

Unified Normalization for Accelerating and Stabilizing Transformers

Qiming Yang

yangqiming5@hikvision.com
Hikvision Research Institute
Hangzhou, China

Kai Zhang

zhangkai23@hikvision.com
Hikvision Research Institute
Hangzhou, China

Chaoxiang Lan

lanchaoxiang@hikvision.com
Hikvision Research Institute
Hangzhou, China

Zhi Yang

yanzhi13@hikvision.com
Hikvision Research Institute
Hangzhou, China

Zheyang Li

lizheyang@hikvision.com
Hikvision Research Institute &
Zhejiang University
Hangzhou, China

Wenming Tan

tanwenming@hikvision.com
Hikvision Research Institute
Hangzhou, China

Jun Xiao

junx@cs.zju.edu.cn
Zhejiang University
Hangzhou, China

Shiliang Pu*

pushiliang.hri@hikvision.com
Hikvision Research Institute
Hangzhou, China

ABSTRACT

Solid results from Transformers have made them prevailing architectures in various natural language and vision tasks. As a default component in Transformers, Layer Normalization (LN) normalizes activations within each token to boost the robustness. However, LN requires on-the-fly statistics calculation in inference as well as division and square root operations, leading to inefficiency on hardware. What is more, replacing LN with other hardware-efficient normalization schemes (e.g., Batch Normalization) results in inferior performance, even collapse in training. We find that this dilemma is caused by abnormal behaviors of activation statistics, including large fluctuations over iterations and extreme outliers across layers. To tackle these issues, we propose Unified Normalization (UN), which can speed up the inference by being fused with other linear operations and achieve comparable performance on par with LN. UN strives to boost performance by calibrating the activation and gradient statistics with a tailored fluctuation smoothing strategy. Meanwhile, an adaptive outlier filtration strategy is applied to avoid collapse in training whose effectiveness is theoretically proved and experimentally verified in this paper. We demonstrate that UN can be an efficient drop-in alternative to LN by conducting extensive experiments on language and vision tasks. Besides, we evaluate the efficiency of our method on GPU. Transformers equipped with UN enjoy about 31% inference

speedup and nearly 18% memory reduction. Code will be released at <https://github.com/hikvision-research/Unified-Normalization>.

CCS CONCEPTS

• **Computing methodologies** → **Computer vision tasks; Machine translation.**

KEYWORDS

neural networks, normalization, transformers

ACM Reference Format:

Qiming Yang, Kai Zhang, Chaoxiang Lan, Zhi Yang, Zheyang Li, Wenming Tan, Jun Xiao, and Shiliang Pu. 2022. Unified Normalization for Accelerating and Stabilizing Transformers. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*, October 10–14, 2022, Lisboa, Portugal. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3503161.3547860>

1 INTRODUCTION

Transformers [38] are initially introduced for Natural Language Processing (NLP) tasks [10, 28]. Since Transformers make few assumptions about the structural bias of input data, these architectures can be universally and flexibly applied in other scenarios, such as multi-modal and speech tasks [20, 40]. The basic modules of Transformers are stackable multi-head self-attention (MHSA) and feed-forward network (FFN) that enable the capture of long-term dependencies between tokens. In these basic modules, Layer Normalization (LN) [2] is chosen as a default component that releases the training process from heavy dependency on mini-batch samples to handle variable-length input. However, LN requires additional computation and memory overheads during inference because of the on-the-fly statistics, as well as division and square root operations. LN thus is inefficient and hardly meets industrial needs [34, 45]. Nevertheless, replacing LN with Batch Normalization (BN) [19] leads to inferior performance in NLP tasks [35]. Shen *et al.* [35] show that the large fluctuations in Transformers within

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

MM '22, October 10–14, 2022, Lisboa, Portugal

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9203-7/22/10...\$15.00

<https://doi.org/10.1145/3503161.3547860>

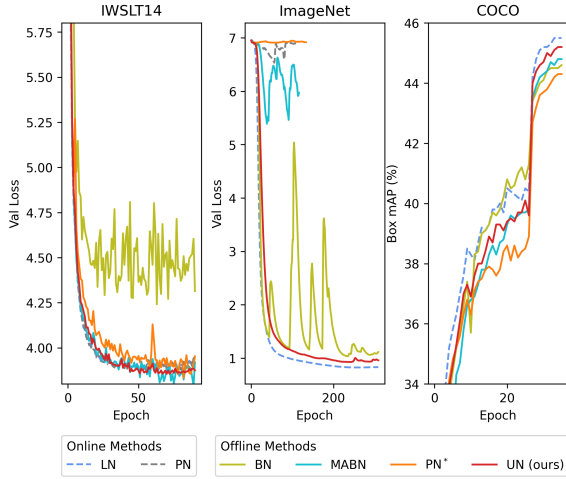


Figure 1: Performance comparison on Transformer (IWSLT14), T2T-ViT-14 (ImageNet), and Swin-T (COCO) during training. The offline methods that could be fused into other linear operations in inference are all plotted with a solid line. PN originally comes with a layer-scale layer, which is removed in PN*.

activation statistics result in poor performance. Recently, Transformers are broadly proliferated to Computer Vision (CV) tasks that have led to a series of breakthroughs in image classification, object detection, instance segmentation, etc., also known as Vision Transformers (ViTs) [4, 7, 12–14, 26, 36, 48, 49]. LN is directly inherited from the original Transformer as an essential component in these ViTs despite the fixed-length inputs. Similar to NLP tasks, significant performance degradation (even collapse in training) will also be triggered by replacing LN with BN in Transformers in CV tasks [6, 33, 46]. Shao *et al.* [33] hold the view that BN is harmful to ViTs that result in performance degradation. In previous works [6, 46], the authors claim that replacing all LN with BN in ViTs leads to convergence problems. Since the naive replacement for LN leads to inferior performance and instability, the deployment of Transformers still suffers from the on-the-fly statistics computation.

There are two main methods to improve the hardware efficiency of normalization in Transformers. 1) Simplifying the computation of online statistics [35, 50] and removing inefficient operations (e.g., square root) [23]. Although the performance is maintained, the dynamic calculation for online statistics still exists. 2) Removing the computation for online statistics that utilizes fixed statistics during inference as an offline method, such as BN [19]. In this way, inference can be sped up by circumventing the redundant statistics computation getting rid of division and square root operations. As there is no free lunch, significant performance drop and convergence problems are reported in these works [35, 45, 46]. MABN [45] and PN [35] utilize moving average strategies to mitigate the fluctuations in Transformers in NLP tasks. However, these mentioned methods are task-specific where the inferior performance and instability still exist in ViTs.

To address the issues above, we dissect the abnormal behaviors of statistics in Transformers. We investigate the activation statistics under moving average strategies in the training of Transformers. We uncover that the fluctuations of the activation statistics are more drastic than that of the gradient statistics during training. Moreover, we find the range of activation statistics task-agnostically keeps increasing along with both the depth and the progress of training, in which the risk of outliers arises. In this sense, extreme outliers are nearly inevitable and continually deteriorate the consistency between activation and gradient statistics. These observations illustrate that the inferior performance and instability (shown in Figure 1) are very like boil down to abnormal behaviors of activation statistics in Transformers.

In this paper, we aim to replace LN with an offline method to promote applications for Transformers in language and vision tasks. We propose Unified Normalization (UN) to accelerate inference in Transformers and achieve comparable performance with LN. Specifically, we design a tailored fluctuation smoothing strategy to deal with the fluctuations of different degree in activation and gradient statistics. At the same time, an adaptive outlier filtration strategy is introduced to ensure stable convergence, where the impact of outliers can be proved to be significantly reduced both in theory and experiments. Extensive experiments demonstrate the effectiveness of UN. In a nutshell, our contributions are as follows:

- We analyze the abnormal behaviors of activation statistics in Transformers and find the large fluctuations and extreme outliers are responsible for inferior performance and instability.
- A tailored fluctuation smoothing strategy is designed to calibrate the activation and gradient statistics and boost the performance.
- An adaptive outlier filtration strategy is introduced to reduce the impact of extreme outliers on the basis of theoretical analysis.
- Extensive evaluations in neural machine translation, image classification, object detection, and instance segmentation illustrate the superiority of our method, which is capable of being a drop-in alternative to LN in Transformers. Furthermore, we show that Transformers equipped with UN gain nearly **18%** memory reduction and over **31%** speedup in inference on GPU.

2 RELATED WORK

2.1 Transformers and Vision Transformers

Transformers [38] initially show surprising capability in sequence modeling and neural machine translation. Owing to the high flexibility, Transformers [3, 10, 15, 47] have become the most recent dominant architectures over various NLP tasks and speech tasks. In 2020, Carion *et al.* [4] propose the first end-to-end Transformer-based detector DETR. Later, ViT [13] is proposed as the very first pure Transformer in CV tasks. The following years have witnessed explosive development of Transformers in CV tasks. T2T [49] introduces token-to-token module that combines adjacent tokens in early stage to model local information. Swin [26] and Swin V2 [25] have applied window-based attention to reduce the overhead of computation in MHSA and achieve state-of-the-art performance.

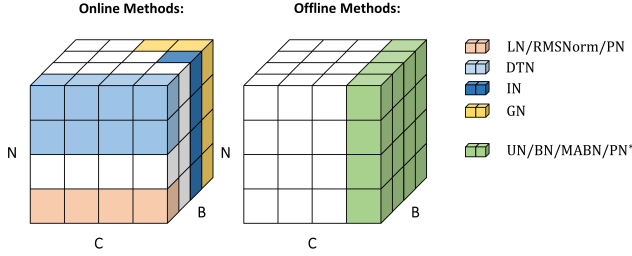


Figure 2: Normalization methods. Each subfigure shows a feature map tensor, where B is the batch axis, N is the number of tokens (or the sequence length) axis, and C is the channel (also known as the embedding size) axis.

Most recently, Ding *et al.* [11] explore attention in both spatial and channel tokens to propose a powerful backbone DaViT for vision tasks. In these aforementioned Transformers and Vision Transformers, LN is the preferred choice for normalization.

In addition, some works [7, 14, 18] originally build Transformers with BN, whilst these works need elaborate design on convolutional operations to stabilize training that have considerably modified the original ViT [13].

2.2 Normalization Methods

Normalization is widely used for stabilizing training and boosting performance in deep neural networks [17]. As illustrated in Figure 2, related normalization methods could be categorized into online methods and offline methods according to whether the inference statistics can be fused or not.

2.2.1 Online Methods. Online methods require the calculation of on-the-fly statistics during training as well as inference. IN [37], GN [41], and LN [2] are representative online methods that calculate statistics in different dimensions as shown in Figure 2. Switchable Normalization [27] learns to switch between different types of normalization by learning their importance weights. It is widely believed that LN is customized for variable-length NLP samples [35]. With the rise of Vision Transformers (ViTs) [13, 26, 39, 49], LN has also become a preferred choice for CV tasks. Lately, DTN [33] exploits the connection within adjacent tokens to improve the performance of LN in ViTs. To make LN more hardware-efficient, some works [23, 44, 50] attempt to reduce the cost of computation in LN. Zhang *et al.* [50] propose a simpler method RMSNorm that scales inputs by the root mean square. However, the inefficient dynamic calculation for online statistics is still not fundamentally removed.

2.2.2 Offline Methods. Offline methods use estimated inference statistics that could be frozen for arbitrary inputs. Only a point-wise add and multiplication are required during inference that enables the fusion of offline methods with adjacent linear operations. In this way, offline methods can be removed entirely from models and achieve efficient inference [45, 46]. However, once these methods cooperate with Transformers, large fluctuations over iterations will lead to performance degradation and even collapse in training [6, 33, 35, 46]. Yao *et al.* [46] find Transformers trained with

Algorithm 1 Fusing Normalization

Input: $\gamma, \beta, \mu, \sigma^2 \in \mathbb{R}^C$ // in Equation 1
 $W \in \mathbb{R}^{C_{out} \times C}, b \in \mathbb{R}^{C_{out}}$ // parameters in the subsequent layer
Output: $W' \in \mathbb{R}^{C_{out} \times C}, b' \in \mathbb{R}^{C_{out}}$ // fused parameters

```

1:  $\tilde{\gamma} = \gamma / \sigma$ 
2:  $\tilde{\beta} = \beta - \tilde{\gamma} \cdot \mu$ 
3:  $b' = b + W \times \tilde{\beta}$  //  $y = W(Norm(x)) + b$  is equivalent to
4:  $W' = W \cdot (\mathbf{1}_{C_{out}} \times \tilde{\gamma}^T)$  //  $y = W'x + b'$ 

```

BN are very unstable and crash irregularly. Chen *et al.* [6] attempt to partially replace LN with BN in FFN for stabilizing the training of Transformers. To mitigate the impact of large fluctuations, MABN [45] leverages exponential moving average statistics in activation statistics, and accordingly uses simple moving average statistics in gradient statistics to estimate gradients. Similarly, Shen *et al.* [35] propose PN* that uses exponential moving average statistics in both activation and gradient statistics. Preceding works aim to improve the efficiency of LN in Transformers but still suffer from inferior performance and instability. Thus, it is valuable for the community to design a more effective and robust method.

3 METHOD

In this section, we describe the design process of Unified Normalization (UN). First, we develop a unified framework for leveraging offline methods. Based on the framework, we next apply a tailored fluctuation smoothing strategy to mitigate the fluctuations and an adaptive outlier filtration strategy for stabilizing training.

3.1 Unified Framework

We develop a unified framework for applying offline methods in Transformers. In this pipeline, the inference statistics are fixed so that they could be fused with other linear operations for speedup.

For a normalization layer, let $X \in \mathbb{R}^{B \times C}$ and $Y \in \mathbb{R}^{B \times C}$ denote the input and output, where B is the batch size and C indicates the number of channels. Note that the number of tokens N , which could be squeezed into B , is omitted in this section for clarity. For arbitrary input in inference, all offline methods perform in a unified manner

$$Y = \gamma \cdot \frac{X - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta. \quad (1)$$

Here, ϵ is a small constant, and $\gamma, \beta \in \mathbb{R}^C$ are learnable parameters. The inference statistics $\mu, \sigma^2 \in \mathbb{R}^C$ are estimated in the training process and independent of inputs. Since the statistics and parameters are fixed during inference, offline normalization can be merged. The pseudo code for fusing offline normalization with adjacent linear operation can be found in Algorithm 1. On the contrary, LN requires calculation for on-the-fly statistics $\mu_{LN} = \mu_{LN}(X), \sigma_{LN}^2 = \sigma_{LN}^2(X)$ that consumes extra computation time.

In forward propagation of training, the normalization procedure is shown as follows,

$$Z_t = \frac{X_t - \hat{\mu}_t}{\sqrt{\hat{\sigma}_t^2 + \epsilon}}, \quad (2)$$

$$Y_t = \gamma \cdot Z_t + \beta. \quad (3)$$

Let Z_t denote the normalized alternative to input X_t at iteration t . The training statistics for normalizing are marked as $\hat{\mu}_t$ and $\hat{\sigma}_t^2$, given by

$$\hat{\mu}_t = \Theta_\mu(\mu_t, \dots, \mu_{t-M+1}), \quad (4)$$

$$\hat{\sigma}_t^2 = \Theta_{\sigma^2}(\sigma_t^2, \dots, \sigma_{t-M+1}^2). \quad (5)$$

Here, $\mu_t, \dots, \mu_{t-M+1}$ and $\sigma_t^2, \dots, \sigma_{t-M+1}^2$ are sequences of recorded statistics from recent M iterations. We consider μ_t and σ_t^2 to be the first-moment and second-moment statistics for current input X_t . In general, the training statistics can be used to update the inference statistics by applying moving averages. In backward propagation, the gradients of loss L pass as:

$$\frac{\partial L}{\partial Z_t} = \gamma \cdot \frac{\partial L}{\partial Y_t}, \quad (6)$$

$$\frac{\partial L}{\partial X_t} = \frac{1}{\sqrt{\hat{\sigma}_t^2 + \epsilon}} \left(\frac{\partial L}{\partial Z_t} - \psi_{\hat{\mu}_t} - Z_t \cdot \psi_{\hat{\sigma}_t^2} \right). \quad (7)$$

Giving gradients $\frac{\partial L}{\partial Y_t}$, $\psi_{\hat{\mu}_t}$ and $\psi_{\hat{\sigma}_t^2}$ indicate the gradient statistics that used for estimating $\frac{\partial L}{\partial X_t}$. In this framework, estimated gradients are gained from averaging functions Θ_{g_μ} and $\Theta_{g_{\sigma^2}}$,

$$\psi_{\hat{\mu}_t} = \Theta_{g_\mu}(g_{\hat{\mu}_t}, \dots, g_{\hat{\mu}_{t-M+1}}), \quad (8)$$

$$\psi_{\hat{\sigma}_t^2} = \Theta_{g_{\sigma^2}}(g_{\hat{\sigma}_t^2}, \dots, g_{\hat{\sigma}_{t-M+1}^2}). \quad (9)$$

The gradients passed from $\hat{\mu}_t$ and $\hat{\sigma}_t^2$ are denoted as $g_{\hat{\mu}_t}$ and $g_{\hat{\sigma}_t^2}$.

Offline Methods in the Unified Framework. With the help of the unified framework, offline methods can be expressed by choosing different statistical objects and averaging functions Θ .

For instance, BN can be formulated by choosing the mean and variance for the first-moment and second-moment statistics, i.e.,

$$\mu_t = \frac{1}{B} \sum_{i=1}^B x_i, \quad \sigma_t^2 = \frac{1}{B} \sum_{i=1}^B (x_i - \mu_t)^2. \quad (10)$$

Then setting averaging functions as:

$$\hat{\mu}_t = \mu_t, \quad \hat{\sigma}_t^2 = \sigma_t^2, \quad \psi_{\hat{\mu}_t} = g_{\hat{\mu}_t}, \quad \psi_{\hat{\sigma}_t^2} = g_{\hat{\sigma}_t^2}. \quad (11)$$

Here, the averaging functions simply focus on statistics of current iteration t and ignore the last $M-1$ statistics in the sequences. During training, the inference statistics are updated as,

$$\mu = \alpha \mu + (1 - \alpha) \hat{\mu}_t, \quad \sigma^2 = \alpha \sigma^2 + (1 - \alpha) \hat{\sigma}_t^2. \quad (12)$$

As illustrated in Equation 11, BN merely focuses on the activations in the current iteration. This makes BN fragile for large fluctuations over iterations.

For MABN [45], the authors reduce the number of statistics for stabilizing training and remove the first-moment statistic. Hence, the quadratic mean is chosen as the second-moment statistic:

$$\sigma_t^2 = \frac{1}{B} \sum_{i=1}^B x_i^2. \quad (13)$$

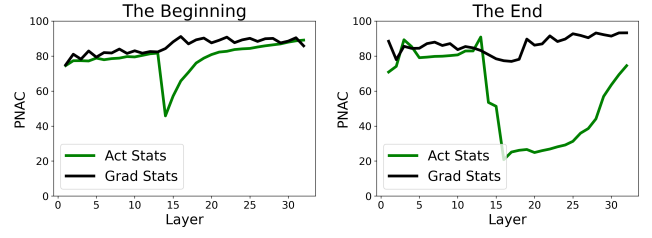


Figure 3: The average PNAC of activation and gradient statistics over iterations in Transformer. A higher PNAC indicates milder fluctuations.

Algorithm 2 Fluctuation Smoothing

Forward Propagation

Input: $X_t \in \mathbb{R}^{B \times C}$

Output: $Y_t \in \mathbb{R}^{B \times C}$

- 1: $\sigma_t^2 = \frac{1}{B} \sum_{i=1}^B x_{t,i}^2$ // mini-batch quadratic mean
- 2: $\hat{\sigma}_t^2 = \sqrt{\prod_{i=0}^{M-1} \sigma_{t-i}^2}$ // geometric mean
- 3: $Z_t = \frac{X_t}{\sqrt{\hat{\sigma}_t^2 + \epsilon}}$ // normalizing
- 4: $Y_t = \gamma \cdot Z_t + \beta$ // re-scaling and shifting
- 5: $\sigma^2 = \alpha \sigma^2 + (1 - \alpha) \hat{\sigma}_t^2$ // updating for inference

Backward Propagation

Input: $\frac{\partial L}{\partial Y_t} \in \mathbb{R}^{B \times C}$

Output: $\frac{\partial L}{\partial X_t} \in \mathbb{R}^{B \times C}$

- 1: $\frac{\partial L}{\partial Z_t} = \gamma \cdot \frac{\partial L}{\partial Y_t}$
 - 2: $g_{\hat{\sigma}_t^2} = \frac{1}{B} \sum_{i=1}^B \frac{\partial L}{\partial z_i} z_i$ // gradients from $\hat{\sigma}_t^2$
 - 3: $\psi_{\hat{\sigma}_t^2} = \alpha \psi_{\hat{\sigma}_{t-1}^2} + (1 - \alpha) \frac{1}{M} \sum_{i=0}^{M-1} g_{\hat{\sigma}_{t-i}^2}$ // estimating gradients
 - 4: $\frac{\partial L}{\partial X_t} = \frac{1}{\sqrt{\hat{\sigma}_t^2 + \epsilon}} \left(\frac{\partial L}{\partial Z_t} - Z_t \cdot \psi_{\hat{\sigma}_t^2} \right)$
-

The averaging functions for MABN are set as follow,

$$\hat{\mu}_t = 0, \quad \hat{\sigma}_t^2 = \text{EMAS}^1, \quad \psi_{\hat{\mu}_t} = 0, \quad \psi_{\hat{\sigma}_t^2} = \text{SMAS}^2. \quad (14)$$

The inference statistics σ^2 for MABN are updated the same as Equation 12. With solely applying EMAS in activation statistics, it is hard for MABN to avoid influence from extreme outliers.

3.2 Fluctuation Smoothing

Analysis with Normality Test. We dive deeper to analyze the abnormal behaviors of activation statistics in Transformers. To investigate the magnitude of fluctuations in activation and gradient statistics, we quantitatively analyze the abnormal behaviors with quadratic mean as the second-moment statistics. For statistics $\sigma_t^2 \in \mathbb{R}^C$, we conduct the normality test[8] on a sequence of $\sigma_t^2, \dots, \sigma_{t-M+1}^2$ to establish whether or not the sequence comes from a normally distributed population, then calculate the percentage of channels held for normality. We define the Percentage of Normality over All Channels (PNAC) which measure the degree of

¹EMAS (Exponential Moving Average Statistics) for K : $K = \eta \cdot K + (1 - \eta) \cdot K_t$

²SMAS (Simple Moving Average Statistics) for K with a window size M : $K = \frac{1}{M} \sum_{i=0}^{M-1} K_{t-i}$

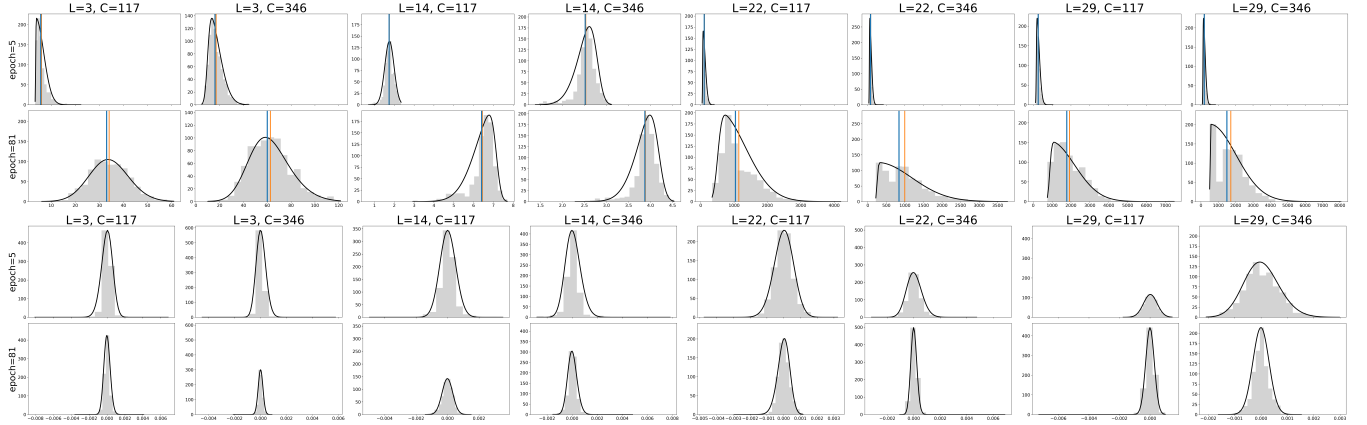


Figure 4: The activation (the 1st and 2nd rows) and gradient (the 3rd and 4th rows) statistics in channel C of normalization layer L. In activation statistics, we show the GM and AM of activation statistics in solid 'blue' and 'orange' lines respectively.

fluctuations in statistics:

$$PNAC = \frac{|\{c_i | \text{Normality Test}(c_i), p > 0.05\}|}{C} \times 100\%, i = 1, \dots, C. \quad (15)$$

The lower PNAC, the larger fluctuations in statistics. In this way, we compute PNAC for each layer averaged over iterations and plot it in Figure 3. At the very beginning of training, both activation and gradient statistics mildly fluctuate over iterations. At the end of the training, there is a significant drop in the PNAC of activation statistics which means large fluctuations exist in activation statistics. On the contrary, we find that there are milder fluctuations in gradient statistics.

Moreover, we next visualize the activation and gradient statistics in different layers and channels, as shown in Figure 4. The range of activation statistics gets larger along the depth and training process. We find that the skewed distribution of activation statistics contains extreme outliers that could impact the arithmetic mean. We thus turn to adopt geometric mean (GM) with less sensitivity to outliers instead of arithmetic mean (AM) to gain a better representation of activation statistics in a skewed distribution. The averaging functions are defined as:

$$\hat{\mu}_t = 0, \quad \hat{\sigma}_t^2 = \sqrt{\prod_{i=0}^{M-1} \sigma_{t-i}^2}, \quad (16)$$

$$\psi_{\hat{\mu}_t} = 0, \quad \psi_{\hat{\sigma}_t^2} = \alpha \psi_{\hat{\sigma}_{t-1}^2} + (1 - \alpha) \frac{1}{M} \sum_{i=0}^{M-1} g_{\hat{\sigma}_{t-i}^2}. \quad (17)$$

By applying quadratic mean as the second-moment statistics, we visualize the GM and AM of activation statistics in Figure 4: in an approximately normal distribution (i.e., mild fluctuations), GM is close to AM; in skewed distribution (i.e., large fluctuations), the extreme outliers greatly influenced AM, while GM is still close to the majority. Owing to the gradient statistics $g_{\hat{\sigma}_t^2}$ are first-moment statistics and do not obey non-negativity constraints, it is unable to use GM directly. Therefore, we utilize AM in gradient statistics that further with a momentum for gradient estimation in backward propagation. Specially, we leverage the quadratic mean as the second-moment statistics in our method to reduce the number of statistics that ensure the stability in applying moving average

strategies, as proved in [45]. Our strategy could be formulated as shown in Algorithm 2. Omitting the impact of updated weights over different iterations, UN is set with moderate window sizes.

3.3 Outlier Filtration

Although the fluctuation smoothing is leveraged to calibrate the activation statistics, extreme outliers are observed and somehow lead to instability in training (as shown in Figure 1). With the moving average strategies (in Equation 4 and 5) applied to activation statistics, it is impossible to calculate the accurate gradients from previous iterations [17]. Once extreme outliers deteriorate the gradient estimation error, the risk of instability increases. Based on the assumption, we attempt to take a step further by introducing an adaptive outlier filtration strategy. More specifically, the main goal of outlier filtration is to decide when to apply the moving average strategies. To identify outliers, we set an adaptive threshold for outlier filtration with the *AM - GM inequality* [1]. Let $\Omega_t = (\sigma_t^2, \sigma_{t-1}^2, \dots, \sigma_{t-M+1}^2)$ denote the M recent activation statistics recorded in forward propagation at iteration t , where $M > 1$, then we have

$$E(\Omega_t) - \Pi(\Omega_t) \leq M \cdot V(\Omega_t^{\frac{1}{2}}), \quad (18)$$

where $\Omega_t^{\frac{1}{2}} = (\sigma_t, \dots, \sigma_{t-M+1})$ and $V(\cdot)$, $E(\cdot)$, $\Pi(\cdot)$ are operators that calculate the variance, arithmetic mean, and geometric mean for input respectively. The extreme outliers will enlarge variances. We thus use the upper bound of the last iteration to detect outliers for the current iteration. Hence, $M \cdot V(\Omega_{t-1}^{\frac{1}{2}})$ can be used as an adaptive threshold for outlier filtration. That is to say, once the mini-batch is deemed to contain extremely large outliers and all the moving average strategies will be dropped in a specific normalization layer,

$$\begin{cases} \hat{\sigma}_t^2 = \sigma_t^2, & \psi_{\hat{\sigma}_t^2} = g_{\hat{\sigma}_t^2} & \text{if } E(\Omega_t) - \Pi(\Omega_t) > M \cdot V(\Omega_{t-1}^{\frac{1}{2}}) \\ \text{Equation (16) and (17)} & & \text{otherwise.} \end{cases} \quad (19)$$

The current statistics σ_t^2 and $g_{\sigma_t^2}$ will be used for forward-propagating and backward-propagating once outliers are found. In Equation 19, the threshold for outlier filtration is independent of the specific input X_t , making this strategy more adaptive to the ever-changing activation statistics during training. Besides, σ^2 and $\psi_{\sigma_t^2}$ are used to update the recorded statistics at iteration t for passivating the outliers in moving average. The rest of the operations are just the same as Algorithm 2.

LEMMA 3.1. Let $A_t = (a_t, a_{t-1}, \dots, a_{t-M+1})$ and $A_{t-1} = (a_{t-1}, a_{t-2}, \dots, a_{t-M})$ are two vectors satisfying $a_i > 0$ and $a_i < a_t, \forall a_i \in A_{t-1}$. $E(\cdot)$, $\Pi(\cdot)$, and $V(\cdot)$ denote the calculation of arithmetic mean, geometric mean, and variance for an arbitrary vector. If $M \cdot V(A_{t-1}^{\frac{1}{2}}) < E(A_t) - \Pi(A_t)$ holds, then $\lambda = \frac{\Pi(A_t)}{a_t} < 1$.

PROOF. From the AM-GM inequality and the lemma condition, we have the following inequality

$$\Pi(A_t) - \Pi(A_{t-1}) < E(A_t) - E(A_{t-1}). \quad (20)$$

Since $E(A_t) - E(A_{t-1}) = \frac{a_t - a_{t-M}}{M}$, then

$$M \cdot (\Pi(A_t) - \Pi(A_{t-1})) < a_t - a_{t-M}. \quad (21)$$

With $a_t > 0$, the above inequality can be transformed to

$$M \cdot \lambda (1 - \sqrt[M]{\frac{a_{t-M}}{a_t}}) < 1 - \frac{a_{t-M}}{a_t}. \quad (22)$$

Therefore,

$$\lambda < \frac{1}{M} \cdot \frac{1 - \frac{a_{t-M}}{a_t}}{1 - \sqrt[M]{\frac{a_{t-M}}{a_t}}} < \frac{1}{M} \cdot M = 1. \quad (23)$$

The last inequality holds because function $f(x) = \frac{1-x}{1-\sqrt[M]{x}}$ is monotonically increasing in $[0, 1)$. \square

COROLLARY 3.2. For a Network using UN without outlier filtration, let $g_{\sigma_t^2}$ denote the ground truth gradient computed based on the chain rule and $\tilde{g}_{\sigma_t^2}$ denote estimated gradient given by a averaging function. If x_t contains an outlier, then

$$\frac{g_{\sigma_t^2}}{\tilde{g}_{\sigma_t^2}} < \frac{1}{M}. \quad (24)$$

PROOF. UN uses geometric mean to estimate the training statistics, which modifies the current mini-batch statistics but detaches from the backward pass. Hence,

$$g_{\sigma_t^2} = \frac{\partial L}{\partial \sigma_t^2} = \frac{\partial L}{\partial \hat{\sigma}_t^2} \cdot \frac{\partial \hat{\sigma}_t^2}{\partial \sigma_t^2} = \tilde{g}_{\sigma_t^2} \frac{\partial \hat{\sigma}_t^2}{\partial \sigma_t^2}. \quad (25)$$

Combined with Lemma 3.1, we have

$$\frac{g_{\sigma_t^2}}{\tilde{g}_{\sigma_t^2}} = \frac{\partial \hat{\sigma}_t^2}{\partial \sigma_t^2} = \frac{1}{M} \cdot \frac{\hat{\sigma}_t^2}{\sigma_t^2} = \frac{1}{M} \cdot \frac{\Pi(\Omega_t)}{\sigma_t^2} < \frac{1}{M}. \quad (26)$$

\square

With the adaptive outlier filtration strategy proposed in this section, we can avoid a catastrophic gradient estimation error. Based on Corollary 3.2, we prove that the gradient estimation error will be shrunk by a factor $1/M$ when an outlier is found.

Table 1: The performance (BLEU [31], higher is better) of Transformers on neural machine translation. 'Offline' indicates a method can be fused in inference. 'NoNorm' means models without normalization. 'FAIL' indicates collapse during training. PN* is PN without a layer-scale layer.

Method	Offline	IWSLT14		WMT14	
		BLEU	Δ	BLEU	Δ
LN [2]	✗	35.3		40.0	
RMSNorm [50]	✗	35.3	0.0	39.8	-0.2
PN [35]	✗	35.3	0.0	39.8	-0.2
NoNorm	/	FAIL	/	32.8	-7.2
BN [19]	✓	31.1	-4.2	35.1	-4.9
MABN [45]	✓	35.4	+0.1	36.5	-3.5
PN* [35]	✓	35.0	-0.3	39.7	-0.3
UN	✓	35.4	+0.1	39.9	-0.1

Table 2: The performance (Top-1 accuracy %) of image classification on ImageNet-1K and CIFAR10/100.

Method	Offline	Swin-T		T2T-ViT-14 [†]					
		ImageNet		ImageNet		CIFAR10		CIFAR100	
		Top1	Δ	Top1	Δ	Top1	Δ	Top1	Δ
LN [2]	✗	81.3		81.5		98.3		88.4	
BN [19]	✓	80.8	-0.5	79.8	-1.7	96.6	-1.7	88.2	-0.2
MABN [45]	✓	80.9	-0.4	FAIL	/	/	/	/	/
PN* [35]	✓	80.9	-0.4	FAIL	/	/	/	/	/
UN	✓	81.0	-0.3	80.9	-0.6	98.3	0.0	88.9	+0.5

[†]: On CIFAR10/100, models are initialized with pre-trained weights from ImageNet-1K. Note that T2T-ViT-14 cooperates with MABN and PN* crash during training on ImageNet, we thus do not report the corresponding results on CIFAR10/100.

4 EXPERIMENTS

4.1 Implementation Details

To put all experiments on an equal footing, we simply replace all LN in corresponding architectures with its drop-in counterparts, without varying the position of the normalization layer or adding extra operators. All models are trained and tested with the same configurations. To simplify the settings, the momentum of UN is set as $\alpha = 0.9$ (the same as BN) to avoid repeatedly tuning hyper-parameter over different tasks. Akin to [45], we set a warming-up step for UN, 4K by default. To show robust results, we report the average performance from 5-run results for small datasets, IWSLT14, CIFAR10, and CIFAR100. See Appendix B for more experimental details.

4.2 Results

4.2.1 Neural Machine Translation. The comparison between online and offline methods is listed in Table 1. Although RMSNorm [50] and PN achieve comparable results with LN, online methods are not able to access efficient deployment on hardware. We report the performance of the Transformers trained without normalization layers (marked as 'NoNorm'). The instability in training and declined performance highlight the necessity of applying normalization. Besides, previous offline methods, such as BN, MABN, and

Table 3: Object detection on COCO val2017 with Faster R-CNN using Swin-T as the backbone. All models are trained with 36 epochs.

Method	Offline	AP ^{Box}	Δ	AP ^{Box} ₅₀	AP ^{Box} ₇₅	AP ^{Box} _s	AP ^{Box} _m	AP ^{Box} _l
LN [2]	✗	45.5		67.5	50.2	31.5	48.8	58.4
BN [19]	✓	44.6	-0.9	66.8	48.9	30.4	48.0	57.8
MABN [45]	✓	44.8	-0.7	66.8	49.0	29.0	47.9	57.3
PN* [35]	✓	44.3	-1.2	66.7	48.4	29.6	47.4	57.2
UN	✓	45.2	-0.3	67.2	49.7	30.1	48.4	58.2

PN* suffer from degradation of performance. Instead, our method outperforms other offline methods and achieves more balanced results that are on par with online methods on both IWSLT14 and WMT14.

4.2.2 Image Classification. Table 2 reports the results of T2T-ViT-14 and Swin-T on ImageNet. In T2T-ViT-14, MABN and PN* that leverage vanilla moving average strategies experience divergence in training. BN appears instability since the early stage of training (shown in Figure 1), leading to degradation in accuracy. UN enjoys stability during training and obtains an improvement of 1.1% over BN. In Swin-T, the top-1 accuracy of BN drops by 0.5%. UN surpasses other offline methods and restores the accuracy to 81.0%. After that, we evaluate UN on downstream classification tasks (CIFAR10/100). The training loss fluctuates dramatically and irregularly in T2T-ViT-14 finetuned with BN, leading to a significant drop in accuracy. UN converges stably without tuning any settings and outperforms LN on top-1 accuracy.

4.2.3 Object Detection and Instance Segmentation. In Table 3, our method restores the performance from other offline methods, with only a slight decrease of 0.3% mAP compared to LN. In Table 4, the result also reveals that our method achieves comparable performance with LN. Here, we show UN surpasses other offline methods that draw strength from the fluctuation smoothing and outlier filtration. With these results, the conclusion easily comes to light that UN could be generalized to other vision tasks more than only image classification.

4.3 Analysis

4.3.1 Ablation Study. In Table 5, we ablate UN to verify the contribution of the basic components. Note that in the fluctuation smoothing of UN, GM, AM, and momentum α are all used as moving average strategies. We remove the basic components one by one, as shown in Exp2-5. In Exp2, the performance deteriorates without any moving strategies applied for activation statistics. In gradient estimation, we conduct a compound moving strategy for gradient statistics, including arithmetic mean and momentum α . In exp3-4, performance degrades after removing any basic part of the compound moving strategy, which implies both of them are contributed to better estimation for gradients. To compare to Exp2, we remove all moving strategies from gradient statistics in exp5. It turns out that models suffer from performance loss without gradient estimation. In this result, we show the fluctuation smoothing proposed in this paper has empowered UN to gain solid performance. Additionally, we also report the ablation study on IWSLT14,

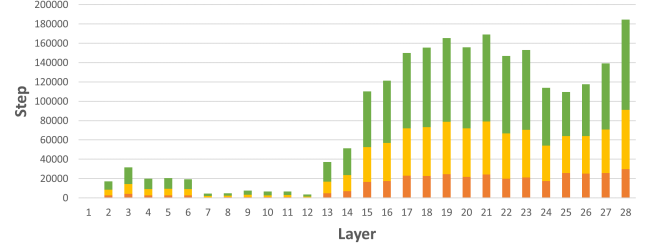


Figure 5: The accumulated steps of dropping moving averages in the outlier filtration during training. We plot the outcomes in red, orange, and green after training 12, 24, and 36 epochs, respectively.

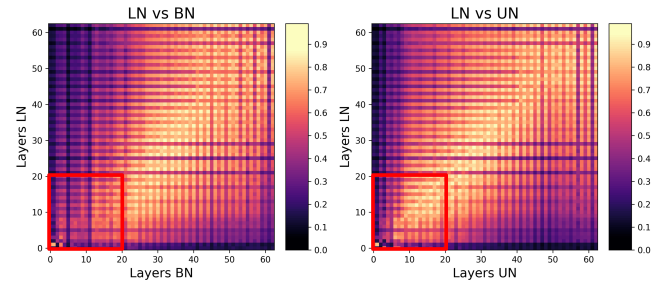


Figure 6: Comparing the similarity of feature maps across T2T-ViT-14 (ImageNet) with shallow layers highlighted in the red box. The similarity is measured with Centered Kernel Alignment (CKA) [29] over layers including normalization layers, MHSA, and FFN in T2T-ViT-14. All LN layers in T2T-ViT-14 are simply replaced with BN and UN.

whose details can be found in Appendix C.1. On IWSLT14, the models can still benefit from the fluctuation smoothing for further improvement.

4.3.2 Effect of the Window Size. We compare different window sizes $M \in \{2, 4, 6, 8, 10\}$ on COCO val2017. Table 6 reports the effect of window sizes in UN. When UN is set with moderate window sizes, the models converge stably and gain competitive performance at the end.

4.3.3 Outlier Filtration. As illustrated in Figure 1, it is easy to see that the outlier filtration stabilizes the training of UN in T2T-ViT while other offline methods tend to crash during training. Table 7 showcases the results with and without outlier filtration. In Transformers, the fluctuations in activation statistics increase along with the depth. Figure 5 shows the accumulated steps of iterations that have found outliers. Outliers tend to emerge from a deeper layer. As the model is trained with more epochs, the percentage increases. The observation implies that Transformers trained with offline methods might get larger fluctuations when scaling up the depth. This result also reveals that the fluctuations also increase along with the training process.

4.3.4 Feature Similarity. With employing Centered Kernel Alignment (CKA), a widely-used representation similarity metric, we can

Table 4: Object detection and semantic segmentation on COCO val2017 with Mask R-CNN using Swin-T as the backbone. All models are trained with 36 epochs.

Method	Offline	AP ^{Box}	Δ	AP ^{Box} ₅₀	AP ^{Box} ₇₅	AP ^{Box} _s	AP ^{Box} _m	AP ^{Box} _l	AP ^{Mask}	Δ	AP ^{Mask} ₅₀	AP ^{Mask} ₇₅	AP ^{Mask} _s	AP ^{Mask} _m	AP ^{Mask} _l
LN [2]	✗	46.0		68.1	50.3	31.2	49.2	60.1	41.6		65.1	44.9	25.9	45.1	56.9
BN [19]	✓	44.9	-1.1	67.2	49.0	29.6	48.4	58.3	40.8	-0.8	64.0	43.8	24.9	44.4	55.3
MABN [45]	✓	45.1	-0.9	67.2	49.6	30.0	48.3	57.7	41.0	-0.6	64.2	44.1	24.9	44.7	55.0
PN* [35]	✓	44.6	-1.4	66.8	48.9	29.1	47.6	57.6	40.7	-0.9	63.7	43.6	24.1	43.8	54.9
UN	✓	45.6	-0.4	67.6	50.4	29.6	49.2	58.9	41.4	-0.2	64.8	44.5	25.2	45.1	55.7

Table 5: Ablation study on COCO val2017. Mask R-CNN with Swin-T is trained for 12 epochs.

Exp	FP	BP	COCO [‡]			
	GM	AM	α	AP ^{Box}	Δ	AP ^{Mask}
1	✓	✓	✓	42.8		39.2
2		✓	✓	42.2	-0.6	38.8
3	✓		✓	42.4	-0.4	38.8
4	✓	✓		42.4	-0.4	39.0
5	✓			42.5	-0.3	39.0

Table 6: The effect of using different window sizes (M) in UN is evaluated on COCO val2017.

M	2	4	6	8	10
AP ^{Box}	42.8	42.8	42.8	42.8	42.5
AP ^{Mask}	39.2	39.3	39.1	39.2	39.2

Table 7: The effect of using the outlier filtration is evaluated on ImageNet and COCO val2017.

Outlier Filtration	T2T-ViT-14	Swin-T			
	ImageNet Top1	ImageNet Top1	Faster RCNN AP ^{Box}	Mask RCNN AP ^{Box}	Mask RCNN AP ^{Mask}
w/	80.9	81.0	42.2	42.8	39.2
w/o	FAIL	80.7	FAIL	FAIL	FAIL

study the internal representations between different models. As depicted in Figure 6, we compare the similarity of intermediate feature maps between models trained with different normalization methods. We mainly focus on the diagonal pixels in the heatmap that indicate the similarity across layers of the same depth. UN shows a more remarkable similarity with LN compared to BN, especially in shallow layers. This result somehow explains the superiority of our method.

4.3.5 Efficiency. Transformers are broadly applied to vision tasks and attempt to achieve efficient deployment. LN comes with an additional overhead of computation and memory that results in inefficient inference. Besides, LN can not be supported on many edge devices, e.g., NXP i.MX Series and TITDA4x. There is still

Table 8: Inference efficiency comparison between LN and UN in Swin-T. MEM (MB) is the maximum allocated memory during inference. TPUT (Img./Sec.) shows the average throughput calculated over 1000 batches. We set a batch size of 512 for ImageNet and 2 for COCO (with Mask R-CNN).

Task	Method	MEM	Reduction	TPUT	Speedup
ImageNet	LN	9978	-	1179.5	-
	UN	8213	17.7%	1547.8	31.2%
COCO	LN	955	-	17.8	-
	UN	897	6.1%	22.1	24.2%

room for Transformers to be further improved to achieve hardware-efficient deployment. In this paper, we focus on improving the normalization layer for Transformers in order to achieve a better trade-off between performance and inference speed. With fusing UN to other linear layers, the division and square root operations are also removed in inference. Experimentally, we test the efficiency on GeForce RTX 3090 with Swin-T as reported in Table 8. For classification, we show that when our method is fused with other linear operations, it gains **about 18%** memory reduction and **over 31%** throughput improvement. For object detection, Mask R-CNN with Swin-T is integrated with other components like FPN and head, whereas LN is solely employed in the backbone. As a result, the increase in speed is limited.

5 CONCLUSION

In this paper, we look at how to deploy Transformers efficiently by replacing LN with an offline method. Previous offline methods suffer from inferior performance and instability due to the large fluctuations and extreme outliers in activation statistics. Based on our analysis, we propose UN that consists of the fluctuation smoothing and the outlier filtration strategies to tackle these challenges. Extensive experiments on NLP and CV tasks demonstrate that our method significantly outperforms previous offline methods. Furthermore, our method provides comparable performance to LN, with a speedup of over 31% in inference. We believe our method will be a general component in Transformers for efficient deployment.

ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. U19B2043).

REFERENCES

- [1] Jesús Munárriz Aldaz. 2012. Sharp bounds for the difference between the arithmetic and geometric means. *Archiv der Mathematik* 99, 4 (2012), 393–399.
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European Conference on Computer Vision*. Springer, 213–229.
- [5] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. 2019. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155* (2019).
- [6] Xinlei Chen, Saining Xie, and Kaiming He. 2021. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9640–9649.
- [7] Zhengsu Chen, Lingxi Xie, Jianwei Niu, et al. 2021. Visformer: The vision-friendly transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 589–598.
- [8] RALPH D'AGOSTINO and Egon S Pearson. 1973. Tests for departure from normality. *Biometrika* 60, 3 (1973), 613–622.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] Mingyu Ding, Bin Xiao, Noel Codella, Ping Luo, Jingdong Wang, and Lu Yuan. 2022. DaViT: Dual Attention Vision Transformers. <https://doi.org/10.48550/ARXIV.2204.03645>
- [12] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, et al. 2021. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *arXiv preprint arXiv:2107.00652* (2021).
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xi-aohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).
- [14] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. 2021. LeViT: a Vision Transformer in ConvNet's Clothing for Faster Inference. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 12259–12269.
- [15] Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100* (2020).
- [16] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*. 2961–2969.
- [17] Lei Huang, Jie Qin, Yi Zhou, Fan Zhu, Li Liu, and Ling Shao. 2020. Normalization techniques in training dnn: Methodology, analysis and application. *arXiv preprint arXiv:2009.12836* (2020).
- [18] Zilong Huang, Yucheng Ben, Guozhong Luo, Pei Cheng, Gang Yu, and Bin Fu. 2021. Shuffle Transformer: Rethinking Spatial Shuffle for Vision Transformer. *arXiv preprint arXiv:2106.03650* (2021).
- [19] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*. PMLR, 448–456.
- [20] Tianyang Lin, Yuxin Wang, Xiangyang Liu, and Xipeng Qiu. 2021. A survey of transformers. *arXiv preprint arXiv:2106.04554* (2021).
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2117–2125.
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*. Springer, 740–755.
- [23] Ye Lin, Yanyang Li, Tengbo Liu, Tong Xiao, Tongran Liu, and Jingbo Zhu. 2020. Towards fully 8-bit integer inference for the transformer model. *arXiv preprint arXiv:2009.08034* (2020).
- [24] Liyuan Liu, Xiaodong Liu, Jianfeng Gao, Weizhu Chen, and Jiawei Han. 2020. Understanding the Difficulty of Training Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*.
- [25] Ze Liu, Han Hu, Yutong Lin, et al. 2021. Swin Transformer V2: Scaling Up Capacity and Resolution. *arXiv preprint arXiv:2111.09883* (2021).
- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10012–10022.
- [27] Ping Luo, Jiamin Ren, Zhanglin Peng, Ruimao Zhang, and Jingyu Li. 2018. Differentiable learning-to-normalize via switchable normalization. *arXiv preprint arXiv:1806.10779* (2018).
- [28] Sachin Mehta, Marjan Ghazvininejad, Srinivasan Iyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2020. Delight: Deep and light-weight transformer. *arXiv preprint arXiv:2008.00623* (2020).
- [29] Thao Nguyen, Maithra Raghu, and Simon Kornblith. 2020. Do wide and deep networks learn the same things? uncovering how neural network representations vary with width and depth. *arXiv preprint arXiv:2010.15327* (2020).
- [30] Myle Ott, Sergey Edunov, Alexei Baevski, et al. 2019. fairseq: A fast, extensible toolkit for sequence modeling. *arXiv preprint arXiv:1904.01038* (2019).
- [31] 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*. 311–318.
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015).
- [33] Wenqi Shao, Yixiao Ge, Zhaoyang Zhang, Xuyuan Xu, Xiaogang Wang, Ying Shan, and Ping Luo. 2021. Dynamic Token Normalization Improves Vision Transformer. *arXiv preprint arXiv:2112.02624* (2021).
- [34] Wenqi Shao, Tianjian Meng, Jingyu Li, et al. 2019. Ssn: Learning sparse switchable normalization via sparsestmax. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 443–451.
- [35] Sheng Shen, Zhewei Yao, Amir Gholami, Michael Mahoney, and Kurt Keutzer. 2020. Powernorm: Rethinking batch normalization in transformers. In *International Conference on Machine Learning*. PMLR, 8741–8751.
- [36] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. 2021. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239* (2021).
- [37] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. 2016. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022* (2016).
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [39] Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. 2021. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122* (2021).
- [40] Shinji Watanabe, Takaaki Hori, Shigeki Karita, Tomoki Hayashi, Jiro Nishitoba, Yuya Unno, Nelson Enrique Yalta Soplin, Jahn Heymann, Matthew Wiesner, Nanxin Chen, Adithya Renduchintala, and Tsubasa Ochiai. 2018. ESPnet: End-to-End Speech Processing Toolkit. In *Proceedings of Interspeech*. 2207–2211. <https://doi.org/10.21437/Interspeech.2018-1456>
- [41] Yuxin Wu and Kaiming He. 2018. Group normalization. In *ECCV*. 3–19.
- [42] Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144* (2016).
- [43] Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tieyan Liu. 2020. On layer normalization in the transformer architecture. In *International Conference on Machine Learning*. PMLR, 10524–10533.
- [44] Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. 2019. Understanding and improving layer normalization. *Advances in Neural Information Processing Systems* 32 (2019).
- [45] Junjie Yan, Ruosi Wan, Xiangyu Zhang, Wei Zhang, Yichen Wei, and Jian Sun. 2020. Towards stabilizing batch statistics in backward propagation of batch normalization. *arXiv preprint arXiv:2001.06838* (2020).
- [46] Zhuliang Yao, Yue Cao, Yutong Lin, Ze Liu, Zheng Zhang, and Han Hu. 2021. Leveraging batch normalization for vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 413–422.
- [47] Zhuoyuan Yao, Di Wu, Xiong Wang, Binbin Zhang, Fan Yu, Chao Yang, Zhendong Peng, Xiaoyu Chen, Lei Xie, and Xin Lei. 2021. WeNet: Production oriented Streaming and Non-streaming End-to-End Speech Recognition Toolkit. In *Proc. Interspeech*. IEEE, Brno, Czech Republic.
- [48] Xiaodong Yu, Dahu Shi, Xing Wei, Ye Ren, Tingqun Ye, and Wenming Tan. 2021. SOIT: Segmenting Objects with Instance-Aware Transformers. *arXiv preprint arXiv:2112.11037* (2021).
- [49] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashe Feng, and Shuicheng Yan. 2021. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 558–567.
- [50] Biao Zhang and Rico Sennrich. 2019. Root mean square layer normalization. *Advances in Neural Information Processing Systems* 32 (2019).

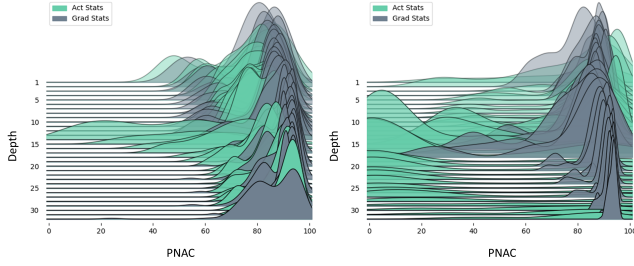


Figure 7: The probability density map of PNAC in different layers (depth) over iterations. Left side: visualization at the beginning of training. Right side: visualization at the end of training. We collect the samples from 150 time steps. The kernel density estimate (KDE) is used for visualization. Note that activation and gradient statistics are plotted in ‘aquamarine’ and ‘gray’, respectively.

A VISUALIZATION

A.1 Large Fluctuations in Activation Statistics

Figure 7 showcases the probability density map of PNAC during training that illustrates a general trend of statistics for all normalization layers in Transformer. At the beginning of training, most channels in activation and gradient statistics hold a high PNAC, which indicates mild fluctuations. After training a couple of epochs, there is a huge shift in activation statistics that a lot of channels gain a much lower PNAC, which means large fluctuations emerge in activation statistics. In this paper, a tailored fluctuation smoothing strategy is utilized to gain a better representation of these ever-changing statistics.

B DETAILS FOR EXPERIMENTS

B.1 Setup for Neural Machine Translation

We evaluate our method with Transformer on two datasets: (1) IWSLT14 De-En (IWSLT14) contains 0.18M sentence pairs; (2) WMT14 En-Fr (WMT14) [42] contains 36M sentence pairs. The setup for preprocessing raw data is the same as [28]. For IWSLT14, we replicate the training and evaluation strategies in [24]. For WMT14, we follow the training and evaluation setup in [28] and average the last 5 checkpoints for the test. Here, all experiments are re-implemented on the code base of Fairseq [30] with pre-normalization [43] setting.

B.2 Setup for Image Classification

In this section, we conduct image classification on ImageNet-1K [9] with two state-of-the-art Vision Transformers, T2T-ViT-14 [49] and Swin-T [26]. ImageNet-1K is a widely-used image classification dataset, which contains 1000 categories, 1.28M training samples, and 50K validation samples. Following the setup in [26, 49], all models are trained from scratch for 300 epochs with a cropped input size of 224×224 . We replace all LN layers in original architectures with BN/MABN/PN and our proposed method UN. After pretraining models on ImageNet, we transfer the models to downstream

Table 9: Ablation of different components of UN evaluated on IWSLT14 with Transformer.

Exp	Method	FP	BP		IWSLT14	
		GM	AM	α	BLEU	Δ
1	UN	✓	✓	✓	35.4	
2			✓	✓	34.5	-0.9
3		✓		✓	35.3	-0.1
4		✓	✓		18.3	-17.1
5		✓			FAIL	/

classification datasets, CIFAR10 and CIFAR100, that focus on general object classification. We follow the training recipe in [49]. All models are fed with a resized 224×224 input and finetuned for 60 epochs.

B.3 Setup for Object Detection and Instance Segmentation

We benchmark our method on COCO [22]. Following the standard setup in [26], object detection is conducted on Faster R-CNN [32] with FPN [21]. For instance segmentation, we evaluate our method with a common framework Mask R-CNN [16]. The setup for training and evaluation are following the original configurations on [26], all models are trained with 36 epochs. The input size is 1333×800 and the total batch size is set as 16. The backbone (Swin-T) is initialized with pretrained weights trained on ImageNet-1K. All experiments are re-implemented based on mmdetection [5].

C EXTRA RESULTS

C.1 Ablation Study on Neural Machine Translation

Table 9, we ablate UN on IWSLT14 to verify the contribution of the basic components. Since Geometric Mean (GM), Arithmetic Mean (AM), and momentum α are all used in the fluctuation smoothing, we try to remove the basic components one by one, as shown in Exp2-5. Without any moving strategies applied in activation statistics, there is a significant drop on BLEU. In backward propagation, we conduct a compound moving strategy for gradient estimation, which consists of arithmetic mean and momentum α . The results exhibit that each one of them is important for the final performance. Especially, the momentum α is greatly helpful for stabilizing the training on IWSLT14. In Exp5, once we remove all moving strategies in BP, the model will fail to converge. The result shows the advantage of the fluctuation smoothing in IWSLT14.

C.2 Effect on α

In UN, we leverage a momentum α for both approximating inference statistics in forward propagation and estimating gradients in backward propagation. We investigate the effect of α as shown in Table 10. By tuning α within a large range of $\{0.6, 0.7, 0.8, 0.9\}$, we find the models still converge stably on COCO with close performance. Choosing α from $\{0.7, 0.8, 0.9\}$ is also fine with IWSLT14. Once we set it with a small ratio, such as $\alpha = 0.6$, the results are

Table 10: The effect of α in UN is evaluated on COCO val2017 (AP) and IWSLT14 (BELU).

α	0.9	0.8	0.7	0.6
AP ^{Box}	42.8	42.6	42.7	42.8
AP ^{Mask}	39.2	39.1	39.2	39.4
BELU	35.4	35.4	35.3	FAIL

task-specific in that the training collapsed on IWSLT14. To some degree, that's similar to what we show in Table 9 (remove the momentum in BP). As a result, we believe that $\alpha = 0.9$ would be a good choice for various tasks while also allowing for a fair comparison with other methods.