

# RESIDUAL: TRANSFORMER WITH DUAL RESIDUAL CONNECTIONS

Anonymous authors

Paper under double-blind review

## ABSTRACT

Transformer networks have become the preferred architecture for many tasks due to their state-of-the-art performance. However, the optimal way to implement residual connections in Transformer, which are essential for effective training, is still debated. Two widely used variants are the Post-Layer-Normalization (Post-LN) and Pre-Layer-Normalization (Pre-LN) Transformers, which apply layer normalization after each residual block’s output or before each residual block’s input, respectively. While both variants enjoy their advantages, they also suffer from severe limitations: Post-LN causes gradient vanishing issue that hinders training deep Transformers, and Pre-LN causes representation collapse issue that limits model capacity. In this paper, we propose ResiDual, a novel Transformer architecture with Pre-Post-LN (PPLN), which fuses the connections in Post-LN and Pre-LN together, and inherits their advantages while avoids their limitations. We conduct both theoretical analyses and empirical experiments to verify the effectiveness of ResiDual. Theoretically, we prove that ResiDual has a lower bound on the gradient to avoid the vanishing issue due to the residual connection from Pre-LN. Moreover, ResiDual also has diverse model representations to avoid the collapse issue due to the residual connection from Post-LN. Empirically, ResiDual outperforms both Post-LN and Pre-LN on several machine translation benchmarks across different network depths and data sizes.

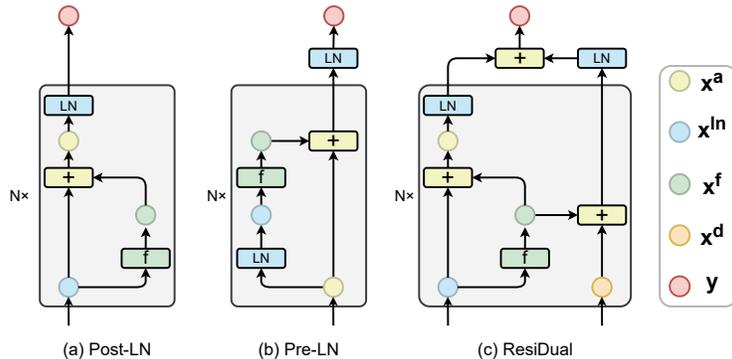


Figure 1: Overview of Post-LN, Pre-LN, and ResiDual. Circles with different colors represent different variables and rectangles represent different operations. See Section 2 for more details.

## 1 INTRODUCTION

Transformer (Vaswani et al., 2017) has emerged as a powerful neural network architecture that has been successfully applied in various AI tasks, including machine translation (Vaswani et al., 2017), language modeling and generation (Radford et al., 2018; 2019; Brown et al., 2020), image recognition (Dosovitskiy et al., 2020), and speech synthesis (Ren et al., 2019). Despite its success, researchers are still exploring ways to further enhance its performance and deepen the understanding of its inner workings (Wang et al., 2019; Katharopoulos et al., 2020; Fedus et al., 2021). Among

them, one area of ongoing research is the study of residual connections in the Transformer architecture (Liu et al., 2020; Xiong et al., 2020; Bachlechner et al., 2021). Two variants of residual connections have been proposed since the introduction of the Transformer, known as Post-LN and Pre-LN. The Post-LN variant applies layer normalization (LN) operations after the output of each residual block. This variant is used in several prominent models such as BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), and ALBERT (Lan et al., 2019). The Pre-LN variant, on the other hand, applies LN operations before the input to each residual block. This variant is used in models such as the GPT series, ViT (Dosovitskiy et al., 2020), and PaLM (Chowdhery et al., 2022).

Although both variants have been widely used, each one has its own drawbacks, which are summarized in Table 1. As shown in Figure 1, the key difference between the two residual variants is how the layer normalization (LN) normalized the outputs of each block. With Post-LN, the output of lower blocks (i.e., the blocks close to input) are normalized multiple times. As a result, the gradient norm decays exponentially with depth and eventually vanishes in the lower layers (Xiong et al., 2020). This problem does not exist in Pre-LN because the gradient can flow directly to each block. However, the Pre-LN architecture has the representation collapse issue (Liu et al., 2020), which will negatively impact the model’s capacity. The representation collapse issue refers to the fact that the hidden representation of higher blocks (i.e., the blocks close to output) will be similar to each other in Pre-LN models. Therefore, the higher blocks will have little contribution to the model capacity.

Several approaches have been proposed to address these problems, which can generally be categorized into three categories. Firstly, some methods aim to modify the architecture, such as DLCL (Wang et al., 2019), NormFormer (Shleifer et al., 2021), RealFormer He et al. (2021), and B2T (Takase et al., 2022), which adds extra components such as aggregations or LNs to stable training. Secondly, some methods add different weights to the residual, such as Admin (Liu et al., 2020), DeepNet (Wang et al., 2022a),  $\tau$ -ResNet (Zhang et al., 2022), and ReZero (Bachlechner et al., 2021). Lastly, some methods use better initialization, such as T-Fixup (Huang et al., 2020), DeepNet (Wang et al., 2022a), and Foundation Transformer (Wang et al., 2022b), to reduce variance and stabilize training.

In this study, we focus on the first category and propose a new architecture for Transformer models to address the drawbacks of both variants while retaining their benefits. Figure 1(c) provides an overview of our method. Our design goal is to maintain the advantages of both variants and avoid their disadvantages by employing two residual connections. In particular, our ResiDual model utilizes a Pre-Post-LN (PPLN) that consists two residuals: one is similar to the Pre-LN to prevent the gradient vanishing issue, while the other one akin to the Post-LN, which sustains representation diversity to avoid the representation collapse issue.

To validate the effectiveness of our proposed method, we conduct both theoretical analysis (Section 3) and empirical study (Section 4) to show that our method can achieve the best of both worlds. From the theoretical perspective, we first show that the gradient vanishing is still a critical problem even using Adam (Kingma & Ba, 2014) optimizer. We also show that ResiDual has a bounded gradient-norm thus do not have such an issue. Furthermore, we study the representation collapse issue and show that ResiDual has the same hidden representation diversity as Post-LN. Therefore, ResiDual do not have the representation collapse issue in Pre-LN.

Empirically, we conduct comprehensive experiments on machine translation tasks, which are among the most representative tasks in natural language processing. Our dataset comprises small-scale (IWLST), mid-scale (WMT), and large-scale (OPUS) datasets. Our experimental results demonstrate that our method outperforms baselines across all three datasets.

In summary, this work makes the following contributions:

Method	Gradient Vanishing	Representation Collapse
Post-LN	☹️	☺️
Pre-LN	☺️	☹️
ResiDual	☺️	☺️

Table 1: Comparison of Post-LN, Pre-LN, and our method. ☺️ means the model does not suffers from the issue and ☹️ means the model has such issue.

- We present ResiDual, a simple yet potent variation of the Transformer architecture, which tackles both the gradient vanishing problem in Post-LN and the representation collapse issue in Pre-LN Transformer models.
- Our theoretical analysis demonstrates that this new design can leverage the strengths of both variants while avoiding their weaknesses.
- Our experimental results provide further evidence of the effectiveness of our approach, as it achieves superior performance compared to both the Post-LN and Pre-LN Transformer models across multiple datasets.

## 2 METHOD

### 2.1 DISADVANTAGES OF POST-LN AND PRE-LN

In this section, we briefly review the architecture of Post-LN and Pre-LN, whose illustrations are available in Figure 1 (a) and (b). We will also discuss the shortcomings of each architecture.

**Gradient Vanishes with Post-LN.** The Post-LN architecture is shown in Figure 1 (a). To be more specific, given a Post-LN Transformer network with  $N$  residual blocks, we assume the input shape is  $n \times d$  where the  $n, d$  denotes the sequence length and embedding size<sup>1</sup>. The variables with vector arrow (e.g.,  $\vec{x} \in \mathbb{R}^{n \times d}$ ) denote the whole sequence and the variables without it (e.g.,  $x \in \mathbb{R}^d$ ) denote an element of the sequence. We use  $\vec{x}^a \in \mathbb{R}^{n \times d}$  denote the tensor after **add** operation and use subscript  $k$  (i.e.  $\vec{x}_k^a$ ) denote the tensor in the  $k$ -th block. We also use  $\vec{x}_k^{ln} \in \mathbb{R}^{n \times d}$  denotes the normalized tensor and  $\vec{x}_k^f \in \mathbb{R}^{n \times d}$  denotes the output of the function  $f_k(\cdot; w_k)$  in the  $k$ -th block. The  $f_k$  can be a self-attention, cross-attention, or feed-forward with parameter  $w_k$ . Using these notations, the Post-LN computation of each element in the  $k$ -th block is

$$x_k^a = x_k^{ln} + x_k^f = x_k^{ln} + f_k(\vec{x}_k^{ln}; w_k); \quad x_{k+1}^{ln} = \text{LN}(x_k^a).$$

Finally, the output  $y$  is computed by  $y = x_{N+1}^{ln} = \text{LN}(x_N^a)$ . Intuitively, the  $x_k^f$  is normalized  $N - k$  times, so does the gradients of  $w_k$ . Therefore, the gradients of lower blocks will be small. From Xiong et al. (2020), we know that for Post-LN Transformer, the gradient norm decreases exponentially from deep layers to shallow layers. Intuitively, such an imbalanced gradients will impede the model training. Therefore, in practise, training tricks such as learning-rate warm-up are necessary to train a Post-LN model.

**Representation Collapses with Pre-LN.** With the same notations, the Pre-LN computation is

$$x_k^{ln} = \text{LN}(x_k^a); \quad x_{k+1}^a = x_k^a + x_k^f = x_k^a + f_k(\vec{x}_k^{ln}; w_k).$$

Similarly, the model output is  $y = \text{LN}(x_{N+1}^a) = \text{LN}(\sum_{k=1}^N x_k^f)$ . Intuitively, as the  $x_k^f$  is only normalized once when computing the  $y$ , neither the forward nor the backward pass are blocked by LN. Thus, Pre-LN do not have the gradient vanish issue. However, it has another issue called

representation collapse. More specifically, Liu et al. (2020) show that the  $\frac{\sqrt{\text{Var}[x_k^f]}}{\sqrt{\text{Var}[x_k^a + x_k^f]}}$  is likely to be smaller for higher blocks (i.e, blocks with larger  $k$ ). This means the output of the later blocks ( $x_k^f$ ) has little contribution to the total variance of  $x_k^a$ . In Section 3.2, we show that the difference between  $x_{k+1}^{ln}$  and  $x_k^{ln}$  (i.e.,  $|x_{k+1}^{ln} - x_k^{ln}|$ ) decays along with  $k$ , which indicates the input of the higher blocks will collapse to similar values. We also show that this issue may limit the capacity of the model.

### 2.2 RESIDUAL

The goal of our model is to take the advantages of both variants and avoid the both disadvantages. To achieve this goal, we use residuals from both variants and the overview of our method is in

<sup>1</sup>We omit the batch dimension that will not affect our analysis.

Figure 1 (c). More specifically, the two residual connections are illustrated in the left and right vertical lines in the Figure. The left one, which is similar to the conventional Post-LN, is

$$\mathbf{x}_k^a = \mathbf{x}_k^{ln} + \mathbf{x}_k^f = \mathbf{x}_k^{ln} + f_k(\vec{\mathbf{x}}_k^{ln}; \mathbf{w}_k); \quad \mathbf{x}_{k+1}^{ln} = \text{LN}(\mathbf{x}_k^a).$$

Meanwhile, the right residual, which is similar to the conventional Pre-LN, is formulated by

$$\mathbf{x}_{k+1}^d = \mathbf{x}_k^d + \mathbf{x}_k^f,$$

where  $\mathbf{x}^d \in \mathbb{R}^{n \times d}$  is the tensor to denote **dual** residual that similar to  $\mathbf{x}^a$  in the Pre-LN that allows the gradients directly flow to each block.

Finally, the output  $y$  is computed by adding the representation of both residuals, which is

$$\mathbf{y} = \mathbf{x}_{N+1}^{ln} + \text{LN}(\mathbf{x}_{N+1}^d).$$

### 2.3 DISCUSSION

In this section, we will only introduce the intuitive understanding of ResiDual and the mathematical analysis is provided in Section 3.

**Avoiding the Gradient Vanishing** In ResiDual, gradient of each block flows from both residual connections. Thus, even if the gradient comes from the Post-LN-like residual vanishes, there will still be gradients from the Pre-LN-like residual. This prevents the gradient vanishing issue. We provide the details of the lower-bound of the gradient norm in Section 3.1.

**Avoiding the Representation Collapse** Our Pre-LN-like residual only affects the model output and does not affect the input to each block. Therefore, the representation capacity is the same as a Post-LN model. Furthermore, because the final output of our model is the sum of two residual connections, the representation of the output will not collapse either. We provide the details of the lower-bound of the representation capacity in Section 3.2.

## 3 THEORETICAL ANALYSIS OF RESIDUAL

In this section, we formally study the gradient vanishing and representation collapse issue. We also prove that our method does not have such issues.

### 3.1 THE GRADIENT VANISHING ISSUE

In order to present the analysis in a concise way, we study a simple setting and make several assumptions. In Transformer, the  $f$  function can be either a feed-forward block or a multi-head attention block. For a feed-forward block,  $f(\mathbf{x}) := \mathbf{W}\mathbf{x}$  where we ignore the layer index. For a multi-head attention block, we have weight matrices  $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V$ . For simplicity, we focus on single-head attention. Similar to Xiong et al. (2020), we initialize  $\mathbf{W}_Q$  to be zero matrices and consequently, the attention is a uniform distribution at initialization and  $f(\mathbf{x}^{(i)}) := \frac{1}{n} \sum_{j=1}^n \mathbf{x}^{(j)} \mathbf{W}_V$  where we drop the layer index and  $\mathbf{x}^{(j)}, j \in [n]$  are the input sequence with length  $n$ . We usually drop the superscript index  $^{(j)}$  for notation simplicity when the context is clear itself. We introduce  $\vec{\mathbf{x}} := \{\mathbf{x}^{(j)}, j \in [n]\}$  and use  $\mathbf{w}$  to denote the collection of parameter matrices in  $f$ .

Based on above assumption, without loss of generality, we further assume that the  $f$  function keeps the norm, i.e.,  $\|f(\mathbf{x})\| = \|\mathbf{x}\|$ . This assumption is asymptotically true when the network width goes to infinity and the initialization variance is properly scaled. We assume that the signal is standardized after layer normalization, i.e.,  $\|\mathbf{x}_k^{ln}\| = \sqrt{d}$  for all  $k \in [N]$ , and that for  $\mathbf{x} \in \mathbb{R}^d$ , the Jacobian matrix through LN satisfies  $\frac{\partial \text{LN}(\mathbf{x})}{\partial \mathbf{x}} \approx \frac{\sqrt{d}}{\|\mathbf{x}\|_2} \mathbf{I}$ . This approximation can be achieved if the mean of  $\mathbf{x}$  is 0 and the variance is  $\frac{1}{d} \|\mathbf{x}\|^2$  while ignoring the gradient back-propagated through mean and variance. The rationale in this assumption is that the error signal (gradients) back-propagating through LN becomes smaller as the norm of the input to the LN gets larger. In the Post-LN Transformer, the scale of the inputs to the layer normalization is independent of  $N$ , and thus the gradients of parameters in the last layer are independent of  $N$ .

**Gradient Norm Estimation for Post and Pre-LN Transformer.** From Xiong et al. (2020), we know that for Post-LN Transformer, the gradient norm of the block  $k$  decreases exponentially as block index  $k$  gets smaller. This indicates that the gradient of the block close to input would be exponentially small for deep transformers. In contrast, for Pre-LN Transformer, the gradient norm of each block is roughly independent with the block index  $k$ .

For completeness, we rephrase the result from Xiong et al. (2020) with our notations and assumptions. We also present the proof in a more accurate way in Appendix.

**Theorem 3.1** (Gradients of the  $k$ -th block in the Post-LN and Pre-LN Transformers). *Given the above assumptions on  $f$  and LN, for the Post-LN Transformer with  $N$  blocks, the gradient of the parameters of the  $k$ -th block satisfies*

$$\left\| \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right\|_F \approx \mathcal{O} \left( (1/2)^{(N-k)/2} e^{\sqrt{N-k}} \right), \quad (1)$$

for the Pre-LN Transformer with  $N$  blocks, the gradient of the parameters of the  $k$ -th block satisfies

$$\left\| \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right\|_F \approx \mathcal{O} \left( \sqrt{\frac{\log(N-k)}{N}} \right), \quad (2)$$

where we ignore the terms irrelevant with  $k, N$ .

**Analysis of Adam** In practice, adaptive optimizers such as Adam are widely used to train Transformer networks. However, the vanished gradients issue cannot be solved by adaptive optimizers and thus we aim to fix the issue in the network architecture. More specifically, we show that the Adam updates is *ill-conditioned* in vanished gradients. More specifically, let the  $\alpha, t, \epsilon, \beta_1, \beta_2$  denote the learning rate, step, smoothing factor, first decay rate and second decay rate, respectively, and the  $\mathbf{w}^{(t)}, \mathbf{g}, \hat{\mathbf{m}}^{(t)}, \hat{\mathbf{v}}^{(t)}$  denote the parameters, gradients, bias-corrected first and second moment estimation at time  $t$ . Meanwhile, we use  $\mathbf{u}(\mathbf{g}^{(t)}) = \alpha \cdot \hat{\mathbf{m}}^{(t)} / (\sqrt{\hat{\mathbf{v}}^{(t)}} + \epsilon)$  denote the Adam update (i.e.,  $\mathbf{w}^{(t)} \leftarrow \mathbf{w}^{(t-1)} - \mathbf{u}(\mathbf{g}^{(t)})$ ) and the full formula is in Appendix B. Because the Adam update is element-wise, we also use  $u(g)$  to denote the scalar function of  $\mathbf{u}(\mathbf{g})$ , which means  $\mathbf{u}(\mathbf{g}) = [u(g_1), u(g_2), \dots, u(g_d)]$ . Then, we will show that, when the gradients vanish, the  $\mathbf{u}(\mathbf{g})$  is sensitive to small perturbation (i.e., ill-conditioned) because of its large condition number.

**Theorem 3.2.** *The Adam update function  $\mathbf{u}(\mathbf{g})$  is ill-conditioned for vanished gradients ( $\mathbf{g} = 0$ ) in early stage ( $t$  is small).*

*Proof.* Considering that the  $\mathbf{u}(\mathbf{g})$  is differentiable, the absolute condition number  $\hat{\kappa}$  for  $\mathbf{u}(\mathbf{g}_t)$  is

$$\hat{\kappa} = \lim_{\delta \rightarrow 0} \sup_{\|\delta \mathbf{g}\| \leq \delta} \frac{\|\mathbf{u}(\mathbf{g} + \delta \mathbf{g}) - \mathbf{u}(\mathbf{g})\|}{\|\delta \mathbf{g}\|} = \|\mathbf{J}(\mathbf{g})\| = \sqrt{\sum_{i=1}^d \left( \frac{\partial u}{\partial g_i} \right)^2}.$$

The full expression of  $\frac{\partial u}{\partial g}$  can be found in Appendix B. In the early stage (i.e.,  $t$  is small), for the vanished gradient ( $g_i = 0$ ), the absolute condition number  $\hat{\kappa}$  is

$$\hat{\kappa} = \alpha \frac{1 - \beta_1}{1 - \beta_1^t} \sqrt{\sum_{i=1}^d \frac{1}{\epsilon + \sqrt{\frac{\beta_2 v_i^{(t-1)}}{1 - \beta_2^t}}}} \approx \frac{\alpha \sqrt{d}}{\epsilon}. \quad (3)$$

For example, in a classic setting where  $d = 1024, \epsilon = 10^{-6}, \alpha = 10^{-4}$ , we have  $\hat{\kappa} = 3200$ , which is a very large number. This tells us that in early stage, the  $\mathbf{u}(\mathbf{g}_t)$  is ill-conditioned.  $\square$

Intuitively, when there is a small noise  $\|\delta \mathbf{g}\| \leq \delta$  added to the gradient  $\mathbf{g}$ , the change of the update  $\|\mathbf{u}(\mathbf{g} + \delta \mathbf{g}) - \mathbf{u}(\mathbf{g})\|$  could be thousand times larger than  $\|\delta \mathbf{g}\|$ . This will make the training

unstable and vulnerable to a small perturbation. This study is also consistent with the empirically findings by Wang et al. (2022a) that the exploding gradients in higher layers is not the root cause of Post-LN training difficulty. Further more, to verify our approximation, we also have simulation in Appendix B.

Moreover, from Equation (3), given a fixed model with width  $d$ , seems there are two possible way to reduce the  $\hat{\kappa}$ : increasing the  $\epsilon$  or decreasing the  $\alpha$ . However, the first one is not viable because a large  $\epsilon$  will make an adaptive optimizer less adaptive. Therefore, in practise, researchers have to reduce the learning-rate  $\alpha$  (e.g., using learning-rate warm-up) to ease this problem.

To conclude, as the gradient vanishing is a critical issue even when the model is trained with adaptive optimizers. As a result, we purpose to solve this problem from the architecture aspect.

### 3.2 THE REPRESENTATION COLLAPSE ISSUE

**The Representation Collapse in Pre-LN** The issue with the representation capability of Pre-LN was initially observed by Liu et al. (2020). In summary, the Pre-LN Transformer’s hidden representation cannot be refined by deeper layers due to the normalization of layer outputs. In this work, we propose a novel analysis approach that directly examines the distribution of hidden state changes, represented by  $|\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln}|$ , and output changes, denoted by  $|\mathbf{y}_N - \mathbf{y}_{N-1}|$ . Our new method offers a straightforward way to obtain quantitative results regarding the convergence rate.

**Theorem 3.3.** *For Pre-LN, assume  $\mathbf{x}_k^f \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$  independently for all  $k \in [N]$ , we have  $\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln} \sim \mathcal{N}(0, \omega_k^2) \mathbf{I}$  where  $\omega_k^2 = \frac{2}{\sqrt{k}(\sqrt{k-1} + \sqrt{k})}$ .*

*Proof.* As  $\mathbf{x}_k^f \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ , we have  $\mathbf{x}_k^a = \sum_{j=1}^{k-1} \mathbf{x}_j^f$  thus  $\mathbf{x}_k^a \sim \mathcal{N}(0, (k-1)\sigma^2 \mathbf{I})$ . For the normalization layer, we approximate its effect as follows,  $\mathbf{x}_k^{ln} = \frac{\mathbf{x}_k^a}{\sqrt{k-1}\sigma}$ . Then we have

$$\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln} = \frac{\mathbf{x}_{k+1}^a}{\sqrt{k}\sigma} - \frac{\mathbf{x}_k^a}{\sqrt{k-1}\sigma} = \frac{\sqrt{k-1} - \sqrt{k}}{\sqrt{k}(k-1)\sigma} \cdot \mathbf{x}_k^a + \frac{1}{\sqrt{k}\sigma} \cdot \mathbf{x}_k^f.$$

We know that  $\frac{\sqrt{k-1} - \sqrt{k}}{\sqrt{k}(k-1)\sigma} \cdot \mathbf{x}_k^a \sim \mathcal{N}(0, \frac{(\sqrt{k-1} - \sqrt{k})^2}{k} \mathbf{I})$  and  $\frac{1}{\sqrt{k}\sigma} \cdot \mathbf{x}_k^f \sim \mathcal{N}(0, \frac{1}{k} \mathbf{I})$ . Because  $\mathbf{x}_k^a$  and  $\mathbf{x}_k^f$  are independent, we have  $\mathbf{a}_{k+1} - \mathbf{a}_k \sim \mathcal{N}(0, \omega_k^2 \mathbf{I})$  and  $\omega_k^2 = \frac{(\sqrt{k-1} - \sqrt{k})^2}{k} + \frac{1}{k} = \frac{2}{\sqrt{k}(\sqrt{k-1} + \sqrt{k})}$ .  $\square$

**Corollary 3.4.** *For each coordinate  $i$  of  $\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln}$ , we have  $\mathbb{E}[|(\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln})_i|] \sim O(\frac{1}{\sqrt{k}})$*

From Corollary 3.4, we can see that the expectation of  $|(\mathbf{a}_{k+1} - \mathbf{a}_k)_i|$  decreases to 0 as  $k$  increases to infinity with rate  $1/\sqrt{k}$ . This means, when the number of layers increases, the inputs to later layers will be similar to each other. Thus, the capability of the later layers are not fully used because they cannot further refine the representations.

**Corollary 3.5.** *When adding an extra layer to a  $N - 1$  layer Pre-LN Transformer, the output difference  $\mathbb{E}[|(\mathbf{y}_N - \mathbf{y}_{N-1})_i|] \sim O(\frac{1}{\sqrt{N}})$  for each coordinate  $i$ .*

The proof of Corollary 3.5 is in Appendix C, it means that adding extra layer in the deep Pre-LN Transformer has little impact on the output. Intuitively, this means the extra layer also cannot refine the model outputs and the model’s capacity is not fully used.

### 3.3 ANALYSIS OF RESIDUAL

**Residual Does Not Suffer From Gradient Vanishing Issue** For the Residual architecture (Figure 1c), we can view it as a mixture of Post-LN Transformer and Pre-LN Transformer. Specifically, in the forward process, Residual Transformer behaves exactly the same as Post-LN except adding a dual branch of normalized sum of all block outputs in the end. In the backward process, the error signal back-propagates through both branches. We can explicitly write down the gradients at block

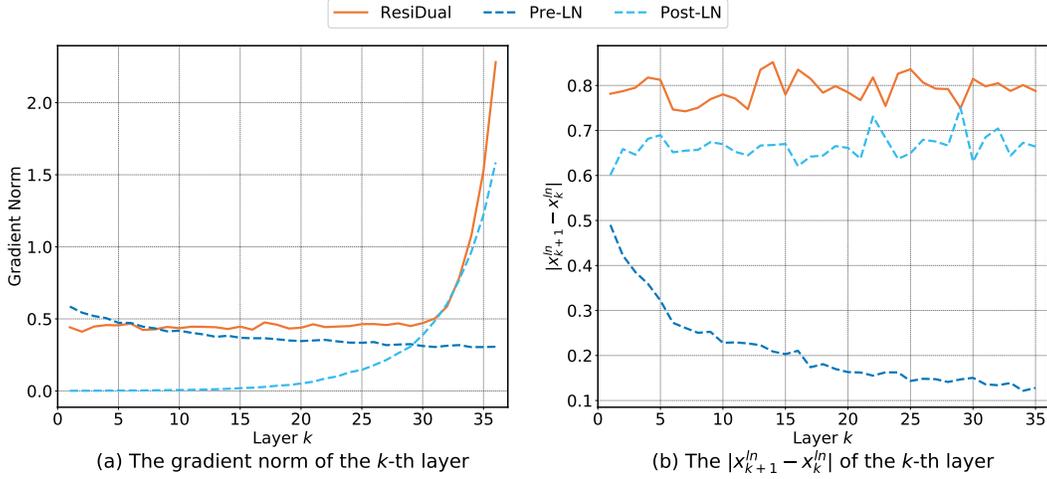


Figure 2: Study of the Gradient norm and hidden representation w.r.t layer  $k$  in each method. Details in Appendix E

$k$  as follows

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} = \left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{post} + \left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{dual}, \quad (4)$$

where  $\left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{post}$  denotes the gradient component from the Post-LN branch and  $\left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{dual}$  denotes the gradient component from the dual branch. Specifically,

$$\left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{post} = \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}} \left( \prod_{l=k}^N \frac{\partial \vec{\mathbf{x}}_{l+1}}{\partial \vec{\mathbf{x}}_l^{ln}} \frac{\partial \vec{\mathbf{x}}_l^{ln}}{\partial \vec{\mathbf{x}}_l} \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k} = \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}} \left( \prod_{l=k}^N \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} \frac{\partial \vec{\mathbf{x}}_l^{ln}}{\partial \vec{\mathbf{x}}_l} \right) \right) \frac{\partial \vec{\mathbf{x}}_k}{\partial \mathbf{w}_k},$$

and

$$\left( \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right)_{dual} = \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}} \left( \prod_{l=k+1}^N \frac{\partial \vec{\mathbf{x}}_{l+1}}{\partial \vec{\mathbf{x}}_l} \right) \frac{\partial \vec{\mathbf{x}}_{k+1}^f}{\partial \mathbf{w}_k} = \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}} \left( \prod_{l=k+1}^N \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} \frac{\partial \vec{\mathbf{x}}_l^{ln}}{\partial \vec{\mathbf{x}}_l} \right) \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k}.$$

We see that when  $k$  is small, the Pre-LN gradient component dominates and when  $k$  is close to  $N$ , the Post-LN gradient component dominates. It is safe to estimate the gradient norm of the  $k$ -th block in ResiDual Transformer as follows,

$$\left\| \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} \right\|_F \approx \max \left\{ \mathcal{O} \left( (1/2)^{(N-k)/2} e^{\sqrt{N-k}} \right), \mathcal{O} \left( \sqrt{\frac{\log(N-k)}{N}} \right) \right\}, \quad (5)$$

where again we ignore the terms irrelevant with  $N, k$ . Therefore, the ResiDual architecture does not suffer gradient vanishing problem. It is worthy to note gradient vanishing problem does not directly relate to inefficient training because in Adam the actual update is rescaled to be normal even if extreme small gradient is obtained. However, the gradient vanishing problem would affect the stability of the Adam optimizer as we argue as follows.

In Figure 2(a), we show the gradient distribution for different methods. We can find that the Post-LN has almost zero gradient for early layers, while the ResiDual (orange line) do not have such an issue. The clearly shows that our method can ensure a lower-bound of the gradient norm. Meanwhile, note that non of these models have the exploding-gradient issue. According to Theorem 3.1, the gradient of last layer (i.e.,  $k = N$ ) is not related to  $N$ .

**ResiDual Does Not Suffer From Representation Collapse Issue.** The Post-LN and ResiDual do not have the representation collapse issue. Formally,

**Theorem 3.6.** *In Post-LN and ResiDual, assume  $\mathbf{x}_k^f \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$  independently for all  $k \in [N]$ , the  $\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln} \sim \mathcal{N}(0, \omega^2)$  where  $\omega$  is not related to  $k$ .*

*Proof.* As  $\mathbf{x}_{k+1}^{ln} = \text{LN}(\mathbf{x}_k^a) = \text{LN}(\mathbf{x}_k^{ln} + \mathbf{x}_k^f)$ , and  $\mathbf{x}_k^{ln} \sim \mathcal{N}(0, \mathbf{I})$ ,  $\mathbf{x}_k^f \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$ , we have

$$\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln} = \frac{\mathbf{x}_k^{ln} + \mathbf{x}_k^f}{\sqrt{1 + \sigma^2}} - \mathbf{x}_k^{ln} = \frac{(1 - \sqrt{1 + \sigma^2})\mathbf{x}_k^{ln} + \mathbf{x}_k^f}{\sqrt{1 + \sigma^2}}.$$

Thus,  $\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln} \sim \mathcal{N}(0, \omega^2)$  where  $\omega^2 = 2 - 2\frac{\sqrt{1+\sigma^2}}{1+\sigma^2}$  and  $\omega$  is not related to  $k$ .  $\square$

**Corollary 3.7.** *When adding an extra layer to a  $N - 1$  layer Pre-LN Transformer, the output difference  $\mathbb{E}[|(\mathbf{y}_N - \mathbf{y}_{N-1})_i|] \geq \sqrt{\frac{2}{\pi}}\omega$  for each coordinate  $i$ .*

The proof of 3.7 is in the supplementary material. From these analyse, we can see that the variance of  $\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln}$  will not decrease when the depth increases, so that later layers can continue refining the hidden representation. Meanwhile, according to Corollary 3.7, the model output can also be refined with a lower bound that not related to depth. In another words, ResiDual can avoid the representation bottleneck of Pre-LN model. To demonstrate this, we also show the  $|\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln}|$  for different architectures in Figure 2(b). As the lines show, our method (orange line) has a consistent value of  $|\mathbf{x}_{k+1}^{ln} - \mathbf{x}_k^{ln}|$ , while the Pre-LN’s value will decrease when the depth is high.

## 4 EXPERIMENTS

### 4.1 EXPERIMENTAL SETTINGS

**Data** We conducted experiments on three datasets: the IWSLT-14 English to German (EN→DE) dataset (Cettolo et al., 2014), the WMT German to English (DE→EN) dataset (Bojar et al., 2014), and the OPUS-100 multilingual dataset (Zhang et al., 2020). More details are in Appendix H.

**Model** Our model is implemented using the FairSeq (Ott et al., 2019) framework with conventional settings as previous works. Notably, our method introduce only negligible parameters to the vanilla Transformer network. Meanwhile, given that the residual connection operations have a relatively small computational cost compared to Attention and FFN layers, the efficiency of our method should not hinder its practical use. We empirically observed about 3% increase in computation cost. Please refer to the Appendix H for hyper-parameters.

### 4.2 EXPERIMENTAL RESULTS ON IWSLT

The experimental results of the IWSLT’14 dataset are presented in Table 2. Two types of models were used: shallow models with 6-layer encoders and 6-layer decoders (E6D6), and deep models with 12-layer encoders and 12-layer decoders (E12D12). We made the following observations:

Firstly, the Post-LN method was successful in converging for E6D6 but not for E12D12. Secondly, the Pre-LN method converged in both depths, but its performance (35.12, 35.18) was inferior to that of the Post-LN E6D6 (35.37) or our E6D6 (35.63). Thirdly, the methods such as DeepNet (Wang et al., 2022a) and Admin (Liu et al., 2020) only showed a slight improvement over the vanilla models, and our method achieved best performance. Especially, in E12D12, we have 0.9-point BLEU gain over the standard Pre-LN model. Our preliminary experiments revealed that increasing the model depth further led to over-fitting issues for all models due to limited data. Therefore, we do not report 18 layer model results on this dataset.

Table 2: Experimental Results on IWSLT.

Method	E6D6	E12D12
Post-LN	35.37	Fail
Pre-LN	35.12	35.18
DeepNet	35.34	35.39
Admin	35.50	35.67
T-Fixup	34.88	35.45
NormFormer	35.14	31.00
<b>ResiDual(Ours)</b>	<b>35.63</b>	<b>36.09</b>

### 4.3 EXPERIMENTAL RESULTS ON WMT

The experimental results on shallow (E6D6) and deep (E18D18) models are presented in Table 3. We only report the average score here and more details can be found in Table 6 and Table 7 in Appendix F. Firstly, we find that the Post-LN model can only converge in the E6D6 setting but not in E18D18 setting. Secondly, the Pre-LN model shows convergence in both E6D6 and E18D18. However, the performance of the Pre-LN model in E18D18 (26.57) is similar to Post-LN model in E6D6 (26.59). Finally, our method achieved the best performance for both shallow and deep models. Particularly, we observed an improvement over the Pre-LN performance by 1.1-point for the E18D18 model.

Table 3: Experimental Results on WMT.

Method	E6D6	E18D18
Pre-LN	26.10	26.57
Post-LN	26.59	Fail
DLCL	26.52	26.90
T-Fixup	26.43	26.94
DeepNet	26.38	27.13
Admin	26.49	26.86
B2T	26.53	27.30
<b>ResiDual</b>	<b>26.85</b>	<b>27.65</b>

#### 4.4 EXPERIMENTAL RESULTS ON OPUS-100

We evaluate our method on the OPUS-100 dataset, which consists of 100 language pairs and 55M parallel sentence pairs. Because we trained single model for both from English (EX) and to English (XE) direction, the total data size is about 110M sentence pairs and approximately 4 billion tokens. Table 4 shows the experimental results. In addition to the original baselines provided by Zhang et al. (2020), we also reproduced the 18-layer encoder and 18-layer decoder model (E18D18). We found that the Post-LN model failed to converge thus only show the Pre-LN results in Table 4. As we can see from the table,

Table 4: Experimental Results on OPUS-100. \*Denotes The Results From Zhang et al. (2020).

Method	#Layers	EX	XE	ALL
Pre-LN*	6	21.4	27.5	24.5
	12	22.9	29.5	26.2
	24	24.0	31.4	27.7
Pre-LN	18	27.9	32.8	30.3
DeepNet	100	29.0	33.2	31.1
<b>ResiDual</b>	18	28.7	33.4	31.0

our method achieves about 0.7 BLEU points over the standard Pre-LN model. The BLEU score is almost identical to a 100-layer DeepNet (Wang et al., 2022a) model, which is about 5 times deeper of our model. This demonstrates that our model can more effectively use deeper layers.

#### 4.5 STUDY OF LEARNING-RATE WARM-UP

One of the objectives of our approach is to facilitate easy and stable training for Transformer models. Therefore, we conducted experiments using different learning rate schedules on the IWSLT dataset. Table 5 presents the results for various models with or without learning-rate warm-up. Further details can be found in the Appendix G. We observe that Post-LN necessitates warm-up for convergence, while Pre-LN and our method are not. This is consistent with our study in Section 3.

Table 5: Study on the Use of Learning-Rate Warm-Up.

Method	Post-LN		Pre-LN		ResiDual	
	Yes	No	Yes	No	Yes	No
E6D6	35.37	Fail	35.12	32.28	35.63	<b>35.76</b>
E12D12	Fail	Fail	35.18	31.82	<b>36.09</b>	35.57

## 5 CONCLUSION

This research is to advance the Transformer architecture and offers an effective strategy for optimizing it with enhanced performance. This paper first examines the limitations of two widely employed variants, and introduces a novel approach, referred to as ResiDual, to mitigate both issues. ResiDual consists two residual connections to circumvent the gradient vanishing and the representation collapse problem. Theoretical analysis and empirical results validates that the suggested model can surmount both challenges while preserving the advantages of each residual connection.

## REPRODUCIBILITY STATEMENT

The complete proof can be found in Appendix A, B, C, and D. The detailed process to build Figure 2 is in Appendix E. Our code is anonymously available at [https://anonymous.4open.science/r/residual\\_review-6F08](https://anonymous.4open.science/r/residual_review-6F08). Meanwhile, you can refer Appendix H for implementation details like data processing scripts and hyper parameters.

## REFERENCES

- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Henry Mao, Gary Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. In *Uncertainty in Artificial Intelligence*, pp. 1352–1361. PMLR, 2021.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. Findings of the 2014 workshop on statistical machine translation. In *Proceedings of the ninth workshop on statistical machine translation*, pp. 12–58, 2014.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, and Marcello Federico. Report on the 11th iwslt evaluation campaign, iwslt 2014. In *Proceedings of the International Workshop on Spoken Language Translation, Hanoi, Vietnam*, volume 57, 2014.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. Palm: Scaling language modeling with pathways, 2022.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, volume 1, pp. 4171–4186, 2018.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale, 2020. URL <https://arxiv.org/abs/2010.11929>.
- William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity, 2021.
- Ruining He, Anirudh Ravula, Bhargav Kanagal, and Joshua Ainslie. Realformer: Transformer likes residual attention. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 2021. doi: 10.18653/v1/2021.findings-acl.81. URL <http://dx.doi.org/10.18653/v1/2021.findings-acl.81>.
- Xiao Shi Huang, Felipe Perez, Jimmy Ba, and Maksims Volkovs. Improving transformer optimization through better initialization. In *International Conference on Machine Learning*, pp. 4475–4483. PMLR, 2020.

- Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rns: Fast autoregressive transformers with linear attention, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations, 2019.
- Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. Understanding the difficulty of training transformers. *arXiv preprint arXiv:2004.08249*, 2020.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. URL [https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language\\_understanding\\_paper.pdf](https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/languageunsupervised/language_understanding_paper.pdf), 2018.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Yi Ren, Yangjun Ruan, Xu Tan, Tao Qin, Sheng Zhao, Zhou Zhao, and Tie-Yan Liu. FastSpeech: Fast, robust and controllable text to speech. *Advances in neural information processing systems*, 32, 2019.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1715–1725, Berlin, Germany, 2016.
- Sam Shleifer, Jason Weston, and Myle Ott. Normformer: Improved transformer pretraining with extra normalization, 2021.
- Sho Takase and Shun Kiyono. Rethinking perturbations in encoder-decoders for fast training. *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021. doi: 10.18653/v1/2021.naacl-main.460. URL <http://dx.doi.org/10.18653/v1/2021.NAACL-MAIN.460>.
- Sho Takase, Shun Kiyono, Sosuke Kobayashi, and Jun Suzuki. On layer normalizations and residual connections in transformers. *arXiv preprint arXiv:2206.00330*, 2022.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, and Furu Wei. Deepnet: Scaling transformers to 1,000 layers. *arXiv preprint arXiv:2203.00555*, 2022a.
- Hongyu Wang, Shuming Ma, Shaohan Huang, Li Dong, Wenhui Wang, Zhiliang Peng, Yu Wu, Payal Bajaj, Saksham Singhal, Alon Benham, Barun Patra, Zhun Liu, Vishrav Chaudhary, Xia Song, and Furu Wei. Foundation transformers, 2022b.
- Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F Wong, and Lidia S Chao. Learning deep transformer models for machine translation. *arXiv preprint arXiv:1906.01787*, 2019.

Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tiejian Liu. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*, pp. 10524–10533. PMLR, 2020.

Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. Improving massively multilingual neural machine translation and zero-shot translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1628–1639, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.148. URL <https://aclanthology.org/2020.acl-main.148>.

Huishuai Zhang, Da Yu, Mingyang Yi, Wei Chen, and Tie-Yan Liu. Stabilize deep resnet with a sharp scaling factor  $\tau$ . *Machine Learning*, 111(9):3359–3392, 2022.

## A PROOF OF THEOREM 3.1

*Proof.* For the Post-LN Transformer, the gradient of the parameters in the  $k$ -th layer (take  $\mathbf{W}_k$  as an example) can be written as

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} &= \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}^{ln}} \frac{\partial \vec{\mathbf{x}}_{N+1}^{ln}}{\partial \vec{\mathbf{x}}_N^a} \left( \prod_{l=k}^{N-1} \frac{\partial \vec{\mathbf{x}}_{l+1}^a}{\partial \vec{\mathbf{x}}_{l+1}^{ln}} \frac{\partial \vec{\mathbf{x}}_{l+1}^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k} \\ &= \frac{\partial \mathcal{L}}{\partial \vec{\mathbf{x}}_{N+1}^{ln}} \frac{\partial \vec{\mathbf{x}}_{N+1}^{ln}}{\partial \vec{\mathbf{x}}_N^a} \left( \prod_{l=k}^{N-1} \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_{l+1}^f}{\partial \vec{\mathbf{x}}_{l+1}^{ln}} \right) \frac{\partial \vec{\mathbf{x}}_{l+1}^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k}. \end{aligned}$$

We care about the spectral norm of the term  $\frac{\partial \vec{\mathbf{x}}_{N+1}^{ln}}{\partial \vec{\mathbf{x}}_N^a} \left( \prod_{l=k}^{N-1} \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_{l+1}^f}{\partial \vec{\mathbf{x}}_{l+1}^{ln}} \right) \frac{\partial \vec{\mathbf{x}}_{l+1}^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right)$ , which varies for different blocks.

For the feedforward layer and attention layer, we respectively have  $l \in [N]$ ,

$$\frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} = \begin{pmatrix} \mathbf{W}_l^T & & \\ & \ddots & \\ & & \mathbf{W}_l^T \end{pmatrix} \quad \text{and} \quad \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} = \begin{pmatrix} \frac{1}{n} \mathbf{W}_{V,l}^T & \cdots & \frac{1}{n} \mathbf{W}_{V,l}^T \\ \vdots & \ddots & \vdots \\ \frac{1}{n} \mathbf{W}_{V,l}^T & \cdots & \frac{1}{n} \mathbf{W}_{V,l}^T \end{pmatrix},$$

based on the setup of the feedforward layer and attention layer at the initialization. For the layer normalization layer, we have

$$\frac{\partial \vec{\mathbf{x}}_{l+1}^{ln}}{\partial \vec{\mathbf{x}}_l^a} = \begin{pmatrix} \frac{\partial \text{LN}(\mathbf{x}_l^{a(1)})}{\partial \mathbf{x}_l^{a(1)}} & & \\ & \ddots & \\ & & \frac{\partial \text{LN}(\mathbf{x}_l^{a(n)})}{\partial \mathbf{x}_l^{a(n)}} \end{pmatrix} = \begin{pmatrix} \frac{\sqrt{d}}{\|\mathbf{x}_l^{a(1)}\|_2} \mathbf{I} & & \\ & \ddots & \\ & & \frac{\sqrt{d}}{\|\mathbf{x}_l^{a(n)}\|_2} \mathbf{I} \end{pmatrix},$$

as we assume on the Jacobian of layer normalization.

We note that  $\mathbf{I} + \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}}$  are block-circulant matrices for all  $l$  and the product of block-circulant matrices is also block-circulant. We know a block-circulant matrix has the following property

$$\left\| \begin{pmatrix} \mathbf{B} & \mathbf{A} & \cdots & \mathbf{A} \\ \mathbf{A} & \mathbf{B} & \cdots & \mathbf{A} \\ \vdots & & \ddots & \vdots \\ \mathbf{A} & \cdots & \mathbf{A} & \mathbf{B} \end{pmatrix} \right\|_2 = \|\mathbf{B} + (n-1)\mathbf{A}\|_2,$$

where  $\mathbf{B}$  and  $\mathbf{A}$  are square matrices and there are  $n-1$   $\mathbf{A}$ s each row. Hence we have

$$\left\| \frac{\partial \vec{\mathbf{x}}_{N+1}^{ln}}{\partial \vec{\mathbf{x}}_N^a} \left( \prod_{l=k}^{N-1} \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_{l+1}^f}{\partial \vec{\mathbf{x}}_{l+1}^{ln}} \right) \frac{\partial \vec{\mathbf{x}}_{l+1}^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right) \right\|_2 = \left( \prod_{l=k}^N \frac{\sqrt{d}}{\|\mathbf{x}_l^{a(i)}\|_2} \right) \left\| \left( \prod_{l=k+1}^N (\mathbf{I} + \mathbf{w}_l^T) \right) \right\|_2,$$

where  $\mathbf{w}_l$  represents either  $\mathbf{W}_{V,l}$  or  $\mathbf{W}_l$ . We know that with high probability,  $\|\mathbf{x}_l^{a(i)}\|_2 \in (1 \pm \epsilon)\sqrt{2d}$  where  $\epsilon$  is a small positive constant, based on the assumption  $\|\mathbf{x}_l^{ln(i)}\|_2 = \sqrt{d}$  and the random initialization of  $\mathbf{w}_l$  for all  $i \in [n]$ . Thus we have a term  $\left(\prod_{l=k}^N \frac{\sqrt{d}}{\|\mathbf{x}_l^{a(i)}\|_2}\right) \approx \mathcal{O}((1/2)^{(N-k)/2})$ . Moreover, based on the random matrix argument Zhang et al. (2022), we have with high probability,

$$\left\| \prod_{l=k}^N (\mathbf{I} + \mathbf{w}_l^T) \right\|_2 \approx \mathcal{O}(e^{\sqrt{N-k}}).$$

Therefore, we have  $\|\frac{\partial \mathcal{L}}{\partial \mathbf{w}_k}\|_F \approx \mathcal{O}((1/2)^{(N-k)/2} \cdot e^{\sqrt{N-k}})$ , which diminishes exponentially as  $N-k$  is large.

On the other hand, we have the bound for Pre-LN transformer as follows.

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} = \frac{\partial \mathcal{L}}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \vec{\mathbf{x}}_{N+1}^a} \left( \prod_{l=k+1}^N \frac{\partial \vec{\mathbf{x}}_{l+1}^a}{\partial \vec{\mathbf{x}}_l^a} \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k} = \frac{\partial \mathcal{L}}{\partial \mathbf{y}} \frac{\partial \mathbf{y}}{\partial \vec{\mathbf{x}}_{N+1}^a} \left( \prod_{l=k+1}^N \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} \frac{\partial \vec{\mathbf{x}}_l^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right) \right) \frac{\partial \vec{\mathbf{x}}_k^f}{\partial \mathbf{w}_k}.$$

We know that with high probability,  $\|\mathbf{x}_l^{a(i)}\|_2 \in (1 \pm \epsilon)\sqrt{ld}$  based on the assumption  $\|\mathbf{x}_l^{ln(i)}\|_2 = \sqrt{d}$  and the random initialization of  $\mathbf{w}_l$  for all  $i \in [n]$ . Hence  $\frac{\partial \mathbf{x}_l^{ln}}{\partial \mathbf{x}_l^a} \approx 1/\sqrt{l}\mathbf{I}$ . Therefore, with high probability, we have

$$\left\| \prod_{l=k+1}^N \left( \mathbf{I} + \frac{\partial \vec{\mathbf{x}}_l^f}{\partial \vec{\mathbf{x}}_l^{ln}} \frac{\partial \vec{\mathbf{x}}_l^{ln}}{\partial \vec{\mathbf{x}}_l^a} \right) \right\|_2 \approx \left\| \prod_{l=k+1}^N \left( \mathbf{I} + \frac{1}{\sqrt{l}} \mathbf{w}_l^T \right) \right\|_2 = \mathcal{O}(\log(N-k)),$$

where the last inequality is based on the argument for the product of random matrices (Zhang et al., 2022). Therefore, by further being aware of  $\|\frac{\partial \mathbf{y}}{\partial \vec{\mathbf{x}}_{N+1}^a}\|_2 \approx 1/\sqrt{N}$ , we have  $\|\frac{\partial \mathcal{L}}{\partial \mathbf{w}_k}\| \approx \mathcal{O}\left(\frac{\log(N-k)}{\sqrt{N}}\right)$ , which scales with  $\sqrt{N}$  inverse proportionally. □

## B STUDY OF ADAM

The Adam update formula is

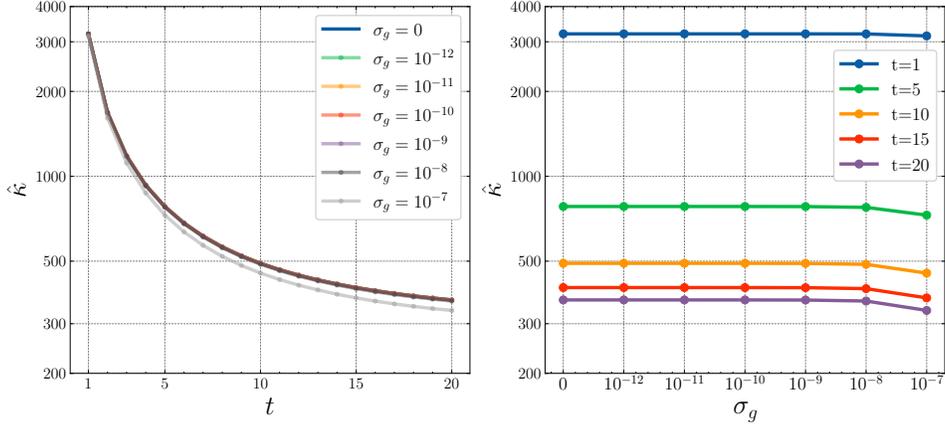
$$\begin{aligned} \mathbf{w}^{(t)} &\leftarrow \mathbf{w}^{(t-1)} - \alpha \cdot \hat{\mathbf{m}}^{(t)} / (\sqrt{\hat{\mathbf{v}}^{(t)}} + \epsilon) \\ \hat{\mathbf{m}}^{(t)} &\leftarrow \mathbf{m}^{(t)} / (1 - \beta_1^t), \quad \hat{\mathbf{v}}^{(t)} \leftarrow \mathbf{v}^{(t)} / (1 - \beta_2^t) \\ \mathbf{m}^{(t)} &\leftarrow \beta_1 \cdot \mathbf{m}^{(t-1)} + (1 - \beta_1) \cdot \mathbf{g} \\ \mathbf{v}^{(t)} &\leftarrow \beta_2 \cdot \mathbf{v}^{(t-1)} + (1 - \beta_2) \cdot \mathbf{g}^2, \end{aligned}$$

The full expression of  $u'(g_{t,i})$  is

$$\begin{aligned} \frac{\partial u}{\partial g} &= \frac{\alpha g \sqrt{\frac{\beta_2 v^{(t-1)} + g^2 \cdot (1 - \beta_2)}{1 - \beta_2^t}} (1 - \beta_2) (\beta_1 m^{(t-1)} + g(1 - \beta_1))}{(1 - \beta_1^t) \left( \epsilon + \sqrt{\frac{\beta_2 v^{(t-1)} + g^2 \cdot (1 - \beta_2)}{1 - \beta_2^t}} \right)^2 (\beta_2 v^{(t-1)} + g^2 \cdot (1 - \beta_2))} \\ &+ \frac{\alpha(1 - \beta_1)}{(1 - \beta_1^t) \left( \epsilon + \sqrt{\frac{\beta_2 v^{(t-1)} + g^2 \cdot (1 - \beta_2)}{1 - \beta_2^t}} \right)} \end{aligned} \quad (6)$$

When the gradient  $g = 0$ , we have

$$\frac{\partial u}{\partial g} = \frac{\alpha(1 - \beta_1)}{(1 - \beta_1^t) \left( \epsilon + \sqrt{\frac{\beta_2 v_{t-1,i}}{1 - \beta_2^t}} \right)}$$

Figure 3: The absolute condition number  $\hat{\kappa}$  w.r.t  $t$  (left) and  $\sigma_g$  (right).

To simulate the Adam update and compute the  $\hat{\kappa}$ , we use the parameter as  $d = 1024$ ,  $\epsilon = 10^{-6}$ ,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.98$ ,  $\alpha = 10^{-4}$ . Then for each step, we random first sample  $\mathbf{g} \sim \mathcal{N}(0, \sigma_g^2 \mathbf{I})$ , can compute the  $\hat{\kappa}$  based on full Equation (6). Finally, we update the Adam momentum with its update rules. In Figure 3, we show the simulated results by sampling  $\mathbf{g} \sim \mathcal{N}(0, \sigma_g^2 \mathbf{I})$  where  $\sigma_g$  ranges from 0 to  $10^{-7}$ . In the left plot, we show how  $\hat{\kappa}$  change w.r.t  $t$  for different  $\sigma_g$ . We can find that our estimation of  $\hat{\kappa}$  is accurate as most of lines are overlapped. Besides, it also show that even after 20 steps update, the  $\hat{\kappa}$  is still greater than 300. As many lines overlapped in the left plot, on the right side, we show a zoomed-in view by selecting five timestamp and show the  $\hat{\kappa}$  w.r.t to  $\sigma_g$ . It is clear that  $\hat{\kappa}$  is large when  $\sigma_g$  is small.

## C PROOF OF COROLLARY 3.5

Given two Pre-LN Transformer with  $N - 1$  and  $N$  layers, we denote their output as  $\mathbf{y}_{N-1}$  and  $\mathbf{y}_N$ , respectively. Then we have

$$\begin{aligned} \mathbf{y}_{N-1} &= \text{LN}(\mathbf{x}_{N-1}^a + \mathbf{x}_{N-1}^f) = \mathbf{x}_N^{ln} \\ \mathbf{y}_N &= \text{LN}(\mathbf{x}_N^a + \mathbf{x}_N^f) = \mathbf{x}_{N+1}^{ln} \\ |\mathbf{y}_N - \mathbf{y}_{N-1}| &= |\mathbf{x}_{N+1}^{ln} - \mathbf{x}_N^{ln}| \end{aligned}$$

From Corollary 3.4, we can approve that  $\mathbb{E}[|(\mathbf{y}_N - \mathbf{y}_{N-1})_i|] \sim O(\frac{1}{\sqrt{N}})$ .

## D PROOF OF COROLLARY 3.7

*Proof.*

$$\begin{aligned} \mathbf{y}_N - \mathbf{y}_{N-1} &= (\mathbf{x}_N^{ln} + \text{LN}(\mathbf{x}_{N+1}^d)) - (\mathbf{x}_{N-1}^{ln} + \text{LN}(\mathbf{x}_N^d)) \\ &= (\mathbf{x}_N^{ln} - \mathbf{x}_{N-1}^{ln}) + (\text{LN}(\mathbf{x}_{N+1}^d) - \text{LN}(\mathbf{x}_N^d)) \end{aligned}$$

The  $\text{LN}(\mathbf{x}_{N+1}^d) - \text{LN}(\mathbf{x}_N^d)$  is also a zero-mean Gaussian distribution, which can be denoted as  $\mathcal{N}(0, \hat{\omega}_N^2)$ . Then we have  $\mathbf{y}_N - \mathbf{y}_{N-1} \sim \mathcal{N}(0, \omega^2 + \hat{\omega}_N^2)$ . Therefore,

$$\begin{aligned} \mathbb{E}[|(\mathbf{y}_N - \mathbf{y}_{N-1})_i|] &= \sqrt{\frac{2}{\pi}} \sqrt{\omega^2 + \hat{\omega}_N^2} \\ &\geq \sqrt{\frac{2}{\pi}} \omega \end{aligned}$$

□

## E HOW TO GET THE FIGURE 2

We created Transformer networks with 36 layers using different architectures: Pre-LN, Post-LN, and ResiDual. The embedding dimension, FFN dimension, and the number of attention heads were set to 256, 1024, and 4, respectively. We used the fairseq framework without modifying any initialization settings. We then created a dummy mini-batch with 16 sentences and 20 tokens per sentence. Next, we conducted a forward and backward process to collect the gradient norm and  $x_k^{\text{in}}$  of each layer. We hope this clarifies how we generated Figure 2, and we are happy to provide more details if needed.

## F FULL RESULTS ON WMT DATASET

The full results on WMT dataset is in Table 6 and Table 7. The baseline results are cited from Takase et al. (2022).

Method	2010	2011	2012	2013	2014	2015	2016	Average
Pre-LN	24.03	21.77	22.08	25.63	26.27	29.07	33.84	26.10
Post-LN	24.27	22.06	22.43	26.11	27.13	29.70	34.40	26.59
DLCL	23.94	22.00	22.24	26.11	27.37	29.71	34.26	26.52
T-Fixup	24.09	21.98	22.04	25.96	26.92	29.45	34.56	26.43
DeepNet	24.08	21.76	22.09	25.90	26.85	29.62	34.39	26.38
Admin	24.32	21.79	22.17	26.26	27.14	29.61	34.12	26.49
B2T	24.12	21.93	22.29	26.31	26.84	29.48	<b>34.73</b>	26.53
<b>ResiDual(Ours)</b>	<b>24.42</b>	<b>22.20</b>	<b>22.66</b>	<b>26.64</b>	<b>27.23</b>	<b>30.22</b>	34.55	<b>26.85</b>

Table 6: Experimental Results on WMT with E6D6 models.

Method	2010	2011	2012	2013	2014	2015	2016	Average
Pre-LN	24.07	21.98	22.4	26.28	27.36	29.74	34.16	26.57
Post-LN				Fail				
DLCL	24.20	22.51	22.83	26.59	27.97	30.24	33.98	26.90
T-Fixup	24.45	22.29	22.76	26.57	27.71	30.13	34.69	26.94
DeepNet	24.70	22.40	22.92	26.85	28.21	30.60	34.25	27.13
Admin	24.56	22.17	22.62	26.48	27.99	30.35	33.88	26.86
B2T	24.62	22.51	22.86	26.74	28.48	30.99	34.93	27.30
<b>ResiDual(Ours)</b>	<b>24.85</b>	<b>22.76</b>	<b>23.18</b>	<b>27.60</b>	<b>28.79</b>	<b>31.12</b>	<b>35.24</b>	<b>27.65</b>

Table 7: Experimental Results on WMT E18D18 models.

## G FULL RESULTS ON LEARNING-RATE WARM-UP

The full results on learning-rate warm-up is in Table 8.

## H IMPLEMENTATION DETAILS

### H.1 DATA PROCESSING

The IWSLT-14 EN→DE dataset is relatively small, with only 140k sentence pairs. We followed the scripts in FairSeq (Ott et al., 2019) to preprocess the data. The WMT DE→EN dataset is larger, with 4M sentence pairs. We followed the preprocessing steps outlined in Takase & Kiyono (2021) by tokenizing the data with Moses tokenizer and then processing it with BPE (Sennrich et al., 2016). The model was trained on the WMT-14 training set and evaluated on the test set from years 2010 to 2016, following Takase & Kiyono (2021). The OPUS-100 dataset is a large-scale multilingual dataset containing 100 languages and approximately 55M sentence pairs. We used the script from Zhang et al.

Method	Learning-Rate Scheduler		E6D6	E12D12
	Warm-up	Decay Formula		
Post-LN	Yes	Inverse Square Root	35.37	Fail
Post-LN	No	Inverse Square Root	Fail	Fail
Post-LN	No	Linear	Fail	Fail
Pre-LN	Yes	Inverse Square Root	35.12	35.18
Pre-LN	No	Inverse Square Root	32.28	31.82
Pre-LN	No	Linear	32.26	31.85
ResiDual(Ours)	Yes	Inverse Square Root	35.63	<b>36.09</b>
ResiDual(Ours)	No	Inverse Square Root	35.76	35.57
ResiDual(Ours)	No	Linear	<b>35.96</b>	35.72

Table 8: Experimental Results on IWSLT with different learning-rate scheduler.

(2020) to tokenize the data and used SentencePiece (Kudo & Richardson, 2018) to segment the tokens. All data processing scripts are available in the Appendix. Because we train our model for both to and from English directions, the total training data is about 110M.

The data processing scripts are

- IWSLT: <https://github.com/facebookresearch/fairseq/blob/main/examples/translation/prepare-iwslt14.sh>
- WMT: <https://github.com/facebookresearch/fairseq/blob/main/examples/translation/prepare-wmt14en2de.sh>
- OPUS-100: <https://github.com/bzhangGo/zero>

## H.2 HYPER-PARAMETERS

The training hyper-parameters are in Table 9, 10, and 11.

Parameter	Value
Dropout	0.3
Embedding dim	256
FFN dim	1024
Attention heads	4
Encoder layers	6/12
Decoder layers	6/12
Learning rate	$5 * 10^{-4}$
Learning rate scheduler	inverse sqrt
Warm-up steps	4000
Label smoothing	0.1
Weight decay	0.0001
Gradient clipping	0
Adam $\beta$	0.9, 0.98
Max update steps	300k

Table 9: Hyper-parameters of IWSLT training

## H.3 IMPLEMENTATION TRICK ON FP16 TRAINING

In ResiDual, sometimes the  $\mathbf{x}_k^d$  will exceed the value range that can be expressed by FP16 and may cause training error. When this happens, a simple numeric trick is to downscale  $\mathbf{x}_k^d$  to make it within the FP16 scope. This will not affect the results because  $\text{LN}(\mathbf{x}_{N+1}^d) = \text{LN}(\eta \cdot \mathbf{x}_{N+1}^d)$  for any  $\eta > 0$ . We did not observe such an issue in FP32 training.

Parameter	Value
Dropout	0.3
Embedding dim	512
FFN dim	2048
Attention heads	8
Encoder layers	6/18
Decoder layers	6/18
Learning rate	$1 * 10^{-3}$
Learning rate scheduler	inverse sqrt
Warm-up steps	4000
Label smoothing	0.1
Weight decay	0.0001
Gradient clipping	0
Adam $\beta$	0.9, 0.98
Max update steps	500k

Table 10: Hyper-parameters of WMT training

Parameter	Value
Dropout	0.1
Embedding dim	512
FFN dim	2048
Attention heads	8
Encoder layers	18
Decoder layers	18
Learning rate	$1 * 10^{-3}$
Learning rate scheduler	inverse sqrt
Warm-up steps	4000
Label smoothing	0.1
Weight decay	0.0001
Gradient clipping	0
Adam $\beta$	0.9, 0.98
Max update steps	100k

Table 11: Hyper-parameters of OPUS-100 training