

Article

Not peer-reviewed version

---

# From Transformer to Transponder: Introducing Contextual Modulation Training for Residual Learning in LLMs

---

[Yingtao Zhang](#) , [Wenqi Gu](#) , Wen Hu , Jianguo Li , [Carlo Vittorio Cannistraci](#) \*

Posted Date: 30 September 2025

doi: 10.20944/preprints202506.0120.v2

Keywords: deep learning; residual connection; modulation training; contextual scaling



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

## Article

# From Transformer to TRANSPONDER: Introducing Contextual Modulation Training for Residual Learning in LLMs

Yingtao Zhang <sup>1,2</sup>, Wenqi Gu <sup>1,3</sup>, Wen Hu <sup>4</sup>, Jianguo Li <sup>4</sup> and Carlo Vittorio Cannistraci <sup>1,2,3,\*</sup>

<sup>1</sup> Center for Complex Network Intelligence (CCNI)

<sup>2</sup> Department of Computer Science

<sup>3</sup> Department of Biomedical Engineering, Tsinghua University

<sup>4</sup> Ant Group

\* Correspondence: kalokagathos.agon@gmail.com

## Abstract

Transformers are the backbone of state-of-the-art systems across language, vision, and multimodal learning tasks, yet the relevance scale of their functional blocks (self-attention and feed-forward networks) is typically constant across inputs and depth. This static design neglects context-sensitive regulation of information flow through residual pathways. We introduce the *contextual modulator*: a lightweight, input-aware mechanism that can scale the outputs of linear sublayers within a block or the entire block output at token- and channel-level granularity. The modulator is implemented via compact parametric functions and adds negligible parameter overhead. Building on this idea, we propose TRANSPONDER, which integrates contextual modulators throughout Transformer blocks to endow functional residual architectures with fine-grained, input-adaptive control. TRANSPONDER provides evident improvement over six other scaling or normalization methods across LLaMA backbones ranging from 60M to 250M parameters, yielding consistent perplexity reductions with < 1% additional parameters. Analysis reveals depth-, module-, and token-specific scaling patterns, indicating that learned modulators act as input-adaptive regulators of residual information flow. TRANSPONDER provides a simple, general mechanism to augment Transformer-based models with context-sensitive modulators, providing robust and significant performance improvements without substantial architectural changes.

**Keywords:** deep learning; residual connection; modulation training; contextual scaling

## 1. Introduction

Transformer architectures deliver state-of-the-art performance across vision, language, and multimodal tasks [1–4]. In standard designs, information flows through a sequence of linear and nonlinear components—query/key/value projections, attention output projections, and the two feed-forward networks (FFN) linear maps—interleaved with residual connections. After training, the scaling by which these components contribute to the residual stream is effectively static for all the input tokens. Several mechanisms attempt to manipulate residual or path scaling via global reparameterizations (e.g., ReZero [5], DeepNorm/DeepNet [6]) or depth-aware normalization [7,8]. Yet these approaches assign fixed gains in spite of the current token, channel, or depth position.

These designs clash with the core goal of representation learning: the relevance of a component's output is input-dependent. Rare tokens may call for amplifying specific heads; bursty channels may require attenuation; deeper layers may benefit from different mixing strengths than shallow ones. When scaling is static, the model must implicitly encode such selectivity inside the functional maps themselves—mixing representation (what to compute) with control (how much to pass through).

To address this limitation, we introduce TRANSPONDER, whose core principle is simple: **pair each Transformer component with a contextual, input-aware modulator that explicitly controls its contribution at inference**. Specifically, for a component with output  $\mathbf{Y} = \mathbf{F}(\mathbf{X})$ , we attach a lightweight controller  $g(\mathbf{X}) \in [0, 1]$  and compute  $\tilde{\mathbf{Y}} = \mathbf{F}(\mathbf{X}) \odot g(\mathbf{X})$ . This decouples representation from control, turning residual mixing from a static heuristic into a principled, context-sensitive regulation mechanism.

We instantiate this principle with TRANSPONDER, designed to be minimally invasive and broadly compatible:

1. **Input-aware modulation.** Modulators calculate the scales that multiply functional-path outputs at use time, yielding context-conditioned scaling factors for each component in Transformer.
2. **Granularity and placement.** We systematically study where to modulate (projection- vs. path-level) and at what resolution (scalar vs. channel-wise), providing a comprehensive view of practical design.
3. **Low overhead, easy integration.** TRANSPONDER adds  $\leq 1\%$  parameters and integrates into standard Transformer blocks via a fused Triton kernel for each modulator to accelerate inference.

Across large-scale language-modeling benchmarks, TRANSPONDER delivers consistent gains on LLaMA backbones from 60M to 250M parameters, achieving up to 5.8–15.3% relative perplexity reductions over the LLaMA baselines. Extensive ablations spanning placement, granularity, hidden size, and component coverage (attention and FFN) show the robustness of the design. Analysis of learned modulator values uncover depth-, module-, and token-specific patterns that adapt layer-wise contributions to input semantics, providing direct evidence that residual functional transformations benefit from adaptive, context-aware scaling.

## 2. Related Work

### 2.1. Residual Connections and the Transformer Architecture

Residual connections [9] enable very deep models by adding a learned transformation to an identity shortcut, improving gradient flow, preserving signal, and allowing layers to refine rather than reconstruct representations. The Transformer [1] instantiates this principle with a persistent residual stream that carries token-wise state across layers while each block applies a functional path—multi-head self-attention (MHA) followed by a position-wise feed-forward network (FFN). Concretely, each block computes linear projections for queries, keys, and values and an attention output, then applies two FFN linear maps with nonlinearities; the results are mixed back into the residual stream through additive shortcuts. Layer normalization (LayerNorm) is interleaved with these sublayers: the original Post-LN normalizes after each sublayer [1], whereas Pre-LN normalizes before sublayers, improving optimization stability for deeper stacks [10]. This decomposition—identity carryover plus functional updates—yields strong gradient propagation and composability, but it also implicitly fixes the contribution of each subcomponent to the residual stream at inference, motivating methods that explicitly control (or scale) block updates relative to the identity path.

### 2.2. Scaling the Block: Static Reparameterizations and Normalization Placement

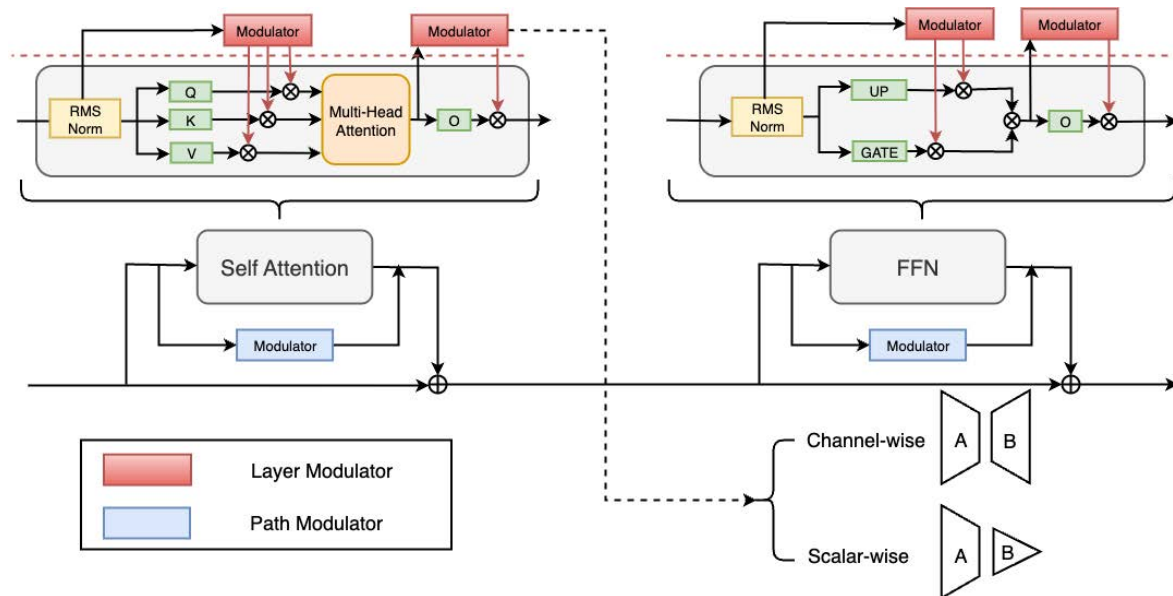
A substantial line of work manipulates the magnitude of each block’s update relative to the residual stream via static, input-agnostic mechanisms. ReZero [5] introduces a zero-initialized, learnable scalar per block that gates the residual branch, effectively starting from an identity mapping and allowing depth to emerge during training. DeepNorm/DeepNet [6] analytically rescales residual and sublayer outputs to maintain stability in very deep Transformers, enabling substantially deeper stacks without divergence. LLaUR [11] generalizes residual/functional mixing with learned coefficients (e.g., RW/LR/PA variants), providing a tunable trade-off between identity flow and the functional path. While these approaches improve training dynamics and depth, their coefficients are typically block-level and input-agnostic at inference time, leaving the strength of updates fixed for any given input.

### 2.3. Context-Conditioned Gating Methods

In convolutional networks, squeeze-and-excitation (SE) [12] introduces feature-wise, input-conditioned scaling by pooling global context and learning channel gates, improving representational efficiency and accuracy. The core idea—context-aware modulation of features—has influenced architectures beyond convolution. Within Transformers, SDPA-gated mechanisms [13] attach the gating mechanism to scaled dot-product attention, commonly parameterized at the head or channels. These gates selectively modulate attention behavior and can stabilize training, but they often incur heavy additional parameters and computation, and primarily target the attention pathway.

## 3. TRANSPONDER - Contextual Modulation Training

We formalize the TRANSPONDER by introducing the lightweight, input-conditioned modulators that scale the functional information flow inside Transformer blocks. The goal is to convert static functional compositions into context-aware modulation while preserving (i) the inductive bias of the base operator (attention/MLP), (ii) training stability, and (iii) negligible parameter/FLOP overhead. We show the main architecture of the TRANSPONDER in Figure 1.



**Figure 1. Illustration of the TRANSPONDER mechanism.** The Figure shows a LLaMA decoder block, consisting of a self-attention module followed by a feed-forward network (FFN). Layer-Modulator (red) dynamically rescales the outputs of key linear transformations (e.g. Q, K, V, O) within both modules, while Path-Modulator (blue) modulates the functional paths directly (e.g. Self-attention, FFN). Each Modulator can operate in two modes: channel-wise control, where the output dimension matches the original feature dimension, and scalar-wise control, where a single scalar is used to regulate the entire function/layer.

### 3.1. Preliminaries: Residual Functional Composition in Transformers

Let  $\mathbf{Y}_{l-1} \in \mathbb{R}^{T \times d}$  denote the hidden states entering layer  $l$  (sequence length  $T$ , width  $d$ ), and let  $\mathcal{F}_l$  denote a parameterized sublayer (self-attention or FFN). The standard residual formulation composes identity and functional streams additively:

$$\mathbf{Y}_l = \mathbf{Y}_{l-1} + \mathcal{F}_l(\mathbf{Y}_{l-1}). \quad (1)$$

While Equation (1) ensures stable signal and gradient propagation, the contribution of  $\mathcal{F}_l$  is typically static at inference time: once trained, the magnitude with which  $\mathcal{F}_l$  perturbs the skip path does not depend on the current token, channel, or context. Prior work improves optimization via block-wise scalars (e.g., ReZero [5]), depth-aware reparameterizations (e.g., DeepNorm [6]), or by mixing identity

and functional streams with learned but input-agnostic coefficients (e.g., LAuReL [11]). Our aim is to endow Equation (1) with contextual modulators applied to, or inside,  $\mathcal{F}_l$ .

### 3.2. Definition: TRANSPONDER Scheme

A *Modulator*  $\mathcal{G}_l$  produces a contextual scaling signal that modulates a sublayer at its point of use. For a generic operator  $\mathcal{F}_l$  acting on  $\mathbf{X}$ ,

$$\mathbf{Y} = \mathcal{F}_l(\mathbf{X}) \odot \mathcal{G}_l(\mathbf{X}), \quad (2)$$

where  $\odot$  denotes broadcasted Hadamard multiplication. In a residual block,

$$\mathbf{Y}_l = \mathbf{Y}_{l-1} + \left[ \mathcal{F}_l(\mathbf{Y}_{l-1}) \odot \mathcal{G}_l(\mathbf{Y}_{l-1}) \right]. \quad (3)$$

Modulator parameterization.

We parameterize the modulator  $\mathcal{G}_l$  with a compact bottleneck applied to each input activation  $\mathbf{X}$ , followed by a bounded nonlinearity:

$$\mathcal{G}_l(\mathbf{X}) = \tilde{\sigma}_{\alpha_l}(\mathbf{B}_l \phi(\mathbf{A}_l \mathbf{X})), \quad \mathbf{A}_l \in \mathbb{R}^{r \times d_{\text{in}}}, \quad \mathbf{B}_l \in \mathbb{R}^{d_{\text{out}} \times r}, \quad (4)$$

where  $d_{\text{in}}$  is the input width of the summary,  $d_{\text{out}}$  matches the modulation resolution (one single scalar or channel-wise).  $\phi$  is a pointwise nonlinearity (e.g., Sigmoid), and  $\tilde{\sigma}_{\alpha}$  is a learnable curvature sigmoid,

$$\sigma_{\alpha}(x) = \frac{1}{1 + \exp(-\alpha x)}, \quad \alpha > 0, \quad \tilde{\sigma}_{\alpha}(x) = 2\sigma_{\alpha}(x). \quad (5)$$

The base gate  $\sigma_{\alpha}$  maps to  $[0, 1]$  and interpolates between soft scaling (small  $\alpha$ ) and near-binary gating (large  $\alpha$ ). We use the calibrated gate  $\tilde{\sigma}_{\alpha} \in [0, 2]$  so that  $\tilde{\sigma}_{\alpha}(0) = 1$ , preserving the expected residual magnitude at initialization.

### 3.3. Granularity: Where and at What Resolution to Modulate?

We distinguish *where* the Modulator is applied (inter-layer) from *how finely* it is parameterized (intra-layer).

Inter-layer granularity.

- **Projection-level modulation.** Apply a modulator after each linear projection inside a functional path:

$$\underbrace{\mathbf{W}\mathbf{X}}_{\text{projection}} \mapsto (\mathbf{W}\mathbf{X}) \odot \mathcal{G}_l(\mathbf{X}). \quad (6)$$

This affords precise control (e.g., on  $Q, K, V$  and MLP up/down projections) with minimal additional computation.

- **Path-level modulation.** Modulate the entire functional path (e.g., self-attention or FFN):

$$\mathbf{Y}_l = \mathbf{Y}_{l-1} + \mathcal{F}_l(\mathbf{Y}_{l-1}) \odot \mathcal{G}_l(\mathbf{Y}_{l-1}). \quad (7)$$

Intra-layer granularity.

Let  $d_{\text{out}}$  be the dimension of  $\mathcal{F}_l$ 's output at the modulation site. We instantiate:

- **Layer-wise (scalar) modulation:**  $\mathcal{G}_l(\mathbf{X}) \in \mathbb{R}$  via  $\mathbf{B}_l \in \mathbb{R}^{1 \times r}$ . A single scale uniformly modulates the sublayer; parameter and compute overheads are negligible.
- **Channel-wise modulation:**  $\mathcal{G}_l(\mathbf{X}) \in \mathbb{R}^{d_{\text{out}}}$  via  $\mathbf{B}_l \in \mathbb{R}^{d_{\text{out}} \times r}$ , enabling feature-specific modulating with two low-rank matrices.



### 3.4. Fused Projection for Projection-Level TRANSPONDER

For each linear layer, the dominant cost is the matrix product  $\mathbf{X}\mathbf{W}_*$ . When the summary  $u_*(\mathbf{X})$  is token-wise, we also need  $\mathbf{X}\mathbf{A}_*^\top$ . Instead of launching two kernels, we concatenate along the output dimension and perform a single tiled GEMM:

$$\begin{bmatrix} \mathbf{Z}_* & \mathbf{U}_* \end{bmatrix} = \mathbf{X} \underbrace{\begin{bmatrix} \mathbf{W}_* & \mathbf{A}_*^\top \end{bmatrix}}_{\mathbf{C}_* \in \mathbb{R}^{d_{\text{in}} \times (d_{\text{out}}+r)}}, \quad \mathbf{Z}_* \in \mathbb{R}^{B \times T \times d_{\text{out}}}, \quad \mathbf{U}_* \in \mathbb{R}^{B \times T \times r}. \quad (8)$$

Here  $\mathbf{A}_* \in \mathbb{R}^{r \times d_{\text{in}}}$  is the rank- $r$  bottleneck,  $\mathbf{W}_* \in \mathbb{R}^{d_{\text{in}} \times d_{\text{out}}}$  is the projection,  $B$  is batch size, and  $T$  is sequence length. After the fused GEMM, we compute the modulator logits and modulate in-register per tile:

$$\text{logits}_* = \mathbf{U}_* \mathbf{B}_*^\top, \quad \mathbf{B}_* \in \mathbb{R}^{d_{\text{out}} \times r}, \quad \gamma_* = \tilde{\sigma}_{\alpha_*}(\text{logits}_*) \in \begin{cases} \mathbb{R}^{B \times T \times 1}, & \text{scalar gate,} \\ \mathbb{R}^{B \times T \times d_{\text{out}}}, & \text{channel-wise gate.} \end{cases} \quad (9)$$

Finally, we apply broadcasted modulation to the projected activations:

$$\hat{\mathbf{Z}}_* = \gamma_* \odot \mathbf{Z}_*. \quad (10)$$

This fusion amortizes the summary computation into the main GEMM and keeps the modulator arithmetic on-chip, yielding precise projection-level control at minimal overhead.

## 4. Experiments

### 4.1. Setup

We evaluate TRANSPONDER on standard language modeling with LLaMA backbones [4] ranging from  $\sim 60\text{M}$  to  $\sim 250\text{M}$  parameters, trained on the C4 [14] and OpenWebText [15] corpus. All models are optimized with Adam and a base learning rate of  $3 \times 10^{-3}$  for Openwebtext dataset and  $1 \times 10^{-3}$  for C4 dataset. All the experiments are trained with bfloat16 for all the activations, weights, and optimization states. The hidden dimension of the contextual modulator is set to  $r = 8$  in the main comparisons to emphasize gains from contextual modulation rather than capacity. For fairness, we keep training data, tokenization, optimizer settings, and learning-rate schedules identical across baselines and TRANSPONDER. We report the perplexity (the lower the better) as the standard language-modeling metrics and measure parameter overhead relative to the corresponding backbone. We initialize  $A, B$  of the Transponder with Kaiming uniform [16] and zero biases; set  $\alpha=1$  initially. Detailed hyperparameter settings can be found in Table A1.

### 4.2. Baselines

We compare against widely used residual/normalization designs and recent residual-scaling or layer-wise contextual control methods:

- **DeepNorm/DeepNet** [6]: rescales the residual branch with depth-dependent constants and tailored initialization to stabilize training of very deep Transformers.
- **LayerNorm Scaling** [8]: multiplies each layer’s normalized output by a depth-aware factor (e.g.,  $\propto 1/\sqrt{L}$ ) to curb variance growth across residual connections.
- **LAuREL** [11]: learns a lightweight mixer that blends the skip and transformed paths (LR/PA variants), improving layer-wise signal routing with minimal extra parameters.
- **SDPA-Gate** [13]: applies a per-head sigmoid gate to the scaled dot-product attention output to modulate contribution and mitigate attention-sink effects.

To match SDPA-Gate’s mechanism more broadly, we also implement an **ALL-Gate** variant that applies the same gating scheme to all linear sublayers, including the MLP projections, providing a stronger control baseline.

#### 4.3. Modulator Variants Validation

We instantiate five different variants of the contextual modulator in the framework of the TRANSPONDER to probe where and at what resolution modulation helps most:

- **Modulator-path-scalar**: a single scalar Modulator modulates the entire functional path (attention output or MLP output) per block.
- **Modulator-path-channel**: a channel-wise Modulator modulates the entire functional path (attention output or MLP output) output channels.
- **Modulator-layer-scalar**: a single scalar Modulator per linear projection.
- **Modulator-layer-channel**: a channel-wise Modulator per linear projection.
- **Modulator-layer-channel-scalar (TRANSPONDER)**: a channel-wise Modulator per linear projection and a single scalar output Modulator per linear projection.

Table 1 shows two consistent trends across model scales. (i) The functional-path modulator does not aid training and can even destabilize optimization, whereas the layer-wise modulator consistently improves both training dynamics and final performance. (ii) With a layer-wise modulator in place, both channel-wise and scalar modulators further benefit TRANSPONDER: the scalar provides coarse, global control, while the channel-wise variant offers finer-grained adjustments. Using both together yields an additional, robust gain with only a modest  $\sim 1\%$  parameter overhead.

**Table 1.** Validation perplexity (PPL, lower is better) of different variants of contextual modulators on Openwebtext. “Param  $\uparrow$ ” denotes the relative parameter increase vs. LLaMA baseline. The best performance is highlighted in bold.

Model	60M		130M		250M	
	PPL	Param $\uparrow$	PPL	Param $\uparrow$	PPL	Param $\uparrow$
Modulator-path-scalar	62.91	$< 1\%$	18.29	$< 1\%$	1088	$< 1\%$
Modulator-path-channel	23.25	$< 1\%$	18.91	$< 1\%$	16.14	$< 1\%$
Modulator-layer-scalar	23.15	$< 1\%$	17.76	$< 1\%$	15.21	$< 1\%$
Modulator-layer-channel	23.08	1%	17.56	1%	15.17	1%
Modulator-layer-channel-scalar	<b>22.50</b>	1%	<b>17.45</b>	1%	<b>14.93</b>	1%

Given its consistently best validation perplexity under comparable budgets and its favorable integration cost, we adopt **Modulator-layer-channel-scalar** as the default TRANSPONDER configuration for all subsequent comparisons against baseline methods.

#### 4.4. Main Results for Language Modeling.

Table 2 compares validation perplexity (PPL; lower is better) across three LLaMA model scales. From the results, TRANSPONDER attains the best PPL wherever reported with only  $\sim 1\%$  parameter overhead. Compared to the LLaMA baseline, TRANSPONDER reduces PPL on OpenWebText by **15.3%** at 60M (26.56 $\rightarrow$ 22.50), **9.4%** at 130M (19.27 $\rightarrow$ 17.45), and **13.6%** at 250M (17.28 $\rightarrow$ 14.93); on C4 the reductions are **5.8%** at 60M (30.31 $\rightarrow$ 28.55), **17.4%** at 130M (26.73 $\rightarrow$ 22.09), and **14.6%** at 250M (21.92 $\rightarrow$ 18.72). These gains persist through 250M under larger token budgets, indicating strong large-scale behavior.

Capacity-heavy gates (SDPA-Gate, ALL-Gate) improve PPL but require 15–80% extra parameters. In contrast, TRANSPONDER matches or exceeds these improvements with only 1% overhead. Among minimal/zero-overhead baselines, LayerNorm Scaling is competitive yet still trails.

Several baselines that attempt to learn the mixing between the residual path and functional blocks exhibit instability at some scales (e.g., DeepNet and LLaMA-PA explode on C4-130M, 143.65 and 106.94 PPL; LLaMA-PA diverges on C4-130M with 1355 PPL and collapses on OpenWebText-250M to 257 PPL), whereas TRANSPONDER delivers consistent improvements across all reported settings.

Under comparable budgets, TRANSPONDER achieves state-of-the-art PPL with minimal overhead, strong large-scale behavior, and improved robustness, establishing it as our default configuration for subsequent comparisons.

**Table 2.** Validation perplexity (PPL, lower is better) on C4 and Openwebtext. “Param ↑” denotes the relative parameter increase vs. LLaMA.

Training tokens		60M		130M		250M	
		1.2B		2.2B		3.9B	
Dataset	Model	PPL	Param ↑	PPL	Param ↑	PPL	Param ↑
Openwebtext	LLaMA (baseline)	26.56	-	19.27	-	17.28	-
	DeepNet	23.78	< 1%	18.74	< 1%	16.53	< 1%
	LAuReL-LR	23.81	< 1%	18.19	< 1%	16.75	< 1%
	LAuReL-PA	23.37	< 1%	18.22	< 1%	257	< 1%
	LayerNorm Scaling	23.31	0%	18.28	0%	16.16	0%
	SDPA-Gate	23.19	15%	18.03	22%	15.44	23%
	ALL-Gate	22.96	44%	17.91	65%	14.97	80%
	TRANSPONDER	<b>22.50</b>	1%	<b>17.45</b>	1%	<b>14.93</b>	1%
C4	LLaMA (baseline)	30.31	-	26.73	-	21.92	-
	DeepNet	30.18	< 1%	143.65	< 1%	21.72	< 1%
	LAuReL-LR	30.05	< 1%	106.94	< 1%	39.14	< 1%
	LAuReL-PA	29.50	< 1%	1355	< 1%	20.88	< 1%
	LayerNorm Scaling	29.77	0%	25.76	0%	20.35	0%
	SDPA-Gate	29.53	15%	23.18	22%	18.95	23%
	ALL-Gate	28.98	44%	22.92	65%	18.82	80%
	TRANSPONDER	<b>28.55</b>	1%	<b>22.09</b>	1%	<b>18.72</b>	1%

#### 4.5. Stability with Post-LayerNorm

Beyond the Pre-LN setting, we evaluate TRANSPONDER under the more delicate Post-LN regime. As shown in Table 3, vanilla LLaMA with Post-LN is unstable at scale: PPL explodes at 250M (1409.79), and even at 130M performance degrades to 26.95 in comparison to the Pre-LN LLaMA. In contrast, TRANSPONDER stabilizes training and improves accuracy in both cases, yielding a **4.6%** PPL reduction at 130M (26.95→25.71) and preventing divergence at 250M, where it achieves **20.28** PPL instead of catastrophic failure. While Post-LN remains slightly behind strong Pre-LN baselines at a similar scale, these results indicate that TRANSPONDER substantially enlarges the viable training regime for Post-LN models with only  $\sim 1\%$  overhead, mitigating the well-known optimization fragility of Post-LN Transformers.

**Table 3.** Validation perplexity of LLaMA 130M and 250M on C4 with Post-Layer Norm.

	130M	250M
LLaMA w. Post-LN	26.95	1409.79
TRANSPONDER w. Post-LN	<b>25.71</b>	<b>20.28</b>

#### 4.6. Ablations and Sensitivity Tests

##### Contextual vs. Non-Contextual.

To assess the contribution of contextual modulation, we compare TRANSPONDER with and without contextual signals. In the non-contextual variant, for each layer  $l$ , we replace the contextual pathway with learnable parameters—a layer-wise scalar  $\beta_1^{(l)}$  and a channel-wise vector  $\beta_2^{(l)}$ —and set

$$\mathcal{G}_l = \tilde{\sigma}_{\alpha_l}(\beta_1^{(l)} \cdot \beta_2^{(l)}),$$

where  $\beta_1^{(l)}$  applies uniform scaling and  $\beta_2^{(l)}$  provides per-channel modulation.



As shown in Table 4, the non-contextual variant already improves over the LLaMA baseline at certain scales, confirming the utility of the proposed modulator design. However, adding contextual control consistently delivers much larger gains across all model sizes, establishing it as the key driver of performance. This result directly supports our initial motivation: each component benefits from input-aware, contextual control rather than relying solely on static modulation.

**Table 4.** Validation perplexity (PPL, lower is better) on Openwebtext comparing non-contextual and contextual variants.

	60M	130M	250M
Original LLaMA	26.56	19.27	17.28
TRANSPONDER w/o contextual control	23.40	19.29	16.29
TRANSPONDER w/o learnable sigmoid	23.56	17.70	<b>14.93</b>
TRANSPONDER	<b>22.50</b>	<b>17.45</b>	<b>14.93</b>

Sigmoid and Learnable Sigmoid.

A key innovation of TRANSPONDER is the introduction of a learnable sigmoid. This learnable non-linear modulation mechanism allows each modulator to dynamically adjust its sensitivity to the input, either amplifying or attenuating modulation strength as needed. As shown in Table 4, incorporating the learnable sigmoid consistently improves (except for 250M that is comparable)perplexity across model sizes, validating our design intuition that adaptive nonlinearity provides a crucial layer of flexibility for effective modulation.

Hidden dimensions.

We further examine the effect of the hidden dimension  $r$  in the contextual modulator. Table 5 shows that even with an extremely compact setting ( $r = 2$ ), TRANSPONDER already delivers large gains over the LLaMA baseline. Scaling up the hidden size from 8 to 32 provides marginal yet consistent improvements. These findings suggest that the modulator does not require a large intermediate capacity to capture input-dependent scaling, underscoring the efficiency of our design: lightweight modulators suffice to model contextual dependencies while adding negligible parameter overhead.

**Table 5.** Ablation study on the hidden dimension  $r$  of both the channel-wise and scalar-wise modulators for the TRANSPONDER on C4 dataset. The reported metric is Perplexity. The best performance is highlighted in bold.

Hidden Size	LLaMA60M	LLaMA130M
Baseline	30.31	26.07
2	28.91	23.01
4	28.65	22.64
8	28.55	22.09
16	<b>28.41</b>	22.12
32	28.73	<b>22.07</b>

Contribution of functional components.

We next investigate which functional modules benefit most from modulation by selectively applying the contextual modulators to individual subcomponents (Table 6). Applying modulators to either the self-attention or MLP block alone yields clear improvements, while restricting modulation to partial components (e.g., only\_first, only\_last, only\_qk) provides only limited gains.

**Table 6.** Ablation study on modulators on different functional modules for the TRANSPONDER on C4 dataset. The reported metric is Perplexity. The best performance is highlighted in bold.

Modules	LLaMA60M	LLaMA130M
Baseline	30.31	26.07
self-attention	29.04	23.21
mlp	29.43	23.36
only_last	29.88	23.97
only_first	29.91	24.11
only_qk	30.01	24.03
w/o up and gate	29.12	22.93
all	<b>28.55</b>	<b>22.09</b>

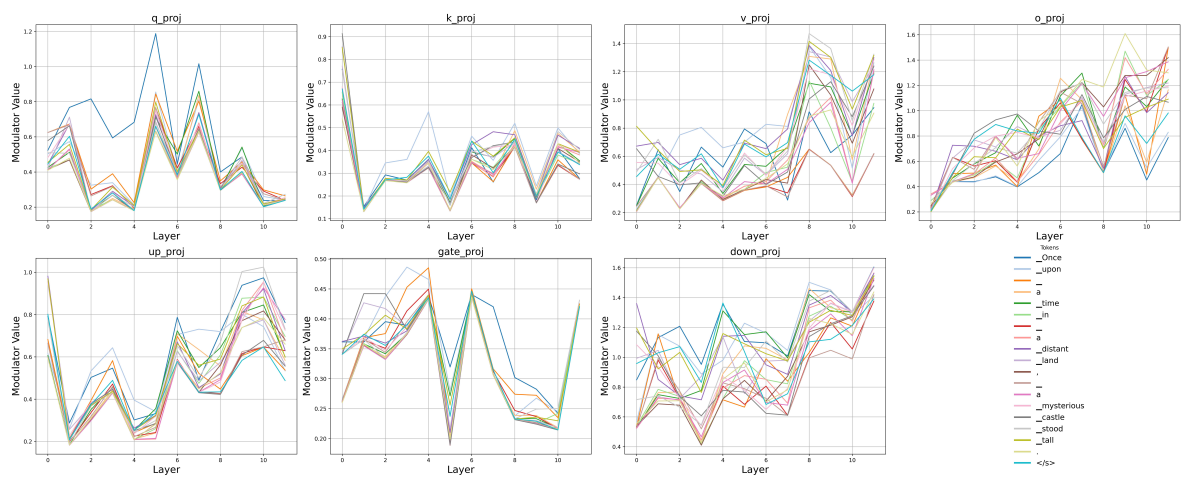
We further evaluate the effect of removing the up- and gate-projections. In the original LLaMA design, the gate projection serves as a channel-wise modulator to the up-projection, though implemented as a full-rank matrix. Interestingly, excluding these projections leads to strong performance, second only to the full “all” configuration. This suggests that the gate-projections already play a key role in channel-wise contextual modulation.

Overall, the best results are consistently achieved when all components are modulated jointly (“all”), yielding the lowest PPL on both LLaMA-60M and LLaMA-130M. This confirms that full-path functional modulation is necessary to provide comprehensive control over residual transformations, enabling robust and consistent improvements across scales.

5. Contextual Modulator Values Analysis

The results in Section 4 suggest that, beyond static architectural scaling (e.g., fixed layer-normalization schemes), context-/input-sensitive modulation is a powerful way to improve optimization and expressivity in deep language models. TRANSPONDER learns, for each token and sub-layer (self-attention and FFN), both a layer scalar and channel-wise modulation. The question is whether these modulators truly encode context rather than acting as static rescalers.

Figure 2 (layer-scalar, per module) shows clear depth- and module-specific structure with token-dependent variation. The traces for different tokens diverge at many layers, indicating context sensitivity. We observe systematic trends across modules: o\_proj and down\_proj steadily strengthen with depth, consistent with greater late-layer amplification; v\_proj exhibits a pronounced surge in upper layers; up\_proj follows a U-shaped pattern (early attenuation, late amplification); gate\_proj remains in a tighter band with a mid-depth peak; and q\_proj / k\_proj show early suppression followed by recovery.



**Figure 2.** Token-wise modulator values across layers and modules in a LLaMA-130M model trained with TRANSPONDER. We report the values of the scalar modulator as the representative. Each subplot corresponds to a specific linear projection of the decoder layer. The plotted curves represent different tokens from the input sequence.

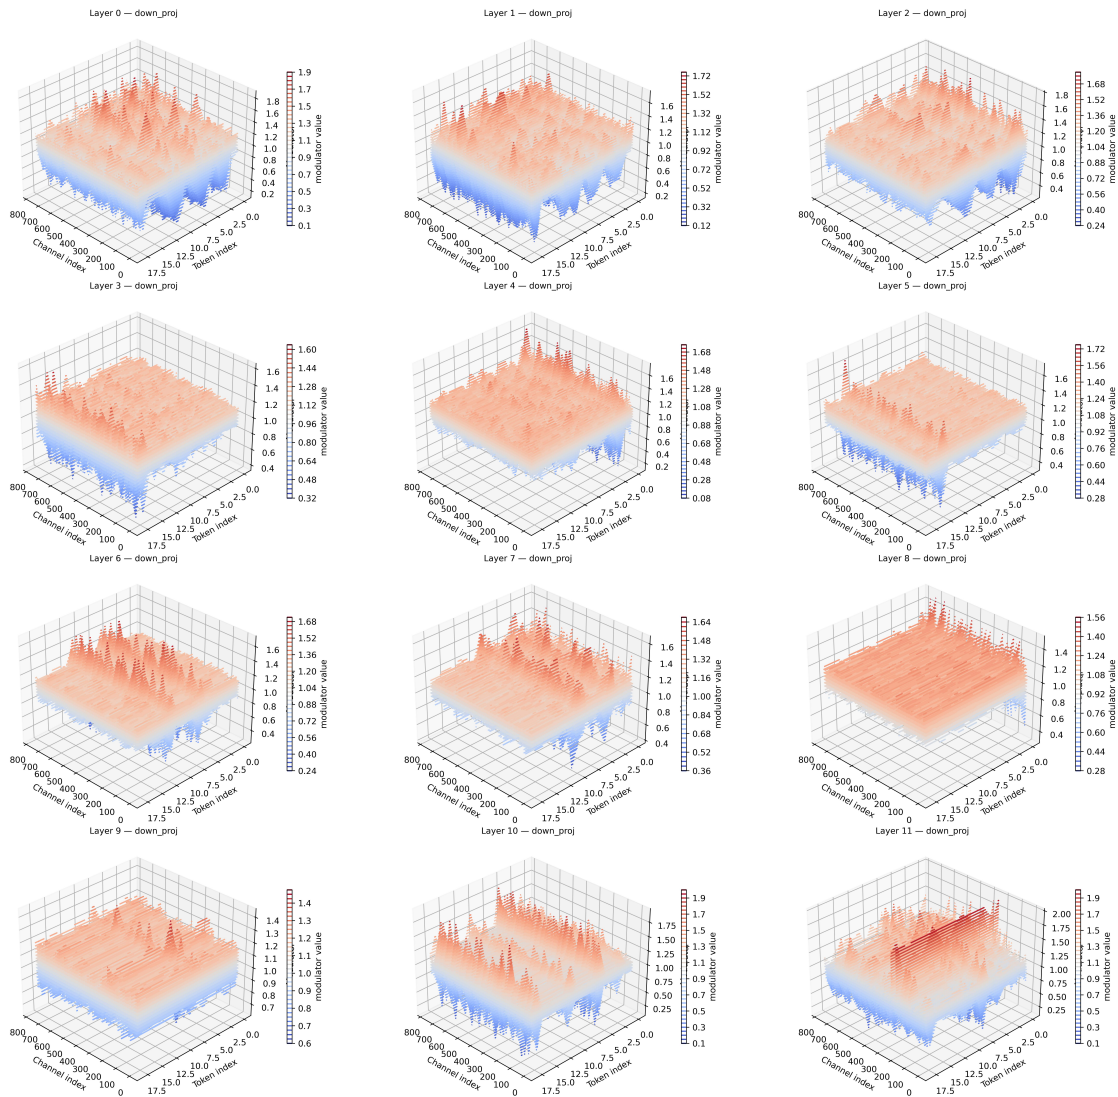
Figure A1 further confirms the effectiveness of the channel-wise modulation. Early layers are relatively flat, but mid–late layers develop structured ridges over both tokens and channels, with increasing contrast across depth. In particular, we observe a similar trend at some specific layers. For instance, at the 5th layer, one token exhibits very large modulator values across all channels, whereas in the 11th layer, one channel displays very large modulation values across the tokens.

Taken together, the two views indicate that TRANSPONDER learns structured, context-aware gains that evolve with depth and specialize by module and channel—precisely the behavior we would expect if the modulators were encoding useful contextual signals rather than acting as fixed, global scalars.

6. Conclusions

In this article, we proposed TRANSPONDER, a simple but effective way to modulate the sub-layers in Transformers *context-aware*. By pairing key projections and paths with compact modulators that produce scalar- and channel-wise gates, TRANSPONDER explicitly controls each component’s contribution at inference, separating representation from control. The design is minimally invasive— $\leq 1\%$  parameter overhead—and broadly compatible with standard decoder blocks. Empirically, TRANSPONDER consistently improves language modeling quality across LLaMA backbones (60M–250M) on OpenWebText and C4. Under comparable budgets, it surpasses stronger, capacity-heavy gating schemes while remaining far lighter, and it outperforms static depth-aware scaling methods. Ablations map a practical design space. We find (i) layer-wise placement is preferable to coarse functional-path-only control; (ii) combining channel-wise modulator with a single scalar modulator per layer yields the best accuracy-cost trade-off; (iii) modulating all the layers is crucial. Modulator value analyses further show structured, token- and depth-dependent behaviors, providing direct evidence that adaptive scaling learns semantically meaningful control. The limitation of this work can be found in Appendix B.

Appendix A.



**Figure A1.** Channel-wise modulator values on down\_proj across layers and tokens in a LLaMA-130M model trained with TRANSPONDER. The X and Y are the channel index and the token index, while the Z-axis are the channel-wise modulator values.

**Table A1.** Hyperparameters of LLaMA-60M, LLaMA-130M, and LLaMA-250M on OpenWebText and C4.

Hyper-parameter	LLaMA-60M	LLaMA-130M	LLaMA-250M
Embedding Dimension	512	768	1024
Feed-forward Dimension	1376	2048	2560
Global Batch Size	512	512	512
Sequence Length	256	256	256
Training Steps	10000	20000	40000
Learning Rate	3e-3	3e-3 (1e-3 for C4)	3e-3 (1e-3 for C4)
Warmup Steps	1000	2000	4000
Learning rate Decay Method	noam	noam	noam
Optimizer	Adam	Adam	Adam
Layer Number	8	12	24
Head Number	8	12	16

## Appendix B. Limitation

Our study is compute-limited: we did not scale beyond  $\sim 1\text{B}$  parameters or train for insufficient tokens because of the limits of the computational resources. In this regime, several zero-shot classification benchmarks yield Matthews correlation coefficient (MCC) scores near 0, which we attribute to undertraining rather than an inherent limitation of the current investigation of TRANSPONDER. Future work should evaluate TRANSPONDER at larger scales and with substantially longer training runs (more steps and tokens), and reassess zero-shot (and few-shot) generalization under those settings.

## Appendix C. Claim of the LLM usage

We used LLM-based tools to improve the language and flow; the principles, core logic, and innovations are entirely the authors'.

## Appendix D. Reproduction Statement

All experiments were conducted using the Hugging Face pretraining framework with data parallelism over  $8 \times$  NVIDIA A100 (40 GB) GPUs. For DeepNorm and LayerNorm scaling, we report the results from [8], as we adopt identical hyperparameter settings. To ensure reproducibility, we include our code and step-by-step instructions in the supplementary materials. Because LAuReL has not released its source code, we reimplemented the method based on our understanding of the paper. Detailed training hyperparameters are provided in Table A1.

## References

1. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Advances in neural information processing systems* **2017**, *30*.
2. Devlin, J.; Chang, M.W.; Lee, K.; Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, 2019, [arXiv:cs.CL/1810.04805].
3. Radford, A.; Narasimhan, K.; Salimans, T.; Sutskever, I. Improving Language Understanding by Generative Pre-Training. [https://cdn.openai.com/research-covers/language-unsupervised/language\\_understanding\\_paper.pdf](https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf), 2018.
4. 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**.
5. Bachlechner, T.; Majumder, B.P.; Mao, H.; Cottrell, G.; McAuley, J. Rezero is all you need: Fast convergence at large depth. In *Proceedings of the Uncertainty in Artificial Intelligence*. PMLR, 2021, pp. 1352–1361.
6. Wang, H.; Ma, S.; Dong, L.; Huang, S.; Zhang, D.; Wei, F. Deepnet: Scaling transformers to 1,000 layers. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **2024**.
7. Li, P.; Yin, L.; Liu, S. Mix-In: Unleashing the power of deeper layers by combining pre-In and post-In. *arXiv preprint arXiv:2412.13795* **2024**.
8. Sun, W.; Song, X.; Li, P.; Yin, L.; Zheng, Y.; Liu, S. The Curse of Depth in Large Language Models. *arXiv preprint arXiv:2502.05795* **2025**.
9. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
10. Xiong, R.; Yang, Y.; He, D.; Zheng, K.; Zheng, S.; Xing, C.; Zhang, H.; Lan, Y.; Wang, L.; Liu, T.Y. On Layer Normalization in the Transformer Architecture, 2020, [arXiv:cs.LG/2002.04745].
11. Menghani, G.; Kumar, R.; Kumar, S. LAUREL: Learned Augmented Residual Layer. *arXiv preprint arXiv:2411.07501* **2024**.
12. Hu, J.; Shen, L.; Albanie, S.; Sun, G.; Wu, E. Squeeze-and-Excitation Networks, 2019, [arXiv:cs.CV/1709.01507].
13. Qiu, Z.; Wang, Z.; Zheng, B.; Huang, Z.; Wen, K.; Yang, S.; Men, R.; Yu, L.; Huang, F.; Huang, S.; et al. Gated Attention for Large Language Models: Non-linearity, Sparsity, and Attention-Sink-Free, 2025, [arXiv:cs.CL/2505.06708].
14. Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; Liu, P.J. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of Machine Learning Research* **2020**, *21*, 1–67.



15. Gokaslan, A.; Cohen, V. OpenWebText Corpus. <http://Skylion007.github.io/OpenWebTextCorpus>, 2019.
16. He, K.; Zhang, X.; Ren, S.; Sun, J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification, 2015, [[arXiv:cs.CV/1502.01852](https://arxiv.org/abs/1502.01852)].

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.