

Rethinking Skip Connection with Layer Normalization

Fenglin Liu^{1*}, Xuancheng Ren^{2*}, Zhiyuan Zhang², Xu Sun², Yuexian Zou^{1,3†}

¹ADSPLAB, School of ECE, Peking University, Shenzhen, China

²MOE Key Laboratory of Computational Linguistics, School of EECS, Peking University

³Peng Cheng Laboratory, Shenzhen, China

{fenglinliu98, renxc, zzy1210, xusun, zouyx}@pku.edu.cn

Abstract

Skip connection is a widely-used technique to improve the performance and the convergence of deep neural networks, which is believed to relieve the difficulty in optimization due to non-linearity by propagating a linear component through the neural network layers. However, from another point of view, it can also be seen as a modulating mechanism between the input and the output, with the input scaled by a pre-defined value one. In this work, we investigate how the scale factors in the effectiveness of the skip connection and reveal that a trivial adjustment of the scale will lead to *spurious* gradient exploding or vanishing in line with the deepness of the models, which could be addressed by normalization, in particular, layer normalization, which induces consistent improvements over the plain skip connection. Inspired by the findings, we further propose to adaptively adjust the scale of the input by recursively applying skip connection with layer normalization, which promotes the performance substantially and generalizes well across diverse tasks including both machine translation and image classification datasets.

1 Introduction

There is a surge of research interest in solving the optimization problem of deep neural networks, which presents a considerable number of novel difficulties that are in need of practical solutions. Batch Normalization (Ioffe and Szegedy, 2015), Layer Normalization (Ba et al., 2016) and Skip Connection (He et al., 2016a) are widely-used techniques to facilitate the optimization of deep neural networks, which prove to be effective in multiple contexts (Szegedy et al., 2016; Vaswani et al., 2017).

Among them, Batch Normalization (BN) and Layer Normalization (LN) make progress by normalizing and rescaling the intermediate hidden representations to deal with the so-called internal co-variate shift in mini-batch training and collective optimization of multiple correlated features. Skip Connection bypasses the gradient exploding or vanishing problem and tries to solve the model optimization problem from the perspective of information transfer. It enables the delivery and integration of information by adding an identity mapping from the input of the neural network to the output, which may ease the optimization and allow the error signal to pass through the non-linearities.

Recently, several efforts have been made to deal with the optimization problem by further exploiting the existing solutions, such as exploring the implementation of skip connection (Srivastava et al., 2015) and combining the skip connection with other normalization methods (He et al., 2016b; Vaswani et al., 2017). Some of the proposals are illustrated in Figure 1. Most of these variations could be formulated as or coarsely equivalent to

$$\mathbf{y} = \mathcal{G}(\lambda\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathbf{W})), \quad (1)$$

where \mathbf{x} denotes the input of the block, i.e., the skip connection or the shortcut, \mathcal{F} denotes the non-linear transformation induced by the neural network parameterized by \mathbf{W} , \mathcal{G} denotes the normalization function,

*Equal contribution.

†Corresponding Author.

This work is licensed under a Creative Commons Attribution 4.0 International License. License details: <http://creativecommons.org/licenses/by/4.0/>.

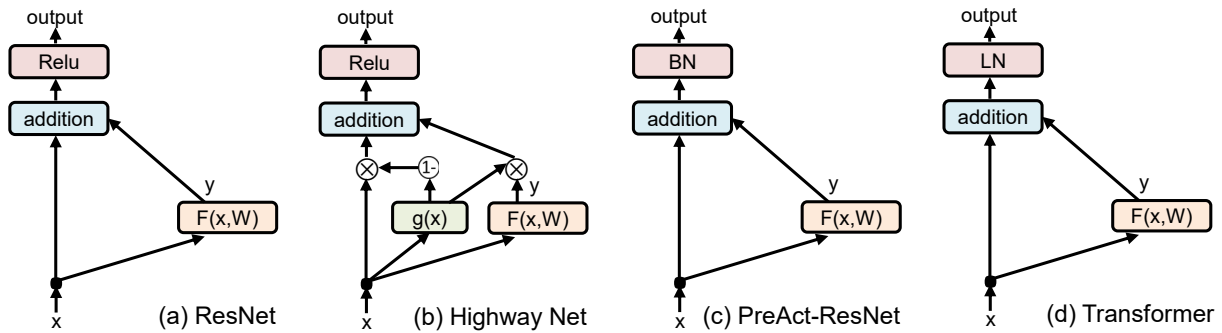


Figure 1: Illustrations of commonly-used skip connections: (a) Conventional Residual Unit (He et al., 2016a); (b) Residual Unit in Highway Net (Srivastava et al., 2015); (c) Residual Unit in PreAct-ResNet (He et al., 2016b), where BN represents batch normalization (Ioffe and Szegedy, 2015); (d) Residual Unit in Transformer (Vaswani et al., 2017), where LN represents layer normalization (Ba et al., 2016).

y denotes the output of the residual block, and λ denotes the modulating factor that controls the relative importance of the skip connection or the shortcut.¹ A residual model consists of a stack of such blocks.

A missing piece from the existing work is how would the residual block perform if \mathcal{G} is realized as layer normalization and λ does not equal one. Although Transformer (Vaswani et al., 2017) has shown the effectiveness of the combination of layer normalization and skip connection, it is intuitive that the modulating factor λ may not always be one, especially when \mathcal{F} is not well trained and not sufficient in representation ability. In addition, the intuition is supported by Srivastava et al. (2015) who found that if using a gate mechanism to determine the balance between the two components, the skip component wins in almost all situations. However, in the study of He et al. (2016b)², which discussed the combination of batch normalization, learned λ , and convolutional neural networks, they found empirically that the best performance is achieved when λ is fixed as one, suggesting the shortcut and the transformed input should be of equal importance.

Therefore, in order to explore whether the optimal ratio of x to $\mathcal{F}(x, \mathcal{W})$ is one in the residual block, and to extend the understandings of the residual block formed as $y = \mathcal{G}(\lambda x + \mathcal{F}(x, \mathcal{W}))$, we propose three different implementations of the residual block in combination with layer normalization. The first implementation, which we name *Expanded Skip Connection*, simply expands the skip connection and sets λ larger than one, following the findings of Srivastava et al. (2015). The second variation, named as *Expanded Skip Connection with Layer Normalization*, includes the layer normalization after the expanded skip connection, since layer normalization is observed to be helpful in facilitating the optimization of skip connection as in Vaswani et al. (2017). The final proposal, *Recursive Skip Connection with Layer Normalization*, is a novel combination that does not fit in the general form of the residual block, which takes the advantages of skip connection and layer normalization in a recursive manner, so that the input signal is amplified yet with layer normalization stabilizing the optimization. The illustrations of those combinations are shown in Figure 2. Our experiments on diverse tasks, including machine translation and image classification with strong baselines, i.e., Transformer and ResNet, conclude the following:

- The expanded skip connection without any normalization indeed causes serve gradient malformation, leading to unsatisfactory learning of the neural networks. Layer normalization helps to deal with the optimization problem introduced by the expanded skip connection to a certain extent.
- The proposed recursive skip connection with layer normalization further facilities the optimization by separating the expanded skip connection into multiple stages to better incorporate the effect of the transformed input. It achieves comfortable margin on datasets from different domains in extensive experiments to determine the proper granularity of the skip connection.

¹The terms *skip connection* and *shortcut* are used interchangeably in this work.

²They move the batch normalization applied to $\mathcal{F}(x, \mathcal{W})$ forward and apply it to x instead. From another point of view, it is similar to applying the batch normalization to y of the previous block.

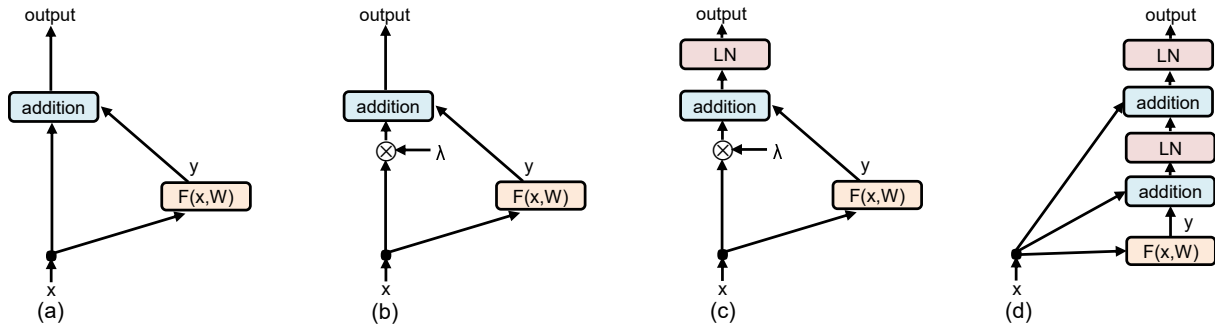


Figure 2: Various combination of skip connection and layer normalization: (a) Original skip connection, (b) Expanded skip connection, (c) Expanded skip connection with layer normalization, and (d) Recursive skip connection with layer normalization ($\lambda = 2$ as an example), where LN represents layer normalization.

- The experimental results on the WMT-2014 EN-DE machine translation dataset using Transformer further proves the effectiveness and the efficiency of the recursive architecture, which helps a model to perform even better than the model three times larger.

2 Related Work

In recent years, a large number of skip connection methods have been used for neural networks (He et al., 2016a; Srivastava et al., 2015; He et al., 2016b; Vaswani et al., 2017; Szegedy et al., 2017; Wang et al., 2019; Nguyen and Salazar, 2019; Xu et al., 2019; Bachlechner et al., 2020; Liu et al., 2020; Shen et al., 2020; Xiong et al., 2020). For the skip connection, there are two main types of problems. The first is how should the skip connection be incorporated into the existing neural network structures so that the best improvements can be achieved. The second is how should the neural network with skip connections be optimized so that its representation capability could be fully mined. Therefore, we categorized the related work into *On the Connection Problem* and *On the Optimization Problem*.

On the connection problem Highway network (Srivastava et al., 2015) built a highway connection from the input to the output, similar to the skip connection. However, a transform gate was proposed to control the balance of the input x and the transformed input $\mathcal{F}(x, \mathcal{W})$, instead of using identity mapping to combine the input and the transformed input. He et al. (2016b) designed five types of skip connections and discussed the possible skip connections in detail. Based on their theory and experiments, they argued that it is likely for a model to perform the best when the skip connection is directly combined into the output without any additional scaling such that the skip and the transformed input have equal contribution to the output. The reason is that with scaling, the gradient of the skip suffers from the gradient exploding or vanishing problem, which hinders the deep neural network from efficient optimization.

On the optimization problem Since the early work has established that it is sufficient to directly combine skip connection into forward propagation of neural networks without any scaling (He et al., 2016b), the succeeding studies on the optimization problem mostly followed the simple yet effective architecture of the skip connection. In computer vision task, PreAct-ResNet (He et al., 2016b) demonstrated that it is helpful to move the batch normalization applied to $\mathcal{F}(x, \mathcal{W})$ forward and apply it to x instead. From another point of view, it is similar to applying the batch normalization to y of the previous block. In natural language processing tasks, Transformer (Vaswani et al., 2017) made use of skip connection and layer normalization extensively in its architecture, and from our experiments, the layer normalization seems essential to the overall performance of the model and can help the optimization of the non-linear transformation to some extent.

In all, we propose to use the form $y = \mathcal{G}(\lambda x + \mathcal{F}(x, \mathcal{W}))$ to summarize the combination of normalization and skip connection in existing work. The form is in accordance with the findings that the normalization \mathcal{G} should be placed after the addition of the skip and it could be beneficial to elevate the portion of the input information with λ . Moreover, this work distinguishes itself in that (a) we mainly

focus on the combination of layer normalization and skip connection, which is less investigated but seems promising, (b) we rethink the common practice that the skip need not scale in consideration of the help of layer normalization, (c) we conduct experiments on both representative computer vision and natural language processing tasks, while existing work only provides result in a single modality, and most importantly (d) we escape from the general form of the residual block that generalizes all previous work and propose a novel recursive construction of the residual block with layer normalization that outperforms all the variants of the general form that are examined in this work .

3 Architecture

In this section, we describe three skip connection constructions, which are illustrated in Figure 2. They are based on the integration of more input information and application of the layer normalization.

Expanded Skip Connection (xSkip) The expanded skip connection is defined similar to the constant scaling method in He et al. (2016b) as

$$\mathbf{y} = \lambda \mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}), \quad (2)$$

where \mathbf{x} and \mathbf{y} are the input and the output of the residual block, respectively. \mathcal{F} denotes the weighted neural network layer, and λ is a modulating scalar. Naturally, this construction adjusts the importance of the skip considering that the neural network layer may come with different representation ability and optimization difficulty. However, it should be noted that in this work λ is fixed to isolate the effect of scaling. While it is possible that a learned λ may better capture the balance between the two components, the learning of λ becomes another variable. Following Srivastava et al. (2015), we mainly experiment with $\lambda > 1$.

Expanded Skip Connection with Layer Normalization (xSkip+LN) Motivated by Transformer, which combines skip connection with layer normalization, we further examine the effect of layer normalization on expanded skip connection, which takes the form of

$$\mathbf{y} = \text{LN}(\lambda \mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})). \quad (3)$$

Layer normalization may help mitigate the gradient malformation in optimization caused by the modulating factor. We use the parameterized layer normalization as suggested in Ba et al. (2016), because different input features naturally have different importance. Different from batch normalization which acts in “sample space”, layer normalization acts in “feature space”. Considering the special case that $\mathcal{F}(\mathbf{x}, \mathcal{W})$ is ill-learned and outputs random values ϵ and the examples $\mathbf{X} = \{\mathbf{x}_i\}$ in the mini-batch all contain the same input feature f of the same value v , batch normalization simply ignores the input feature f which results in uneventful optimization according to mini-batch statistics, while layer normalization still recognizes the input feature f and its value can be properly adjusted according to the other input features. The same may also apply to the situation where $\lambda \mathbf{x}$ is dominant compared to $\mathcal{F}(\mathbf{x}, \mathcal{W})$ and the transformed input falls into the range comparable to random noise. In such situations, which may happen when the neural network is hard to optimize, layer normalization could still help the learning of the shortcut, while batch normalization might fail.

Recursive Skip Connection with Layer Normalization (rSkip+LN) Another way to stabilize the gradient is to keep $\lambda = 1$ each time but repeatedly add the shortcut with layer normalization, such that more input information is also modeled. It is defined recursively as

$$\mathbf{y}_\lambda = \text{LN}(\mathbf{x} + \mathbf{y}_{\lambda-1}) \quad \text{and} \quad \mathbf{y}_1 = \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})), \quad (4)$$

where λ should be an integer no smaller than one. For example, when $\lambda = 1$, it regresses to the block used in Transformer and conforms to the findings that the skip need not scale. A further example is when $\lambda = 2$, $\mathbf{y}_2 = \text{LN}(\mathbf{x} + \mathbf{y}_1) = \text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})))$, which is shown in Figure 2. Through recursive skip connection with layer normalization, the model is encouraged to use layer normalization

multiple times to improve optimization, and can incorporate more information of \mathbf{x} by the skip connection. Moreover, the model may gain more expressive power than simply incorporating the scaled skip at once, since each recursive step essentially constructs a different feature distribution and the recursive structure can learn the adaptive ratio of \mathbf{x} to $\mathcal{F}(\mathbf{x}, \mathcal{W})$. The ratio between the input \mathbf{x} and the residual $\mathcal{F}(\mathbf{x}, \mathcal{W})$ is $\mathbf{x}/\mathcal{F}(\mathbf{x}, \mathcal{W}) = \sigma_1/w_1 + 1$, where σ_1 is the standard deviation of $\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$ and w_1 is the gain parameters of the inner layer normalization, which penetrate through the normalization and effectively act as the learnable ratio. It is also important to note that the ratio is also dependent on the current layer distribution, which could serve as a regularization component. For any λ larger than 2, the general form of the ratio is $1 + \sum_{i=1}^{\lambda} \prod_{j=1}^i \sigma_j/w_j$. This derivation, which is detailed in Appendix A, also indicates that setting $\lambda = 2$ should be sufficient for learning the shortcut-residual ratio and adjusting the ratio based on instant data distribution. Yet with larger λ , the distribution of the higher order combination of \mathbf{x} and $\mathcal{F}(\mathbf{x})$ will be considered. The recursive definition introduces gradual and smooth combination of the input and the residual up to the specified order λ which could deal with different distributions better.

4 Empirical Study

In this section, we examine the effect of the discussed combinations on two representative baselines through empirical studies. First, we conduct exploratory experiments on datasets that are relatively small to gain insights of the combinations, such as the choice of λ and the effectiveness of the constructions. Then, we discuss the observed results and conduct further analysis. Finally, we validate our findings on larger and more difficult datasets.

4.1 Tasks and Settings

For ease of introduction, brief descriptions of the tasks and corresponding models are summarized together in the following. As the combinations to be examined only relates to the skip connection and we keep the inner structure of the residual block untouched, please refer to the cited work for more detailed description of the baselines.

4.1.1 Image Classification

Datasets CIFAR-10 and CIFAR-100 (Krizhevsky, 2009) are two datasets for colored image classification. Both contain 50K images for training and 10K images for testing. The former has 10 classes, while the latter has 100 classes and is considered more difficult. We report the classification error rates as the evaluation of the performance.

Baselines and settings We conduct experiments using the ResNet with pre-activation (PreAct-ResNet) as proposed in He et al. (2016b). For experiments on CIFAR-10, we use ResNet-110 for exploration because of its deep structures. It contains 54 two-layer residual blocks, which makes it a challenging baseline for optimization. We validate the proposal on CIFAR-100, using a variety of ResNet structures, such as ResNet-20, ResNet-32, ResNet-44, ResNet-56, and ResNet-110, to ensure the generality of the results. We mostly adopt the hyperparameters from He et al. (2016b) to keep the number of variables controlled and comparable, except that we use a weight decay rate of 0.0002.

4.1.2 Machine Translation

Datasets IWSLT 2015 English-Vietnamese dataset (EN-VI) (Cettolo et al., 2015) consisting of 133K sentence pairs, WMT 2014 English-German dataset (EN-DE) consisting of 4.5M sentence pairs and WMT 2014 English-French dataset (EN-FR) consisting of 36M sentence pairs are for machine translation. Following previous work (Luong and Manning, 2015), we use `tst2012` and `tst2013` as development set and test set, respectively for EN-VI dataset. For EN-DE, we use `newstest2013` for development and `newstest2014` for testing. For EN-FR, we validate on `newstest2012+2013` and test on `newstest2014`. The evaluation metric is BLEU (Papineni et al., 2002) against the reference translations.

Baselines and settings Transformer (Vaswani et al., 2017) is a representative model for machine translation, which includes skip connection and layer normalization in the model structure. Transformer follows the encoder-decoder framework, except that it substitutes the conventional recurrent layers with

Method	Architecture	λ	\mathcal{G}	Error (%)
Baseline	$\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	1	-	6.37
1xSkip	$\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	1	-	<u>6.31</u>
2xSkip	$2\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	2	-	8.41
3xSkip	$3\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	3	-	9.23
1xSkip+LN	$\text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	1	LN	7.72
2xSkip+LN	$\text{LN}(2\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	2	LN	6.29
3xSkip+LN	$\text{LN}(3\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	3	LN	<u>6.07</u>
1rSkip+LN	$\text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	1	LN	7.72
2rSkip+LN	$\text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})))$	2	LN	6.02
3rSkip+LN	$\text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))))$	3	LN	<u>6.55</u>

Table 1: Results on CIFAR-10 using PreAct-ResNet-110. Average of 5 runs. The lower the better. The underlined numbers denote the best results under the same comparable settings. λ xSkip, λ xSkip+LN, and λ rSkip+LN stand for expanded skip connection, expanded skip connection with LN, and recursive skip connection with LN, respectively. Please note that 1xSkip is equivalent to Baseline and 1rSkip+LN is equivalent to 1xSkip+LN. It is clear that simply expanding the skip connection hurts the performance and using layer normalization can mitigate the problem.

the self-attention based blocks, i.e., $\mathcal{F}(\mathbf{x}, \mathcal{W})$ consists of a dot-product based attention layer and two feed-forward layers. Although Transformer is not as deep as ResNet-110, it is heavily parameterized and the commonly-used six-layer Transformer-Base structure embodies 65M learnable parameters, which also poses challenge for optimization. We base our experiments on the tensor2tensor (Vaswani et al., 2018) code. For experiments on IWSLT-2015, we set batch size to 4096 and train on single GPU, as it is relatively small. For experiments on WMT-2014 that is much larger, due to the complication caused by the special batching algorithm and the distributed learning, we find setting batch size to 16,384 tokens in total on 3 GPUs can reproduce the reported results on 8 GPUs using batch size of 32,768 tokens in total (Vaswani et al., 2017). However, the training steps required are increased from 100,000 steps to 550,000 steps. The rest of hyperparameters are the same as Vaswani et al. (2017).

4.2 Exploratory Results

In this section, we evaluate the effect of the proposed constructions and report the exploratory results on CIFAR-10 using ResNet and IWSLT-2015 using Transformer.

The main results on CIFAR-10 using PreAct-ResNet-110 are summarized in Table 1. We experiment with different constructions and different λ s. 1xSkip is our reimplement of the ResNet-110 with pre-activation and outperforms the reported result in He et al. (2016b) by 0.06. This margin could approximately serve as the boundary deciding whether the improvement is caused by random variations or experimental settings. As we can see, for skip connection without normalization, it is indeed the best to add the input without scaling and expanding the skip results in significant loss in accuracy. However, with the application of layer normalization, the expanded skip can also be sufficiently optimized and surpass the baseline, which is suggested by He et al. (2016b) by extensive structure search. Moreover, with the recursive design, the performance can also be maintained and yet further improved, which reaches the best result in our experiments when using $\lambda = 2$. It is also interesting to see that in combinations with layer normalization, the λ that results in the best accuracy is also different. It is intuitive that the recursive structure adjusts the effective ratio of the skip and the transformed input more sophisticatedly than the plain expanding structure so that different λ may be required.

Table 2 shows the results on IWSLT-2015 using Transformer and the results about layer normalization are similar. With the increasing of the weight of the skip and layer normalization, the performance of the model improves. Since the models incorporated expanded skip connection with layer normalization have the same number of parameters, regardless of λ , thus embodying the same representation ability, the improved performance should mainly be attributed to the enhanced shortcut information. Moreover, the recursive version brings significant improvements upon the simply expanding version, pushing the

Method	Architecture	λ	\mathcal{G}	BLEU
Baseline	$\text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	1	-	30.31
1xSkip	$\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	1	-	28.98
1xSkip+LN	$\text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	1	LN	30.31
2xSkip+LN	$\text{LN}(2\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	2	LN	30.99
3xSkip+LN	$\text{LN}(3\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	3	LN	30.92
1rSkip+LN	$\text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	1	LN	30.31
2rSkip+LN	$\text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})))$	2	LN	31.45
3rSkip+LN	$\text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))))$	3	LN	31.10

Table 2: Results on EN-VI machine translation dataset using Transformer. The higher the better. λ xSkip, λ xSkip+LN, and λ rSkip+LN are defined similarly. Please note that Baseline, 1xSkip+LN, and 1rSkip+LN are equivalent. It is clear that without layer normalization, Transformer experiences significantly performance degradation. By expanding the shortcut, substantially better results can be achieved. The best improvements of over 1.1 BLEU scores are brought by the proposed recursive skip connection.

BLEU score past 31 and reaching a margin of 1.1 BLEU scores at most, demonstrating the advantage of gradually adding the expanded skip to ease the optimization. A notable difference from the results on PreAct-ResNet-110 is that the plain skip connection version does not perform better than the one with layer normalization (30.31 BLEU vs. 28.98 BLEU on Transformer and 92.28 accuracy vs. 93.69 accuracy on PreAct-ResNet-110). We find that the reason is likely to be that in the residual block of PreAct-ResNet-110, there are already two instances of batch normalization before each convolutional layer to facilitate its optimization, while in the Transformer block no other normalization technique exists.

In summary, our experiments show that feeding more input signals to the layer output should be helpful to learning of the overall model. Previous concerns on the optimization difficulty introduced by the expanded skip connection could be alleviated by the usage of layer normalization. Especially, the proposed recursive skip connection with layer normalization brings the most improvements by dividing the expanded skip and iteratively adding the pieces to the output, which can further stabilize the gradient on the input and help the training of the weighted layer.

4.3 Discussions

However, questions remain whether the improvements can also be achieved by contracting skip connection and whether batch normalization can achieve the same effect which is more commonly in convolutional neural networks. We conduct the following analysis to answer these questions. Furthermore, we demonstrate that layer normalization adaptively adjusts the gradients of the scaled skip, solving the gradient malformation caused by the modulating factor λ .

Contracting skip connection We experiment two different ways to relatively tune down the skip ratio, i.e., using smaller λ and expanding $\mathcal{F}(\mathbf{x}, \mathcal{W})$. The results are shown in Table 3. As we can see, if the shortcut is forced to play the minor role, the deep neural network fails to be optimized properly and shows catastrophic performance. In contrast, the expanded skip causes less harm to the overall performance (cf. Table 1), indicating the importance of linear component in the optimization of deep neural networks. It is consistent with the findings of Srivastava et al. (2015) and He et al. (2016b). Moreover, the results also show that applying layer normalization under this situation further deteriorates the performance.

Using batch normalization Since batch normalization has been widely used in CNNs, we replace layer normalization with batch normalization to see the difference. However, comparing Table 1 and Table 4, it can be seen that batch normalization is less effective than layer normalization in combination with skip connection. The proposed recursive structure can help batch normalization deal with the expanded skip as each time the ratio is still kept as one. As we explained in Section 3, batch normalization makes use of mini-batch statistics which may prevent the plainly expanded skip from coming into effect.

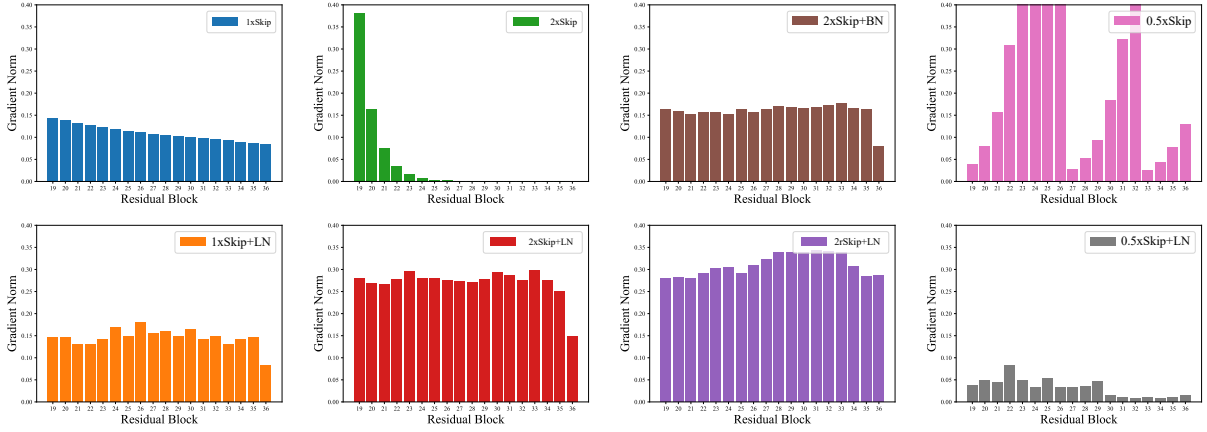


Figure 3: Gradient norm of the output of residual blocks of various constructions on CIFAR-10 using PreAct-ResNet-110. We show the gradient norm of the output \mathbf{y} of the residual blocks of middle depth in the model averaged over 10,000 random examples. By definition, \mathbf{y} equals \mathbf{x} of the following block.

Architecture	Error (%)
$\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	6.31
$0.5\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W})$	18.61
$\text{LN}(0.5\mathbf{x} + \mathcal{F}(\mathbf{x}, \mathcal{W}))$	29.88
$\text{LN}(\mathbf{x} + 2 * \mathcal{F}(\mathbf{x}, \mathcal{W}))$	24.98
$\text{LN}(\mathbf{x} + 3 * \mathcal{F}(\mathbf{x}, \mathcal{W}))$	61.02

Table 3: Results on CIFAR-10 with contracted skip connection. Lowering the skip ratio leads to catastrophic results and LN further aggravates the performance.

ResNet-110	λ	\mathcal{G}	Error (%)
1xSkip	1	-	6.31
1xSkip+BN	1	BN	8.23
2xSkip+BN	2	BN	6.43
3xSkip+BN	3	BN	6.50
1rSkip+BN	1	BN	8.23
2rSkip+BN	2	BN	6.77
3rSkip+BN	3	BN	6.68

Table 4: Results on CIFAR-10 with BN. Compared to LN, BN is less effective dealing with the expanded skip connection.

Transformer	# Blocks	BLEU
1xSkip+LN	6	30.31
2rSkip+LN	6	31.45
1xSkip+LN	2	30.04
2rSkip+LN	2	30.49

Table 5: Results on Transformer of fewer blocks. As we can see, with the proposed combination of recursive skip connection, the same performance could be achieved by fewer parameters.

Parameter efficiency A byproduct of the proposed structure is that it could be possible to use fewer parameters to reach the same performance, especially for Transformer, which is heavily parameterized, as the expanded skip connection and layer normalization both helps to ease the optimization. Fewer parameter may translate into lower time cost. We conduct experiments on the IWSLT-2015 EN-VI machine translation dataset using Transformer with two blocks instead of six blocks used in the main experiments. The results are shown in Table 5. Surprisingly, a Transformer of two blocks does not perform substantially worse, and with the proposed rSkip+LN, the reduction could be easily made up. A interesting phenomenon is that the proposed method gains larger margin on deeper models, which also attests that proposed combination helps the optimization of the model, so that the performance of deeper model could be further exploited.

Better Optimization To see how the proposed combinations help the optimization of the deep neural models, in Figure 3, we show the averaged gradient norm of the output of each residual block in pretrained ResNet-110 on CIFAR-10 based on 10,000 random training examples. Due to limited space, we only show the representative blocks at the “middle” of the model, meaning those blocks have the same filter size but different from the blocks lower and the blocks higher. As we can see, the original construction 1xSkip shows slight gradient “explosion”, while simply scaling the skip causes severe gradient malformation (e.g., 2xSkip and 0.5xSkip). With normalization techniques such as layer normalization and batch normalization, the gradient malformation is generally mitigated. However, the magnitude of the gradient is different, with layer normalization more accurately reflecting the modulating factor λ (e.g. 1xSkip and 1xSkip+LN have similar values, 2xSkip+LN is approximately twice 1xSkip, etc.), which is intuitive since batch normalization is based on mini-batch statistics unaware of the feature space. The recursive structure

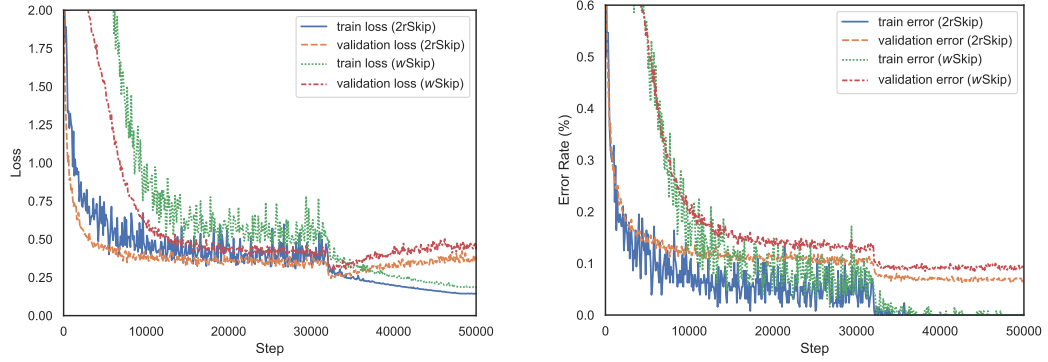


Figure 4: Learning curve for 2rSkip+LN and w Skip initialized as 1 on CIFAR-10 using ResNet-110. It can be observed that w Skip has difficulty in start training.

2rSkip+LN is similar to 2xSkip+LN, except that its norm is larger. Especially, note the gradient at the Block 36, which receives gradient from the higher block that have different filter size; xSkip+LN and xSkip+BN seem to fail normalize its gradient properly, while rSkip+LN successfully elevates its scale, suggesting it could mitigate the difference between blocks of different nature as well. The last but not the least, layer normalization does not adjust the effective scale of the skip connection in forward propagation such that the effective scale becomes 1 and regresses to the plain residual structure. The learned effective scale is 2.82 and 2.64 for 2xSkip+LN and 2rSkip+LN, respectively, averaged over the shown blocks.

Learning λ To automatically adjust the ratio between the skip and the residual branch, a trivial extension is to replace the fixed scalar λ with a learnable scalar or vector. We also try this method on CIFAR-10 using ResNet-110. Concretely, we adopt the form of $y = \text{LN}(w \cdot x + \mathcal{F}(x))$ (w Skip+LN), as using a parameter vector should be more general and a vector covers all the possible solutions of a scalar. However, the findings are not in favor of this approach over the fixed scalar version. When the w is initialized as 1, the model experiences significantly performance degradation. The averaged error rate over 5 runs is 8.66 and the final norm of a sampled $w \in \mathbb{R}^{32}$ at middle layers is 1.34, suggesting very small scale for each dimension of the skip connection (0.24 on average). The small scale of the skip connection indicates that the model relies more on the residual branch that is parameterized and optimized for the training data, entailing a higher overfitting risk. It can be confirmed from Figure 4, where we show the learning curve of w Skip compared to 2rSkip (+LN is omitted for conciseness). w Skip and 2rSkip both achieve low training loss, but w Skip has difficulty in starting learning and has higher validation loss. We also try different initialization strategy for w Skip, e.g., initializing w as 2 since we find 2 is a good empirical value for rSkip+LN. We observe a performance of 6.45 in terms of error rate, which is indeed better but still worse than rSkip+LN. Moreover, such strategy would make w Skip no better than using pre-determined fixed scalar values, as a proper initialization value still needs to be searched.

4.4 Validation Results

We verify the proposed combination of recursive skip connection with layer normalization on more difficult datasets, i.e., CIFAR-100, WMT-2014 EN-DE and EN-FR datasets. Considering the training budget, we only validate the performance of 2rSkip+LN based on the exploratory results.

The results are shown in Table 6 and Table 7, respectively, which demonstrate the effectiveness of the proposed recursive combination that brings consistent improvements regardless of the model and the dataset. The results are in accordance with our findings in exploratory experiments as well. For CIFAR-100, the best improvement of 1.02 absolute error rates is achieved upon the baseline that is the deepest and should be the most difficult to optimize. The relative improvements across models of different depth also suggest that the representation ability of the deep models still needs to be developed and transferred into real performance elevation, while the proposal shows potential. The results on EN-VI, EN-DE and EN-FR datasets are also encouraging. With approximately one third of the parameters, the Transformer-Base model could be improved to compete with the Transformer-Big model on EN-VI and

Dataset	Method	ResNet-20	ResNet-32	ResNet-44	ResNet-56	ResNet-110
CIFAR-100	Baseline	33.38	32.00	30.50	29.81	28.71
	+ proposal	32.99 (+0.39)	31.39 (+0.61)	29.85 (+0.65)	28.98 (+0.83)	27.69 (+1.02)

Table 6: Results on CIFAR-100 using ResNet and the proposed 2rSkip+LN skip connection construction. Average of 5 runs. The numbers in parentheses denote the improvements. The proposal brings consistent improvements. The improvements are more evident for deeper and complex models.

Method	# Parameters	EN-VI	EN-DE	EN-FR
Transformer-Big (Vaswani et al., 2017)	213M	-	28.4	41.0
Deep Representations (Dou et al., 2018)	356M	-	29.2	-
Evolved Transformer (So et al., 2019)	218M	-	29.8	41.3
Fixup (Zhang et al., 2019)	-	-	29.3	-
AdaNorm (Xu et al., 2019)	-	31.4	28.5	-
ScaleNorm (Nguyen and Salazar, 2019)	-	32.8	27.6	-
Admin (Liu et al., 2020)	-	-	29.0	43.8
Transformer-Base	65M	30.3	27.9	38.2
+ proposal	+0.03M	<u>31.5 (+1.2)</u>	<u>28.6 (+0.7)</u>	<u>39.0 (+0.8)</u>
Transformer-Big	213M	31.4	28.5	41.0
+ proposal	+0.06M	<u>32.2 (+0.8)</u>	<u>29.0 (+0.5)</u>	<u>41.6 (+0.6)</u>

Table 7: Results on EN-VI, EN-DE and EN-FR using Transformer and the proposed 2rSkip+LN skip connection construction. The numbers of parameters are reported for the EN-DE models. As we can see, the proposal brings consistent substantial improvements.

EN-DE, with the help of the refined skip connection architecture. It again demonstrates the potential of the proposal in realizing the capability of the models by inducing more efficient optimization.

5 Conclusion and Future Work

In this work, we focus on exploring the combination of skip connection and layer normalization to build a novel skip connection architecture able to incorporate reasonable skip information at output level and solve the gradient explosion problem in practice. According to the empirical studies, we find that expanding skip information should be helpful to the optimization of the model, which is hindered by the gradient malformation in practice, and layer normalization does better than batch normalization in mitigating such problem. The proposed Recursive Skip Connection with Layer Normalization guarantees the stability of the shortcut gradient by recursively dividing the expanded skip connection and adding it to the output iteratively with layer normalization. The refined skip connection architecture may enable more efficient optimization. Validation on CIFAR-100, WMT-2014 EN-DE and EN-FR proves the findings and the effectiveness of the proposed recursive skip connection construction with best improvements over 1.0 absolute point, demonstrating its generalization ability to a wide range of existing systems.

Unfortunately, a limitation of our investigation is that λ is selected as integer and only a small number of values are used, which may not be of the best granularity for expanding the skip connection. A more desirable solution would be to adaptively determine the value of λ , maybe in terms of the gradient distribution during the learning so that more efficient optimization could be realized. We will further seek to address this limitation in future work.

Acknowledgments

This work is partly supported by AOTO-PKUSZ Joint Research Center for Artificial Intelligence on Scene Cognition Technology Innovation, and Beijing Academy of Artificial Intelligence (BAAI). We thank all the anonymous reviewers for their constructive comments and suggestions. Yuexian Zou is the corresponding author of this paper.

References

- Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *CoRR*, abs/1607.06450.
- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W. Cottrell, and Julian J. McAuley. 2020. Rezero is all you need: Fast convergence at large depth. *CoRR*, abs/2003.04887.
- Mauro Cettolo, Jan Niehues, Sebastian Stüker, Luisa Bentivogli, Roldano Cattoni, and Marcello Federico. 2015. The iwslt 2015 evaluation campaign. In *International Workshop on Spoken Language Translation 2015, IWSLT 2015*.
- Zi-Yi Dou, Zhaopeng Tu, Xing Wang, Shuming Shi, and Tong Zhang. 2018. Exploiting deep representations for neural machine translation. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4253–4262. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016a. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016b. Identity mappings in deep residual networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV*, volume 9908 of *Lecture Notes in Computer Science*, pages 630–645. Springer.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In Francis R. Bach and David M. Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pages 448–456. JMLR.org.
- Alex Krizhevsky. 2009. Learning multiple layers of features from tiny images. Technical report, Computer Science Department, University of Toronto.
- Fenglin Liu, Xuancheng Ren, Yuanxin Liu, Houfeng Wang, and Xu Sun. 2018. simnet: Stepwise image-topic merging network for generating detailed and comprehensive image captions. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 137–149. Association for Computational Linguistics.
- Fenglin Liu, Yuanxin Liu, Xuancheng Ren, Xiaodong He, and Xu Sun. 2019a. Aligning visual regions and textual concepts for semantic-grounded image representations. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 6847–6857.
- Fenglin Liu, Xuancheng Ren, Yuanxin Liu, Kai Lei, and Xu Sun. 2019b. Exploring and distilling cross-modal information for image captioning. In Sarit Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 5095–5101. ijcai.org.
- Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. 2020. Understanding the difficulty of training transformers. *CoRR*, abs/2004.08249.
- Minh-Thang Luong and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domain. In *International Workshop on Spoken Language Translation 2015, IWSLT 2015*.
- Toan Q. Nguyen and Julian Salazar. 2019. Transformers without tears: Improving the normalization of self-attention. In *International Workshop on Spoken Language Translation 2019, IWSLT 2019*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA.*, pages 311–318. ACL.
- Sheng Shen, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2020. Powernorm: Rethinking batch normalization in transformers. *CoRR*, abs/2003.07845.

- David R. So, Quoc V. Le, and Chen Liang. 2019. The evolved transformer. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 5877–5886. PMLR.
- Rupesh Kumar Srivastava, Klaus Greff, and Jürgen Schmidhuber. 2015. Highway networks. In *ICML 2015 Deep Learning Workshop*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 2818–2826. IEEE Computer Society.
- Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In Satinder P. Singh and Shaul Markovitch, editors, *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 4278–4284. AAAI Press.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 6000–6010.
- Ashish Vaswani, Samy Bengio, Eugene Brevdo, François Chollet, Aidan N. Gomez, Stephan Gouws, Llion Jones, Lukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and Jakob Uszkoreit. 2018. Tensor2tensor for neural machine translation. In Colin Cherry and Graham Neubig, editors, *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas, AMTA 2018, Boston, MA, USA, March 17-21, 2018 - Volume 1: Research Papers*, pages 193–199. Association for Machine Translation in the Americas.
- Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F. Wong, and Lidia S. Chao. 2019. Learning deep transformer models for machine translation. In Anna Korhonen, David R. Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 1810–1822. Association for Computational Linguistics.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. On layer normalization in the transformer architecture. *CoRR*, abs/2002.04745.
- Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. 2019. Understanding and improving layer normalization. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pages 4383–4393.
- Hongyi Zhang, Yann N. Dauphin, and Tengyu Ma. 2019. Fixup initialization: Residual learning without normalization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.

A On Adaptive Ratio of Recursive Skip Connection with Layer Normalization

To see how the recursive structure can learn the ratio, we take $\lambda = 2$ as an example and unroll the formula:

$$\begin{aligned}
 \mathbf{y}_2 &= \text{LN}(\mathbf{x} + \text{LN}(\mathbf{x} + \mathcal{F}(\mathbf{x}))) \\
 &= \mathbf{w}_2 \cdot \left(\frac{\mathbf{x} + \left(\mathbf{w}_1 \cdot \left(\frac{\mathbf{x} + (\mathcal{F}(\mathbf{x})) - \mu_1}{\sigma_1} \right) + \mathbf{b}_1 \right) - \mu_2}{\sigma_2} \right) + \mathbf{b}_2 \\
 &= \left(\frac{\mathbf{w}_2}{\sigma_2} + \frac{\mathbf{w}_2 \cdot \mathbf{w}_1}{\sigma_2 \sigma_1} \right) \mathbf{x} + \frac{\mathbf{w}_2 \cdot \mathbf{w}_1}{\sigma_2 \sigma_1} \mathcal{F}(\mathbf{x}) + C,
 \end{aligned} \tag{5}$$

where \cdot denotes entrywise-product, $\mathbf{w}_1, \mathbf{w}_2, \mathbf{b}_1, \mathbf{b}_2$ are the gain and the bias parameters of the affine transformation in layer normalization, and $\sigma_1, \sigma_2, \mu_1, \mu_2$ are the standard deviation and the mean of

$x + \mathcal{F}(x)$ and $x + y_1$, respectively. As we can see, the ratio between the input and the residual is

$$x : \mathcal{F}(x) := \frac{w_2}{\sigma_2} + \frac{w_2 \cdot w_1}{\sigma_2 \sigma_1} : \frac{w_2 \cdot w_1}{\sigma_2 \sigma_1} = \frac{\sigma_1}{w_1} + 1. \quad (6)$$

Because the recursive definition, the inner gain parameters of layer normalization could penetrate through the normalization and effectively act as the learnable ratio. It is also important to note that the ratio is also dependent on the current layer distribution, which could serve as a regularization component. For any λ larger than 2, the general form of the ratio is $1 + \sum_{i=1}^{\lambda} \prod_{j=1}^i \sigma_j / w_j$. It should be reminded that setting $\lambda = 2$ is sufficient for automatically learning the shortcut-residual ratio and dynamically adjusting the ratio based on instant data distribution. Yet with larger λ , the distribution of the higher order combination of x and $\mathcal{F}(x)$ will be considered. The recursive definition introduces gradual and smooth combination of the input and the residual up to the specified order λ which could deal with different distributions better. Our experimental results also demonstrate that the proposal performs better than the scaled shortcut.