arXiv:2611.30634v2 [cs.LG] 24 Nov 2026

Rewirable Circuits in Transformers with Loom Attention

Da Xiao [1] **Qingye Meng** [2] **Shengping Li** [2] **Xingyuan Yuan** [2]

[1]Beijing University of Posts and Telecommunications [2]ColorfulClouds Technology

xiaoda99@bupt.edu.cn **Code:** https://github.com/Caiyun-AI/LoomFormer

Abstract

We propose loom attention, which allows a Transformer to dynamically rewire its circuits at every token by interweaving a head's vertical attention over the outputs of earlier heads (to obtain its query, key and value) with its horizontal attention over tokens. We instantiate vertical attention efficiently with normalized and gated linear attention over a per-token matrix memory, which stores the compressed head outputs and serves as a second residual stream. This yields LoomFormer, a dual-stream architecture that splits control and content between the vector and matrix streams. Across model sizes and training token budgets in language modeling and downstream tasks, LoomFormer matches Transformers trained with over 3× compute and recent residual-stream baselines (HC/mHC, MUDDFormer, AttnRes) trained with ~1.7–2.8×. For example, trained on the same 300B tokens with only 2% extra FLOPs, LoomPythia-2.8B achieves ~1.6× the average downstream gain of mHC-Pythia-2.8B over Pythia-2.8B, and matches (0-shot) or surpasses (5-shot) Pythia-12B.

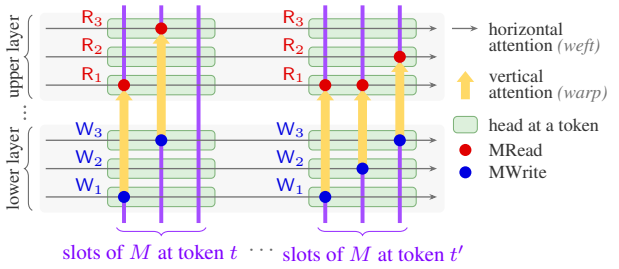

Figure 1. Conceptual sketch of loom attention. At each token, each head chooses which slots of that token's matrix memory M to read and write (i.e., vertical attention, every head does both; only one is shown per layer) and which tokens to attend to (horizontal attention), so a different circuit is woven for each token.

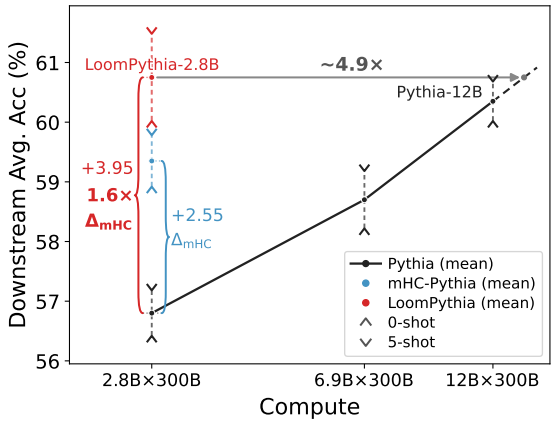

Figure 2. Downstream accuracy averaged over 10 tasks.

1. Introduction

Transformers solve tasks by composing attention heads across layers into circuits (Elhage et al., 2020; Olsson et al., 2022; Wang et al., 2023; Merullo et al., 2024): an earlier head writes a feature into a subspace of the residual stream, and a later head reads from that subspace into its query, key or value. Induction heads, for example, read the output of a previous-token head through their keys (K-composition) (Olsson et al., 2022).

The composition between two heads, however, is determined by the product of their output and input projections, a *virtual weight* that does not change after training (Elhage et al., 2020). Attention makes the *horizontal* flow of information across tokens input-dependent: a head chooses *which tokens* to attend to. In contrast, the *vertical* flow across layers is fixed: a head cannot choose *which heads* to read from. A model that needs different circuits for different inputs therefore keeps all of them wired at once, lets them compete for the residual stream, and can gate them only indirectly, e.g., by sending a head's attention to a sink token (Xiao et al., 2024b). We ask whether the wiring itself can be chosen per token, i.e., whether circuits can be made *rewirable*.

Recent work has enriched the vertical flow. Dense connections let a layer aggregate the outputs of all earlier layers (Pagliardini et al., 2024); MUDDFormer (Xiao et al., 2025) and Attention Residuals (AttnRes) (Kimi Team, 2026a) make this aggregation input-dependent, turning it into a depth-wise (i.e., vertical) attention (separate for the query,

key, value and residual inputs in MUDDFormer); and Hyper-Connections (HC) (Zhu et al., 2025) and its manifold-constrained variant mHC (Xie et al., 2025) maintain several residual streams with learnable or dynamic mixing. In all of them, however, the vertical channel meets horizontal attention only through the hidden state of a layer, in which the outputs of all heads have already been summed and mixed together: a head can neither single out a particular earlier head nor be singled out by a later one. The vertical and horizontal flows run side by side, never woven together.

To make circuits rewirable, we propose *loom attention*, in which a head chooses which earlier heads to read from just as it chooses which tokens to attend to. By attending vertically over the outputs of these heads at the same token, the head obtains its query, key and value to perform the standard horizontal attention over tokens, and its output is in turn attended to vertically by later heads. Vertical and horizontal attention are thus interwoven, each feeding the other at every head, like the warp and weft of a loom (Figure 1).

Naive vertical attention over the outputs of all earlier heads is, however, expensive. We therefore instantiate it with normalized and gated linear attention over a per-token matrix memory, which stores the compressed head outputs and serves as a second residual stream. We call the resulting dual-stream architecture *LoomFormer*, after a drawloom—a loom in which a separate stored pattern selects which warp threads to lift at each pick, rather than repeating one fixed pattern. [1] Here the ordinary vector stream plays the pattern, carrying control—which memory slots each head writes and reads (weaving a different circuit per token) as well as which tokens to attend to—while the matrix stream carries the content moved across tokens. The two streams are coupled through attention outputs and FFN writes.

Extensive experiments on language modeling and downstream tasks across model sizes and training token budgets show that LoomFormer adds only 2.0–3.7% FLOPs while matching Transformers trained with over 3× compute and recent residual-stream baselines (HC/mHC, MUDDFormer, AttnRes) trained with ~1.7–2.8×. For example, trained on the same 300B tokens, LoomPythia-2.8B achieves ~1.6× the average downstream gain of mHC-Pythia-2.8B over Pythia-2.8B, and matches (0-shot) or surpasses (5-shot) Pythia-12B (Figure 2). We also analyze the trained 2.8B models to understand how and why LoomFormer works—e.g., by rewiring head outputs to inputs per token, it keeps over three times as many deep-layer heads off attention sinks as Pythia and mHC.

[1] The automated form of the drawloom, the Jacquard loom, encoded this pattern on punched cards and inspired early programmable computers such as Babbage's Analytical Engine.

2. Rewirable Circuits

Writers, readers and circuits. Attention heads communicate through the residual stream (Elhage et al., 2020). Consider a Transformer with L layers of n heads of dim d, and model dim $D = nd$. We call the output of a head a *writer* and each of its inputs (query, key and value) a *reader*, so a head has one writer and three readers. Writer W_i adds its output $o_i \in \mathbb{R}^d$ to the residual stream through its block $W_i^O \in \mathbb{R}^{d\times D}$ of the output projection, and reader R_j of a later head reads the stream through its input projection $W_j^s \in \mathbb{R}^{D\times d}$ ($s \in \{Q, K, V\}$), so W_i reaches R_j through the *virtual weight* $\bar{C}_{ij} = W_i^O W_j^s \in \mathbb{R}^{d\times d}$: the two communicate only if W_i writes into a subspace that R_j reads. Heads composed in this way form *circuits* whose nodes are the writers and readers.

A task may need different circuits at different tokens, i.e., different pairs of writers and readers to be connected (Figure 3). To describe how the wiring can change, we give it a dynamic part besides the static virtual weights: at a token, reader R_j receives

$$u_j = \textstyle\sum_i o_i \, c_{ij} \bar{C}_{ij}, \tag{1}$$

where the per-token scalars c_{ij} form the *wiring matrix* $C \in \mathbb{R}^{w\times r}$ from the w earlier writers to the r readers of one layer, and $c_{ij}\bar{C}_{ij}$ is the *effective wiring* from W_i to R_j. The two are connected at the token if their effective wiring is nonzero.[2] A *wiring pattern* $P \subseteq [w] \times [r]$ is a set of writer-reader pairs, realized when each reader in P is connected to exactly its writers in P.

Definition 2.1. An architecture with fixed virtual weights is *rewirable w.r.t.* a set of wiring patterns if every pattern in the set is realized by some wiring matrix C that it admits at a token.

Transformers are node-gated. In a Transformer, pre-LN only rescales all writers of a token by a common factor (Theorem A.1 in Section A.1). Within a token, nodes can thus only be gated through whole heads: a head that attends to a sink token (Xiao et al., 2024b) outputs nearly nothing, turning off both its writer and, since what its readers receive has no effect, its readers.

Definition 2.2. An architecture is *node-gated* if, at every token, its wiring matrix has the form

$$C = g\, g'^{\top}, \tag{2}$$

with gate vectors $g \in \mathbb{R}^w$ and $g' \in \mathbb{R}^r$ whose nonzero entries mark the writers and readers that are on.

[2] In trained models entries are rarely exactly zero; we count those below a threshold as zero (Section A.3).

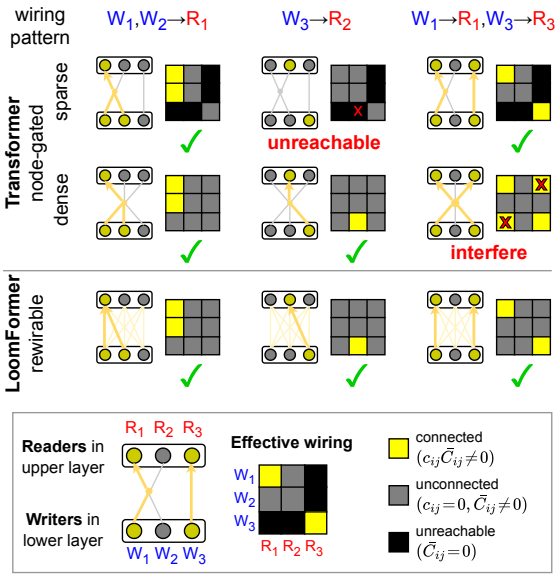

Figure 3. Effective wiring for three wiring patterns. Transformers (top two rows) can only gate whole heads, i.e., rows and columns, on and off, so they connect only rectangles within the fixed support of their virtual weights. LoomFormer (bottom row) connects exactly the required pairs.

We call the rank of C the *wiring rank*; a node-gated architecture has wiring rank at most one, so existing connections can be gated on and off, but not rewired.

Proposition 2.3. *A Transformer with pre-LN is node-gated: at any token, it connects precisely the pairs with $\bar{C}_{ij} \neq 0$ whose writer and reader are both on, a rectangle within the fixed support of the virtual weights.*

Gating a writer or a reader acts on its whole row or column of C, so a wiring pattern is realized only if it is such a rectangle, as in the first column of Figure 3; other patterns fail in one of two ways. With *dense* virtual weights, where heads share subspaces, every pair can be connected, but a non-rectangular pattern such as $\{\mathsf{W}_1 \to \mathsf{R}_1, \mathsf{W}_3 \to \mathsf{R}_3\}$ (last column) also connects $\mathsf{W}_1 \to \mathsf{R}_3$ and $\mathsf{W}_3 \to \mathsf{R}_1$: the two circuits *interfere*. With *sparse* virtual weights, where groups of heads use separate subspaces, circuits in different groups do not interfere, but pairs across groups such as $\mathsf{W}_3 \to \mathsf{R}_2$ (second column) are *unreachable*.

The price of rewiring by node gating. Can a node-gated architecture overcome this by spending more heads? Suppose the role of each of $r \geq 2$ readers may be played by several heads, gated on and off per token, and at every token each reader must receive exactly one of w writers, chosen independently of the other readers; the last pattern in Figure 3 is the smallest instance. Let $N(w, t)$ denote the minimum size of a ground set carrying a t-cover-free family of w sets, i.e., w sets none of which is covered by the union of t others (Kautz & Singleton, 1964).

Theorem 2.4. *To realize all such wiring patterns, a node-gated architecture needs at least $N(w, r-1)$ heads per reader, and this number suffices, where $N(w, 1) \approx \log_2 w$, $N(w, t) = \Omega(t^2 \log w / \log t)$ for $2 \leq t \leq \sqrt{w}$ and $N(w, t) = \Theta(w)$ for $t > \sqrt{w}$.*

Intuitively, a head receives every writer that is on and has a nonzero virtual weight to it, so a reader must hold in reserve a head that sees its chosen writer and none of the up to $r-1$ writers chosen by the other readers (Theorem A.2 in Section A.2; bounds in Section A.3). With the $r = 3n$ readers of a layer choosing among its $w = nl$ earlier writers, $r - 1 > \sqrt{w}$ whenever $l \leq 9n - 6$, so at practical sizes node gating needs $\Theta(w)$ heads per reader.

Naive loom attention makes Transformers rewirable. To rewire with one head per reader, the wiring matrix must be chosen entry by entry at each token. The most direct way is to let every reader attend over the outputs of all earlier heads of the same token, with attention weights generated from the token itself:

$$u_j = \textstyle\sum_i \alpha_{ij}(x)\, o_i, \qquad \alpha_{\cdot j}(x) = \mathrm{softmax}(x W_j^{\alpha}), \quad (3)$$

where $W_j^{\alpha} \in \mathbb{R}^{D \times nl}$ for a reader in layer l, which has nl earlier writers. We call this *naive loom attention*: its virtual weights are identities and all the wiring is left to $C = [\alpha_{ij}]$, unlike a Transformer, which wires heads through fixed virtual weights and only gates them through a rank-one C. With one head per reader, it is rewirable w.r.t. all wiring patterns (cf. Figure 3, bottom row), but expensive: over the $3n$ readers of all L layers, generating the attention weights takes about $\frac{3}{2} n^2 L^2 D$ parameters and multiply-adds per token, growing quadratically in the total number of heads nL and rivaling the whole model at practical sizes (Section A.4).

3. Method

3.1. From Naive to Affordable Loom Attention

As linear attention compresses a growing key-value cache into a fixed-size state (Katharopoulos et al., 2020; Schlag et al., 2021), we compress the head outputs of each token along depth into a matrix memory $M \in \mathbb{R}^{d_k \times d}$ of d_k slots of width d. Writer W_i writes its output to M with a key $k_i^m \in \mathbb{R}^{d_k}$, and reader R_j reads with a query $q_j^m \in \mathbb{R}^{d_k}$, so a reader attends over d_k slots instead of all earlier writers, at a cost independent of depth. Since M accumulates writes layer by layer, it is a second residual stream, the *matrix stream*, besides the original *vector stream* (Figure 4 left).

We call the interface through which a head reads from or writes to a stream a *port*. The two streams have dual ports:

$$\begin{aligned} &\text{vector stream: } u_j = x W_j^s, \quad \Delta x = o_i W_i^O \\ &\text{matrix stream: } u_j = q_j^m M, \quad \Delta M = k_i^m \otimes o_i \end{aligned} \quad (4)$$

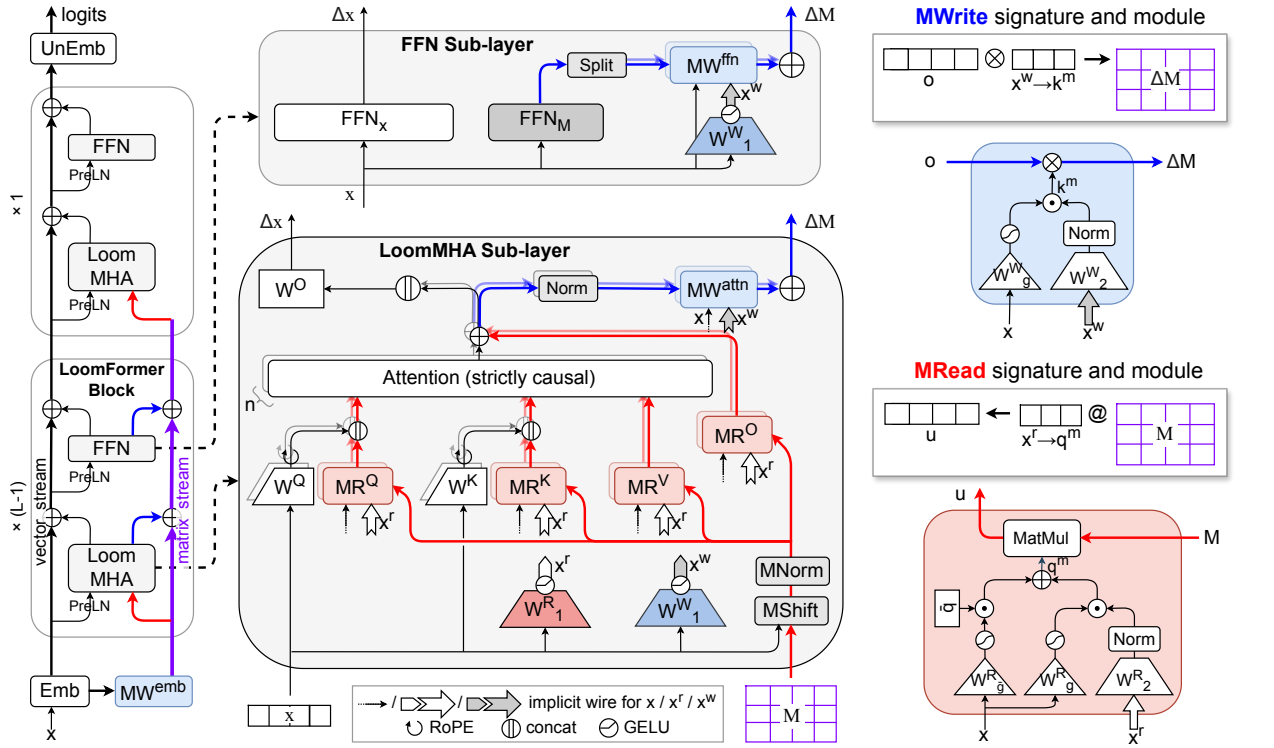

Figure 4. LoomFormer architecture: the model (left), the LoomMHA and FFN sub-layers of a block (middle), and the MRead (MR) and MWrite (MW) modules (right). Vertical attention: reads from M are in red and writes to M in blue.

where $k_i^m \otimes o_i = (k_i^m)^\top o_i$ is the outer product. Data and ports swap shapes: a port shrinks from a $D \times d$ matrix to a vector of length d_k, small enough to be generated per token from the vector stream, which we keep as the *control path*.

3.2. Architecture

LoomFormer keeps the backbone of the Transformer block and changes how sub-layers interface with the two streams: the vector stream carries control, i.e., where to attend and how to wire, and the matrix stream carries the content that is moved (Figure 4). We describe the architecture bottom up, from ports to sub-layers to the model (pseudocode in Section C), with x denoting a sub-layer's pre-LN input.

MRead and MWrite. MRead and MWrite realize the matrix-stream ports of Eq. (4). $\mathrm{MRead}(M, x, x^r)$ generates a query $q^m \in \mathbb{R}^{d_k}$ to read $u \in \mathbb{R}^d$ from M, where $x^r = \mathrm{GELU}(xW_1^R) \in \mathbb{R}^{D_r}$ with $W_1^R \in \mathbb{R}^{D \times D_r}$ ($D_r \ll D$) is a low-rank feature shared by all MRead modules of a sub-layer to keep port generation cheap:

$$\begin{aligned} q^m &= \mathrm{Norm}(x^r W_2^R) \odot \sigma(xW_g^R) + \bar{q} \odot \sigma(xW_{\bar{g}}^R) \\ u &= q^m M \end{aligned} \tag{5}$$

Here x^r is projected by $W_2^R \in \mathbb{R}^{D_r \times d_k}$, normalized with RMSNorm and gated to get a dynamic query, which is summed with a static query $\bar{q}$ (also gated) to obtain $q^m \in \mathbb{R}^{d_k}$. The two scalar gates are computed from x by projecting to scalars with $W_g^R, W_{\bar{g}}^R \in \mathbb{R}^{D \times 1}$ and then applying sigmoid. Similarly, with a shared $x^w = \mathrm{GELU}(xW_1^W) \in \mathbb{R}^{D_r}$, $\mathrm{MWrite}(o, x, x^w)$ generates a key $k^m \in \mathbb{R}^{d_k}$ and writes a vector $o \in \mathbb{R}^d$ into M as their outer product:

$$\begin{aligned} k^m &= \mathrm{Norm}(x^w W_2^W) \odot \sigma(xW_g^W) \\ \Delta M &= k^m \otimes o \end{aligned} \tag{6}$$

Normalization and gating separate a port's direction from its magnitude: the normalized vector selects which slots to read or write, and the gate alone sets how much.

LoomMHA. To access the previous token's memory without horizontal attention, MShift first mixes $M_{[t]}$ with $M_{[t-1]}$ via a per-slot gate $z_{[t]} = \sigma(x_{[t]} W_g^S) \in \mathbb{R}^{d_k}$ ($W_g^S \in \mathbb{R}^{D \times d_k}$), and MNorm RMS-normalizes the matrix into $\widetilde{M}_{[t]}$. Each head then reads four vectors from $\widetilde{M}$ with separate MRead modules: query, key, value and output bypass o^m:

$$\begin{aligned} \widetilde{M}_{[t]} &= \mathrm{MNorm}\left(z_{[t]} \odot M_{[t-1]} + (1 - z_{[t]}) \odot M_{[t]}\right) \\ \tilde{q}, \tilde{k}, v, o^m &= \mathrm{MRead}^{Q,K,V,O}(\widetilde{M}, x, x^r) \\ q, k &= \mathrm{RoPE}(xW^{Q,K}) \,\|\, (\tilde{q}, \tilde{k}) \end{aligned} \tag{7}$$

where $W^Q, W^K \in \mathbb{R}^{D \times d/4}$ project x to $d/4$ dimensions, to which rotary position embedding (RoPE) (Su et al., 2024)

is applied; the results are concatenated ($\|$) with $\tilde{q}$ and $\tilde{k}$ read from M. The value is read from M alone: the vector stream supplies only ports and RoPE queries/keys for slot and token routing, while all content moved by attention comes from M. Unlike MHA, LoomMHA has two outputs: the projected head outputs for the vector stream, and the writes of all heads $i \in \{1, \dots, n\}$ summed into ΔM:

$$\begin{aligned} o_i &= o_i^m + \mathrm{Attention}(q_i, k_i, v_i) \\ \Delta M &= \textstyle\sum_{i=1}^{n} \mathrm{MWrite}_i(\mathrm{Norm}(o_i), x, x^w) \\ \mathrm{LoomMHA}(x, M) &= \left(\mathrm{Concat}_{i=1}^{n}[o_i] W^O,\ \Delta M \right) \end{aligned} \tag{8}$$

where q_i, k_i, v_i, o_i^m are given by Eq. (7) for head i, and Attention is scaled dot-product attention over tokens with a strictly causal mask that excludes the current token: a token's own content reaches the head's output through the bypass o_i^m instead. LoomMHA is where vertical and horizontal attention are interwoven: the vertical reads from M feed the horizontal attention of head i, and its output o_i is written back to M for later layers.

FFN with MWrite. The feed-forward network (FFN) sublayer reads only x and, like LoomMHA, returns a pair $(\mathrm{FFN}_x(x), \Delta M)$, where FFN_x is a standard FFN module. A second, smaller FFN_M computes the values to write, which are split into n head-sized vectors (split_i) and written to M with the sub-layer's MWrite modules. This write is the only channel through which content of the vector stream enters M, since attention moves only content read from M.

$$\Delta M = \textstyle\sum_{i=1}^{n} \mathrm{MWrite}_i(\mathrm{split}_i(\mathrm{FFN}_M(x)), x, x^w) \tag{9}$$

LoomFormer. The embedding initializes both residual streams; M_0 is written from the embedding x_0 in the same way as the FFN writes:

$$\begin{aligned} x_0 &= \mathrm{Embedding}(X) \\ M_0 &= \textstyle\sum_{i=1}^{n} \mathrm{MWrite}_i(\mathrm{split}_i(x_0 W^E), x_0, x_0^w) \end{aligned} \tag{10}$$

where $W^E \in \mathbb{R}^{D \times D}$ is a value projection and $x_0^w = \mathrm{GELU}(x_0 W_1^W)$. Each block l then applies the two sublayers with pre-LN, updating both streams residually (pairs are added element-wise):

$$\begin{aligned} (x_l', M_l') &= (x_l, M_l) + \mathrm{LoomMHA}(\mathrm{LN}_1(x_l), M_l) \\ (x_{l+1}, M_{l+1}) &= (x_l', M_l') + \mathrm{FFN}(\mathrm{LN}_2(x_l')) \end{aligned} \tag{11}$$

for $l = 0, \dots, L-1$. Finally, x_L is normalized and projected to produce the final output. As M_L is not used, the last block does not write to M. We typically set $d_k = 2n$ slots (so M has size $2D$) and $D_r = nd_k/2 = n^2$ (a $2\times$ compression bottleneck relative to nd_k), matching the Transformer's parameter budget by adjusting FFN widths at an approximately 2:1 ratio between FFN_x and the FFN_M write branch.

3.3. Rewiring Capacity

On the matrix stream without MShift, reader R_j receives $u_j = q_j^m M = \sum_i \langle q_j^m, k_i^m \rangle o_i$ by Eq. (4). The matrix stream thus keeps the division of roles of naive loom attention: its virtual weights are identities, and its wiring matrix is $C = KQ^\top$, of rank at most d_k, where the rows of K and Q are the keys k_i^m of the writers and the queries q_j^m of the readers.

Proposition 3.1. *(i) With ports fixed to (K, Q), LoomFormer is a standard Transformer $\mathrm{Tfm}_{K,Q}$ with residual stream $(x, \mathrm{vec}(M)) \in \mathbb{R}^{D+d_k d}$, whose virtual weights absorb $KQ^\top$, hence node-gated. (ii) With ports generated per token, LoomFormer is rewirable with one head per reader w.r.t. the wiring patterns with at most d_k writers, whatever the total number of writers w.*

In (ii), only the writers in the pattern are on, each taking a slot through a one-hot key while each reader's query selects the slots of its writers (Section A.6). The $N(w, r-1)$ heads per reader that node gating needs (Theorem 2.4) are thus replaced by a per-token assignment of d_k slots; when more than d_k writers are on, slots are shared and writers interfere, as in linear attention. By (i), this rewiring capacity comes from generating ports per token, not from the matrix memory itself. Together, (i) and (ii) give a model-level view of LoomFormer's expressiveness: within each token, LoomFormer wires heads as $\mathrm{Tfm}_{K(x),Q(x)}$, selecting per token one Transformer from the continuum $\{\mathrm{Tfm}_{K,Q}\}$ much as a mixture of experts selects experts (Fedus et al., 2022), except that the experts are wirings, share all other weights, and cost only d_k numbers per port to select.

4. Connections to Prior Work

Table 1 compares LoomFormer with residual-stream mechanisms (derivations in Section A.7; extended related work in Section B). In DenseFormer (Pagliardini et al., 2024), HC (Zhu et al., 2025), mHC (Xie et al., 2025), AttnRes (Kimi Team, 2026a) and MUDDFormer (Xiao et al., 2025), the ports are shared by all heads of a layer: apart from node gating and the fixed projections, their wiring depends on a writer only through its layer and on a reader only through its layer (and Q/K/V input in MUDDFormer). Their wiring rank is therefore 1, or 3 in MUDDFormer, one per Q/K/V input, whereas LoomFormer generates its ports per head and per token and reaches $\min(d_k, 3n)$, i.e., $2n$ by default, with a smaller state than HC ($2D + D$ against $4D$).

Linear attention. Linear attention (Katharopoulos et al., 2020; Schlag et al., 2021) and its gated and delta-rule variants (Dao & Gu, 2024; Yang et al., 2024a;b; 2025; Peng et al., 2024) also read and write a fixed-size matrix state per head with keys and queries from the vector stream, but

Table 1. LoomFormer among residual-stream mechanisms. Ports k^m, q^m are indexed by heads i, j or, when shared across a layer, by layers l_i, l_j (and Q/K/V input s), and marked (x) or (M) when computed from the vector stream or from the memory itself; $\mathrm{softmax}_{l_i}$ normalizes over writer layers. Typical sizes assume $d_k = 4$ in HC and $d_k = 2n$ in LoomFormer ($+D$ for its vector stream).

	Effective wiring	Wiring rank	Granularity	Memory: slots × width = typical size	Residual stream
Transformer	$\bar{C}_{ij}$	1	–	$1 \times D = D$	vector
DenseFormer	$\langle q^m_{l_j}, k^m_{l_i} \rangle \bar{C}_{ij}$	1	layer	$l \times D = lD$	vector
HC / mHC	$\langle q^m_{l_j}(M), k^m_{l_i}(M) \rangle \bar{C}_{ij}$	1	layer	$d_k \times D = 4D$	matrix
AttnRes (full)	$\mathrm{softmax}_{l_i} \langle q^m_{l_j}, k^m_{l_i}(M) \rangle \bar{C}_{ij}$	1	layer	$2l \times D = 2lD$	none
MUDDFormer	$\langle q^{m,s}_{l_j}(x), k^m_{l_i} \rangle \bar{C}_{ij}$	3	layer × QKV	$l \times D = lD$	vector
LoomFormer	$\langle q^{m,s}_{j}(x), k^m_{i}(x) \rangle I$	$\min(d_k, 3n) = 2n$	head × QKV	$d_k \times d = 2D$ $(+D)$	matrix + vector

differ in axis and sharing. Their state accumulates across tokens and its read is the head output, replacing horizontal attention, whereas M accumulates along depth and its reads feed the query, key and value of horizontal attention. Their state is also private to one head across tokens, so heads still compose through the static ports of the vector stream, whereas M is shared by all heads of all layers at a token, as a channel through which heads read each other. The two are thus orthogonal and can be combined.

5. Experiments

We train all models on the Pile dataset (Gao et al., 2020) in JAX, initializing dynamic port gates near zero so training starts smoothly from static ports (Section D.1). We evaluate scaling laws (Section 5.1), large-scale pretraining and downstream performance (Section 5.2), mechanisms of the trained models (Section 5.3), computational efficiency (Section 5.4) and ablations (Section 5.5).

5.1. Scaling Laws

Settings. We train three model scales—Medium (432M), Large (778M) and XL (1.4B)—with sequence length $T = 4096$, batch size 0.5M tokens (128 sequences), and token budgets proportional to non-embedding parameters following Hoffmann et al. (2022) (architectural dimensions and optimizer hyperparameters in Section D.1).

Baselines. We compare LoomFormer against recent residual-stream baselines on the same Transformer++ backbone (Touvron et al., 2023) (RoPE (Su et al., 2024), SwiGLU FFN (Shazeer, 2020), RMSNorm, no linear bias): DHC (dynamic HC with expansion rate 4; Zhu et al., 2025), mHC (Xie et al., 2025), AttnRes (full) (Kimi Team, 2026a), and MUDDFormer (Xiao et al., 2025) (details in Section D.2).

Results. Across the Medium, Large and XL scales (Figure 5), mHC, AttnRes and DHC improve over Transformer++, and MUDDFormer improves further, while LoomFormer achieves the lowest loss at every scale. Interpolating log-linearly along each baseline curve across Medium and Large, LoomFormer matches Transformer++ trained with 3.19–3.25× compute (3.10× net at Large after the 3.1% FLOP overhead in Table 3) and all four residual-stream baselines at 1.73–2.76×. At XL, where LoomFormer lies below the largest model of every baseline, it lowers validation loss from 2.101 to 1.979 ($\Delta L = 0.122$, 1.8× MUDDFormer's 0.068) at 2.8% extra FLOPs (Table 3).

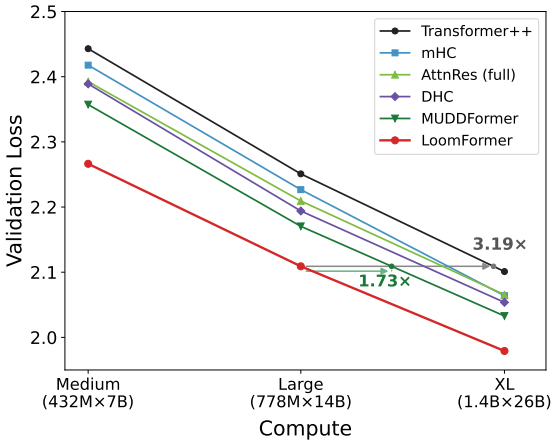

Figure 5. Scaling curves on Pile validation loss against training compute (total parameters × training tokens).

5.2. Large-Scale Pretraining and Evaluation

Settings. To test whether LoomFormer's gains persist under extended training, we compare against the Pythia suite (Biderman et al., 2023) on 300B Pile tokens (over 5× the Chinchilla budget of Section 5.1; Hoffmann et al., 2022). Specifically, we train LoomPythia-2.8B and mHC-Pythia-2.8B—the residual-stream baseline that also keeps a depth-independent state size (Table 1)—and compare them with Pythia 2.8B–12B. Except for using serial instead of parallel sub-layers[3], both LoomPythia and mHC-Pythia use the same architecture choices and training hyperparameters as

[3]mHC requires serial sub-layers by design, and parallel vs. serial FFN placement is quality-neutral for standard Transformers, as shown by Chowdhery et al. (2023) and our own experiments.

Table 2. Zero/five-shot downstream evaluation results (Δacc is over Pythia-2.8B; best in bold within each block).

Model	Pile ppl↓	FLAN ppl↓		LAM BADA	PIQA	Wino Grande	ARC -E	ARC -C	SciQ	BoolQ	Hella Swag	RACE -M	RACE -H	Avg acc↑ /Δacc
Pythia-2.8B	6.63	8.16		64.7	73.9	59.4	64.4	29.5	88.2	64.5	45.4	38.1	34.9	56.3
Pythia-6.9B	6.29	7.85		67.3	75.2	60.9	67.3	31.3	89.7	63.7	48.0	40.6	37.0	58.1
mHC-Pythia-2.8B	6.24	7.35	*0-shot*	68.5	74.4	62.8	67.1	30.7	90.5	66.9	47.5	**42.3**	37.0	58.8/+2.5
Pythia-12B	**6.01**	7.26		70.5	**76.0**	**63.9**	**70.2**	**31.8**	90.2	**67.4**	**50.3**	40.6	38.3	**59.9**
LoomPythia-2.8B	6.02	**6.94**		**71.8**	74.7	63.7	68.5	31.7	**92.4**	65.8	49.1	**42.3**	**39.0**	**59.9**/+3.6
Pythia-2.8B	-	-		60.5	73.6	60.6	67.3	32.3	94.3	65.6	45.1	38.4	35.6	57.3
Pythia-6.9B	-	-		63.8	75.5	63.7	70.2	35.6	95.1	65.7	48.1	39.0	36.5	59.3
mHC-Pythia-2.8B	-	-	*5-shot*	62.4	74.9	64.1	70.8	35.5	95.6	69.4	47.6	40.3	38.2	59.9/+2.6
Pythia-12B	-	-		67.3	**76.0**	64.2	**71.0**	**36.5**	95.3	68.0	**50.3**	40.1	38.8	60.8
LoomPythia-2.8B	-	-		**69.0**	75.7	**65.7**	70.8	35.1	**96.0**	**71.2**	49.5	**43.1**	**40.1**	**61.6**/+4.3

Pythia (Biderman et al., 2023, Appendix E). We evaluate zero- and five-shot accuracy across 10 standard benchmarks.

Results. As shown in Table 2 and Figure 2, LoomPythia-2.8B improves average accuracy over Pythia-2.8B by +3.6 in 0-shot and +4.3 in 5-shot (+3.95 on their mean), achieving 1.6× the gain of mHC-Pythia-2.8B (+2.5 / +2.6, mean +2.55). Its 0-shot average (59.9) matches Pythia-12B (59.9), its 5-shot average (61.6) exceeds Pythia-12B (60.8), and its 0/5-shot mean (60.75) matches a log-linear extrapolation of Pythia at ~4.9× compute. On pretraining metrics, LoomPythia-2.8B reaches Pile ppl 6.02, on par with Pythia-12B (6.01) and well ahead of mHC-Pythia-2.8B (6.24).

As an out-of-domain evaluation, Table 2 also reports target-span perplexities on a randomly sampled subset of the FLAN Collection (Longpre et al., 2023), which features instruction following, chain-of-thought and in-context few-shot data. The advantage of LoomPythia on FLAN is even larger than on Pile: FLAN ppl 6.94 is best in the table, beating both Pythia-12B (7.26) and mHC-Pythia-2.8B (7.35). Whereas mHC's gain barely changes from 0-shot to 5-shot (+2.5 vs. +2.6), LoomPythia's widens from +3.6 to +4.3, indicating stronger in-context learning that Section 5.3 traces to per-token rewiring freeing deep heads from attention sinks.

5.3. Mechanistic Analysis

We analyze Pythia-2.8B, mHC-Pythia-2.8B and LoomPythia-2.8B (Section 5.2) on held-out Pile validation sequences; metric definitions and additional visualizations are in Section H.

Per-token circuit rewiring. After subtracting each pair's token-mean static component, the dynamic wiring matrix on the $3n = 96$ Q/K/V readers of layer l has rank 0 under fixed ports (Theorem 3.1(i)), at most 1 in Transformer and mHC, at most 3 in MUDDFormer, and up to $\min(d_k, 3n) = 64$ in LoomFormer (Table 1). In Figure 6(a), LoomPythia-2.8B's dynamic effective rank rises with depth to 41.0 (26.4 mean), placing 71.0% of dynamic wiring energy (90.7% at layer 32) outside rank 1, and averages 14.3 (20.9 at layer 32) within a single Q, K or V input where MUDDFormer also has rank ≤ 1. Figure 6(b) directly tests Theorem 3.1(i) vs. (ii) in $\mathbb{R}^{d_k}$: whereas static or scalar-gated ports have zero cross-token cosine distance, LoomPythia's read ports $\tilde{q}, \tilde{k}, v, o^m$ average 0.33–0.46 (up to 0.68) and its attention and FFN_M write ports average 0.36 and 0.56 (up to 0.75; Section H traces an in-context circuit rewiring across adjacent tokens).

Head utilization and within-head specialization. In a node-gated model (Section 2, Theorem 2.4), heads whose

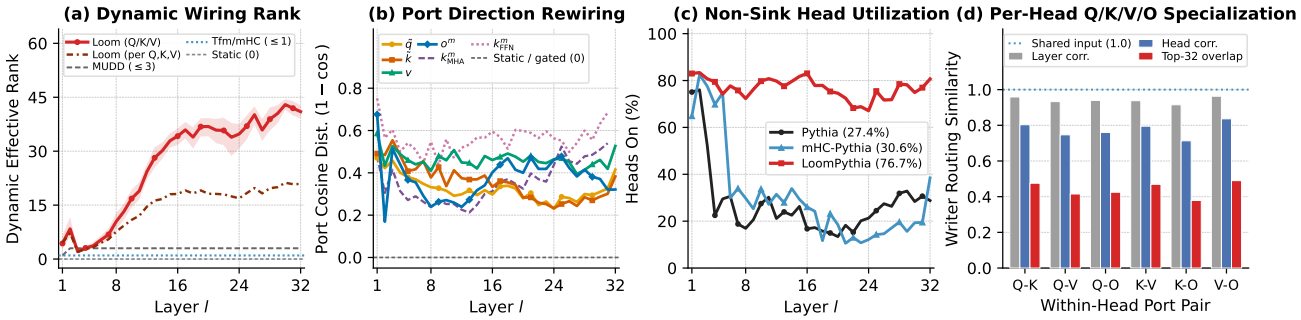

Figure 6. Mechanistic comparison of 2.8B models across layers $l = 1, \ldots, 32$ on Pile validation sequences (definitions in Section H). (a) Energy effective rank of the token-centered dynamic Q/K/V wiring matrix (band: 10th–90th percentiles) and per single Q, K or V input, vs. the bounds of Table 1. (b) Cross-token cosine distance ($1 - \cos$) of read ports ($\tilde{q}, \tilde{k}, v, o^m$) and write ports ($k^m_{\mathrm{MHA}}, k^m_{\mathrm{FFN}}$) in $\mathbb{R}^{d_k}$ (static/gated = 0). (c) Percentage of heads that are on (non-sink; legend: 32-layer mean). (d) Within-head writer routing similarity across port pairs in layers 3–32.

Table 3. Analytical computational overhead (+FLOPs) and measured relative throughput of LoomFormer vs. matched-size Transformer at sequence length T (absolute speeds in Section F).

Model	+FLOPs (%)	Train (%) v5p	Prefill (%) bs 1 / 8	Decode (%) bs 1 / 16 / 64
Medium	3.70	90.5	87.4 / 90.4	91.7 / 89.2 / 92.3
Large	3.13	89.2	85.3 / 86.2	85.4 / 87.8 / 84.9
XL	2.77	92.3	83.8 / 85.3	86.9 / 84.6 / 86.6
2.8B	1.95	92.4	87.5 / 90.5	86.4 / 87.6 / 88.9

fixed inputs are not needed at a token turn themselves off by attending to sink tokens (Vig & Belinkov, 2019; Xiao et al., 2024b). Counting non-sink heads as on (Xiao et al., 2025), Figure 6(c) shows that in layers 17–32, only 22.9% of heads are on in Pythia-2.8B and 17.7% in mHC-Pythia-2.8B—whose four streams still merge into one shared input per layer—whereas LoomPythia-2.8B keeps 74.4% on (3.3× and 4.2×), more than doubling content-token attention mass and accounting for its 5-shot and FLAN gains (Section 5.2). Finally, Figure 6(d) verifies head × QKV specialization (Table 1): although routing looks correlated at the layer level (0.92–0.96), per-writer correlation within the same head drops to 0.71–0.84 and top-32 writer overlap is 0.38–0.49, so over half of the top writers read by a head's $\tilde{q}, \tilde{k}, v, o^m$ are distinct.

5.4. Computational Complexity and Efficiency

Analytical computational overhead. Under a matched parameter budget, LoomFormer adds $\sim \frac{1}{d(1+\rho/12)} + \frac{1}{8(1+12/\rho)}$ relative forward FLOPs from reading/writing M and wider Q/K ($d_k = 2n$, $\rho = T/D$; derived in Section E), totaling 1.95%–3.70% across our four scales (Table 3).

Wall-clock training and inference throughput. Table 3 reports relative throughput for JAX training on TPU v5p pods (with fused Pallas kernels) and BF16 PyTorch inference on an NVIDIA A800 GPU (with a fused Triton kernel and `torch.compile`; Section F). LoomFormer sustains 89%–92% of Transformer training speed and 84%–92% of inference speed across prefill and decode batch sizes; the gap beyond the 2.0%–3.7% FLOPs reflects HBM I/O on M, lower arithmetic intensity of per-token M contractions, and smaller split projections and norms, and is far outweighed by the over 3× compute savings (Sections 5.1 and 5.2).

5.5. Ablation Studies

We ablate LoomFormer at the Medium scale (432M, 7B tokens; Section 5.1), reallocating saved parameters to FFN_x to match the parameter budget, and report Pile validation loss and the retained gain $\mathrm{Gain} = (L_{\mathrm{Tfm++}} - L)/(L_{\mathrm{Tfm++}} - L_{\mathrm{Loom}})$ (Table 4; full suite in Section G).

Dynamic rewiring. Replacing dynamic k^m, dynamic q^m, or both with learned static vectors (dropping their gates) tests Theorem 3.1. Fixing both reduces LoomFormer to $\mathrm{Tfm}_{K,Q}$ and cuts the gain to 56.0%—still ahead of layer-wise dynamic baselines (Figure 5)—so the static head-level dual-stream backbone and per-token rewiring contribute roughly equal halves. Fixing q^m hurts far more than fixing k^m (75.4% vs. 90.3% retained): writers can use stable slots if readers select among them per token.

Table 4. Ablations of LoomFormer at the Medium scale.

Variant	Loss↓	ΔLoss	Gain
LoomFormer (full)	**2.266**	—	**100.0%**
Transformer++ (baseline)	2.443	+0.177	0.0%
Dynamic rewiring (Theorem 3.1)			
Static write keys (k^m)	2.283	+0.017	90.3%
Static read queries (q^m)	2.310	+0.043	75.4%
Static read & write (q^m, k^m)	2.344	+0.078	56.0%
Port normalization & gating (Eqs. 5–6)			
− dynamic port gates	2.277	+0.011	93.7%
− dynamic port Norm	2.289	+0.023	87.1%
− dynamic port gates & Norm	2.283	+0.017	90.3%
Sub-layer components (Section 3.2)			
− FFN_M write	2.300	+0.034	80.7%
− output bypass o^m	2.299	+0.032	81.8%
− vector-stream RoPE $W^{Q,K}$	2.286	+0.020	88.7%
− MShift	2.277	+0.010	94.2%

Port normalization and gating. Removing either the dynamic-branch gates or port RMSNorm hurts, yet removing both beats removing RMSNorm alone (90.3% vs. 87.1%): an unnormalized projection already varies in norm, making a separate gate redundant. Pairing RMSNorm with gating works best by separating slot direction from read/write magnitude.

Sub-layer components. Removing FFN_M (Eq. 9) hurts most, as the sole post-embedding path from the vector stream into M, followed closely by removing the bypass o^m and restoring self-attention to the current token (Eq. 8). Applying RoPE to $d/4$ of the M-read coordinates instead of $xW^{Q,K}$ (Eq. 7) also hurts, supporting separate streams

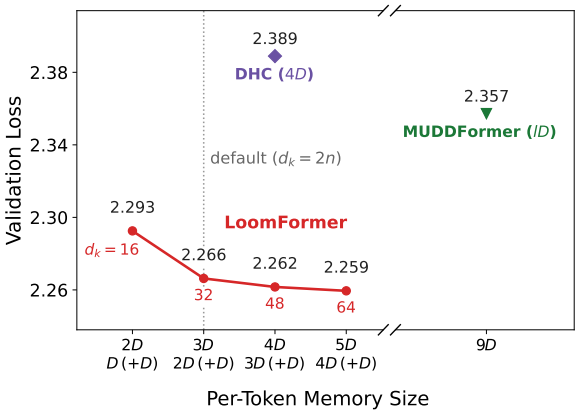

Figure 7. Validation loss vs. per-token memory size and matrix slots d_k at the Medium scale ($L = 18$).

for positions and content, whereas removing MShift still retains 94.2% of the gain.

Memory capacity. As we vary the slot count d_k (Figure 7), validation loss drops sharply (-0.026) from $d_k = 16$ to 32 with diminishing returns beyond, placing our default $d_k = 2n$ at the knee given the linear growth of MRead/MWrite FLOPs (Section E) and I/O with d_k. Even at $d_k = 16$ ($2D$ total state), LoomFormer already outperforms DHC ($4D$) and MUDDFormer (lD, averaging $9D$ across $L = 18$ layers), showing that head-level rewiring uses state memory far more effectively than layer-wise mechanisms.

6. Conclusion

Viewing Transformer circuits through their dynamic wiring matrix shows that standard Transformers and layer-wise residual-stream mechanisms are restricted to low-rank gating, requiring many heads per reader to avoid circuit interference. Loom attention lifts this bottleneck by interweaving vertical linear attention over a per-token matrix memory with horizontal attention over tokens, making circuits rewirable per token at one head per reader and 2.0%–3.7% extra forward FLOPs. Across model scales and 300B-token pretraining, LoomFormer matches Transformers trained with over 3× compute and recent residual-stream baselines with ~1.7–2.8×, while keeping over three quarters of its attention heads on and rewiring its ports from token to token. Because the matrix memory operates along depth at each token, combining loom attention with linear or sparse attention along the sequence is a natural direction for future foundation models.

References

Biderman, S., Schoelkopf, H., Anthony, Q. G., Bradley, H., O'Brien, K., Hallahan, E., Khan, M. A., Purohit, S., Prashanth, U. S., Raff, E., et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning (ICML)*, pp. 2397–2430. PMLR, 2023.

Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.

Brandon, W., Mishra, M., Nrusimha, A., Panda, R., and Kelly, J. R. Reducing transformer key-value cache size with cross-layer attention. *arXiv preprint arXiv:2405.12981*, 2024.

Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H. W., Sutton, C., Gehrmann, S., et al. PaLM: Scaling language modeling with pathways. *Journal of Machine Learning Research*, 24(240):1–113, 2023.

Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.

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. *arXiv preprint arXiv:1803.05457*, 2018.

Conmy, A., Mavor-Parker, A., Lynch, A., Heimersheim, S., and Garriga-Alonso, A. Towards automated circuit discovery for mechanistic interpretability. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, pp. 16318–16352, 2023.

Dai, D., Deng, C., Zhao, C., Xu, R., Gao, H., Chen, D., Li, J., Zeng, W., Yu, X., Wu, Y., et al. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*, 2024.

Dao, T. and Gu, A. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. In *Proceedings of the Forty-First International Conference on Machine Learning (ICML)*, 2024.

D'yachkov, A. G. and Rykov, V. V. Bounds on the length of disjunctive codes. *Problemy Peredachi Informatsii*, 18 (3):7–13, 1982.

Elhage, N., Neel, N., Olsson, C., et al. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2020. URL https://transformer-circuits.pub/2021/framework/index.html.

Elhage, N., Hume, T., Olsson, C., Schiefer, N., Henighan, T., Kravec, S., Hatfield-Dodds, Z., Lasenby, R., Drain, D., Chen, C., et al. Toy models of superposition. *Transformer Circuits Thread*, 2022.

ElNokrashy, M., AlKhamissi, B., and Diab, M. Depth-wise attention (dwatt): A layer fusion method for data-efficient classification. *arXiv preprint arXiv:2209.15168*, 2022.

Erdős, P., Frankl, P., and Füredi, Z. Families of finite sets in which no set is covered by the union of r others. *Israel Journal of Mathematics*, 51:79–89, 1985.

Fang, Y., Cai, Y., Chen, J., Zhao, J., Tian, G., and Li, G. Cross-layer retrospective retrieving via layer attention. *arXiv preprint arXiv:2302.03985*, 2023.

Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple

and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022.

Füredi, Z. On r-cover-free families. *Journal of Combinatorial Theory, Series A*, 73(1):172–173, 1996.

Gao, L., Biderman, S., Black, S., Golding, L., Hoppe, T., Foster, C., Phang, J., He, H., Thite, A., Nabeshima, N., et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.

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. A framework for few-shot language model evaluation, 12 2023. URL https://zenodo.org/records/10256836.

Gu, A. and Dao, T. Mamba: Linear-time sequence modeling with selective state spaces. In *Proceedings of the First Conference on Language Modeling*, 2024.

Hanna, M., Pezzelle, S., and Belinkov, Y. Have faith in faithfulness: Going beyond circuit overlap when finding model mechanisms. In *Proceedings of Conference on Language Modeling (COLM)*, 2024.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 770–778, 2016.

Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L. A., Welbl, J., Clark, A., et al. An empirical analysis of compute-optimal large language model training. *Advances in Neural Information Processing Systems*, 35: 30016–30030, 2022.

Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pp. 4700–4708, 2017.

Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are RNNs: Fast autoregressive transformers with linear attention. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, 2020.

Kautz, W. H. and Singleton, R. C. Nonrandom binary superimposed codes. *IEEE Transactions on Information Theory*, 10(4):363–377, 1964.

Kimi Team. Attention residuals. *arXiv preprint arXiv:2603.15031*, 2026a.

Kimi Team. Kimi K3: Open frontier intelligence. *arXiv preprint arXiv:2607.24653*, 2026b.

Lai, G., Xie, Q., Liu, H., Yang, Y., and Hovy, E. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*, 2017.

Leviathan, Y., Kalman, M., and Matias, Y. Selective attention improves transformer. *arXiv preprint arXiv:2410.02703*, 2024.

Liu, A., Feng, B., Wang, B., Wang, B., Liu, B., Zhao, C., Dengr, C., Ruan, C., Dai, D., Guo, D., et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model. *arXiv preprint arXiv:2405.04434*, 2024.

Longpre, S., Hou, L., Vu, T., Webson, A., Chung, H. W., Tay, Y., Zhou, D., Le, Q. V., Zoph, B., Wei, J., et al. The flan collection: Designing data and methods for effective instruction tuning. *arXiv preprint arXiv:2301.13688*, 2023.

Mak, B. and Flanigan, J. Residual matrix transformers: Scaling the size of the residual stream. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, pp. 42712–42729, 2025.

Merullo, J., Eickhoff, C., and Pavlick, E. Talking heads: Understanding inter-layer communication in transformer language models. *arXiv preprint arXiv:2406.09519*, 2024.

Ni, R., Xiao, D., Meng, Q., Li, X., Zheng, S., and Liang, H. Benchmarking and understanding compositional relational reasoning of llms. In *Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI)*, 2025.

Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, A., et al. In-context learning and induction heads. *Transformer Circuits Thread*, 2022.

Pagliardini, M., Mohtashami, A., Fleuret, F., and Jaggi, M. Denseformer: Enhancing information flow in transformers via depth weighted averaging. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

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. *arXiv preprint arXiv:1606.06031*, 2016.

Peng, B., Goldstein, D., Anthony, Q., Albalak, A., Alcaide, E., Biderman, S., Cheah, E., Du, X., Ferdinan, T., Hou, H., et al. Eagle and finch: Rwkv with matrix-valued states and dynamic recurrence. *arXiv preprint arXiv:2404.05892*, 2024.

Ruszinkó, M. On the upper bound of the size of the r-cover-free families. *Journal of Combinatorial Theory, Series A*, 66(2):302–310, 1994.

Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.

Schlag, I., Irie, K., and Schmidhuber, J. Linear transformers are secretly fast weight programmers. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, 2021.

Shazeer, N. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.

Sperner, E. Ein satz über untermengen einer endlichen menge. *Mathematische Zeitschrift*, 27:544–548, 1928.

Srivastava, R. K., Greff, K., and Schmidhuber, J. Training very deep networks. *Advances in neural information processing systems*, 28, 2015.

Su, J., Ahmed, M., Lu, Y., Pan, S., Bo, W., and Liu, Y. Roformer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024.

Tay, Y., Dehghani, M., Aribandi, V., Gupta, J., Pham, P. M., Qin, Z., Bahri, D., Juan, D.-C., and Metzler, D. Omninet: Omnidirectional representations from transformers. In *International Conference on Machine Learning (ICML)*, pp. 10193–10202. PMLR, 2021.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

Vig, J. and Belinkov, Y. Analyzing the structure of attention in a transformer language model. *arXiv preprint arXiv:1906.04284*, 2019.

Wang, K., Variengien, A., Conmy, A., Shlegeris, B., and Steinhardt, J. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR)*, 2023.

Wang, K., Xia, X., Liu, J., Yi, Z., and He, T. Strengthening layer interaction via dynamic layer attention. *arXiv preprint arXiv:2406.13392*, 2024.

Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., and Chao, L. S. Learning deep transformer models for machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2019.

Welbl, J., Liu, N. F., and Gardner, M. Crowdsourcing multiple choice science questions. *arXiv preprint arXiv:1707.06209*, 2017.

Xiao, D., Meng, Q., Li, S., and Yuan, X. Improving transformers with dynamically composable multi-head attention. 2024a.

Xiao, D., Meng, Q., Li, S., and Yuan, X. MUDDFormer: Breaking residual bottlenecks in transformers via multiway dynamic dense connections. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, 2025.

Xiao, G., Tian, Y., Chen, B., Han, S., and Lewis, M. Efficient streaming language models with attention sinks. In *The Twelfth International Conference on Learning Representations (ICLR)*, 2024b.

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.

Yang, S., Wang, B., Shen, Y., Panda, R., and Kim, Y. Gated linear attention transformers with hardware-efficient training. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pp. 56501–56523, 2024a.

Yang, S., Wang, B., Zhang, Y., Shen, Y., and Kim, Y. Parallelizing linear transformers with the delta rule over sequence length. In *Proceedings of the Thirty-Eighth Annual Conference on Neural Information Processing Systems*, 2024b.

Yang, S., Kautz, J., and Hatamizadeh, A. Gated delta networks: Improving Mamba2 with delta rule. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

Ye, T., Dong, L., Xia, Y., Sun, Y., Zhu, Y., Huang, G., and Wei, F. Differential transformer. 2024.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

Zhu, D., Huang, H., Huang, Z., Zeng, Y., Mao, Y., Wu, B., Min, Q., and Zhou, X. Hyper-connections. In *Proceedings of the Thirteenth International Conference on Learning Representations (ICLR)*, 2025.

A. Theory of Rewirable Circuits

A.1. Normalization and proof of Proposition 2.3

Lemma A.1. *In a Transformer with pre-LN (LayerNorm or RMSNorm), the contribution of earlier writers to reader R_j's input $u_j = \mathrm{LN}(x)W_j^s$ is $\lambda(x)\sum_i o_i \bar{C}_{ij}$, where the scale $\lambda(x) > 0$ is shared by all writers and the virtual weights $\bar{C}_{ij}$ are fixed.*

Proof. Write $x = \sum_i o_i W_i^O + x_{\text{other}}$, where x_{other} sums the embedding and earlier FFN outputs. LayerNorm computes $\mathrm{LN}(x) = \lambda(x)\,(x - \mu(x)\mathbf{1}) \odot \gamma + \beta$ with $\lambda(x) = \sqrt{D}/\|x - \mu(x)\mathbf{1}\| > 0$ (putting $\lambda(x) = 0$ on the degenerate set $x \in \mathrm{span}(\mathbf{1})$). Since centering is the fixed linear map $x - \mu(x)\mathbf{1} = x\big(I - \frac{1}{D}\mathbf{1}\mathbf{1}^\top\big)$,

$$u_j = \mathrm{LN}(x)W_j^s = \lambda(x)\sum_i o_i \bar{C}_{ij} + \lambda(x)\, x_{\text{other}}\big(I - \tfrac{1}{D}\mathbf{1}\mathbf{1}^\top\big)\,\mathrm{diag}(\gamma)W_j^s + \beta W_j^s,$$

with fixed virtual weights $\bar{C}_{ij} = W_i^O\big(I - \frac{1}{D}\mathbf{1}\mathbf{1}^\top\big)\,\mathrm{diag}(\gamma)W_j^s$ absorbing the linear maps of the normalization (any projection bias adds to βW_j^s). For RMSNorm, centering is omitted, $\beta = 0$ and $\lambda(x) = \sqrt{D}/\|x\|$. □

Proof of Proposition 2.3. By Theorem A.1, pre-LN rescales all writers of a token by the same $\lambda(x) > 0$. Within a token, what depends on the input is only this scale and which heads are on: let $g_i = \lambda(x)$ if the head of W_i is on and 0 otherwise, and $g'_j = 1$ if the head of R_j is on and 0 otherwise. Then $C = g\,g'^\top$, which is Eq. (2), and since $\lambda(x) > 0$, W_i and R_j are connected ($g_i g'_j \bar{C}_{ij} \neq 0$) if and only if both are on and $\bar{C}_{ij} \neq 0$. □

A.2. Full statement and proof of Theorem 2.4

Setting. The role of reader R_j ($j \in [r]$, $r \geq 2$) is played by a set H_j of heads, so a reader is here a role rather than one input of one head; by Eq. (2), head $h \in H_j$ *receives* W_i when its effective wiring $g_i g'_h \bar{C}_{ih}$ is nonzero. A wiring task assigns to each reader a writer $\pi(j) \in [w]$, and the tasks are all maps $\pi : [r] \to [w]$. Reader R_j *receives exactly* $\mathsf{W}_{\pi(j)}$ if some $h \in H_j$ receives $\mathsf{W}_{\pi(j)}$ and no $h \in H_j$ receives any W_i with $i \neq \pi(j)$. We require every head in H_j that is on to be free of other writers because query and key inputs feed a softmax, so interference at one head cannot be cancelled downstream. Every writer in $\pi([r])$ must therefore be on ($g_{\pi(j)} \neq 0$), and turning others on can only add interference, so we may assume the writers that are on are $\pi([r])$. For a fixed reader R_j, let $B_i = \{h \in H_j : \bar{C}_{ih} \neq 0\}$. A family $\{B_i\}_{i=1}^w$ is t-cover-free if $B_i \not\subseteq \bigcup_{i' \in J} B_{i'}$ for every $i \in [w]$ and every $J \subseteq [w] \setminus \{i\}$ with $|J| \leq t$.

Theorem A.2 (full version of Theorem 2.4)**.** *For $r \geq 2$, a node-gated architecture with w writers realizes all wiring tasks $\pi : [r] \to [w]$ if and only if, for every reader, $\{B_i\}_{i=1}^w$ is an $(r-1)$-cover-free family. Hence each reader needs at least $N(w, r-1)$ heads, and this number suffices.*

Proof. If. Given π, turn on only the writers $\pi([r])$. For reader R_j, let $J = \pi([r]) \setminus \{\pi(j)\}$, so $|J| \leq r - 1$. Since $\{B_i\}$ is $(r-1)$-cover-free, there is a head $h \in B_{\pi(j)} \setminus \bigcup_{i \in J} B_i$; turn on h and no other head of H_j. Then h receives $\mathsf{W}_{\pi(j)}$ and none of the other writers that are on.

Only if. Fix a reader R_j, a writer i and a set $J \subseteq [w] \setminus \{i\}$ with $|J| \leq r - 1$. Choose π with $\pi(j) = i$ that maps the other $r - 1$ readers onto J, which is possible since $r - 1 \geq 1$, $|J| \leq r - 1$ and different readers may share a writer (if $J = \emptyset$, map them to i). Realizing π requires a head $h \in H_j$ that is on with $\bar{C}_{ih} \neq 0$ and $\bar{C}_{i'h} = 0$ for all $i' \in J$, i.e., $h \in B_i \setminus \bigcup_{i' \in J} B_{i'}$. Hence $B_i \not\subseteq \bigcup_{i' \in J} B_{i'}$, and the family is $(r-1)$-cover-free.

Counting. The sets B_i are subsets of H_j, so $|H_j| \geq N(w, r-1)$. Conversely, given an $(r-1)$-cover-free family of w sets on a ground set of size $N(w, r-1)$ for each reader, create one head per point and set $\bar{C}_{ih} = I_d$ if $h \in B_i$ and 0 otherwise; this is realized in $D = \min\{w, rN(w, r-1)\}\,d$ dimensions by assigning orthogonal d-dimensional subspaces either to the w writers or to the $rN(w, r-1)$ heads. □

Remarks. (i) Attention sinks gate heads off only approximately; under the threshold of Section 2, the counting argument is unchanged. (ii) Sufficiency at $N(w, r-1)$ heads per reader holds in a Transformer when each of the r readers belongs to a separate head (with its other two inputs fixed); when $r = 3n$ counts all Q, K and V inputs of n heads, coupling Q, K and V inside the same physical head (and sink gate g'_h) can only require more heads, while the lower bound $N(w, r-1)$ still applies. Likewise, if several heads of a value reader share identical attention weights, their outputs add linearly and

interference could in principle cancel, whereas our clean-head requirement is necessary for query and key inputs. (iii) The model excludes routing through FFNs, discussed in Section A.5.

A.3. Bounds on $N(w,t)$ and remarks on Theorem 2.4

Bounds. A 1-cover-free family is an antichain, so by Sperner's theorem (Sperner, 1928) $N(w,1)$ is the smallest N with $\binom{N}{\lfloor N/2 \rfloor} \geq w$, about $\log_2 w$. For $2 \leq t \leq \sqrt{w}$, $N(w,t) = \Omega(t^2 \log w / \log t)$ (D'yachkov & Rykov, 1982; Ruszinkó, 1994; Füredi, 1996) and $N(w,t) = \mathcal{O}(\min\{t^2 \log w, w\})$ (Erdős et al., 1985); for $t > \sqrt{w}$, Bassalygo's bound $N(w,t) \geq \min\{\binom{t+2}{2}, w\}$ (D'yachkov & Rykov, 1982) together with the singleton construction $N(w,t) \leq w$ gives $N(w,t) = \Theta(w)$, reaching $N(w,t) = w$ (one head per writer) once $\binom{t+2}{2} \geq w$.

Group testing. Theorem A.2 is an instance of non-adaptive group testing: the heads of a reader are tests that must isolate the selected writer from up to $r-1$ other writers that are on.

Dense and sparse layouts. With a single head per reader ($|H_j| = 1$), the two layouts of Figure 3 are degenerate codes: $B_1 = \cdots = B_w = H_j$ (dense), which is not 1-cover-free and makes circuits interfere, and $B_i = \emptyset$ for some writer i (sparse), which leaves (i,j) unreachable.

Superposition. At layer scale ($r = 3n$, $N(w, r-1) = \Theta(w)$), separating w writers as above takes $\Theta(wd)$ dimensions (wd with orthogonal writer subspaces). With nL writers sharing $D = nd$ dimensions, writers across layers must overlap in superposition (Elhage et al., 2022), so virtual weights are rarely exactly zero and in practice part of the price is paid as interference.

A.4. Cost of naive loom attention

Indexing layers by $l = 0, \ldots, L-1$ as in Section 3.2, every reader of layer l attends over nl earlier head writers. Over the $3n$ readers of all L layers, the projections W^α of Eq. (3) take $\sum_{l=0}^{L-1} 3n \cdot nl \cdot D = \frac{3}{2} n^2 L(L-1)D \approx \frac{3}{2} n^2 L^2 D$ parameters, and as many multiply-adds per token. This grows quadratically in the total number of heads nL and rivals the $12LD^2$ of the whole model once $n^2 L \approx 8D$ (i.e., $nL \approx 8d$), e.g., 70% of it for our XL configuration ($L = 28$, $n = 20$, $D = 1920$). Factorizing the weights into keys and queries removes W^α, but still leaves $\approx \frac{3}{2} n^2 L^2$ attention weights per token to compute and store for backpropagation, plus $\approx \frac{3}{2} nL^2 D$ multiply-adds to aggregate the o_i's.

A.5. Routing through FFNs

An intermediate FFN between the writers and the reader layer can route any pattern $P \subseteq [w] \times [r]$ by copying o_i into a d-dimensional subspace private to R_j whenever $(i,j) \in P$. Concretely, suppose the residual stream $x \in \mathbb{R}^D$ carries binary indicators $p_{ij} \in \{0,1\}$ of P ($p_{ij} = 1$ iff $(i,j) \in P$) and the writer outputs satisfy $\|o_i\|_\infty \leq \Lambda$ (ignoring pre-LN rescaling). Since $\mathrm{ReLU}\big(z - \Lambda(1-p)\big) - \mathrm{ReLU}\big(-z - \Lambda(1-p)\big)$ equals z when $p = 1$ and 0 when $p = 0$ for every $z \in [-\Lambda, \Lambda]$, a two-layer ReLU FFN can compute the gated copy $p_{ij} o_i \in \mathbb{R}^d$ coordinate-wise for each pair (i,j) and sum over i into R_j's private subspace. However, this construction uses $2wrd$ hidden units—hence $4wrdD = \mathcal{O}(wrdD)$ parameters, or $12nlD^2$ when routing $w = nl$ writers to $r = 3n$ readers (n times the $12lD^2$ parameters of all l preceding blocks)—requires an intervening FFN step, and needs disjoint subspaces for the w writers, r readers and wr control coordinates ($D \geq (w+r)d + wr$). We give this construction only as an upper bound and do not prove a lower bound on FFN routing.

A.6. Proof of Proposition 3.1

Part (i). Flatten $M \in \mathbb{R}^{d_k \times d}$ row by row into $\mathrm{vec}(M) \in \mathbb{R}^{d_k d}$, so that entry (a,b) of M is entry $(a-1)d + b$ of $\mathrm{vec}(M)$. For a static read port $q \in \mathbb{R}^{d_k}$, $(qM)_b = \sum_a q_a M_{ab} = \big(\mathrm{vec}(M)\,(q^\top \otimes I_d)\big)_b$. For a static write port $k \in \mathbb{R}^{d_k}$, $\mathrm{vec}(k^\top o)_{(a-1)d+b} = k_a o_b = \big(o\,(k \otimes I_d)\big)_{(a-1)d+b}$. Hence every read and write of M is a fixed linear map between $\mathrm{vec}(M)$ and the head, i.e., a Transformer projection with Kronecker structure, $q^\top \otimes I_d$ for reads and $k \otimes I_d$ for writes, and the composition weight between W_i and R_j is the constant $\langle q_j^m, k_i^m \rangle$. The coupling $KQ^\top$ is therefore fixed and belongs to the virtual weights, $\bar{C}_{ij} = \langle q_j^m, k_i^m \rangle I_d$; only gating heads through attention sinks (together with normalization scales) depends on the input, as in Eq. (2), so $\mathrm{Tfm}_{K,Q}$ is node-gated. In the full architecture (Section 3.2), MNorm is

Table 5. Ports of residual-stream mechanisms in the common form of Eq. (12).

	Write port k_i^m	Read port q_j^m	Update Φ_l
Transformer	1	1	1
DenseFormer	e_{l_i}	static	I
HC / mHC	$\beta_{l_i}(M)$	$\alpha_{l_j}(M)$	$A_l(M)$
AttnRes (full)	$\mathrm{Norm}(v_{l_i})$	static, softmax	I
MUDDFormer	e_{l_i}	per Q/K/V, from x	I
LoomFormer	per head, from x	per head and Q/K/V, from x	I

RMSNorm over $\mathrm{vec}(M)$ and rescales all matrix-stream writers by one shared factor (Theorem A.1), while the bypass o_i^m and $(\mathrm{FFN}_x, \mathrm{FFN}_M)$ become a fixed linear skip and a standard FFN on $(x, \mathrm{vec}(M))$. The per-head $\mathrm{Norm}(o_i)$ on LoomMHA writes (Eq. 8) applies a fixed diagonal gain (absorbed into the write projection) and rescales writer W_i by $\sqrt{d}/\|o_i\|$, and even if the scalar port gates $\sigma(xW_g^W), \sigma(xW_{\bar{g}}^R)$ (Eqs. 5–6) are retained on static port directions, they multiply whole rows or columns of C, leaving $C = g\, g'^{\top}$ of rank at most one (cf. Section G). With ports generated per token, the computation at each token and layer coincides with that of $\mathrm{Tfm}_{K(x),Q(x)}$. The coincidence holds within a token; across tokens, attention mixes values computed by different members, so the model as a whole is not any single $\mathrm{Tfm}_{K,Q}$.

Part (ii). Let S_W be the writers in a wiring pattern $P \subseteq [w] \times [r]$, with $|S_W| \le d_k$. Turn on only S_W, pick an injection $\iota : S_W \to [d_k]$, give each $i \in S_W$ the write key $k_i^m = e_{\iota(i)}$ ($k_i^m = 0$ for $i \notin S_W$), and give reader R_j the read query $q_j^m = \sum_{i \in S_W :\, (i,j) \in P} e_{\iota(i)}$. Then $\langle q_j^m, k_i^m \rangle = 1$ if $(i,j) \in P$ and 0 otherwise for all $i \in [w]$, so $c_{ij} \neq 0$ if and only if $(i,j) \in P$; since the virtual weights of the matrix stream are identities, the connected pairs are P, with one head per reader and d_k independent of w. The RMSNorm of the actual ports rescales each key and each query by a constant, and the sigmoid gates set their magnitudes (turning off writers outside S_W), so the pattern is realized up to per-reader scales. In the full model, computing ι and P from the vector stream is the role of the port generators on the control path.

A.7. Wiring matrices of residual-stream mechanisms

This section derives the entries of Table 1. As in Section 2, we consider the wiring within a token from the earlier head writers ($o_i \in \mathbb{R}^d$) to the query, key and value readers (u_j) of one later layer. Other reads of the history—the FFN inputs of HC/mHC and AttnRes, the residual input of MUDDFormer and the bypass o^m of LoomFormer—feed no head input and are not readers. We mark a port by (x) if it is computed from the vector (residual) stream and by (M) if it is computed from the memory itself. Per-token scales such as pre-LN are absorbed into the wiring matrix as in Theorem A.1, and the node-gating factors g, g' of Theorem 2.2 are common to all mechanisms, so we omit them from the tables. For LoomFormer we consider the matrix stream only and ignore MShift, which mixes neighbouring tokens (the $d/4$-dimensional RoPE slice from x belongs to the node-gated vector stream of Theorem 2.3). Ports are row vectors, as in Section 3.1.

A common form. Each mechanism keeps at a token a memory M of d_k slots of width d_v. Writer W_i adds $k_i^m \otimes (o_i W_i^O)$ to M, reader R_j receives $u_j = (q_j^m M)\, W_j^s$, and between layers M may be updated linearly, $M \leftarrow \Phi_l M$; here k_i^m and q_j^m are the write and read ports, and $W_i^O = W_j^s = I_d$ in the matrix stream of LoomFormer. Composing these steps gives

$$u_j = \sum_i \left(q_j^m \, \Phi_{l_i l_j} \, (k_i^m)^{\top} \right) o_i W_i^O W_j^s, \tag{12}$$

where $\Phi_{l_i l_j} = \Phi_{l_j - 1} \cdots \Phi_{l_i}$ is the product of the updates between the layer l_i of W_i and the layer l_j of R_j: the wiring matrix has entries $c_{ij} = q_j^m \Phi_{l_i l_j} (k_i^m)^{\top}$ and the virtual weights are $\bar{C}_{ij} = W_i^O W_j^s$. This representation is not unique: a change of slot basis by an invertible G maps $k_i^m \mapsto k_i^m G^{\top}$, $q_j^m \mapsto q_j^m G^{-1}$ and $\Phi_l \mapsto G \Phi_l G^{-1}$ without changing c_{ij}. For mechanisms that keep the outputs of all earlier layers, we use the basis in which every stored output has its own slot, the write port is e_{l_i} and $\Phi_l = I$; residual copies then become part of the read port. Transitions $\Phi_{l_i l_j}$ can also be absorbed into the ports without adding slots: for a fixed reader layer l_j, setting $\tilde{k}_i^m = k_i^m \Phi_{l_i l_j}^{\top}$ writes $c_{ij} = \langle q_j^m, \tilde{k}_i^m \rangle$ even when Φ_l is singular (and when every Φ_l is invertible, setting $\Phi_l^{\times} = \Phi_{l-1} \cdots \Phi_1$ makes this split independent of l_j via $\tilde{q}_j^m = q_j^m \Phi_{l_j}^{\times}$ and $\tilde{k}_i^m = k_i^m \Phi_{l_i}^{\times -\top}$). Table 5 lists the parameterization of each mechanism, justified below; Table 1 omits factors equal to 1 or I.

Transformer. One slot of width D with $k_i^m = q_j^m = \Phi = 1$ gives $c_{ij} = 1$ and the virtual weights $\bar{C}_{ij} = W_i^O W_j^s$; with the node gates, this is Eq. (2).

DenseFormer. The input of each block is a static weighted sum of the embedding and the outputs of all earlier blocks, and the block adds its own attention and FFN outputs to this input through its internal residual (Pagliardini et al., 2024). Unfolding these residual copies, the input of block l' is a static combination of the outputs of all earlier blocks, so $c_{ij}\bar{C}_{ij} = \langle q_{l_j}^m, k_{l_i}^m \rangle W_i^O W_j^s$ with a static read port, so only node gating makes the wiring matrix input-dependent. DenseFormer keeps one slot per block, about l slots at layer l (size lD).

HC and mHC. HC keeps d_k parallel streams of width D (we use $d_k = 4$, the main configuration of Zhu et al. (2025)). Each sub-layer reads αM and updates the streams as $M \leftarrow AM + \beta^\top h$, with α, β and A all computed from M in the dynamic version (Zhu et al., 2025), and A projected onto the Birkhoff polytope of doubly stochastic matrices in mHC (Xie et al., 2025). Heads read and write through the attention sub-layer, with ports α_l and β_l, while the FFN and subsequent attention stream-mixing matrices between attention sub-layers l_i and l_j compose into $\Phi_{l_i l_j} = A_{l_i l_j}$. In the common form, $c_{ij} = \alpha_{l_j} A_{l_i l_j} \beta_{l_i}^\top$—the unrolled weight of Eq. (27) of Zhu et al. (2025), which Kimi Team (2026a) interpret as depth-wise linear attention with queries α and keys β. Since all factors are computed from the streams M, absorbing $A_{l_i l_j}$ into the write port as above (without requiring A_l to be invertible) gives the entry $\langle q_{l_j}^m(M), k_{l_i}^m(M) \rangle$ of Table 1, with d_k-dimensional ports and state size $d_k D = 4D$.

AttnRes. Full AttnRes keeps the embedding and the outputs v_m of all earlier sub-layers, without a residual connection, and gives sub-layer m' the input $\sum_m \mathrm{softmax}_m \langle w_{m'}, \mathrm{Norm}(v_m) \rangle v_m$ with a learned query $w_{m'}$ (Kimi Team, 2026a). In the slot basis of Eq. (12), writers of layer l_i write to their attention sub-layer's slot (e_{l_i}, while FFN and embedding slots hold no head writer) and heads of layer l_j read linearly with the softmax weight vector over all stored slots; Tables 1 and 5 instead display the underlying softmax query $q_{l_j}^m = w_{l_j}$ (static) and key $k_{l_i}^m(M) = \mathrm{Norm}(v_{l_i})$ (from M), normalized by $\mathrm{softmax}$ across all stored sub-layer slots. Full AttnRes keeps one slot per sub-layer, about $2l$ slots at layer l (size $2lD$); its block variant sums the outputs within each block into one slot.

MUDDFormer. MUDDFormer keeps the embedding and the outputs of all earlier blocks, and gives the next block four inputs—for its queries, keys, values and residual—each a weighted sum of these outputs with weights generated from the latest hidden state (Xiao et al., 2025). The residual input carries dynamic combinations of earlier outputs into the next block output; unfolding these recursions expresses the effective read port $q_{l_j}^{m,s}(x)$ of input $s \in \{Q, K, V\}$ as a function of the hidden states along the vector stream x, giving $c_{ij}\bar{C}_{ij} = \langle q_{l_j}^{m,s}(x), k_{l_i}^m \rangle W_i^O W_j^s$. Like DenseFormer, MUDDFormer keeps about l slots at layer l (size lD).

LoomFormer. The matrix stream keeps d_k slots of width d that accumulate writes ($\Phi_l = I$) and passes content without fixed projections: absorbing the per-token RMS scales of $\mathrm{Norm}(o_i)$ and MNorm into k_i^m and $q_j^{m,s}$ (and the learned diagonal gain of $\mathrm{Norm}(o_i)$ into the writer output o_i) leaves $\bar{C}_{ij} = I_d$. Indexing a reader by its head $j \in [n]$ and input $s \in \{Q, K, V\}$ as in Table 1 gives $c_{ij} = \langle q_j^{m,s}(x), k_i^m(x) \rangle$, with both ports generated per head and per token from the vector stream (Theorem 3.1(ii)). With $d_k = 2n$, M has size $d_k d = 2D$, plus D for the vector stream.

Wiring ranks. Node gating multiplies C entrywise by $g\, g'^\top$—equivalently $C \mapsto \mathrm{diag}(g)\, C\, \mathrm{diag}(g')$—which cannot increase its rank. All readers of a layer share one read port in DenseFormer, HC/mHC and AttnRes, so C has a single distinct column and rank at most 1; MUDDFormer has one read port per Q/K/V input, hence at most three distinct columns and rank at most 3 (reached once at least three earlier layers are present with linearly independent Q/K/V read ports). In LoomFormer every reader has its own query, so the rank is that of $KQ^\top$ restricted to the $3n$ readers of the layer, at most $\min(w, d_k, 3n)$—hence $\min(d_k, 3n)$ once $w \geq \min(d_k, 3n)$, reached whenever the writer keys and reader queries span a common $\min(d_k, 3n)$-dimensional subspace.

Linear attention. The same form describes linear attention along the sequence instead of depth: the writers and readers are one head at different tokens, M is the head's private state, the ports are its keys and queries, and the update Φ distinguishes the variants—I for linear attention (Katharopoulos et al., 2020), a scalar decay in Mamba2 (Dao & Gu, 2024), a gated diagonal in gated linear attention (Yang et al., 2024a), $I - \beta k^\top k$ in DeltaNet (Yang et al., 2024b), and $\gamma(I - \beta k^\top k)$ in Gated DeltaNet (Yang et al., 2025). It differs from LoomFormer in the axis and the sharing of the state (Section 4).

B. Extended Related Work

Cross-layer connections and residual-stream architectures. Beyond standard residual connections (He et al., 2016), early work strengthened depth-wise signal propagation through input-dependent gating in Highway Networks (Srivastava et al., 2015) and all-to-all layer aggregation with static or learnable weights in DenseNet (Huang et al., 2017), DLCL (Wang et al., 2019) and DenseFormer (Pagliardini et al., 2024). Subsequent architectures introduced cross-layer attention over earlier layer representations (ElNokrashy et al., 2022; Fang et al., 2023; Wang et al., 2024; Kimi Team, 2026a), multi-stream residual states with dynamic or manifold-constrained mixing in HC and mHC (Zhu et al., 2025; Xie et al., 2025), and decoupled Q/K/V/residual layer aggregation in MUDDFormer (Xiao et al., 2025). As analyzed in Section 4 and Section A.7, all of these mechanisms operate at the granularity of whole layers (or layer $\times$ Q/K/V), summing the outputs of all heads within a layer before routing and restricting the within-token wiring rank to 1 (or 3).

The Residual Matrix Transformer (RMT) (Mak & Flanigan, 2025) replaces the vector residual stream entirely with a pure matrix residual stream, which heads and FFNs read and write through static key vectors so the stream can be widened cheaply. LoomFormer instead keeps a dual-stream architecture that separates control on the vector stream from content on the matrix stream: the compact vector stream generates per-token ports cheaply, whereas RMT's static ports fall under Theorem 3.1(i) and remain node-gated however wide the matrix is made ($4D$ in RMT's main configuration vs. $2D + D$ in LoomFormer).

Cross-layer token attention and KV sharing. OmniNet (Tay et al., 2021) lets every token attend directly to all tokens across all layers, incurring attention cost quadratic in LT; loom attention instead factors cross-layer routing within each token through a fixed-size matrix memory M in $\mathcal{O}(LTDd_k)$ time and composes it with standard within-layer horizontal attention. Cross-Layer Attention (CLA) (Brandon et al., 2024) shares key-value caches across adjacent layers to reduce inference memory at a modest quality trade-off, whereas loom attention dynamically composes earlier head outputs into distinct per-head queries, keys and values at every layer (Section 5.3).

Mechanistic interpretability and intra-layer advances. Circuit-level interpretability studies (Elhage et al., 2020; Olsson et al., 2022; Conmy et al., 2023; Hanna et al., 2024; Merullo et al., 2024; Ni et al., 2025) show that Transformers implement tasks by composing specific upstream and downstream heads across layers; loom attention turns this implicit, static composition into an explicit, per-token operation over M. Orthogonal intra-layer advances—including improved horizontal attention (Liu et al., 2024; Xiao et al., 2024a; Ye et al., 2024; Leviathan et al., 2024), sequence-axis linear attention and state-space models (Gu & Dao, 2024; Dao & Gu, 2024; Peng et al., 2024; Yang et al., 2024b; 2025), and sparse Mixture-of-Experts FFNs (Fedus et al., 2022; Dai et al., 2024)—modify horizontal mixing or FFN capacity within a layer and can be combined directly with loom attention along depth.

C. PyTorch-Style Pseudocode for LoomFormer

```
# b = batch size, t = seq len, n = num_heads, d = head_dim
# k = num_slots d_k, r = RoPE dim d // 4. M is of shape btkd

def MRead(M, x, xr):
    qm = rearrange(xr @ W_r2, 'bt(nk)->btnk')
    qm = rmsnorm(qm) * sigmoid(rearrange(x @ W_rg, 'btn->btn1'))
    qm += q_bar * sigmoid(rearrange(x @ W_rg_bar, 'btn->btn1')) # nk,btn1->btnk
    u = einsum('btnk,btkd->btnd', qm, M)
    return u

def MWrite(o, x, xw):
    km = rearrange(xw @ W_w2, 'bt(nk)->btnk')
    km = rmsnorm(km) * sigmoid(rearrange(x @ W_wg, 'btn->btn1'))
    dM = einsum('btnk,btnd->btkd', km, o)
    return dM

def MShift(M, x):
    M_prev = shift_right(M, 1, dim=1)  # shift t dim of btkd
```

```
    z = sigmoid(rearrange(x @ W_sg, 'btk->btk1'))
    return z * M_prev + (1 - z) * M

def loom_mha(x, M, write_M=True):
    M = RMSNorm(MShift(M, x), dim=(-2, -1))
    xr, xw = gelu(x @ W_r1), gelu(x @ W_w1)
    q, k, v, o = [MRead[s](M, x, xr) for s in ['q', 'k', 'v', 'o']]
    qr = rope(rearrange(x @ W_q, 'bt(nr)->btnr'))
    kr = rope(rearrange(x @ W_k, 'bt(nr)->btnr'))
    q, k = concat([qr, q], dim=-1), concat([kr, k], dim=-1)

    allow_mask = tril(ones(T, T), diagonal=-1) # strictly causal mask
    o = o + MHA(q, k, v, allow_mask=allow_mask)
    dx = rearrange(o, 'btnd->bt(nd)') @ W_o
    dM = MWrite(RMSNorm(o, dim=-1), x, xw) if write_M else 0.
    return dx, dM

def ffn(x, write_M=True):
    xw = gelu(x @ W_w1)
    dx = _ffn_x(x)
    dM = MWrite(rearrange(_ffn_m(x), 'bt(nd)->btnd'), x, xw) if write_M else 0.
    return dx, dM

def loomformer(input_ids):
    x = embedding(input_ids)
    o = rearrange(x @ W_e, 'bt(nd)->btnd')
    M = MWrite(o, x, gelu(x @ W_w1))

    for l in range(L):
        write_M = l < L - 1
        # pairs are added element-wise: (x, M) + (dx, dM) = (x + dx, M + dM)
        (x, M) = (x, M) + loom_mha[l](ln_1[l](x), M, write_M=write_M)
        (x, M) = (x, M) + ffn[l](ln_2[l](x), write_M=write_M)
    return unembedding(ln_f(x))
```

D. Experimental Configurations, Baselines and Benchmarks

D.1. Model Configurations and Hyperparameters

Table 6 lists the architectural dimensions, peak learning rates and token budgets across all four LoomFormer scales (Sections 5.1 and 5.2). In blocks $0, \ldots, L-2$, FFN_x and the FFN_M write branch (FFN_M plus its MWrite module) follow a 2:1 parameter ratio (rounded to multiples of 128 at 2.8B, giving 2.2:1), while block $L-1$ (last) omits unused writes to M_L and widens FFN_x to match the baseline Transformer's total parameter count. Table 7 lists the training hyperparameters for the scaling-law experiments (Section 5.1; weight decay is excluded on normalization scales and biases, including gate biases), while the 2.8B models (Section 5.2) follow Pythia-2.8B (Biderman et al., 2023, Appendix E). To improve training stability at the XL scale, LoomFormer-XL replaces SwiGLU in FFN_x with Kimi K3's bounded SiTU-GLU activation (Kimi Team, 2026b) with $\beta_1 = 4$, $\beta_2 = 7$.

Port and gate initialization. Each gate projection in Eqs. (5)–(7) includes a learnable bias (omitted in Section 3.2 for brevity) that sets its initial opening. Specifically, the gate projection weights are initialized to zero, with biases setting initial gate values to 0.05 on the dynamic read branch, 0.9 on the static read branch, 0.1 on the write branch, and 0.5 on MShift ($z_{[t]}$). For the static queries $\bar{q} \in \mathbb{R}^{d_k}$, value reads v are initialized from $\mathcal{N}(0, 1/d_k)$, whereas $\tilde{q}, \tilde{k}$ and o^m start at $\bar{q} = 0$ so memory-read Q/K coordinates and the output bypass grow in gradually.

Table 6. Model configurations across scales. $d_k = 2n$; shared port-feature width $D_r = nd_k/2 = n^2$, capped at $D/4 = 640$ at 2.8B.

Model	Params	L	D	D/L	n	d	D_r	Tfm FFN	FFN_x (last)	FFN_M	LR	Tokens
Medium	432M	18	1200	66.7	16	75	256	3200	2510 (3904)	1128	3.0e-4	7B
Large	778M	22	1536	69.8	16	96	256	4096	3317 (5083)	1539	2.5e-4	14B
XL	1.4B	28	1920	68.6	20	96	400	5120	4048 (6277)	1827	2.0e-4	26B
2.8B	2.8B	32	2560	80	32	80	640	10240	7808 (9984)	2944	1.6e-4	300B

Table 7. Training hyperparameters for the scaling-law experiments (Section 5.1).

Hyperparameter	Value	Hyperparameter	Value
Optimizer	AdamW	Sequence length	4096
(β_1, β_2)	$(0.9, 0.95)$	Batch size (seqs)	128
ϵ	10^{-8}	Vocab size	50,432
Weight decay	0.1	Tied embeddings	False
Gradient clipping	1.0	Linear bias	False
LR schedule	Cosine	Normalization	RMSNorm ($\epsilon = 10^{-6}$)
Min LR ratio	0.1	FFN activation	SwiGLU
Warmup fraction	0.01	Positional emb.	RoPE ($\theta = 10{,}000$)
Dropout	0.0	Precision	BF16

D.2. Scaling-Law Baselines

All baselines in Section 5.1 are trained in JAX/MaxText on the identical data order, tokenizer and schedule as LoomFormer:

- **Transformer++**: Llama-style pre-norm Transformer (Touvron et al., 2023) with RoPE (Su et al., 2024), SwiGLU FFN (Shazeer, 2020), RMSNorm and no linear biases.
- **DHC**: Dynamic Hyper-Connections with expansion rate $d_k = 4$ and $\tanh$-bounded dynamic routing following Zhu et al. (2025).
- **mHC**: Manifold-constrained Hyper-Connections ($d_k = 4$ streams) projecting the stream-mixing matrix onto the Birkhoff polytope via Sinkhorn normalization (Xie et al., 2025).
- **AttnRes (full)**: Full Attention Residuals aggregating all preceding sub-layer outputs per sub-layer via learned depth-wise softmax queries (Kimi Team, 2026a).
- **MUDDFormer**: Multiway Dynamic Dense connections generating per-token depth-aggregation weights separately for the query, key, value and residual inputs of each block (Xiao et al., 2025).

D.3. Downstream Evaluation Benchmarks

In addition to the benchmarks used by Pythia (Biderman et al., 2023)—LAMBADA (Paperno et al., 2016), PIQA (Bisk et al., 2020), WinoGrande (Sakaguchi et al., 2021), ARC-Easy/Challenge (Clark et al., 2018) and SciQ (Welbl et al., 2017)—we evaluate BoolQ (Clark et al., 2019) and HellaSwag (Zellers et al., 2019) for commonsense reasoning, and RACE-Middle/High (Lai et al., 2017) for reading comprehension (Table 2), all under LM Evaluation Harness (Gao et al., 2023) in zero-shot and five-shot settings.

E. Details of Complexity Analysis

We derive the forward-FLOP overhead in Table 3 under a matched parameter budget, following the counting convention of DCFormer (Xiao et al., 2024a) and MUDDFormer (Xiao et al., 2025). Let L be the number of layers, T the sequence length, D the model width, n the number of heads, $d = D/n$ the head dimension, and $M \in \mathbb{R}^{d_k \times d}$ the matrix memory. We report FLOPs per sequence at two FLOPs per multiply-add, omitting the shared token embedding and unembedding as well as normalization, nonlinearities, biases and other elementwise operations (such as gate activations and the elementwise mixing in MShift). Causal attention has approximately $T^2/2$ query–key pairs; using a strictly rather than inclusively causal mask changes only the lower-order diagonal terms.

Parameter budget and tokenwise projections. The reference Transformer has $4D^2$ attention and $8D^2$ FFN weights per layer, for a total of $P_0 = 12LD^2$ (applying to both a GELU FFN of hidden width $4D$ and a SwiGLU FFN of hidden width $8D/3$). LoomFormer replaces the full Q/K/V projections with smaller RoPE projections and matrix reads, redistributing the saved parameter budget among the port generators, FFN_x and FFN_M, as well as the initial value projection W^E and embedding MWrite (Eq. 10), so that $P_{\mathrm{Loom}} = P_0$.

Every tokenwise dense matrix W costs $2T|W|$ FLOPs, where $|W|$ is its number of entries. Consequently, to leading order, the combined projection and FFN cost is $2TP_0$ in both models. This includes both stages of the low-rank port generators, all projected gates (including W_g^S in MShift), W^E, FFN_x and FFN_M; a shared projection is counted once. Adding the port-generation or FFN_M FLOPs without subtracting the FFN_x FLOPs that paid for their weights would double-count them. Static ports, learned normalization scales and biases give only lower-order corrections to this dense-weight relation, and for the same reason the bottleneck width D_r affects parameter allocation without entering the leading net FLOP overhead.

Reference Transformer FLOPs. In addition to the tokenwise projections and FFNs, each layer computes $QK^\top$ and attention-weighted values. Each costs $2nd \cdot T^2/2 = DT^2$ FLOPs under the causal approximation. The reference forward cost is therefore

$$F_0 = 2TP_0 + 2LDT^2 = 2LDT(12D + T). \tag{13}$$

Reading and writing the matrix memory. A single MRead contracts a d_k-dimensional port with $M \in \mathbb{R}^{d_k \times d}$, at a cost of $2d_k d$ FLOPs per token. Its dynamic and static ports are summed before this contraction (Eq. 5), so they do not require two separate reads. With four reads per head and n heads, the L layers cost

$$F_{\mathrm{read}} = 4 \cdot 2Tndd_k \cdot L = 8LTDd_k. \tag{14}$$

For each writing sub-layer, MWrite sums n outer products of shape $d_k \times d$, costing $2Tndd_k = 2TDd_k$ FLOPs. There are two such sub-layers per block, attention and FFN_M, except in the last block, which does not write to M. The embedding adds one initial write of the same size (Eq. 10). Hence

$$\begin{aligned} F_{\mathrm{write}} &= 4(L-1)TDd_k + 2TDd_k, \\ F_{\mathrm{M}} &= F_{\mathrm{read}} + F_{\mathrm{write}} = (12L-2)TDd_k \simeq 12LTDd_k. \end{aligned} \tag{15}$$

The last approximation neglects the small endpoint factor $1 - 1/(6L)$ multiplying the memory-overhead term. Since M has a fixed size, the read/write cost per layer does not increase with depth.

Wider query and key. In the reference Transformer, Q, K and V each have d coordinates per head. LoomFormer concatenates $d/4$ RoPE coordinates to the d coordinates of Q/K read from M (Eq. 7), while V retains width d. Thus only the $QK^\top$ contraction is wider; the attention-weighted value contraction is unchanged. The extra cost is

$$\Delta F_{\mathrm{QK}} = 2Ln\frac{d}{4} \cdot \frac{T^2}{2} = \frac{LDT^2}{4}. \tag{16}$$

Total overhead and scaling. Combining Eqs. (13)–(16) gives

$$R_{\Delta\mathrm{FLOPs}} \simeq \frac{12LTDd_k + LDT^2/4}{2LDT(12D+T)} = \frac{d_k}{2D(1+\rho/12)} + \frac{1}{8(1+12/\rho)}, \tag{17}$$

where $\rho = T/D$. With $D = nd$ and $d_k = 2n$, we have $d_k/(2D) = 1/d$, giving the expression in Section 5.4:

$$\boxed{R_{\Delta\mathrm{FLOPs}} \simeq \underbrace{\frac{1}{d(1+\rho/12)}}_{\text{MRead/MWrite}} + \underbrace{\frac{1}{8(1+12/\rho)}}_{\text{wider Q/K}}.} \tag{18}$$

At fixed ρ, the matrix read/write term scales as $1/d$ while the wider-Q/K term is independent of d; as ρ grows, the first term falls from $1/d$ toward 0 and the second rises toward $1/8$, with both independent of L at leading order. In Table 3 ($d = 75, 96, 96, 80$ and $\rho = T/D = 3.41, 2.67, 2.13, 0.80$ for Medium, Large, XL and 2.8B, using $T = 4096$ at

Medium/Large/XL and $T = 2048$ at 2.8B; Table 6), Large, XL and 2.8B have integer $d/4$ and follow Eq. (18) directly (3.13%, 2.77% and 1.95%), whereas Medium rounds its RoPE width $d/4 = 75/4$ to $\lfloor d/4 \rfloor = 18$ in Eq. (16), giving 3.70% instead of Eq. (18)'s unrounded 3.81% (retaining $12L - 2$ in Eq. 15 changes every row by < 0.01 percentage points). For a backbone with a different FFN parameter ratio, one can retain the general denominator $2TP_0 + 2LDT^2$ instead of substituting $P_0 = 12LD^2$. Wall-clock training and inference speeds are detailed in Section F.

F. Training and Inference Throughput Details

Training throughput. Table 8 reports paired training benchmarks in JAX/MaxText on TPU v5p pods at each model's training configuration (Sections 5.1 and 5.2), with both Transformer and LoomFormer using Splash attention (1024-token blocks and fused backward). LoomFormer executes custom Pallas kernels (forward and backward) that fuse MShift, MNorm, port normalization and gating, and the four MReads of all n heads into one per-token $(4n \times d_k) \times (d_k \times d)$ MXU matrix multiplication in VMEM, and likewise fuse MWrite's n-head outer-product accumulation, avoiding intermediate HBM traffic on M. We compare median XPlane device step time under full activation rematerialization on both sides (Table 3). At 2.8B, LoomFormer uses 512-token cross-entropy chunking to fit `v5p-128` HBM whereas Pythia-2.8B does not, making the 2.8B ratio slightly conservative.

Table 8. TPU v5p training step time and relative throughput of LoomFormer (Loom) vs. Transformer (Tfm).

Model	Pod	T	Batch (seqs)	Tfm (ms)	Loom (ms)	Relative speed
Medium	v5p-16	4096	128	1224	1352	90.5%
Large	v5p-32	4096	128	962	1078	89.2%
XL	v5p-32	4096	128	1590	1723	92.3%
2.8B	v5p-128	2048	1024	2988	3233	92.4%

Inference throughput on NVIDIA A800. Table 9 reports BF16 inference throughput on a single NVIDIA A800-80GB PCIe GPU (`torch` 2.14 / `triton` 3.8 / cuDNN 9.24) using trained checkpoints at all four scales, taking the median over five measurements. LoomFormer executes one fused Triton kernel per layer covering the preceding write, the MShift read view, the four MReads, RoPE and cache updates, together with grouped decode GEMVs for the small port projections; prefill uses each model's fastest SDPA backend (FlashAttention for Transformer; cuDNN SDPA for LoomFormer, whose Q/K head width $1.25d$ differs from V width d). Prefill measures one uncached forward pass of $B \times T$ tokens, and decode measures single-token generation steps after a $(T - 64)$-token prompt under `torch.compile(mode="max-autotune")` with CUDA graphs.

Table 9. BF16 inference throughput on one NVIDIA A800-80GB GPU, reported as relative throughput and absolute tokens/s (`Loom / Tfm`); $T = 4096$ for Medium, Large and XL, and $T = 2048$ for 2.8B.

Model	Prefill bs 1	Prefill bs 8	Decode bs 1	Decode bs 16	Decode bs 64
Medium	87.4% (148.9k/170.3k)	90.4% (166.0k/183.5k)	91.7% (565/616)	89.2% (2109/2365)	92.3% (2499/2709)
Large	85.3% (95.3k/111.7k)	86.2% (104.0k/120.6k)	85.4% (438/513)	87.8% (1916/2183)	84.9% (2311/2721)
XL	83.8% (50.8k/60.6k)	85.3% (55.3k/64.9k)	86.9% (272/313)	84.6% (1245/1472)	86.6% (1452/1677)
2.8B	87.5% (30.4k/34.7k)	90.5% (33.1k/36.6k)	86.4% (171/198)	87.6% (1248/1425)	88.9% (1836/2066)

G. Extended Ablation Studies

Table 10 reports the complete 18-variant ablation suite at the Medium scale ($L = 18$, $D = 1200$, $n = 16$, $d = 75$, $d_k = 32$, $D_r = 256$, RoPE dim $\lfloor d/4 \rfloor = 18$, $T = 4096$, 7B tokens), extending Table 4. In full LoomFormer, FFN_M has SwiGLU hidden width 1128, and FFN_x has width 2510 in blocks 0–16 and 3904 in block 17 (which omits unused writes to M_L). Whenever an ablation removes parameters, we widen FFN_x uniformly across blocks 0–16 and absorb any embedding or integer-rounding residual in block 17, matching the 432.11M parameter budget within ± 1.4K ($< 0.0004\%$).

Write-port gating vs. slot selection, and reduction to $\mathrm{Tfm}_{K,Q}$. Fixing write keys k^m to learned static vectors while keeping their dynamic gates $\sigma(xW_g^W)$ preserves nearly the full gain (100.0% for LoomMHA, 99.3% for FFN_M, and 95.6% across all three write sites including the embedding). Dropping the gate on static k^m, however, hurts LoomMHA over three times as much as FFN_M (100.0% $\rightarrow$ 91.1% vs. 99.3% $\rightarrow$ 96.6%): LoomMHA normalizes its head outputs $\mathrm{Norm}(o_i)$ before writing (Eq. 8), leaving the gate as the sole control of write magnitude (and whether a head is on), whereas FFN_M

Table 10. Complete Medium-scale (432M) ablation suite of LoomFormer under a matched parameter budget (432.11M$\pm$1.4K). †Included in Table 4.

Variant	Specification (vs. full LoomFormer)	Loss↓	ΔLoss	Gain
LoomFormer (full)†	—	**2.266**	—	**100.0%**
Transformer++ (baseline)†	Standard MHA + SwiGLU FFN, no matrix stream M	2.443	+0.177	0.0%
Dynamic rewiring & write-port gating (Theorem 3.1)				
Static LoomMHA k^m, gated	Static k^m in LoomMHA; write gate $\sigma(xW_g^W)$ kept	2.266	0.000	100.0%
Static LoomMHA k^m	Static k^m in LoomMHA; write gate dropped ($\equiv 1$)	2.282	+0.016	91.1%
Static FFN_M k^m, gated	Static k^m in FFN_M; write gate $\sigma(xW_g^W)$ kept	2.268	+0.001	99.3%
Static FFN_M k^m	Static k^m in FFN_M; write gate dropped ($\equiv 1$)	2.272	+0.006	96.6%
Static write keys (k^m), gated	Static k^m in emb., LoomMHA, FFN_M; write gates kept	2.274	+0.008	95.6%
Static write keys (k^m)†	Static k^m in emb., LoomMHA, FFN_M; write gates dropped	2.283	+0.017	90.3%
Static read queries (q^m)†	Only static $\bar{q}$ in $\mathrm{MRead}^{Q,K,V,O}$; all read gates dropped	2.310	+0.043	75.4%
Static read & write (q^m, k^m), gated	Static q^m (no read gates), all k^m; write gates kept	2.327	+0.060	65.8%
Static read & write (q^m, k^m)†	Static q^m, all k^m; port gates dropped ($\mathrm{Tfm}_{K,Q}$ + MShift)	2.344	+0.078	56.0%
Static read & write (q^m, k^m), − MShift	Static q^m, all k^m; port gates, MShift dropped (Prop. 3.1(i))	2.348	+0.082	53.6%
Port normalization & gating (Eqs. 5–6)				
− dynamic port gates†	$\sigma(xW_g^R), \sigma(xW_g^W) \equiv 1$; $\sigma(xW_{\tilde{g}}^R)$ and MShift kept	2.277	+0.011	93.7%
− dynamic port Norm†	Drop Norm on dynamic q^m, k^m; $\mathrm{Norm}(o_i)$, MNorm kept	2.289	+0.023	87.1%
− dynamic port gates & Norm†	Drop both $\sigma(xW_g^R), \sigma(xW_g^W)$ and port Norm	2.283	+0.017	90.3%
Sub-layer components (Section 3.2)				
− FFN_M write†	Remove FFN_M and its MWrite (Eq. 9)	2.300	+0.034	80.7%
− output bypass o^m†	Drop MRead^O; restore self-attention to current token (Eq. 8)	2.299	+0.032	81.8%
− vector-stream RoPE $W^{Q,K}$†	Drop $W^{Q,K}$; RoPE on 18 of 75 dims of $\tilde{q}, \tilde{k}$ (Eq. 7)	2.286	+0.020	88.7%
− MShift†	Set $z_{[t]} \equiv 0$ in Eq. (7)	2.277	+0.010	94.2%
− LoomMHA write $\mathrm{Norm}(o_i)$	Drop Norm on o_i in LoomMHA writes (Eq. 8)	2.268	+0.002	98.8%

writes unnormalized slices (Eq. 9) whose norms can still modulate write strength per token. When both q^m and k^m are static, retaining dynamic write gates gives a node-gated circuit on fixed slots (65.8%), dropping those gates leaves $\mathrm{Tfm}_{K,Q}$ with MShift (56.0%), and further removing MShift yields the exact $\mathrm{Tfm}_{K,Q}$ of Theorem 3.1(i) (53.6%)—which has no token-dependent port or shift gates on M, yet still outperforms all layer-wise dynamic baselines at the same scale (Figure 5, where MUDDFormer reaches 48.5% of LoomFormer's gain).

Attention write-value normalization. Removing $\mathrm{Norm}(o_i)$ from LoomMHA writes (Eq. 8) loses only 1.2% of the gain (98.8% retained), as whole-matrix pre-normalization MNorm before reads (Eq. 7) already bounds the scale of M.

H. Mechanistic Analysis Details and Visualizations

Evaluation protocol and linear attribution. The statistics in Section 5.3 and this section are computed on held-out Pile validation sequences of length $T = 2048$ at the 300B-token checkpoints of Pythia-2.8B, mHC-Pythia-2.8B and LoomPythia-2.8B (evaluating token positions $t \geq 64$ to avoid sequence-boundary transients). Because MShift and MNorm act linearly on the accumulated writes at each layer (Equation (7)), absorbing the per-slot shift gate $z_{[t]}$ (or $1 - z_{[t]}$) and the scalar MNorm factor into $c_{ij} = \langle q_j^m, k_i^m \rangle$ decomposes every MRead output as $u_j = \sum_i c_{ij} o_i$ over all earlier writers (embedding, LoomMHA and FFN_M writes at t and $t-1$; 0.4% median RMS difference from BF16 execution across layers). Projecting each term onto u_j gives the additive writer attribution

$$a_{ij} = \frac{c_{ij} \langle o_i, u_j \rangle}{\|u_j\|^2}, \qquad \sum_i a_{ij} = 1. \tag{19}$$

Definitions for Figure 6. In Figure 6(a), to isolate dynamic port rewiring from both static virtual weights (Theorem 3.1(i)) and per-token scalar gate or writer-norm fluctuations, let $\hat{q}_j^m(t) = q_j^m(t)/\|q_j^m(t)\|_2$ and $\hat{k}_i^m(t) = k_i^m(t)/\|k_i^m(t)\|_2 \in \mathbb{S}^{d_k - 1}$ denote the unit port directions of reader j and same-token writer i, weighted by each writer's corpus RMS norm $\bar{\omega}_i = \sqrt{\mathbb{E}_t[\|o_i(t)\|^2]}$: $\widehat{C}_{ij}(t) = \bar{\omega}_i \langle \hat{q}_j^m(t), \hat{k}_i^m(t) \rangle$. We center across tokens, $\Delta\widehat{C}(t) = \widehat{C}(t) - \mathbb{E}_{t'}[\widehat{C}(t')]$, and report the energy effective rank $\exp\big(-\sum_k p_k \log p_k\big)$ with $p_k = \sigma_k^2 / \sum_{k'} \sigma_{k'}^2$ over the $3n = 96$ Q/K/V readers of layer l (or the $n = 32$ readers of a single Q, K or V input); under fixed-port $\mathrm{Tfm}_{K,Q}$ with arbitrary positive scalar gates, $\Delta\widehat{C}(t) \equiv 0$. In

Figure 6(b), for each read port $q_j^m(t) \in \mathbb{R}^{d_k}$ ($\tilde{q}, \tilde{k}, v, o^m$) and write port $k_i^m(t) \in \mathbb{R}^{d_k}$ (LoomMHA and FFN_M), we average $1 - \cos(p(t_1), p(t_2))$ over distinct token pairs $t_1 \neq t_2$ ($t \geq 64$) and over the $n = 32$ heads of layer l (embedding write ports remain near-static at 0.005, providing a sanity check). In Figure 6(c), following Xiao et al. (2025), a head is on at query position t if $\arg\max_s \alpha_{ts}$ falls neither in $\{0, 1\}$ nor on a sink token (<|endoftext|>, ".", " .", "\n", "\n\n"); in layers 17–32, first-two-token sink mass drops from 26.4% (Pythia) and 27.2% (mHC) to 1.5% in LoomPythia while content-token attention mass rises from 22.7% and 31.2% to 65.4%, and masking the self-diagonal $s = t$ in Pythia and mHC to match LoomMHA's strict causal mask (Equation (8)) leaves the gap unchanged (74.4% vs. 21.2% and 15.7%). In Figure 6(d), for each head h in layers 3–32 and each pair among its four ports $\{\tilde{q}_h, \tilde{k}_h, v_h, o_h^m\}$, we compute from token-averaged positive attributions $\bar{a}_{ij}^+ = \mathbb{E}_t[\max(a_{ij}(t), 0)]$ over same-token writers: (i) Pearson correlation after summing attributions by writer group (Emb, each earlier LoomMHA layer, each earlier FFN_M layer), (ii) Pearson correlation across individual head/module writers, and (iii) $|S_A \cap S_B|/32$ for the top-32 writer sets S_A, S_B, reporting the median across heads (in slot space $\mathbb{R}^{d_k}$, same-token cosine similarity between ports of the same head averages 0.01–0.07 outside Q–K and 0.27 on Q–K).

Writer composition and depth-to-depth routing. Averaged over layers 2–32 in LoomPythia-2.8B, positive attribution mass $\max(a_{ij}, 0)$ is shared across LoomMHA writes (53.4%–62.5% across $\tilde{q}, \tilde{k}, v, o^m$), FFN_M writes (27.5%–35.2%) and embedding writes (10.0%–11.7%), while previous-token MShift reads account for 13.7%–15.2% of $\sum_i |a_{ij}|$ (Figure 8(a)). FFN_M reaches its largest share (35.2%) on value reads v, consistent with FFN_M serving as the sole post-embedding channel from the vector stream into M (Equation (9)) and with its large effect in Table 4. Plotted with writer sources on the vertical axis (bottom to top) and reader layers on the horizontal axis, the column-max normalized routing map (Figure 8(b)) combines local relay along the diagonal with persistent horizontal bands from hub writer sources—most prominently the embedding at early layers and writer layers 13–14 and 19 across all subsequent layers.

Adjacent-token rewiring on an in-context prompt. Figure 9 shows the head-writer attribution submatrix $[a_{ij}]$ at hub layer 19 (selecting the top 6 readers per MRead arm and top 18 earlier LoomMHA head writers by peak $|a_{ij}|$ across the two tokens) on `"France: Paris\nJapan: Tokyo\nChina: Beijing\nSpain: Madrid\nGermany:"`, where LoomPythia-2.8B predicts `" Berlin"` top-1 with probability 0.65. At `"Germany"` (Figure 9(a)), writers route selectively to distinct MRead arms—`L13.H11` feeds key reads ($\tilde{k}$: `K14`, `K23`, `K25`) without feeding value or bypass reads, whereas `L13.H12` and `L08.H25` feed value and bypass reads (`V10`, `O09`, `O22`)—forming a non-rectangular within-token pattern (Theorem 2.3) that prepares `"Germany"`'s horizontal keys, values and residual state. At the immediately following delimiter `":"` (Figure 9(b)), LoomMHA's strict causal mask ($s < t$, Equation (8)) means that only $\tilde{q}$ (attending back to `"Germany"` and the few-shot capitals) and o^m affect the `" Berlin"` prediction, whereas $\tilde{k}$ and v at `":"` serve only future tokens. Correspondingly, `L13.H11` and `L13.H12` switch off on layer 19, while layer 19 rewires its causal query reads $\tilde{q}$ onto `L18.H25`, `L15.H32` and `L14.H01` (`Q06`, `Q07`, `Q17`, `Q22`) and its bypass outputs o^m onto `L14.H06`, `L16.H29` and `L14.H19` (`O16`, `O21`); causally freezing layer 19's dynamic read ports at `":"` to those at `"Germany"` alone drops $P($`" Berlin"`$)$ from 0.65 to 0.54 (0.41 when frozen across all layers).

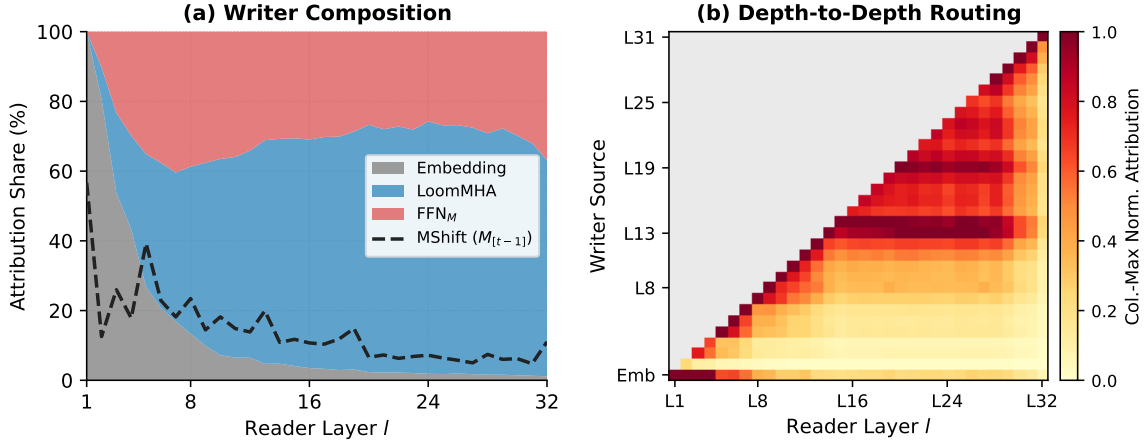

Figure 8. Matrix-stream routing in LoomPythia-2.8B across Pile validation sequences. (a) Share of positive attribution from Embedding, LoomMHA and FFN_M writes, and share of $|a_{ij}|$ from previous-token MShift ($M_{[t-1]}$). (b) Column-max normalized positive attribution from each writer source (vertical, bottom to top) to each reader layer l (horizontal; gray upper triangle: causally unreachable).

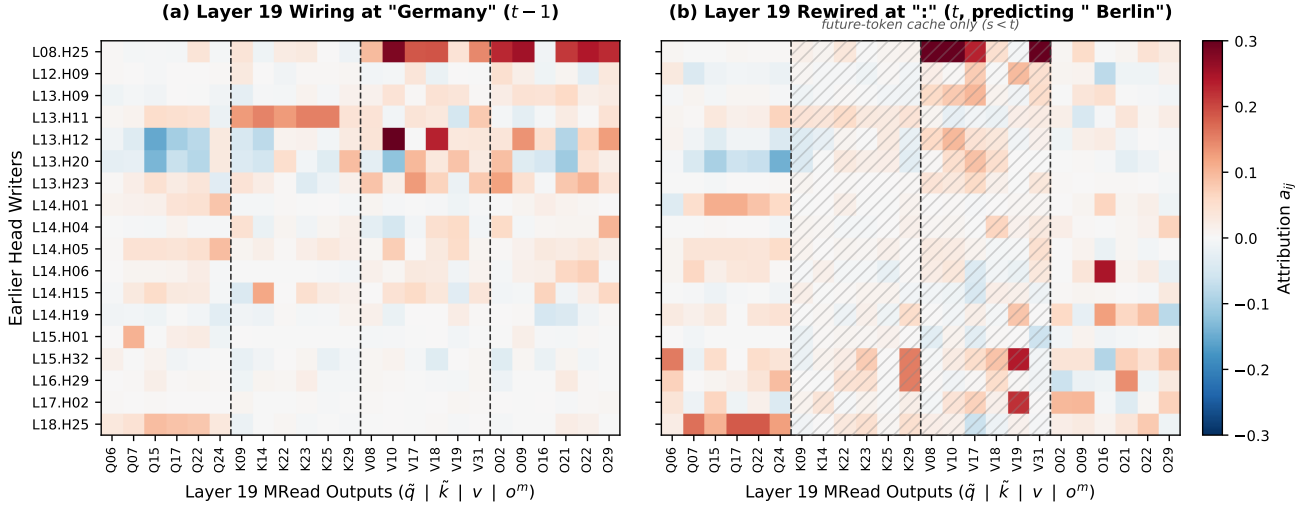

Figure 9. Attribution submatrix $[a_{ij}]$ (rows: top-18 earlier head writers `L`l`.H`h; columns: top-6 layer-19 MRead outputs per arm $\tilde{q}, \tilde{k}, v, o^m$) of LoomPythia-2.8B at two consecutive tokens of an in-context prompt (next-token top-1: `" Berlin"`, $p = 0.65$), showing non-rectangular within-token wiring at `"Germany"` (a) and per-token rewiring onto causal $\tilde{q}$ and o^m pathways at `":"` (b; hatched $\tilde{k}, v$ columns affect only future tokens under $s < t$).