

Residual Stream Duality in Modern Transformer Architectures

Yifan Zhang

yifanzhangresearch@gmail.com

Abstract

Recent work has made clear that the residual pathway is not mere optimization plumbing; it is part of the model’s representational machinery. We agree, but argue that the cleanest way to organize this design space is through a two-axis view of the Transformer. A decoder evolves information along two ordered dimensions: sequence position and layer depth. Self-attention already provides adaptive mixing along the sequence axis, whereas the residual stream usually performs fixed addition along the depth axis. If we fix a token position and treat layer index as the ordered variable, then a causal depth-wise residual attention read is the same local operator as causal short sliding-window attention (ShortSWA), except written over depth rather than over sequence. This is the core residual-stream duality underlying Transformer².

This perspective also clarifies the recent literature. ELC-BERT and DenseFormer already show that learned aggregation over depth can outperform uniform residual accumulation. In contrast, Vertical Attention, DeepCrossAttention (DCA), MUDDFormer, and Attention Residuals move further toward explicit attention-based routing over earlier layers. The key point, however, is that operator-level duality does not imply systems-level symmetry. For large-scale autoregressive models, sequence-axis ShortSWA is usually the more hardware-friendly placement because it reuses token-side sliding-window kernels, KV-cache layouts, and chunked execution. If the goal is instead to change the shortcut itself, Deep Delta Learning (DDL) is the cleaner intervention because it modifies the residual operator directly rather than adding a separate cross-layer retrieval path. Our recommendation is therefore simple: use DDL when the shortcut is the object of interest, and use sequence-axis ShortSWA when the goal is local adaptive mixing.

Project Page: <https://github.com/yifanzhang-pro/residual-stream-duality>

1 Introduction

A modern Transformer evolves information along two ordered axes: sequence position and layer depth. Along the sequence axis, self-attention performs learned, content-dependent mixing. Along the depth axis, the residual stream usually performs uniform addition. The title Transformer² is meant literally: modern Transformer architectures have two ordered directions of information flow, but only one of them is usually equipped with an adaptive attention operator. The main theme of this note is that this asymmetry is conceptually revealing and practically consequential.

That asymmetry has already motivated a broad family of proposals that replace or augment uniform depth aggregation. Earlier examples include ELC-BERT (Charpentier and Samuel, 2023), which feeds each layer a convex combination of earlier layer outputs, and DenseFormer, which inserts a depth-weighted average after each block (Pagliardini et al., 2024). More recent work makes the

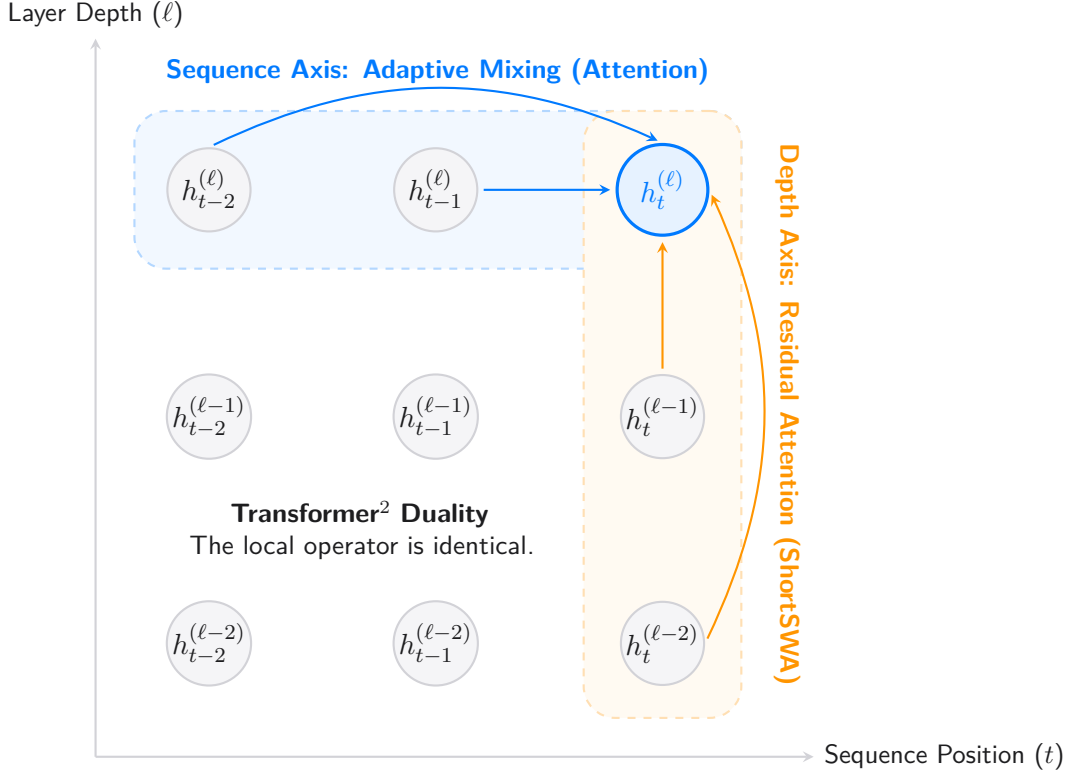


Figure 1 Overview of Residual Stream Duality. A modern Transformer evolves information along two ordered dimensions: sequence position and layer depth. Computing the target state $h_t^{(\ell)}$ via explicit depth-wise residual attention operates identically to causal short sliding-window attention (ShortSWA) on the sequence axis, effectively bridging the representation mechanisms of the two pathways.

cross-depth routing explicitly attention-based, including Vertical Attention, DeepCrossAttention (DCA), MUDDFormer, and Attention Residuals (Kojima et al., 2026; Heddes et al., 2025; Xiao et al., 2025; Chen et al., 2026). Related interventions such as Hyper-Connections and Deep Delta Learning (DDL) further underscore that shortcut design remains an active architectural degree of freedom (Zhu et al., 2024; Zhang et al., 2026). The shared lesson is that the residual pathway participates in representation, not merely optimization.

Our claim is not that all of these proposals are identical. ELC-BERT and DenseFormer sit on the learned-static end of the spectrum; Vertical Attention, DCA, MUDDFormer, and Attention Residuals use more expressive routing modules. But the common object is learned aggregation over the ordered depth axis. The cleanest exact statement applies to explicit depth-wise attention reads: once a token position is fixed, and a layer index is treated as a one-dimensional ordered axis, a truncated residual attention read is precisely causal ShortSWA written over depth. The full-memory variant is simply the full-window limit of the same operator family.

That duality is mathematical, not systems-symmetric. Sequence-axis ShortSWA reuses existing sliding-window attention kernels, token-side KV-cache layouts, and chunked execution strategies. Depth-axis aggregation, by contrast, requires an additional layer-indexed state path: each block needs online access to earlier layer states or block summaries for the same token, and under pipeline parallelism, those states may need to be forwarded, stored, or recomputed. The practical question

is therefore not whether attention can be applied over depth, but whether depth is the right axis on which to place a short adaptive mixer. The thesis of this note is therefore:

- A depth-wise residual attention read is not a new local operator. It is ShortSWA written on the depth axis rather than the sequence axis.
- Learned cross-depth aggregation spans a continuum from static depth weighting (ELC-BERT, DenseFormer) to attention-based routing (Vertical Attention, DCA, MUDDFormer, Attention Residuals). These systems are not identical end-to-end, but they occupy the same design space.
- Once that distinction is explicit, the natural design choice is either to use Deep Delta Learning (Zhang et al., 2026) to improve the shortcut itself or to place ShortSWA directly on the sequence axis, which is usually more hardware-efficient for current training and inference stacks.
- Following Zhang (2025), we view ShortSWA as the successor to ShortConv and, in spirit, the attention-form successor to Canon layers (Allen-Zhu, 2025).

Relation to prior depth-aggregation work. ELC-BERT and DenseFormer are important precursors because they already replace uniform depth accumulation with learned aggregation. ELC-BERT feeds each layer a convex combination of previous layer outputs, while DenseFormer adds a depth-weighted average of current and past representations after each block (Charpentier and Samuel, 2023; Pagliardini et al., 2024). Vertical Attention, DCA, MUDDFormer, and Attention Residuals move further toward attention-based routing over earlier layers (Kojima et al., 2026; Heddes et al., 2025; Xiao et al., 2025; Chen et al., 2026). Our claim is therefore not that these methods are end-to-end identical. It is that, once depth is treated as an ordered axis, they are best compared inside one common design space of learned cross-depth aggregation. DDL, by contrast, attacks a different target: it changes the shortcut update itself rather than adding a separate retrieval path over stored earlier states (Zhang et al., 2026). Hyper-Connections make a related point, that residual design is itself a meaningful architectural degree of freedom, but they do not remove the system’s asymmetry between token-side local mixing and layer-side state management (Zhu et al., 2024).

Relation to ShortConv and Canon layers. ShortConv, Canon layers, and ShortSWA all occupy the same architectural slot: they are local mixers that operate before or alongside a broader global mechanism. ShortConv uses a fixed, small kernel. Canon layers compute learned weighted sums over nearby tokens (Allen-Zhu, 2025). As argued in Zhang (2025), once chunked computation is already part of the implementation, the natural upgrade is ShortSWA: the same local role, but with content-adaptive mixing and a chunk-aligned receptive field. In that sense, ShortSWA is the natural successor to ShortConv and the attention-form successor to Canon layers.

2 Residual Stream Duality

2.1 Preliminaries

Let $H = [H^{(0)}, \dots, H^{(L)}] \in \mathbb{R}^{(L+1) \times T \times d}$ denote the hidden-state stack of an L -block decoder, where $H^{(0)}$ is the input stream, T is the sequence length, and d is the model width. We write $H^{(\ell)} \in \mathbb{R}^{T \times d}$ for the hidden states at depth ℓ . A standard pre-norm Transformer block, for $\ell = 0, \dots, L - 1$, is

$$\begin{aligned}
 U^{(\ell)} &= H^{(\ell)} + \text{Attn}\left(\text{Norm}\left(H^{(\ell)}\right)\right), \\
 H^{(\ell+1)} &= U^{(\ell)} + \text{MLP}\left(\text{Norm}\left(U^{(\ell)}\right)\right).
 \end{aligned}$$

The sequence axis is mixed adaptively by attention, while the depth axis is mixed by fixed addition.

2.2 Depth-wise residual attention is ShortSWA on the depth axis

Fix a token position t and collect its trajectory through depth:

$$X_t = [h_t^{(0)}; h_t^{(1)}; \dots; h_t^{(L)}] \in \mathbb{R}^{(L+1) \times d}.$$

Now consider a causal depth window of size K . Define

$$\mathcal{W}_{t,\ell}^{(K)} = [h_t^{(\max(0, \ell-K+1))}; \dots; h_t^{(\ell)}] \in \mathbb{R}^{K_\ell \times d}, \quad K_\ell = \min(K, \ell + 1).$$

A depth-wise residual attention read at layer ℓ can be written as

$$q_{t,\ell} = W_Q h_t^{(\ell)}, \quad \mathbf{K}_{t,\ell} = \mathcal{W}_{t,\ell}^{(K)} W_K, \quad \mathbf{V}_{t,\ell} = \mathcal{W}_{t,\ell}^{(K)} W_V,$$

$$z_t^{(\ell)} = \text{softmax}\left(\frac{q_{t,\ell} \mathbf{K}_{t,\ell}^\top}{\sqrt{d_k}}\right) \mathbf{V}_{t,\ell}.$$

This is exactly causal ShortSWA applied to the one-dimensional sequence X_t whose index is the layer number:

$$z_t^{(\ell)} = \text{ShortSWA}(X_t; K)_\ell.$$

Hence, after transposing the hidden-state tensor so that depth becomes the ordered axis, truncated depth-wise residual attention and ShortSWA belong to the same operator family. The full-memory residual attention variant is simply the full-window limit $K = \ell + 1$.

This exact equivalence applies whenever the cross-depth retrieval is implemented as an explicit attention read. Simpler learned-weight schemes such as ELC-BERT or DenseFormer belong to the same broader design space but are not literally instances of the QKV operator above (Charpentier and Samuel, 2023; Pagliardini et al., 2024).

2.3 A unified view of learned depth aggregation

The exact equivalence above applies to the explicit depth-wise residual attention read written here. It also suggests a useful taxonomy of nearby methods. ELC-BERT and DenseFormer are learned depth aggregators with parameterized weights over earlier layers, but without a full depth-wise QK attention read (Charpentier and Samuel, 2023; Pagliardini et al., 2024). Vertical Attention, DCA, MUDDFormer, and Attention Residuals are closer to the explicit attention end of the spectrum: Vertical Attention learns inter-layer paths through routing modules, DCA computes attention inputs from mixtures of previous layer outputs, MUDDFormer introduces separate dynamic dense modules for query, key, value, and residual streams, and Attention Residuals presents the read most directly as attention over depth (Kojima et al., 2026; Heddes et al., 2025; Xiao et al., 2025; Chen et al., 2026). These systems are not identical end-to-end architectures; they differ in factorization, parameter sharing, gating, and injection point. What the duality statement contributes is a common coordinate system: once depth is treated as an ordered axis, explicit cross-depth attention is simply local causal attention on that axis, and the broader family can be read as increasingly expressive parameterizations of learned depth aggregation.

2.4 Why the sequence axis is the better placement

Once the equivalence above is explicit, the main design question becomes where to place the short attention primitive when the goal is local adaptive mixing. Our view is that the sequence axis is the better answer:

$$\begin{aligned} S^{(\ell)} &= H^{(\ell)} + \text{ShortSWA}(\text{Norm}(H^{(\ell)}); w), \\ A^{(\ell)} &= S^{(\ell)} + \text{Attn}(\text{Norm}(S^{(\ell)})), \\ H^{(\ell+1)} &= A^{(\ell)} + \text{MLP}(\text{Norm}(A^{(\ell)})). \end{aligned}$$

This preserves the same local-to-global story but places the adaptive local mixer on the axis that modern kernels and inference stacks already optimize.

At autoregressive inference time, sequence-axis ShortSWA can reuse the usual token-side cache layout over the most recent w tokens. In chunked training or inference, the local window can be aligned to the chunk already loaded into SRAM. Under pipeline parallelism, the implementation preserves the standard forward flow of activations between layer partitions rather than introducing an additional layer-indexed state path. Depth-axis attention-style aggregation faces the opposite incentives: each block needs online access to earlier layer states or block summaries for the same token. Methods such as Vertical Attention, DCA, MUDDFormer, and blockwise Attention Residuals differ in how they parameterize or compress this access, but they all live with the same underlying pressure: depth-side routing must manage cross-layer state explicitly (Kojima et al., 2026; Heddes et al., 2025; Xiao et al., 2025; Chen et al., 2026). If the target is instead the shortcut operator itself, we would choose Deep Delta Learning rather than add another cross-depth read, because DDL changes the residual update directly and does not require an explicit stack of earlier layer states (Zhang et al., 2026).

2.5 Recommended block

The resulting recommendation is therefore a clean two-way design fork:

- If the goal is a better shortcut, use Deep Delta Learning (Zhang et al., 2026).
- If the goal is a local content-adaptive mixer, use ShortSWA directly on the sequence axis.

For current large-scale training and inference stacks, we do not see a general systems case for a third option that repackages sequence-local attention as a depth-axis residual mechanism. When the goal is sequence-side local mixing, the resulting local-to-global Transformer block is as follows.

Algorithm 1 Recommended local-to-global block when the goal is sequence-side local mixing

Require: hidden states $H^{(\ell)}$, local window w

- 1: $S^{(\ell)} \leftarrow H^{(\ell)} + \text{ShortSWA}(\text{Norm}(H^{(\ell)}); w)$
 - 2: $A^{(\ell)} \leftarrow S^{(\ell)} + \text{Attn}(\text{Norm}(S^{(\ell)}))$
 - 3: $H^{(\ell+1)} \leftarrow A^{(\ell)} + \text{MLP}(\text{Norm}(A^{(\ell)}))$
-

2.6 Complexity and systems notes

Ignoring head-wise constants, ShortSWA adds a local attention term of roughly $O(Twd)$ per layer, where $w \ll T$. If a block still includes full self-attention, the asymptotic sequence-mixing cost

remains $O(T^2d)$ up to constant factors. The important point is not a new asymptotic regime, but a better hardware placement: the local operation lives on the token axis and can reuse standard sliding-window kernels and KV-cache layouts.

The next estimates apply to explicit attention-style reads over depth, not to lighter learned-weight schemes such as ELC-BERT or DenseFormer. Depth-wise residual attention with a depth window K adds $O(TKd)$ work per block, hence $O(TKLd)$ across an L -block network, together with additional online access to earlier layer states or block summaries. The full-depth variant grows to $O(TL^2d)$ for the score/value interactions. These formulas make the compute overhead visible, but the more consequential issue in practice is systems complexity: one now needs an extra layer-indexed state that must be retained, forwarded, or recomputed, especially when depth windows cross pipeline-stage boundaries. In some deployments, this behaves like a second cache over depth. DDL avoids this depth-axis state-management overhead because it modifies the per-block shortcut rather than attending over stored earlier layer states (Zhang et al., 2026).

3 Conclusion

This note formalizes a simple residual-stream duality. A Transformer evolves information along two ordered axes: sequence position and layer depth. When a token position is fixed, and the layer index is treated as the ordered variable, an explicit depth-wise residual attention read is ShortSWA on the transposed axis: a causal local attention operator over previous layers rather than previous tokens. This view places prior depth-aggregation methods in one design space. ELC-BERT and DenseFormer use learned but mostly static depth aggregation, while Vertical Attention, DCA, MUDDFormer, and Attention Residuals move toward content-dependent cross-depth routing (Charpentier and Samuel, 2023; Pagliardini et al., 2024; Kojima et al., 2026; Heddes et al., 2025; Xiao et al., 2025; Chen et al., 2026). They differ architecturally, but all modify how information is aggregated across depth.

The duality is therefore conceptual, not a mandate to attend over depth by default. If the goal is to improve the shortcut itself, DDL is the more direct intervention. If the goal is local adaptive mixing, sequence-axis ShortSWA is usually the cleaner systems choice, since it matches existing sliding-window kernels, token-side KV caches, and chunked execution. In this sense, ShortSWA remains the natural successor to ShortConv and the content-adaptive analogue of Canon-style local mixing (Zhang, 2025; Allen-Zhu, 2025). Transformer² is best read as an organizing principle: one local attention family, two ordered axes, and a practical preference for sequence placement unless cross-depth retrieval is the object of study.

Acknowledgement

We sincerely thank Xinyu Yang for helpful discussions. We used large language models to assist in polishing the writing of this work.

References

Zeyuan Allen-Zhu. Physics of language models: Part 4.1, architecture design and the magic of canon layers. *arXiv preprint arXiv:2512.17351*, 2025.

- Lucas Georges Gabriel Charpentier and David Samuel. Not all layers are equally as important: Every layer counts bert. In *Proceedings of the BabyLM Challenge at the 27th Conference on Computational Natural Language Learning*, pages 238–252, 2023.
- Guangyu Chen, Yu Zhang, Jianlin Su, Weixin Xu, Siyuan Pan, Yaoyu Wang, Yucheng Wang, Guanduo Chen, Bohong Yin, Yutian Chen, Junjie Yan, Ming Wei, Y. Zhang, Fanqing Meng, Chao Hong, Xiaotong Xie, Shaowei Liu, Enzhe Lu, Yunpeng Tai, Yanru Chen, Xin Men, Haiqing Guo, Y. Charles, Haoyu Lu, Lin Sui, Jinguo Zhu, Zaida Zhou, Weiran He, Weixiao Huang, Xinran Xu, Yuzhi Wang, Guokun Lai, Yulun Du, Yuxin Wu, Zhilin Yang, and Xinyu Zhou. Attention residuals. *GitHub*, 2026. URL <https://github.com/MoonshotAI/Attention-Residuals>.
- Mike Heddes, Adel Javanmard, Kyriakos Axiotis, Gang Fu, MohammadHossein Bateni, and Vahab Mirrokni. Deepcrossattention: Supercharging transformer residual connections. *arXiv preprint arXiv:2502.06785*, 2025.
- Takeshi Kojima, Yusuke Iwasawa, Rio Yokota, Yusuke Miyao, Jun Suzuki, and Yutaka Matsuo. Vertical attention: Automatic exploration of inter-layer connections in transformer-based language models. *Openreview*, 2026.
- Matteo Pagliardini, Amirkeivan Mohtashami, Francois Fleuret, and Martin Jaggi. Denseformer: Enhancing information flow in transformers via depth weighted averaging. *Advances in neural information processing systems*, 37:136479–136508, 2024.
- Da Xiao, Qingye Meng, Shengping Li, and Xingyuan Yuan. Muddformer: Breaking residual bottlenecks in transformers via multiway dynamic dense connections. *arXiv preprint arXiv:2502.12170*, 2025.
- Yifan Zhang. Rethinking swa: Why short sliding window attention will replace shortconv in modern architectures. <https://github.com/yifanzhang-pro/Rethinking-SWA>, December 2025. Blog post and project page.
- Yifan Zhang, Yifeng Liu, Mengdi Wang, and Quanquan Gu. Deep delta learning. *arXiv preprint arXiv:2601.00417*, 2026.
- Defa Zhu, Hongzhi Huang, Zihao Huang, Yutao Zeng, Yunyao Mao, Banggu Wu, Qiyang Min, and Xun Zhou. Hyper-connections. *arXiv preprint arXiv:2409.19606*, 2024.

Appendix