arXiv:2611.22426v1 [cs.LG] 17 Nov 2026

INFUSELOOP: CROSS-TRAJECTORY LATENT INJECTION FOR SCALABLE PARALLEL LOOPED TRANSFORMERS

Zhaofeng Sun [1 *] **Lanxiang Hu** [2] **Yu-Yang Qian** [3] **Jiaqi Leng** [2] **Yuanbo Yang** [2] **Chi Zhang** [1] **Youyou Lu** [1]
Peng Zhao [3] **Hao Zhang** [2]

ABSTRACT

Parallel Looped Transformers (PLTs) use otherwise idle decoding compute by overlapping recurrent iterations. However, vanilla PLTs scale poorly with depth: they exhibit training-loss spikes, and their quality gains shrink as more loops are added. The benefit of additional loops depends on the information available to each iteration, so we examine whether the existing computational dependencies limit effective depth scaling. We find that only the preceding in-trajectory state is passed across loop boundaries, even though completed cross-trajectory representations are already available. Based on this finding, we propose InfuseLoop, a cross-trajectory latent injection mechanism. It aggregates causally available hidden states from later prediction trajectories through lightweight attention residuals. This design preserves token-parallel training and concurrent loop execution, and requires only $O(Rd)$ additional hidden-state storage. Experiments show that InfuseLoop eliminates training-loss spikes and improves validation-loss scaling up to eight loops. At eight loops, InfuseLoop reduces perplexity by 3.4–17.1% relative to vanilla PLTs and by 10.5–37.3% relative to the non-looped baseline. It also improves average downstream accuracy by 0.84 and 1.99 percentage points over these two baselines. InfuseLoop incurs only 0.3% decoding time over vanilla PLTs and 23.3% over the non-looped baseline, while keeping the same logical KV-cache footprint as vanilla PLTs.

1 INTRODUCTION

Scaling computational depth can improve model quality, but adding layers to a standard Transformer also increases its parameter count, raising the memory and bandwidth costs of deployment (Vaswani et al., 2023). This motivates using additional computation to improve model quality without increasing model size. Looped Transformers (LTs) achieve this by repeatedly applying shared Transformer blocks, so that effective depth grows while the parameter count stays fixed (Zhu et al., 2025; Geiping et al., 2025). However, because LTs execute loops sequentially, they incur substantial inference overhead: both decoding latency and KV-cache memory grow roughly linearly with the number of loop iterations. This overhead limits how far recurrent depth can be scaled in practice.

Parallel Looped Transformers (PLTs) address this bottleneck by restructuring the dependencies across token positions and loop iterations, which allows different loop stages from multiple prediction trajectories to run concurrently (Wu et al., 2025). Combined with cross-loop KV-cache sharing, this parallel schedule keeps decoding latency and cache storage close to those of a non-looped model when decoding is memory-bound. PLTs therefore make it possible to improve model quality by using decoding compute that would otherwise be idle.

However, efficient execution alone does not guarantee effective depth scaling in PLTs. Prior work (Yang et al., 2026) reports that performance degrades beyond two loops and attributes this to diminishing refinement gains. Our experiments reveal a further problem: vanilla PLTs show training-loss spikes with as few as four loops. Examining the computational graph, we find a gap between the states that are available and the states that the recurrent connections actually use. At each loop boundary, a vanilla PLT passes only the previous output of the same prediction trajectory to the next loop. Yet its overlapping execution schedule also produces completed states on other trajectories, and these states are causally available to the receiving loop. They provide additional representations of the observed context, but the loop-to-loop connections give no direct access to them. We therefore ask: can aggregating these available cross-trajectory representations make additional parallel computation more effective?

In this work, we introduce **InfuseLoop**, which is based on a simple intuition: earlier loops can perform more useful com-

[*]Work done during an internship at UC San Diego. [1]Tsinghua University [2]University of California, San Diego [3]Nanjing University. Correspondence to: Hao Zhang <haz094@ucsd.edu>.

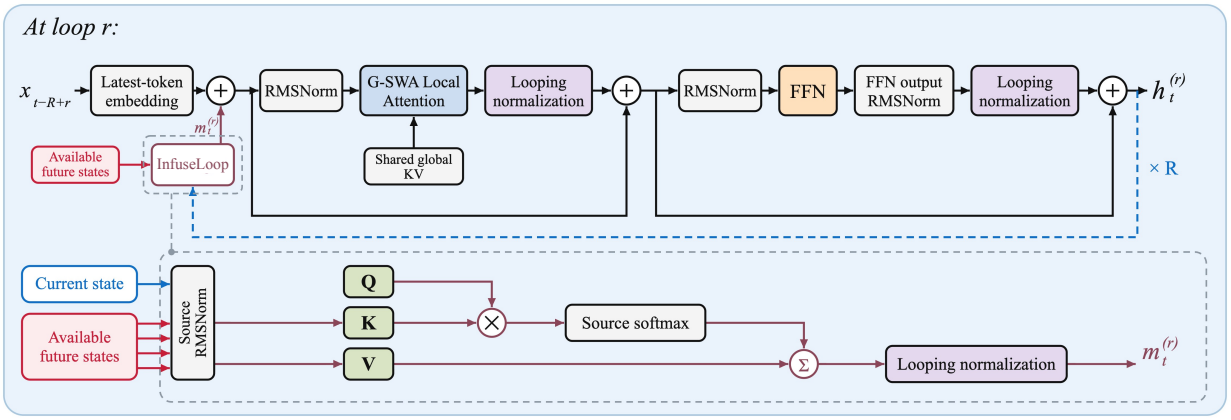


Figure 1. **Overview of InfuseLoop.** Each loop combines the latest token embedding with the current state and causally available future states through attention residual aggregation. The enlarged view details the attention-based aggregation of current latent state and future states.

putation if they can directly combine more of the contextual representations that are already available. As illustrated in Figure 2, InfuseLoop injects hidden states from shallower loops of later prediction trajectories into deeper loops of earlier trajectories. These additional connections allow earlier loop stages to use representations beyond their own recurrent state, so that the computation before the final prediction is used more effectively. Because every injected state is computed before it is used, InfuseLoop expands information flow across trajectories while preserving the parallel execution structure of PLTs.

To the best of our knowledge, InfuseLoop is the first PLT architecture to train stably and continue improving in quality up to eight loops. At eight loops, our model achieves lower validation loss than vanilla PLTs. Compared with vanilla PLTs and a non-looped baseline, it reduces perplexity by 3.4–17.1% and 10.5–37.3% and improves average downstream accuracy by 0.8 and 2.0 percentage points, respectively. InfuseLoop achieves these gains with less than 1% additional decoding time over vanilla PLTs while keeping the same logical KV-cache footprint. These results show that cross-trajectory information flow is an effective way to improve depth scaling in PLTs. It allows additional recurrent computation to improve quality while preserving efficient parallel decoding across loops.

In summary, our contributions are as follows:

- **Computational graph analysis of PLTs.** We find that, at each loop boundary, a vanilla PLT passes only the preceding in-trajectory state to the next loop, so completed cross-trajectory states are never used as direct recurrent inputs. We further show that selected shallower-loop states from the previous cycle can be incorporated into deeper-loop updates without breaking autoregressive causality, token-parallel training, or concurrent loop execution.
- **InfuseLoop: cross-trajectory aggregation.** We propose InfuseLoop, which aggregates causally available hidden states from future positions through lightweight attention residuals. InfuseLoop reduces training-loss spikes and improves model quality while preserving token-parallel training and concurrent loop execution, and it requires only $O(Rd)$ additional hidden-state storage.
- **Depth-scaling and system-level evaluation.** We evaluate depth scaling up to eight loops and characterize the quality–latency trade-offs, execution costs, and cache-memory overhead of InfuseLoop. The results show that InfuseLoop scales recurrent computation while remaining efficient at the system level.

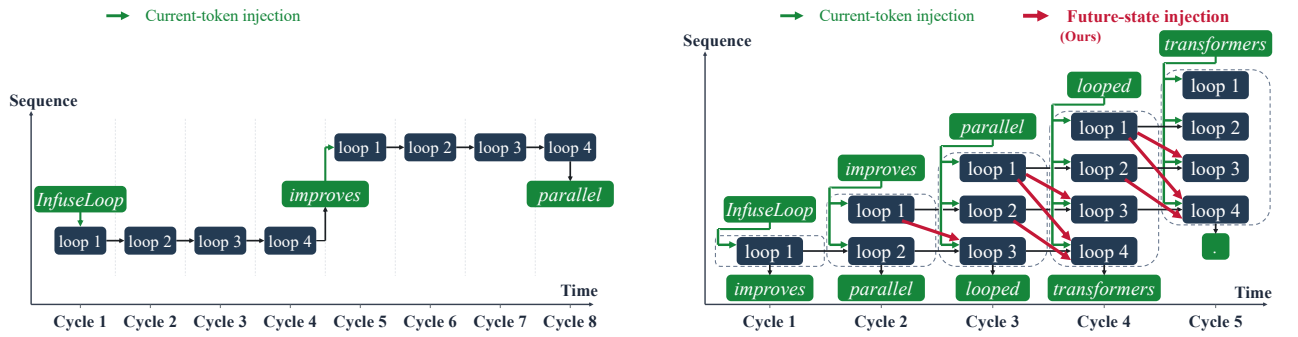
2 PRELIMINARY

2.1 Looped Transformers Depth Scaling

Looped Transformers (LTs) (Dehghani et al., 2019) decouple computational depth from parameter count by repeatedly applying shared Transformer blocks. Let V be the vocabulary and $x_t \in V$ the input token at position t. Consider a standard full-stack LT with K Transformer blocks $f_{\theta_1}, \ldots, f_{\theta_K}$, whose composition is $F_\theta = f_{\theta_K} \circ \cdots \circ f_{\theta_1}$. Starting from $h_t^{(0)} = \mathrm{Embed}(x_t)$, the model applies this backbone for R loops:

$$\begin{aligned} h_t^{(r)} &= F_\theta\left(h_t^{(r-1)}\right), \qquad r = 1, \ldots, R, \\ x_{t+1} &= \mathrm{LM_Head}\left(h_t^{(R)}\right). \end{aligned} \tag{1}$$

Because the parameters $\theta = \{\theta_1, \ldots, \theta_K\}$ are shared across loops, the model performs $K \times R$ Transformer-block applications with only K distinct blocks.



(a) Looped Transformer

(b) InfuseLoop (Ours)

Figure 2. **Execution schedules and information flow.** (a) Looped Transformers execute successive loops serially before generating the next token. (b) InfuseLoop overlaps computation across prediction positions and loop iterations, and introduces future-state injection (red) on top of vanilla PLT.

Empirical studies show that recurrent computation can improve model quality across architectures and training stages, from pretrained dense and mixture-of-experts (MoE) language models (Geiping et al., 2025; Wang et al., 2026b) to instruction-tuned models (Zhu et al., 2025). Realizing these gains, however, requires addressing the optimization and stability problems caused by reusing parameters across loops. Existing approaches address these problems with loop-aware residual scaling and initialization (Li et al., 2026b; Wang et al., 2026a), constrained input injection (Prairie et al., 2026), and variable-depth training with intermediate supervision (Geiping et al., 2025; Zhu et al., 2025). These advances make recurrent computation more effective, but its sequential execution still increases autoregressive decoding latency, which motivates parallel looping schemes.

2.2 Parallel Looped Transformers

Despite these benefits, LTs are difficult to deploy because each prediction requires R sequential evaluations of the shared backbone F_θ. As a result, autoregressive decoding latency grows roughly linearly with the number of loops, which limits the practical benefit of greater computational depth.

One way to reduce this sequential cost is to overlap the loop stages of different predictions. This overlap is possible only if each loop receives its recurrent input without waiting for another loop that runs at the same timestep. Parallel Looped Transformers (PLTs) enable this execution pattern by restructuring the dependencies across prediction trajectories (Wu et al., 2025). As illustrated in Figure 2, a trajectory starts its computation from earlier context and incorporates newly available tokens as it moves through the loops. As a result, different trajectories can run different loop stages at the same time.

Consider a trajectory whose final loop runs at timestep t and predicts x_{t+1}. Note that t indexes the last timestep of the trajectory, not its first. The trajectory computes

$$\begin{aligned} h_t^{(1)} &= F_\theta(\mathrm{Embed}(x_{t-R+1}))\,, \\ h_t^{(r)} &= F_\theta\Big(h_t^{(r-1)} + \mathrm{Embed}(x_{t-R+r})\Big)\,, \quad r = 2, \dots, R, \\ x_{t+1} &= \mathrm{LM_Head}\Big(h_t^{(R)}\Big)\,. \end{aligned} \tag{2}$$

Loop r runs at timestep $t - R + r$ and incorporates the token available at that step, and the final state $h_t^{(R)}$ predicts x_{t+1}. Therefore, unlike in a standard LT, a trajectory starts computing before x_t, the token immediately preceding its prediction target, is available. This allows early loops of later trajectories to run at the same time as deeper loops of earlier trajectories. For example, the following states can be computed in parallel at timestep t:

$$\begin{aligned} h_{t+R-1}^{(1)} &= F_\theta(\mathrm{Embed}(x_t))\,, \\ h_{t+R-2}^{(2)} &= F_\theta\Big(h_{t+R-2}^{(1)} + \mathrm{Embed}(x_t)\Big)\,, \\ &\vdots \\ h_t^{(R)} &= F_\theta\Big(h_t^{(R-1)} + \mathrm{Embed}(x_t)\Big)\,, \end{aligned}$$

because each depends only on x_t and on states computed at earlier timesteps.

Although LTs and PLTs share weights across loops, their hidden representations differ from loop to loop, so standard implementations keep a separate KV cache for each loop. This increases the per-token cache footprint by a factor of R, which can increase memory traffic and limit serving batch sizes. KV-cache sharing can reduce this overhead. For example, gated sliding-window attention (G-SWA) (Wu et al., 2025) shares the global KV cache of the first loop with all later loops, while each loop keeps its own bounded local window. Beyond looped models, YOCO (Sun et al.,

2024) builds global KV representations in a self-decoder and reuses them across cross-decoder layers. These designs show that cache reuse can limit the memory cost of repeated attention, which complements methods that reduce the cost of sequential execution.

However, efficient execution does not guarantee effective depth scaling. Prior work reports diminishing refinement gains and even performance degradation beyond two loops (Yang et al., 2026). We therefore focus on enabling PLTs to benefit from more than two loops ($R > 2$).

2.3 In-trajectory Aggregation and Cross-trajectory Aggregation

Residual connections give later computation direct access to intermediate representations (He et al., 2015). In recurrent Transformers, this kind of routing can be viewed along two complementary axes: *in-trajectory aggregation*, which combines representations across layers or loop iterations within the same prediction trajectory (Team et al., 2026; He et al., 2026; Li et al., 2026a); and *cross-trajectory aggregation*, which passes hidden representations between prediction positions (Wang et al., 2026c; Cai et al., 2026).

PLTs bring these two axes together through their overlapping execution schedule. With R loops, each decoding cycle produces up to R latent states, each at a different prediction position and recurrent depth. These states are different transformations of the observed context, not additional unobserved input tokens. Although attention already passes information across token positions, the explicit loop-output connections in a vanilla PLT carry only the previous state of the same trajectory. An important design question is therefore how to use these simultaneously available states without introducing dependencies that would serialize loop execution. In this work, we study how directly aggregating completed loop outputs across trajectories affects model quality and depth scaling, while preserving the execution schedule of PLTs.

3 INFUSELOOP

We propose InfuseLoop, which gives recurrent computation direct access to completed representations from overlapping prediction trajectories. Section 3.1 analyzes which states are available and describes how we select cross-trajectory sources. Section 3.2 presents the aggregation architecture. Sections 3.3 and 3.4 describe training and inference, and explain how the same design supports token-parallel training, concurrent loop execution, and bounded additional state storage.

3.1 InfuseLoop Computational Graph

Equation (2) describes a trajectory that predicts x_{t+1} through R evaluations of the backbone. The first $R-1$ evaluations run before x_t is available, so x_t is injected only in the final loop. At each loop transition, a vanilla PLT passes only the previous output of the same trajectory to the next loop. Its overlapping execution schedule also produces completed states on other trajectories, but the recurrent update does not take them as inputs. This raises a natural question: can access to these additional representations improve recurrent computation without changing the parallel execution schedule?

InfuseLoop adds direct connections from later prediction trajectories to earlier ones, using only states that have already been computed. For example, $h^{(1)}_{t+1}$ is computed at cycle $t-R+2$, one cycle before $h^{(3)}_t$, so it can be aggregated with $h^{(2)}_t$ to form the input to the third loop. Although $h^{(1)}_{t+1}$ belongs to a later trajectory, it encodes only the observed context and has no access to x_{t+1}. More generally, these connections do not make x_t available before cycle t; they only enlarge the set of contextual representations that recurrent updates can use before the final loop conditions on x_t. Our experiments evaluate how this change affects optimization behavior and predictive quality.

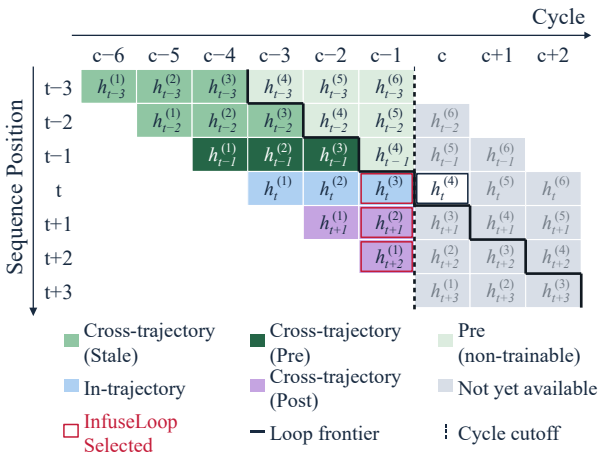


Figure 3. **State availability and source selection** for $h^{(4)}_t$ in a six-loop model ($R = 6$). $h^{(4)}_t$ is computed at cycle $c = \tau(t, 4)$, and the outlined states are the sources selected from cycle $c-1$.

Figure 3 illustrates which states are available when computing $h^{(4)}_t$. Following Section 2, let $\tau(t, r) = t-R+r$ denote the execution cycle of $h^{(r)}_t$. We impose two constraints on an aggregation source $h^{(j)}_s$. First, $\tau(s, j) < \tau(t, r)$: the source must be completed in an earlier cycle, so that no new dependencies arise between loops running in the same cycle. Second, $j < r$: the source must come from a shallower loop, so that it is available under loop-ordered training, in

which all sequence positions of a loop are processed in parallel. Together, these constraints define the admissible source set

$$\mathcal{A}_{t,r} = \left\{ h_s^{(j)} : 1 \le j < r,\ s + j < t + r \right\}. \quad (3)$$

Relative to the receiving trajectory t, admissible sources fall into three groups: In-trajectory States from the same trajectory ($s = t$), Pre States from earlier trajectories ($s < t$), and Post States from later trajectories ($s > t$). These groups are defined by trajectory order, not by execution time. The figure also marks older Pre States (Stale) and Pre States with $j \ge r$ (non-trainable). The latter are causally available but are excluded, because they violate the loop order required for token-parallel training.

Rather than aggregating the full admissible set, whose size grows with sequence length, InfuseLoop selects all shallower-loop outputs from the immediately preceding cycle:

$$\mathcal{S}_{t,r} = \{ h_{t+r-j-1}^{(j)} : 1 \le j < r \}. \quad (4)$$

For example, $\mathcal{S}_{t,4} = \{h_t^{(3)}, h_{t+1}^{(2)}, h_{t+2}^{(1)}\}$. Each selected state satisfies $\tau(t + r - j - 1, j) = \tau(t, r) - 1$ and $j < r$, and therefore belongs to $\mathcal{A}_{t,r}$. The source with $j = r - 1$ is the in-trajectory base state $h_t^{(r-1)}$, and the remaining $r - 2$ sources are completed outputs of later trajectories. No Pre State is selected, because every admissible Pre State was computed at least two cycles earlier. Because all sources come from the previous cycle, a rolling buffer of $O(R)$ hidden vectors suffices for all receiving loops. Appendix A proves that the selected states are causally available, that the selection is compatible with the training and inference schedules, and that the storage bound holds.

3.2 Architecture Design

Inspired by Attention Residuals (Team et al., 2026), we aggregate the selected sources with a learned pseudo-query $w_r \in \mathbb{R}^d$ for each receiving loop r. The keys and values are both obtained by applying parameter-free RMS normalization to the source states:

$$\begin{aligned} k_{t,j\to r} &= v_{t,j\to r} = \mathrm{RMSNorm}\left(h_{t+r-j-1}^{(j)} \right), \\ e_{t,j\to r} &= \frac{w_r^\top k_{t,j\to r}}{\sqrt{d}} + b_{r,j}, \\ \alpha_{t,j\to r} &= \frac{\exp(e_{t,j\to r})}{\sum_{i=1}^{r-1} \exp(e_{t,i\to r})}, \end{aligned} \quad (5)$$

where $b_{r,j}$ is a learned bias for source j. The aggregated residual $m_t^{(r)}$ then replaces $h_t^{(r-1)}$ as the recurrent input:

$$\begin{aligned} m_t^{(r)} &= \rho_r \sum_{j=1}^{r-1} \alpha_{t,j\to r} v_{t,j\to r}, \\ h_t^{(r)} &= F_\theta\left(\mathrm{Embed}(x_{t-R+r}) + m_t^{(r)} \right), \quad r = 2, \dots, R, \end{aligned} \quad (6)$$

where ρ_r is a learned aggregation scale. The first loop is unchanged. At $r = 2$, the only source is the in-trajectory state $h_t^{(1)}$; Post States are included from $r = 3$ onward. Sources that cross a document boundary are masked before the softmax, and an empty source set yields a zero residual.

We keep the G-SWA design of PLT (Wu et al., 2025). In each backbone layer, the first loop provides global KV representations that are shared with all later loops. Each later loop combines attention over these shared representations with its own sliding-window attention through a learned gate. For clarity, we omit the causal masks and KV-cache arguments from F_θ.

To stabilize recurrent computation, we apply loop-aware residual scaling to the attention and FFN updates (Li et al., 2026b; He et al., 2026). Following the normalization design of looped MoE models (He et al., 2026), we also apply RMSNorm to the FFN output before the residual addition. Together with the RMS normalization of the sources and the learned aggregation scale ρ_r, these choices control the magnitudes of both the backbone updates and the cross-trajectory residuals. Section 4.4 evaluates the contributions of these components separately from that of source selection.

3.3 Training

Algorithm 1 Token-parallel training of InfuseLoop

Require: Input sequence $X = (x_1, \dots, x_L)$,
next-token targets Y,
number of loops R, backbone F_θ

$\mathbf{E} \leftarrow \mathrm{Embed}(X)$ {Token embeddings}
$\mathbf{H}^{(1)} \leftarrow F_\theta(\mathbf{E})$ {First-loop states}
for $r = 2, \dots, R$ **do**
  for $j = 1, \dots, r-1$ **do**
    $\widetilde{\mathbf{H}}^{(j)} \leftarrow \mathrm{ShiftRight}(\mathbf{H}^{(j)}, 1)$ {Previous-cycle states}
  end for
  $\mathbf{M}^{(r)} \leftarrow \mathrm{Agg}_r(\{\widetilde{\mathbf{H}}^{(j)}\}_{j<r})$ {Aggregate sources}
  $\mathbf{H}^{(r)} \leftarrow F_\theta(\mathbf{E} + \mathbf{M}^{(r)})$ {All positions in parallel}
end for
$\mathcal{L} \leftarrow \mathrm{CrossEntropy}(\mathrm{LM_Head}(\mathbf{H}^{(R)}), Y)$
Backpropagate $\mathcal{L}$ and update parameters

During training, loops are processed sequentially, while all sequence positions within a loop are computed in parallel. This order is valid because Equation (4) restricts every selected source to a shallower loop ($j < r$), so all sources for loop r are computed before loop r begins. In Algorithm 1, uppercase bold symbols denote sequence-level tensors whose rows are ordered by execution cycle. Shifting these tensors right by one position therefore gives the previous-cycle states required by Equation (4). As a result,

training computes exactly the same computational graph as inference, without approximating any of its dependencies.

We apply the next-token objective to the outputs of the final loop and backpropagate through all recurrent and aggregation paths. Appendix B.2 describes the G-SWA masks and how global and local KV representations are handled during training.

3.4 Efficient Inference

Algorithm 2 describes one decoding cycle, using lowercase symbols for individual states. Each loop builds its input from the current token and the cached outputs of shallower loops from the previous cycle, so aggregation introduces no dependencies on outputs of the current cycle. The backbone follows the layer-synchronous G-SWA schedule of PLT: at each layer, the states of all active loops are batched together, and the global KV representations of the first loop are shared with the other loops. The output of the final loop predicts the next token, and the cached hidden states are updated only after all outputs of the current cycle have been computed.

Keeping one output per loop requires only $O(Rd)$ additional hidden-state storage per sequence, independent of sequence length, whereas keeping all outputs over L cycles would require $O(RLd)$. This bound excludes model parameters, temporary buffers, and the KV cache, whose size is determined by the G-SWA design of PLT (Wu et al., 2025).

For a batch of B requests with the same number of loops, each steady-state cycle passes $B \times R$ loop states to the shared backbone, which allows batched kernels and CUDA Graph execution. Section 4.3.1 evaluates the decoding cost and cache footprint.

Algorithm 2 One decoding cycle of InfuseLoop

Require: Current token x_c,
previous-cycle outputs $p_1, \ldots, p_R$,
number of loops R, backbone F_θ

$e_c \leftarrow \mathrm{Embed}(x_c)$
$z_1 \leftarrow e_c$ {First-loop input}
for $r = 2, \ldots, R$ **do**
  $m_r \leftarrow \mathrm{Agg}_r(p_1, \ldots, p_{r-1})$ {Aggregate sources}
  $z_r \leftarrow e_c + m_r$ {Loop input}
end for
$u_{1:R} \leftarrow \mathrm{ParallelForward}_\theta(z_{1:R})$ {All loops in parallel}
Generate x_{c+1} from $\mathrm{LM_Head}(u_R)$
$p_{1:R} \leftarrow u_{1:R}$ {Cache sources for the next cycle}
return x_{c+1} and $p_{1:R}$

4 Experiments

4.1 Experimental Setup

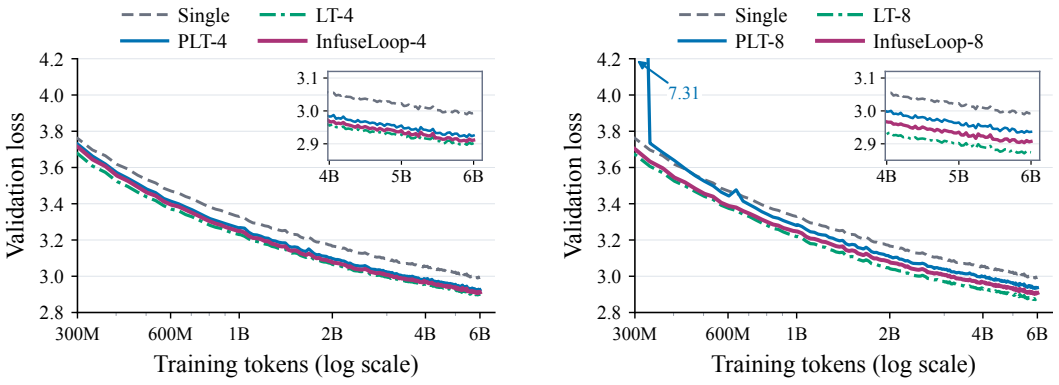
Models and Data. We study models with about 0.6B parameters, each with a shared backbone of 28 Transformer layers and a hidden size of $d = 1{,}024$. We vary the number of loops while keeping the backbone fixed, so additional loops add recurrent computation but no new untied Transformer layers. This setup isolates how each architecture uses the shared backbone across repeated evaluations. Appendix B.1 gives the full model configurations.

All models are trained on FineWeb-Edu (Penedo et al., 2024) with the same tokenizer and data splits, using 6B tokens for the main experiments and 1B tokens for the component ablations.

We conduct all benchmark evaluations using `lm-evaluation-harness` (Gao et al., 2024). For language modeling, we report word perplexity on WikiText-2 (Merity et al., 2017) and final-word perplexity on LAMBADA (Paperno et al., 2016). For downstream tasks, we report accuracy on MMLU (Hendrycks et al., 2021), HellaSwag (Zellers et al., 2019), ARC-Easy (ARC-E) and ARC-Challenge (ARC-C) (Clark et al., 2018), PIQA (Bisk et al., 2020), WinoGrande (WG) (Sakaguchi et al., 2020), and OpenBookQA (OBQA) (Mihaylov et al., 2018). MMLU is evaluated in the five-shot setting, and all other tasks in the zero-shot setting. We use length-normalized accuracy (`acc_norm`) for HellaSwag, ARC-E, ARC-C, and OBQA, and standard accuracy (`acc`) for MMLU, PIQA, and WG. *Avg.* is the unweighted mean of the seven downstream accuracies and does not include perplexity.

Training Settings and Implementation Details. We train all models on NVIDIA GB200 and H200 GPUs with a sequence length of 1,024 and a global batch size of 128 sequences. We use AdamW with a learning rate of 3×10^{-4}, and all other optimization settings are shared across the compared models. Within each comparison, all models are trained on the same number of tokens, so training compute grows with the number of loops. Unless ablated, all looped models use the looping normalization described in Section 3.2.

Baselines. We compare InfuseLoop with three baselines that use the same backbone: Single, a standard LT, and a vanilla PLT. Single applies the backbone once, LT runs its loops sequentially, and PLT overlaps computation across prediction trajectories. All looped models use G-SWA with a window size of $w = 64$ during decoding. The component ablations change either the normalization or the source selection, while keeping the rest of the architecture and the training settings fixed.



(a) Four loops

(b) Eight loops

Figure 4. Validation loss during pretraining with (a) $R = 4$ and (b) $R = 8$ loops.

4.2 Results

We compare the looped architectures at $R = 4$ and $R = 8$, using Single as the non-looped reference.

Figure 4 shows validation loss during pretraining. Vanilla PLT shows pronounced loss spikes, whereas InfuseLoop removes these spikes and converges more smoothly at both depths. These results suggest that changing the restrictive loop-to-loop information flow of vanilla PLT is an effective way to stabilize parallel recurrent training.

Table 1 reports validation loss, held-out perplexity, and downstream accuracy. At both depths, InfuseLoop outperforms vanilla PLT and Single in validation loss, both perplexity metrics, and average accuracy, and it closes **52–61%** of the validation-loss gap between vanilla PLT and LT. Across both depths and both perplexity benchmarks, its perplexity is **1.4–28.4%** lower than that of vanilla PLT. It also improves average downstream accuracy over vanilla PLT by **0.96** percentage points at $R = 4$ and by **0.84** percentage points at $R = 8$. LT still achieves the highest average accuracy, but at eight loops InfuseLoop trails it by only 0.59 percentage points, recovering about **59%** of the accuracy gap between vanilla PLT and LT. Thus, the benefits of cross-trajectory aggregation extend beyond training stability to predictive quality.

4.3 InfuseLoop Efficiency and Scalability

4.3.1 Inference Efficiency

We benchmark decoding on a single NVIDIA GB200 GPU using CUDA Graphs, the `vLLM` FA4 attention backend (Kwon et al., 2023), and custom Triton kernels for G-SWA. Linear and attention operands use FP16, while residuals, reductions, and output logits use FP32. After warmup and graph capture, we measure 64 consecutive decoding cycles and report the median throughput over three runs. Component timings are collected separately from short CUDA Graph execution traces.

Memory footprint. Table 2 reports the logical cache footprint per cached token with $w = 64$. It includes the shared global KV cache, the loop-specific local windows, and the $R - 1$ latent vectors used as aggregation sources. KV entries are stored in 16-bit precision and latent vectors in FP32. The footprint is reported in decimal KB and excludes model weights, temporary buffers, and allocator padding. Because the global KV cache is shared, no loop needs its own full-context cache, and the local windows and latent vectors have a fixed size per sequence. Their relative overhead therefore decreases as the context grows, falling to **5.5%** for InfuseLoop-8 at 8,192 tokens. Because every term scales with the number of sequences, the per-token footprint does not depend on batch size.

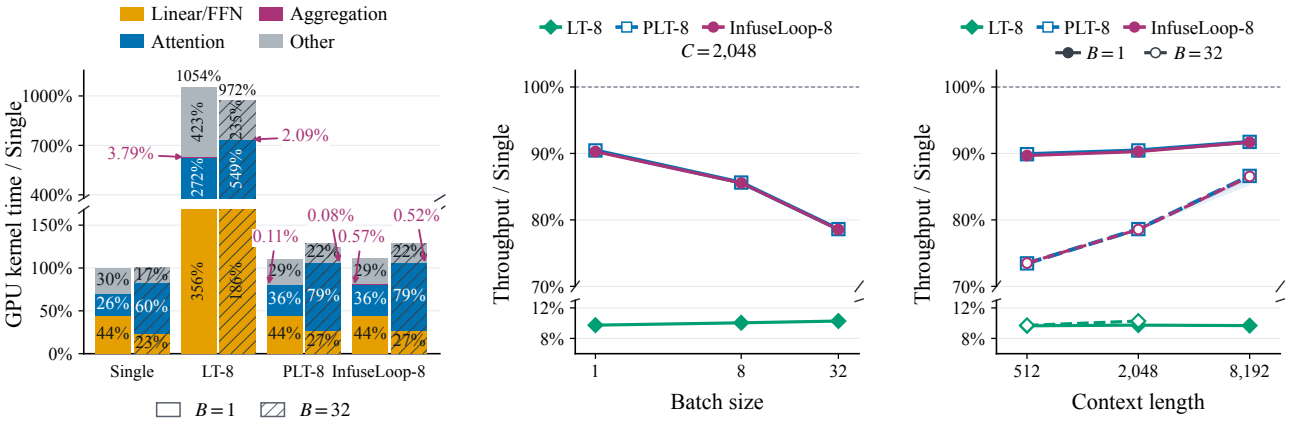
Arithmetic intensity. We estimate arithmetic intensity as the FLOPs of the dominant projection and attention operations divided by the distinct bytes of packed weights and KV inputs, assuming that each element is read once. At $B = 32$, $R = 8$, and a 4,096-token context, InfuseLoop has an estimated intensity of **26.8 FLOPs/byte**. This is far below the GB200 ridge point of **312.5 FLOPs/byte**, which is the 2.5 PFLOP/s dense FP16 peak divided by the 8 TB/s HBM bandwidth. Decoding therefore remains memory-bound even with eight loops. Parallel execution reuses the backbone weights across loops, and G-SWA shares the global KV cache across loop queries. Additional loops can therefore use idle compute without a proportional increase in memory traffic, so latency grows modestly rather than by a factor of R.

Execution cost breakdown. Figure 5(a) breaks down GPU kernel time at $B = 1$ and $B = 32$ with a 2,048-token initial context, normalized to the total kernel time of Single at the same batch size. Cross-trajectory aggregation accounts for

Table 1. Validation loss, perplexity, and downstream accuracy (%) after training on 6B tokens.

Loops	Model	Val. loss ↓	Perplexity ↓		Downstream tasks ↑							
			Wiki2	LAMBADA	MMLU	HellaSwag	ARC-E	ARC-C	PIQA	WG	OBQA	Avg.
$R=1$	Single	2.9941	39.72	67.02	24.58	36.44	49.92	28.58	63.38	51.70	33.00	41.09
$R=4$	LT	2.9022	34.78	41.07	24.35	40.17	54.34	30.55	66.43	54.70	33.80	43.48
	PLT	2.9252	36.19	62.96	**24.71**	39.24	51.68	28.58	65.78	51.70	32.60	42.04
	InfuseLoop	**2.9112**	**35.70**	**45.08**	24.49	**39.49**	**53.75**	**31.40**	**65.89**	**52.96**	**33.00**	**43.00**
$R=8$	LT	2.8775	33.78	35.04	23.99	41.88	55.77	31.91	65.51	54.38	32.20	43.66
	PLT	2.9365	36.82	50.69	24.50	39.25	51.39	29.61	65.72	52.33	**32.80**	42.23
	InfuseLoop	**2.9058**	**35.55**	**42.04**	**24.96**	**40.00**	**53.20**	**31.43**	**66.16**	**53.35**	32.40	**43.07**

Note: Within each loop group, underlined values are the best overall, and **bold** values are the best among non-LT models.



(a) Execution breakdown (b) Batch-size sensitivity (c) Context-length sensitivity

Figure 5. **Decoding cost and workload sensitivity.** (a) GPU kernel time with a 2,048-token initial context, normalized to the total kernel time of Single at the same batch size. (b, c) Throughput relative to Single across batch sizes and context lengths.

Table 2. **Logical cache footprint** per token (KB) with window size $w = 64$. Parentheses show the overhead relative to Single.

	Context length		
Model	512	2,048	8,192
Single	114.69	114.69	114.69
InfuseLoop-4	157.72 (+37.5%)	125.45 (+9.4%)	117.38 (+2.3%)
InfuseLoop-8	215.10 (+87.5%)	139.79 (+21.9%)	120.96 (+5.5%)

less than **0.6%** of this reference time. Because loops are batched, linear and FFN costs stay close to those of Single, and most of the additional time comes from attention.

Batch and context sensitivity. Figure 5(b) varies the batch size with a 2,048-token initial context, and panel (c) varies the context length at $B = 1$ and $B = 32$. The context grows during each measurement window, and all throughput ratios are relative to Single at the same batch size and context length. InfuseLoop closely tracks vanilla PLT in all settings, as expected from its small aggregation cost. The throughput gap to Single widens with larger batches but narrows with longer contexts. At an 8,192-token context, InfuseLoop retains **86.5%** of the throughput of Single even at $B = 32$. This trend is consistent with the fixed costs of local windows and aggregation becoming smaller relative to global attention as the context grows.

4.3.2 *Depth Scalability*

Figure 6(a) compares models trained separately with $R = 4$ and $R = 8$, using the same backbone and 6B-token budget, with Single as the $R = 1$ reference. The horizontal axis is logarithmic, so successive doublings of depth are equally spaced. Because the data budget is fixed, this comparison measures the quality gained from additional training compute.

Increasing vanilla PLT from four to eight loops worsens validation loss, even though it doubles the number of backbone applications. In contrast, InfuseLoop reduces validation

loss from four to eight loops and outperforms vanilla PLT at both depths, although its average downstream accuracy improves only slightly. LT still achieves the best quality, but InfuseLoop narrows the gap without changing the parallel execution schedule.

4.3.3 *Performance–Latency Pareto Frontier*

Figure 6(b) shows performance–latency Pareto frontiers for validation loss and average downstream accuracy. We measure decode-only wall-clock time at $B = 32$ with 2,048-token prompts and 2,048 output tokens per sequence, timing the 2,047 cached steps after prefill. We use the setup in Section 4.3.1 and report the median of three runs, excluding model loading, warmup, and graph capture. The loss and accuracy frontiers are computed separately from measured configurations only. A configuration is dominated if another is at least as good and at least as fast, and strictly better in at least one of the two.

InfuseLoop configurations lie on the measured Pareto frontier in the low-latency regime. InfuseLoop-4 dominates PLT-8: it achieves lower validation loss and higher average accuracy at lower decoding latency. InfuseLoop-8 further reduces validation loss at a modest latency cost. LT-8 achieves higher quality, but its decoding latency is about **7.7**$\times$ that of InfuseLoop-8.

LT remains nondominated in this fixed-backbone comparison, but parameter efficiency does not imply latency efficiency (Wu et al., 2025). A larger non-looped model is an alternative use of the same inference resources. For example, at matched inference FLOPs, a 3.9B non-looped model matches or exceeds a three-loop 1.4B model on all 12 benchmarks in a controlled comparison (Marchenko et al., 2026). InfuseLoop addresses this deployment cost: it obtains quality gains from additional depth with a fixed backbone and without the large latency increase of serial looping.

4.4 Ablations

Table 3 evaluates source selection and normalization at $R = 4$ and $R = 8$, using the same 0.6B backbone trained on 1B tokens.

Source selection. Following the categories in Section 3.1, we compare four alternative source sets: Base only, FullState, In-trajectory, and Pre. All variants keep the in-trajectory base state $h_t^{(r-1)}$ and use the same normalization and aggregation. InfuseLoop achieves the lowest validation loss and the highest average accuracy among these variants at both depths. The Base-only and In-trajectory results show that information from the same trajectory alone does not recover the benefit of cross-trajectory aggregation.

FullState excludes stale Pre States but still aggregates $O(R^2)$ candidates for deep receiving loops, compared with $O(R)$ previous-cycle sources for InfuseLoop; all variants use the same number of backbone evaluations. Its higher loss suggests that a larger source set does not necessarily improve quality at this training budget, possibly because denser aggregation makes optimization harder. Pre also underperforms, especially at $R = 8$. Its most recent eligible states are at least two cycles old, compared with one cycle for InfuseLoop, which suggests that staleness limits their usefulness.

Table 3. **Component ablations.** All configurations are trained on 1B tokens.

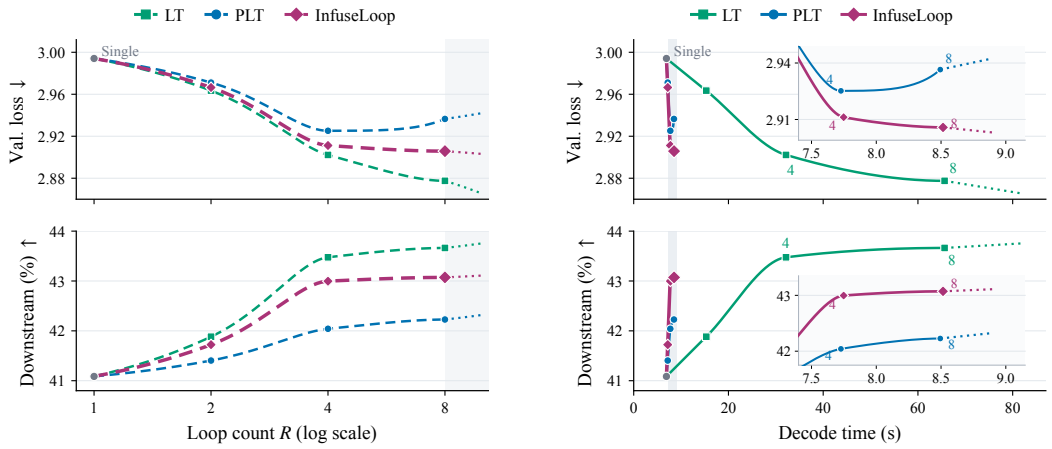
	$R = 4$		$R = 8$	
Configuration	Loss ↓	Avg. ↑	Loss ↓	Avg. ↑
InfuseLoop	3.2482	38.26	3.2436	38.35
Source selection				
Base only	3.2660	38.04	3.2799	37.59
FullState	3.3975	37.43	3.4244	37.21
In-trajectory	3.2735	38.19	3.2839	37.87
Pre	3.2820	37.81	3.3707	37.41
Normalization				
w/o source RMSNorm	3.2883	38.02	3.3550	37.59
w/o looping normalization	3.2488	38.33	3.2434	38.17

Normalization. Removing source RMSNorm worsens both validation loss and average accuracy, especially at $R = 8$. Removing looping normalization, which consists of the attention and FFN residual scaling and the FFN output RMSNorm, while keeping source RMSNorm, leaves validation loss nearly unchanged and changes average accuracy only slightly ($+0.07$ at $R = 4$ and -0.18 at $R = 8$). Thus, at the 1B-token scale, source normalization affects final quality more than the tested backbone normalization.

5 Related Work

5.1 Efficient Looped Transformers

Efficient inference with looped models requires limiting both repeated computation and the state stored for each loop. Parameter-efficient designs combine layer tying with low-rank, depth-specific adaptation (Bae et al., 2025a), while cache-compression methods use the structure shared across loops through latent representations and quantized residuals (O'Neill & Reid, 2026; Kim et al., 2026). Adaptive methods reduce redundant computation by choosing the recursion depth for each token (Bae et al., 2025b) or by applying sparse updates across loops (Yang & Liu, 2026). Other execution schemes increase parallelism through depth-wise batching and pipelining across tokens (Bae et al., 2025a; Wu et al., 2025), or by predicting the inputs of later loops and correcting their outputs afterward (Song et al., 2026).



(a) Depth scaling

(b) Performance–latency Pareto frontier

Figure 6. **Depth scaling and performance–latency Pareto frontiers.** Top: validation loss; bottom: average downstream accuracy. Interpolated $R = 2$ points and projected $R > 8$ extensions are visual guides only; Pareto membership is computed from measured configurations.

InfuseLoop keeps the pipelined schedule of PLT and improves information flow across trajectories by aggregating completed loop outputs.

The regularity of execution affects how well these designs map to batched inference. Adaptive depths can produce uneven workloads across requests, which require dedicated batching and load-balancing mechanisms (Bae et al., 2025a;b). In contrast, InfuseLoop keeps the regular execution pattern of PLT for requests with the same R, so each cycle has static shapes and supports CUDA Graph execution (Section 3.4).

5.2 Exploiting Memory-Bound Decoding

When decoding is memory-bound, parallel computation can amortize the cost of reading model weights and caches. Speculative decoding drafts several tokens with a cheaper model and verifies them in a single batched pass of the target model (Leviathan et al., 2023; Chen et al., 2023). Auxiliary prediction heads and lightweight block-diffusion drafters further reduce drafting cost and increase token-level parallelism (Cai et al., 2024; Chen et al., 2026).

InfuseLoop instead uses this spare compute for recurrent computation that improves model quality, while keeping loop execution parallel. It is therefore potentially complementary to draft-and-verify decoding. However, because both approaches draw on the same spare compute, their combined benefit depends on the workload and the available compute. For example, a lightweight block-diffusion drafter could propose token blocks for a InfuseLoop target model (Chen et al., 2026). Such a combination would require verification that accounts for the loop pipeline, as well as consistent handling of cached recurrent states when drafts are rejected. We leave its implementation and end-to-end evaluation to future work.

6 Limitations

Because of our training compute budget, we evaluate only dense models with about 0.6B parameters. Scaling InfuseLoop to larger dense and mixture-of-experts (MoE) models remains unexplored.

Although InfuseLoop improves depth scaling over vanilla PLT, its quality gains from additional loops are still smaller than those of sequential LTs at the evaluated depths. Further study is needed to identify the architectural and optimization factors behind this gap.

7 Conclusions

We presented InfuseLoop, an efficient architecture for scaling parallel recurrent computation. By injecting causally available hidden states from later prediction trajectories, InfuseLoop improves training stability and model quality while preserving token-parallel training and concurrent loop execution. Its lightweight aggregation and $O(Rd)$ hidden-state storage add little decoding overhead over vanilla PLT.

Our results suggest that cross-trajectory aggregation is important for deep recurrent models: additional computation becomes more effective when representations from later trajectories can reach the loops of earlier ones. More broadly, InfuseLoop illustrates how redesigning these information paths can improve depth scaling while retaining the execution efficiency of parallel looping.

REFERENCES

Bae, S., Fisch, A., Harutyunyan, H., Ji, Z., Kim, S., and Schuster, T. Relaxed recursive transformers: Effective parameter sharing with layer-wise LoRA. In *International Conference on Learning Representations*, 2025a. URL https://openreview.net/forum?id=WwpYSOkkCt.

Bae, S., Kim, Y., Bayat, R., Kim, S., Ha, J., Schuster, T., Fisch, A., Harutyunyan, H., Ji, Z., Courville, A., and Yun, S.-Y. Mixture-of-recursions: Learning dynamic recursive depths for adaptive token-level computation. In *Advances in Neural Information Processing Systems*, 2025b. URL https://arxiv.org/abs/2507.10524.

Bisk, Y., Zellers, R., Le Bras, R., Gao, J., and Choi, Y. PIQA: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 7432–7439, 2020. doi: 10.1609/aaai.v34i05.6239. URL https://arxiv.org/abs/1911.11641.

Cai, T., Li, Y., Geng, Z., Peng, H., Lee, J. D., Chen, D., and Dao, T. Medusa: Simple LLM inference acceleration framework with multiple decoding heads. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 5209–5235. PMLR, 2024. URL https://proceedings.mlr.press/v235/cai24b.html.

Cai, Z., Zhu, X., Dong, Y., He, Y., and Arora, S. T^2MLR: Transformer with Temporal Middle-Layer Recurrence, 2026. URL https://arxiv.org/abs/2607.15178.

Chen, C., Borgeaud, S., Irving, G., Lespiau, J.-B., Sifre, L., and Jumper, J. Accelerating large language model decoding with speculative sampling, 2023. URL https://arxiv.org/abs/2302.01318.

Chen, J., Liang, Y., and Liu, Z. DFlash: Block diffusion for flash speculative decoding, 2026. URL https://arxiv.org/abs/2602.06036. Accepted at ICML 2026.

Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try ARC, the AI2 reasoning challenge, 2018. URL https://arxiv.org/abs/1803.05457.

Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Łukasz Kaiser. Universal transformers, 2019. URL https://arxiv.org/abs/1807.03819.

Gao, L., Tow, J., Abbasi, B., Biderman, S., Black, S., DiPofi, A., Foster, C., Golding, L., Hsu, J., Le Noac'h, A., Li, H., McDonell, K., Muennighoff, N., Ociepa, C., Phang, J., Reynolds, L., Schoelkopf, H., Skowron, A., Sutawika, L., Tang, E., Thite, A., Wang, B., Wang, K., and Zou, A. The language model evaluation harness, 2024. URL https://zenodo.org/records/12608602.

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, 2025. URL https://arxiv.org/abs/2502.05171.

He, D., Li, P., Chang, D., Meng, Q., Yin, L., and Liu, S. Looping beyond twice: A scalable recipe for looped mixture-of-experts, 2026. URL https://arxiv.org/abs/2610.01153.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition, 2015. URL https://arxiv.org/abs/1512.03385.

Hendrycks, D., Burns, C., Basart, S., Zou, A., Mazeika, M., Song, D., and Steinhardt, J. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021. URL https://arxiv.org/abs/2009.03300.

Kim, H., Lee, J., Cho, S., Han, D., Han, I., and Kim, S. ResidualQuant: KV cache quantization for looped transformers with 2-bit residuals, 2026. URL https://arxiv.org/abs/2610.10381.

Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J. E., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with pagedattention, 2023. URL https://arxiv.org/abs/2309.06180.

Leviathan, Y., Kalman, M., and Matias, Y. Fast inference from transformers via speculative decoding. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 19274–19286. PMLR, 2023. URL https://proceedings.mlr.press/v202/leviathan23a.html.

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, 2026a. URL https://arxiv.org/abs/2610.11570.

Li, S., Zhang, Y., Guo, J., Gu, Q., and Wang, M. Deeploop: Depth scaling for looped transformers, 2026b. URL https://arxiv.org/abs/2607.13491.

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, 2026. URL https://arxiv.org/abs/2610.00673.

Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. In *International Conference on Learning Representations*, 2017. URL https://arxiv.org/abs/1609.07843.

Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2381–2391. Association for Computational Linguistics, 2018. doi: 10.18653/v1/D18-1260. URL https://aclanthology.org/D18-1260/.

O'Neill, J. and Reid, F. Looped latent attention: Cross-loop KV compression for looped transformers, 2026. URL https://arxiv.org/abs/2607.15456.

Paperno, D., Kruszewski, G., Lazaridou, A., Pham, Q. N., Bernardi, R., Pezzelle, S., Baroni, M., Boleda, G., and Fernández, R. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1525–1534. Association for Computational Linguistics, 2016. doi: 10.18653/v1/P16-1144. URL https://aclanthology.org/P16-1144/.

Penedo, G., Kydlíček, H., allal, L. B., Lozhkov, A., Mitchell, M., Raffel, C., Werra, L. V., and Wolf, T. The fineweb datasets: Decanting the web for the finest text data at scale, 2024. URL https://arxiv.org/abs/2406.17557.

Prairie, H., Novack, Z., Berg-Kirkpatrick, T., and Fu, D. Y. Parcae: Scaling laws for stable looped language models, 2026. URL https://arxiv.org/abs/2604.12946.

Sakaguchi, K., Le Bras, R., Bhagavatula, C., and Choi, Y. WinoGrande: An adversarial Winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 8732–8740, 2020. doi: 10.1609/aaai.v34i05.6399. URL https://arxiv.org/abs/1907.10641.

Song, F., Zeng, B., Hao, Y., Song, S., Li, H., Yang, Y., Wang, Z., Liu, Y., He, Z., and Lin, Z. LoopFold: Parallelizing recurrent depth in looped transformers, 2026. URL https://arxiv.org/abs/2610.35556.

Sun, Y., Dong, L., Zhu, Y., Huang, S., Wang, W., Ma, S., Zhang, Q., Wang, J., and Wei, F. You only cache once: Decoder-decoder architectures for language models, 2024. URL https://arxiv.org/abs/2405.05254.

Team, K., 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.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. Attention is all you need, 2023. URL https://arxiv.org/abs/1706.03762.

Wang, S., Li, B., Zhu, T., Zhang, G., Huang, W., Yan, S., and Li, J. Loop-μp: Residual scaling and hyperparameter transfer in looped transformers, 2026a. URL https://arxiv.org/abs/2606.18524.

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, 2026b. URL https://arxiv.org/abs/2609.01343.

Wang, X., Cai, Z., Zhan, Z., Dong, H., Fan, Y., de Rosa, G., Pearce, T., and Langford, J. Full-bandwidth transformer, 2026c. URL https://arxiv.org/abs/2608.08888.

Wu, B., Chen, M., Luo, X., Yan, S., Yu, Q., Xia, F., Zhang, T., Zhan, H., Zhong, Z., Zhou, X., Qiao, S., and Bin, X. Parallel loop transformer for efficient test-time computation scaling, 2025. URL https://arxiv.org/abs/2510.24824.

Yang, J., Guo, S., Zhang, W., Zheng, T., Du, Y., Li, H.-S., Wu, J., Song, Y., Xing, Y., Cai, Q., Huang, Z., Hao, C., Tao, R., Liu, X., Zhao, W. X., Tang, M., Lv, W., Zhou, M., and Dai, B. Loopcoder-v2: Only loop once for efficient test-time computation scaling, 2026. URL https://arxiv.org/abs/2606.18023.

Yang, W. and Liu, S. FlashLoop: Fast and memory-efficient looped transformers via lazy updates, 2026. URL https://arxiv.org/abs/2609.29812.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800. Association for Computational Linguistics, 2019. doi: 10.18653/v1/P19-1472. URL https://aclanthology.org/P19-1472/.

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.

AI USAGE DISCLOSURE

The authors used AI-assisted tools, including Codex and Claude Code, to improve productivity during research and manuscript preparation. These tools primarily assisted with code implementation, debugging, data processing, experimental automation, visualization, and language refinement under the authors' supervision and verification. All research ideas, methodological innovations, experimental designs, scientific analyses, and conclusions were developed and determined by the authors. AI tools served solely as implementation and productivity aids. The authors reviewed and validated AI-assisted outputs and take full responsibility for the accuracy and integrity of this work.

A CAUSAL ROUTING AND PARALLEL EXECUTION

This appendix establishes the causal availability, execution compatibility, and auxiliary storage requirements of InfuseLoop's source selection. We use the notation of Section 3.1: $h_t^{(r)}$ is the output of loop r on the trajectory that predicts x_{t+1}, and its execution cycle is $\tau(t,r) = t - R + r$. The source characterization below assumes steady-state execution. Unavailable source states are masked, and an empty source set produces a zero residual.

A.1 Availability and Source Selection

The admissible source set satisfies both the execution-cycle and training-loop ordering constraints:

$$\mathcal{A}_{t,r} = \left\{ h_s^{(j)} : 1 \le j < r,\ s + j < t + r \right\}. \tag{7}$$

InfuseLoop restricts this set to the immediately preceding cycle:

$$\mathcal{S}_{t,r} = \left\{ h_{t+r-j-1}^{(j)} : 1 \le j < r \right\}. \tag{8}$$

Lemma 1 (Previous-cycle source characterization)**.** *For a receiver $h_t^{(r)}$ with $r \ge 2$, $\mathcal{S}_{t,r}$ consists exactly of the admissible sources produced at cycle $\tau(t,r)-1$. It contains one in-trajectory state and $r-2$ Post states. Any admissible Pre state is at least two cycles old.*

Proof. Let $c = \tau(t,r)$. A state $h_s^{(j)}$ is produced at cycle $c-1$ if and only if

$$s - R + j = t - R + r - 1, \tag{9}$$

or equivalently $s = t + r - j - 1$. For every $1 \le j < r$, this source belongs to $\mathcal{A}_{t,r}$, since $s + j = t + r - 1 < t + r$. Moreover, $s = t$ when $j = r - 1$, while $s > t$ for $1 \le j \le r - 2$. Thus, the $r - 1$ selected sources comprise the base $h_t^{(r-1)}$ and $r - 2$ completed Post states.

For any Pre source with $s < t$ and $j < r$,

$$\tau(t,r) - \tau(s,j) = (t - s) + (r - j) \ge 2. \tag{10}$$

Hence no admissible Pre state was produced in the immediately preceding cycle. □

This characterization explains the temporal distinction between the selected Post sources and admissible Pre sources; it does not determine their relative predictive utility. At $r = 2$, selection contains only the in-trajectory base state, and cross-trajectory aggregation begins at $r = 3$.

A.2 Autoregressive Causality

Lemma 2 (Prefix causality)**.** *For fixed model parameters, a state $h_t^{(r)}$ produced at cycle $c = \tau(t,r)$ depends only on input tokens observed by cycle c. Future-state aggregation does not introduce dependence on $x_{c+1}, x_{c+2}, \ldots$.*

Proof. We proceed by induction over execution cycles, and within each cycle over Transformer layers. Initially, unavailable source states are masked. At cycle c, Lemma 1 shows that every aggregation source was produced at cycle $c-1$. By the induction hypothesis, these states depend only on tokens already observed by that cycle. Combining them with $\mathrm{Embed}(x_c)$ therefore produces inputs depending only on the observed prefix $x_{\le c}$.

At each backbone layer, causal G-SWA reads first-loop global KV entries and loop-specific local entries only from positions at or before c. The current first-loop KV entry is computed from that loop's layer input, which is itself causal by induction over layers. Attention, normalization, and FFN operations therefore preserve dependence only on the observed prefix. This establishes the claim for every current-cycle loop output. □

Consequently, a larger trajectory index denotes a later prediction, rather than access to a later input token. In particular, for $R \ge 3$, the connection $h_{t+1}^{(1)} \to h_t^{(3)}$ reads a source computed at $t - R + 2$ before a receiver at $t - R + 3$; it does not provide x_t before cycle t.

A.3 Training and Inference Compatibility

Theorem 1 (Equivalent evaluation orders)**.** *Conditioned on the same input sequence and using the same parameters, attention masks, and initial states, loop-ordered token-parallel training and cycle-ordered layer-synchronous inference evaluate the same InfuseLoop dependency graph. The routing does not require an approximation to the training dependencies or an additional serial loop stage within a decoding cycle.*

Proof. Consider a receiving state $h_t^{(r)}$ at cycle $c = \tau(t,r)$. By Lemma 1, each selected source $h_{t+r-j-1}^{(j)}$ was produced at cycle $c-1$. In Algorithm 1, sequence-level tensors are ordered by execution cycle. A one-position right shift therefore aligns the loop-j output from cycle $c-1$ with the receiving position at cycle c. Since $j < r$, the source tensor is already available when training processes loop r. Aggregation acts independently at each position, and causal attention evaluates all sequence positions within each backbone layer in parallel.

In Algorithm 2, p_j is the cached output of loop j from cycle $c-1$. The states $p_1, \ldots, p_{r-1}$ thus supply exactly the

same sources as the shifted training tensors. All current loop inputs can consequently be prepared without reading current-cycle loop outputs. The layer-synchronous backbone schedule then computes the same layer operations: the first loop provides each layer's global KV entries, and the remaining loops use the same causal global and local attention masks. Current-layer projections may be batched before attention; they do not require another loop to complete its full backbone evaluation.

Thus, both schedules respect every aggregation and backbone dependency and apply the same operations to the same inputs. Matching masks and initial states extend the argument to positions with unavailable history. The resulting states agree in exact arithmetic; different kernel evaluation orders can introduce ordinary floating-point differences. □

The statement compares evaluation on a fixed token prefix. It does not remove the usual distinction between teacher-forced training and autoregressive generation. Training retains gradients through every selected source, so token parallelism requires neither stop-gradient operations nor truncated routing paths.

A.4 Auxiliary State Storage

At cycle c, define the retained source buffer as

$$\mathcal{B}_{c-1} = \left\{ h^{(j)}_{c+R-j-1} : 1 \le j \le R \right\}. \tag{11}$$

Every receiving loop reads a subset of this buffer:

$$\mathcal{S}_{c+R-r,r} = \left\{ h^{(j)}_{c+R-j-1} : 1 \le j < r \right\}. \tag{12}$$

The previous-cycle final-loop output is not read by aggregation; retaining all R outputs gives a uniform cache update matching Algorithm 2. After the current cycle, the R new loop outputs replace the retained states. For hidden dimension d and batch size B, this requires $O(BRd)$ persistent auxiliary storage, independent of the generated sequence length. Separate old and new buffers change only the constant factor. This bound excludes KV caches, model weights, training activations, and temporary kernel workspaces.

There are $\sum_{r=2}^{R}(r-1) = R(R-1)/2$ logical source–receiver pairs per sequence and cycle. Their scoring and weighted reduction require $O(R^2 d)$ arithmetic, but no additional backbone evaluations and no persistent $R \times R \times d$ source tensor. This distinction separates aggregation work from retained-state storage.

Scope of the structural results. These results establish how cross-trajectory aggregation can be executed causally with the training and inference orders used in PLT. They do not imply an expressivity separation or a guaranteed improvement in optimization or prediction. The effects of source selection on training stability, validation loss, and downstream performance are assessed by the experiments and ablations.

B Training Details

B.1 Model Configuration

All models follow the backbone configuration of `Qwen3-0.6B` and use its tokenizer. Table 4 summarizes the shared architectural parameters. The input embedding and output projection share weights. Looped variants reuse the same 28-layer backbone across iterations, increasing computational depth without replicating backbone parameters. The total parameter count remains approximately 0.6B, with small additional parameter overhead from architecture-specific aggregation modules.

Table 4. Backbone configuration and parameter counts.

Configuration	Value
Backbone configuration	`Qwen3-0.6B`
Transformer layers (K)	28
Hidden dimension (d)	1,024
FFN intermediate dimension	3,072
Query heads	16
KV heads	8
Head dimension	128
Vocabulary size	151,936
FFN activation	SwiGLU
Positional encoding	RoPE
RoPE base	10^6
RMSNorm ϵ	10^{-6}
Tied input/output embeddings	Yes
Training sequence length	1,024
Base total parameters	596,136,384
Embedding parameters	155,582,464
Base non-embedding parameters	440,553,920
InfuseLoop-4 total parameters	596,139,465
InfuseLoop-8 total parameters	596,143,587

Note: Base counts exclude cross-loop aggregation parameters. The tied embedding and output projection are counted once.

B.2 G-SWA During Training

Following PLT (Wu et al., 2025), we use gated sliding-window attention (G-SWA) to combine shared global context with loop-specific local representations. Training tensors are aligned by execution cycle: row c at loop r corresponds to the computation that produces $h^{(r)}_{c+R-r}$. For backbone layer $\ell = 1, \ldots, K$, let $\mathbf{Q}^{(r)}_\ell$, $\mathbf{K}^{(r)}_\ell$, and $\mathbf{V}^{(r)}_\ell$ denote the sequence-level query, key, and value tensors at loop r. We suppress batch and attention-head indices below.

Shared global and loop-specific local KV. The first loop performs causal self-attention and constructs a separate

global KV bank at each layer:

$$\left(\mathbf{K}_\ell^{\mathrm{g}}, \mathbf{V}_\ell^{\mathrm{g}}\right) = \left(\mathbf{K}_\ell^{(1)}, \mathbf{V}_\ell^{(1)}\right). \tag{13}$$

These tensors are constructed once per training forward pass and reused by all subsequent loops at the corresponding layer. Each later loop computes its own queries and local KV representations. For $r > 1$, the two attention branches are

$$\begin{aligned} \mathbf{A}_{\ell,\mathrm{g}}^{(r)} &= \mathrm{Attn}\Big(\mathbf{Q}_\ell^{(r)}, \mathbf{K}_\ell^{\mathrm{g}}, \mathbf{V}_\ell^{\mathrm{g}}; \mathbf{M}^{\mathrm{g}}\Big), \\ \mathbf{A}_{\ell,\mathrm{loc}}^{(r)} &= \mathrm{Attn}\Big(\mathbf{Q}_\ell^{(r)}, \mathbf{K}_\ell^{(r)}, \mathbf{V}_\ell^{(r)}; \mathbf{M}^{\mathrm{loc}}\Big), \end{aligned} \tag{14}$$

where Attn denotes scaled dot-product attention with an additive mask. The two branches normalize their attention distributions independently.

Causal masking and gated combination. Let D_c identify the document segment containing position c. For valid query and key positions, the masks are

$$\begin{aligned} M_{c,s}^{\mathrm{g}} &= \begin{cases} 0, & s \le c,\ D_s = D_c, \\ -\infty, & \text{otherwise}, \end{cases} \\ M_{c,s}^{\mathrm{loc}} &= \begin{cases} 0, & 0 \le c - s < w,\ D_s = D_c, \\ -\infty, & \text{otherwise}. \end{cases} \end{aligned} \tag{15}$$

Thus, global attention accesses the observed prefix within the same document, while local attention accesses the current position and its preceding $w - 1$ positions. We use $w = 64$. Queries and keys use the corresponding within-document rotary positions consistently across loops.

For each later loop, a query-dependent sigmoid gate combines the two outputs:

$$\mathbf{A}_\ell^{(r)} = \left(1 - \mathbf{G}_\ell^{(r)}\right) \odot \mathbf{A}_{\ell,\mathrm{g}}^{(r)} + \mathbf{G}_\ell^{(r)} \odot \mathbf{A}_{\ell,\mathrm{loc}}^{(r)}. \tag{16}$$

The gate assigns one scalar to each sequence position and attention head, broadcast across value features. It is computed from normalized queries before rotary position encoding, with gate parameters shared across loops. The first loop uses only the global branch.

Token-parallel execution and gradients. Training first evaluates the first loop over the entire teacher-forced sequence, retaining its per-layer global KV tensors. Subsequent loops are evaluated sequentially, with all sequence positions processed in parallel within each layer. The attention masks enforce the logical cache visibility directly on these sequence-level tensors. The shared global KV tensors remain connected to the computation graph, allowing gradients from all subsequent uses to accumulate into the first-loop representations. The final-loop next-token loss is backpropagated through both attention branches and all recurrent and Foresight aggregation paths.

Consistency with decoding. During decoding, each layer projects the active loop inputs, adds the current first-loop KV to the shared global bank, and adds each later loop's KV to its local window before evaluating attention. The current-cycle entries are therefore visible under the same masks used during training. This ordering requires layer-level coordination while preserving concurrent execution across loops. For identical input tokens, loop-ordered training and cycle-ordered decoding evaluate the same attention dependencies. The bounded local caches apply to incremental decoding; training retains the differentiable activations required for backpropagation.