

# “BNN - BN = ?”: Training Binary Neural Networks without Batch Normalization

Tianlong Chen<sup>1</sup>, Zhenyu Zhang<sup>2</sup>, Xu Ouyang<sup>3</sup>, Zechun Liu<sup>4</sup>, Zhiqiang Shen<sup>4</sup>, Zhangyang Wang<sup>1</sup>  
<sup>1</sup>University of Texas at Austin, <sup>2</sup>University of Science and Technology of China,  
<sup>3</sup>Cornell University, <sup>4</sup>Carnegie Mellon University

{tianlong.chen, atlaswang}@utexas.edu, zzy19969@mail.ustc.edu.cn,  
xo28@cornell.edu, {zechunl, zhiqians}@andrew.cmu.edu

## Abstract

Batch normalization (BN) is a key facilitator and considered essential for state-of-the-art binary neural networks (BNN). However, the BN layer is costly to calculate and is typically implemented with non-binary parameters, leaving a hurdle for the efficient implementation of BNN training. It also introduces undesirable dependence between samples within each batch. Inspired by the latest advance on Batch Normalization Free (BN-Free) training [7], we extend their framework to training BNNs, and for the first time demonstrate that BNs can be completely removed from BNN training and inference regimes. By plugging in and customizing techniques including adaptive gradient clipping, scale weight standardization, and specialized bottleneck block, a BN-free BNN is capable of maintaining competitive accuracy compared to its BN-based counterpart. Extensive experiments validate the effectiveness of our proposal across diverse BNN backbones and datasets. For example, after removing BNs from the state-of-the-art ReActNets [38], it can still be trained with our proposed methodology to achieve 92.08%, 68.34%, and 68.0% accuracy on CIFAR-10, CIFAR-100, and ImageNet respectively, with marginal performance drop (0.23% ~ 0.44% on CIFAR and 1.40% on ImageNet). Codes and pre-trained models are available at: [https://github.com/VITA-Group/BNN\\_NoBN](https://github.com/VITA-Group/BNN_NoBN).

## 1. Introduction

Despite widespread success [59, 24, 53, 52, 40, 11, 45], state-of-the-art deep networks usually have hundreds of millions of parameters [7, 33, 47], and suffer from burdensome computational cost. It is questionable how practical they are when it comes to deployment on real-world resource-constrained platforms, e.g., FPGA, ASICs, and mobile devices. Binary neural network (BNN) [14, 13, 50, 51, 73] are therefore proposed for the efficiency purpose. It takes only 1-bit with two discrete values, i.e.,  $\{-1, 1\}$  to represent networks’ weights and activations, leading to sig-

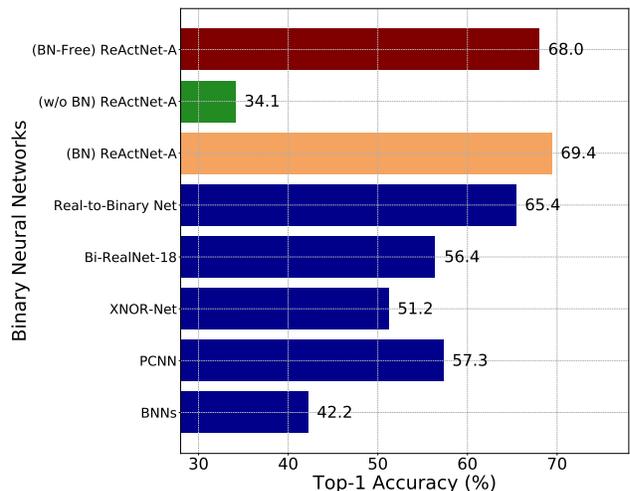


Figure 1. Top-1 accuracies of different binary neural networks (BNN) evaluated on ImageNet. Blue bars denote previous BNN methods. The orange bar presents the existing state-of-the-art (SOTA) method, i.e., ReActNet [38]. The green bar shows the performance of ReActNet if naively dropping Batch Normalization (BN) modules. The red bar indicates our proposed *Batch Normalization Free* (BN-Free) binary neural network, which reaches competitive performance compared to its counterpart with BN.

nificantly accelerated and energy-efficient inference as the 1-bit convolution operation can be efficiently implemented with XNOR and Bitcount operations [51].

Despite these appeals, BNNs are notoriously difficult to train, and undergo performance degradation. Particularly, [55] shows that the Batch Normalization (BN) [31] is critical to train BNNs successfully, allowing for stable training under larger learning rates, from both theoretical and empirical perspectives. Unfortunately, the batch normalization implementation [31] hinges on high precision values to compute the sum of squares, square-root and reciprocal. Therefore, it comes as no surprise that most BNNs [51, 73, 39, 38] kept BN layers in full precision during training, and some used reduced-precision such as

8-bit [4]. Although BNs can be absorbed into the BNN weights (e.g., scaling factors) post-training, their presence becomes a bottleneck for BNNs training efficiency on hardware [65]. Moreover, BNs often account for a substantial fraction of run-time, are hard to accelerate [17], and incurs memory overhead [9]. That applies to both inference and the feedforward stage of a training pass. Besides, BN causes discrepant behaviors between the network training and inference stages [58], which may break down the independence assumption between samples within each batch.

We hence anticipate finding an alternative to eliminate the unwanted properties of BNs in BNNs, while maintaining competitive performance. Motivated by the recent advance [6, 7], we propose Batch Normalization Free (BN-Free) binary neural networks. Specifically, we leverage the adaptive gradient clipping to constraining BNN’s gradient distribution and mitigate gradient explosion due to removing BNs [55]. Then, the scaled weight standardization and specialized bottleneck block [7] are integrated for preserving the variances and preventing the mean shifts of activations. Our contributions are outlined as follows:

- We provide the first proof-of-concept study that general BNNs can be successfully trained without BNs but maintain competitive performance.
- We introduce adaptive gradient clipping, scaled weight standardization, and specialized block to BNNs, and show these techniques can be easily plugged in various BNN backbones to make them BN-Free.
- Comprehensive experiments validate the effectiveness of our proposed mechanisms. For example, BN-Free ReActNet achieves 92.08%, 68.34%, and 68.0% accuracy on CIFAR-10, CIFAR-100, and ImageNet respectively, with only marginal performance drops compared to state-of-the-arts.

## 2. Related Work

**Binary neural networks.** Numerous model compression and acceleration algorithms have been proposed to reduce the latency of models while maintaining comparable accuracy performance. General model compression approaches fall under multiple forms [12]: pruning [21, 63], quantization [64, 56, 16], knowledge distillation [27, 44], as well as their compositions [61, 69, 71].

A Binary Neural Network (BNN) [13, 14, 34, 73, 51, 14, 36, 48, 44, 29, 72, 39, 25, 37, 10, 57, 20, 66] represents the most extreme form of model quantization as it quantizes weights in convolution layers to only 1 bit, enjoying great speed-up compared with its full-precision counterpart. [50] roughly divides previous BNN literature into two categories: (i) native BNN [13, 14, 34] which directly applies binarization to a full-precision model by a

pre-defined binarization function. Straight-through estimator (STE) [26, 5] is usually adopted to enable the back-propagation in binarized models [13]. (ii) optimization-based BNNs techniques, including minimizing the quantization error [73, 51, 14, 36], improving the network loss function [48, 44, 29, 72], and reducing the gradient error [39, 25, 37, 10].

However, such aggressive quantization usually results in severe accuracy decline. To tackle this limitation, [14] proposes an end-to-end gradient back-propagation framework for training the discrete binary weights and activations, establishing great successes on small datasets, such as MNIST [60] and CIFAR10 [62], while still has unsatisfactory performance on large datasets like ImageNet [38]. Follow-up researches [8, 19, 42, 39, 38] devote themselves to build state-of-the-art (SOTA) accuracies on ImageNet. Among these works, ReActNet [38] proposes generalized activation functions and a distributional loss, reaching the superior performance which reduces the gap to its full-precision counterpart within 3.0% accuracy on ImageNet. Note that, all mentioned SOTA BNNs are not sustained without batch normalization.

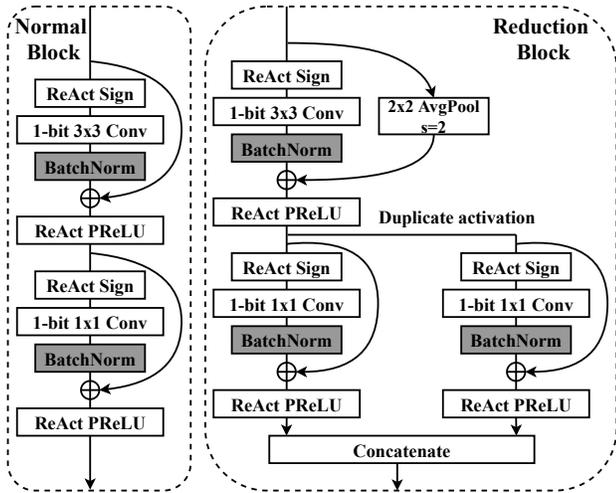
### **Batch normalization and normalization-free networks.**

Batch normalization (BN) [31] is a well-known and widely used technique to stabilize model training. It also plays a critical role in the BNN training, as evidenced by [55]. However, both normalization is an expensive computational primitive [17], and its inefficiency is further amplified in low bits precision context which hinders the deployment of BNN to resource-limited hardware [65].

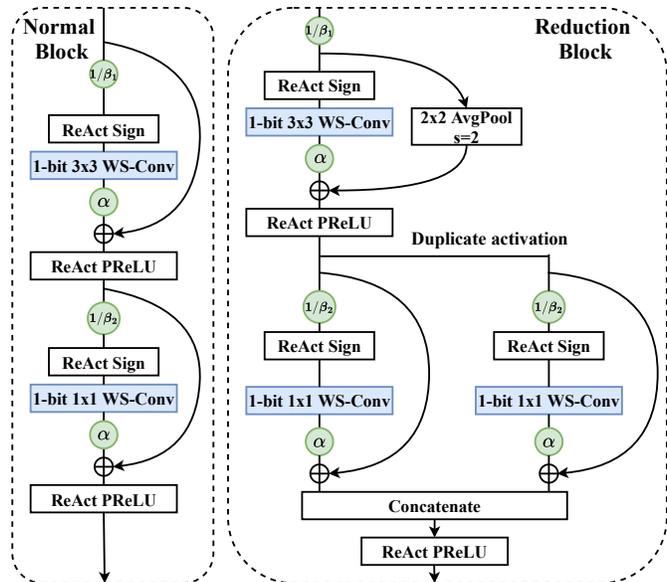
To seek a simple and effective alternative for batch normalization, various studies [70, 2, 22, 49, 6, 7] are proposed. [70] introduces an initialization and rescaling rule (i.e., fixed-update initialization) to stabilizes the training of very deep models in place of BN. [2, 22] share similar observations that appropriately initializing weights and scaling residual modules benefit avoiding the gradient exploding and vanishing, leading to a stabilized training. Another promising substitution is weight standardization [49], which subtracts the mean from weights and divides weights by their standard deviation. [6] proposes a modified variant, i.e., scaled weight standardization, to suppresses the quickly enlarging of the mean in hidden activations. Recently, [7] proposes Adaptive Gradient Clipping (AGC) to enable the larger batch size training of normalization-free networks, and to overcome the instabilities from eliminating BN.

## 3. Technical Approach

In this section, we present the detailed normalization-free methodologies for binary neural networks (BNN) in Section 3.1 and adopt BNN backbone architecture in Sec-



(a) Baseline Network Block of ReActNet



(b) Proposed BN-Free Network Block

Figure 2. The architecture overview of baseline network block (a) and proposed BN-Free network block (b). The baseline network blocks are inherited from the recent state-of-the-art (SOTA) BNN framework, i.e., ReActNet [38], which are modified from MobileNetV1 [30] and have the same configuration of channel and layer numbers. For the reduction block, [38] duplicates the input activation and concatenate the outputs to increase the channel number, which is also maintained in our proposed BN-Free network block. The most important thing is that all original Batch Normalization modules are *removed*, replaced by scaling factors (e.g.,  $\alpha$ ,  $1/\beta_1$ ,  $1/\beta_2$ ) and adjusted convolutional layers with scaled weight standardization (i.e., WS-Conv).

tion 3.2. Before that, we briefly list the main Batch Normalization benefits from previous literature.

**Understanding Batch Normalization.** The Batch Normalization (BN) can (i) reduce the scale of hidden activations on the residual branches [15, 3, 22, 68], and maintains well-behaved gradients early in training; (ii) eliminates mean-shift by enforcing the mean activation of each channel to zero across the current training batch [15, 32, 6]; (iii) serves as an implicit regularization [41] and enhances the models’ generalization [28]; (iv) enables large-batch training [18] and smoothens the loss landscapes [55].

Removing batch normalization directly usually leads to an inferior performance [31, 7]. It is further aggravated in training the binary neural network, due to its challenge regime with discrete values of variables [55]. Particularly, [55] provides both theoretical and empirical analyses to demonstrate the critical role of BN is to alleviate exploding gradients in the case of binary neural networks, which motivates us to introduce adaptive gradient clipping to establish the framework of BN-Free BNN.

### 3.1. Normalization-free Training Methodology

**Adaptive gradient clipping (AGC).** Gradient clipping is typically adopted to constrain the norm of gradients [46], leading to stabilized training [43]. Recently [7] proposes

adaptive gradient clipping (AGC) to ameliorate the NF-ResNets [6]’s performance, which clips gradients based on the unit-wise ratios of gradient norms to parameter norms. It can be described as follows:

$$G_i^l \rightarrow \begin{cases} \lambda \frac{\|W_i^l\|_F^*}{\|G_i^l\|_F} G_i^l & \text{if } \frac{\|G_i^l\|_F}{\|W_i^l\|_F^*} > \lambda \\ G_i^l & \text{otherwise.} \end{cases} \quad (1)$$

Where  $G_i^l$  denotes the  $i$ th row of gradient matrix  $G^l$ ; similarly,  $W_i^l$  is the  $i$ th row of weight matrix  $W^l$ ;  $l$  is the layer index of the considered network;  $\|W_i^l\|_F^* = \max\{\|W_i^l\|_F, \epsilon\}$ ,  $\epsilon = 10^{-3}$  and  $\|\cdot\|_F$  is the Frobenius norm.

The clipping threshold  $\lambda$  is a crucial hyperparameter, which is usually tuned by a grid search. Equipped with AGC, BNN training tends to have a constrained gradient distribution as evidenced in Figure 8, avoiding the gradient explosion issue.

**Scaled weight standardization.** To deal with the mean-shift in the hidden activation distributions caused by removing BN, we also introduced the Scaled Weight Standardization from [6]. Specifically, we modify all convolutional layers in BNN backbones as follows:

$$\hat{W}_{i,j} = \gamma \cdot \frac{W_{i,j} - \mu_i}{\sqrt{N} \sigma_i} \quad (2)$$

where  $\mu_i = (1/N)\sum_j W_{i,j}$ ,  $\sigma_i^2 = (1/N)\sum_j (W_{i,j} - \mu_i)^2$ ,  $N$  is the fan-in, and  $W_{i,j}$  is the corresponding standardized weights.  $\gamma$  is a fixed scalar for variance preserving, and has diverse values for different adopted activation functions [6]. For example,  $\gamma = \sqrt{2/(1 - (1/\pi))}$  for the ReLU activation function [1]. We name the modified convolutional layer as *WS-Conv* for simplicity. Note that, such WS-Conv has consistent performance between training and inference, mitigating the discrepancy behaviour of the batch normalization [7] and leading to a hardware-friendly implementation of BN-Free binary neural networks.

**Specialized bottleneck block.** For the batch normalization benefits preserving purpose, we inherit the specialized bottlenecks block from [6, 7] that applies input/output normalization with hand-crafted scaling factor (e.g.,  $\alpha, \beta$ ). As shown in Figure 2, we utilize  $x_{i_0}$  and  $x_{i_1}$  to present the input of the  $i_{\text{th}}$  BN-Free block and activation after the ReAct PReLU function. In order to normalize the input variance,  $\beta_1 = \sqrt{\text{Var}(x_{i_0})}$  is adopted before the 3x3 WS-Conv operation. We then multiply it with a scalar  $\alpha$  and feed it to the ReAct PReLU. Similarly, we divide the obtain activation  $x_{i_1}$  with  $\beta_2 = \sqrt{\text{Var}(x_{i_1})}$  and multiply it with  $\alpha$ . Blessed by the variance preserving design [6, 7], the output variance of the  $i_{\text{th}}$  BN-Free block is  $\text{Var}(x_{i_1}) + \alpha^2$ . Note that,  $\beta_1$  and  $\beta_2$  are usually the expected empirical standard deviation of the corresponding activation at initialization [7].

### 3.2. The Backbone Architecture of BNN

**Generalized activation functions.** [51, 67, 8, 38] advocate that enforcing binary neural networks to learn similar distribution as full-precision (i.e., real-valued or 32 bits) networks plays a significant role in the final achievable performance of BNN. Specifically, XNOR-Net [51] pursues close logits distribution as real-valued ones by calculating analytical real-valued scaling factors and multiplying them with the activations. [67, 8] introduce further improvements by learning these factors through back-propagation. ReActNet [38] explores an orthogonal perspective that mimics the activation distribution from a pre-trained full precision model. However, it is challenging for binary neural networks with a highly limited capacity to learn appropriate activation distribution, since even small variations to their activation distribution can substantially affect the feature representations in BNNs [38].

To tackle this issue, [38] proposes the generalized activation functions with learnable parameters, for sign and PReLU [23] functions, which are termed as RSign and RReLU respectively. Such learnable parameters enable the adaptive reshape and shift of BNNs’ activation to match the desired distributions. Following [38]’s definition, we introduce adopted activation functions.

$$\text{(RSign)} \quad x_i^b = h(x_i^r) = \begin{cases} +1, & \text{if } x_i^r > \alpha_i \\ -1, & \text{if } x_i^r \leq \alpha_i \end{cases} \quad (3)$$

where  $x_i^r$  is full-precision input of the RSign function  $h(\cdot)$  on the  $i$ th channel,  $x_i^b$  is the binary output and  $\alpha_i$  is a learnable coefficient controlling the threshold. The superscripts  $b$  and  $r$  above  $x_i$  denote the corresponding binary and full-precision values.

$$\text{(RReLU)} \quad f(x_i) = \begin{cases} x_i - \gamma_i + \zeta_i, & \text{if } x_i > \gamma_i \\ \beta_i \times (x_i - \gamma_i) + \zeta_i, & \text{if } x_i \leq \gamma_i \end{cases} \quad (4)$$

In RReLU function  $f(\cdot)$ ,  $x_i$  is the input in the  $i$ th channel,  $\gamma_i$  and  $\zeta_i$  are learnable shifts, and  $\beta_i$  is a learnable coefficient determines the slope of the negative half.

Meanwhile, we also use the default setting that adding parameter-free shortcuts to blocks, similar to [39] and [38]. As shown in Figure 2, our proposed *Batch Normalization free* (BN-Free) network block maintains the duplication of input activation from [38], replaces by scaling factors (e.g.,  $\alpha, 1/\beta_1, 1/\beta_2$ ) and adjusts convolutional layers with scaled weight standardization (i.e., WS-Conv).

**Distillation loss functions.** To establish the state-of-the-art BNN results, we also introduce the distribution loss function [38] to enforce the similarity of distributions between full-precision networks and binary neural networks. It can also be regarded as a knowledge distillation technique. Specifically, the formulation is depicted as follows:

$$\mathcal{L}_{\text{Dis}} = -\frac{1}{n} \sum_c \sum_{i=1}^n \rho_c^{\mathcal{R}}(X_i) \times \log\left(\frac{\rho_c^{\mathcal{B}}(X_i)}{\rho_c^{\mathcal{R}}(X_i)}\right) \quad (5)$$

where  $\mathcal{L}_{\text{Dis}}$  is the Kullback–Leibler (KL) divergence,  $X_i$  is the input image,  $c$  represents classes and  $n$  denotes the batch size.  $\rho_c^{\mathcal{R}}$  is the softmax output of the full-precision (i.e., real-valued) model and  $\rho_c^{\mathcal{B}}$  is the softmax output of the binary neural network. With the assistance of introduced distribution loss, BNN is capable of imitating the prediction distribution from full-precision models, leading to a superior performance. In the implementation, the full-precision NFNet [6, 7] is utilized, which is also a BN-Free network.

## 4. Experiments

### 4.1. Setup

We conduct experiments on three binary models, i.e., XNOR-Net [51], Bi-RealNet [39], and ReActNet [38] with two widely used backbones, i.e., ResNet-18 [24] and MobileNetV1 [30]. Meanwhile, we evaluate their BN-free

counterparts and report the performance on three representative classification datasets, i.e., CIFAR-10 [35], CIFAR-100 [35], and ILSVRC12 ImageNet [54].

**Implementation details on ImageNet.** We use the ImageNet dataset with 1000 classes. There are 1,281,167 images for training and 50,000 images for validation. Considering the superior performance of ReActNet [37] on the ImageNet classification task, we apply our BN-Free network design on ReActNet-18 and ReActNet-A, which are the modifications of ResNet-18 and MobileNetv1 respectively. We also adopt the adaptive gradient clipping (AGC) [7] in the back-propagation when training our BN-Free BNNs with the upper bound value set to 0.02.

When training the model, We follow the original two-step training strategy [38], where we only binarize the activations and train the network from scratch in the first step, then we fine-tune the network with both binary activations and weights in the second step. In both steps, we train the network for 120 epochs with the Adam optimizer and an initial learning rate of  $5 \times 10^{-4}$ , which follows a linear decreasing scheduler to zero. The weight decay is set to  $1 \times 10^{-5}$  for the first step and 0 for the second. Besides, the data augmentation method we used in our experiments follows [30], which contains random cropping, lighting, and random horizontal flipping. The input resolution is 224 and the top-1 accuracy on the validation set will be reported in the following section.

**Implementation details on CIFAR-10 and CIFAR-100.** Both CIFAR-10 and CIFAR-100 contain 50,000 training images and 10,000 testing images from 10 and 100 classes respectively. To comprehensively investigate the effectiveness of BN-Free networks, We conduct the classification experiments with four binary networks: XNORNet-18, Bi-RealNet-18, ReActNet-18, and ReActNet-A on CIFAR-10 and CIFAR-100. The first three networks all have a modified ResNet-18 backbone while ReActNet-A is constructed on MobileNetv1. We follow the two-step training strategy consistent with the ImageNet experiments and train the network for 256 epochs in each step. The upper bound of clipping value in AGC is set to 0.001 by default, according to the grid searching in Section 4.3. Additionally, other training hyperparameters remain the same as those in the ImageNet experiments. Differently, we use only random cropping and horizontal flipping for data augmentation.

## 4.2. Comparison to State-of-the-art networks

We begin by investigating the performance of batch normalization free BNN (BN-Free BNNs). For each network, we apply the Scaled Weight Standardization [7] to all convolution layers and replace the basic blocks with our BN-Free blocks after removing all batch normalization (BN)

layers. We report the accuracy between its three variants: the baseline network with BN, the network without BN, and the BN-Free network.

**Results on ImageNet.** We first evaluate our proposed BN-Free (BF) BNN on ImageNet. Specifically, the BN-Free versions of ReActNet-18 and ReActNet-A are constructed to compare with other existing state-of-the-art BNNs (with BN). Top-1 accuracies are collected in Table 1 and Figure 3.

Table 1. Comparison of the top-1 accuracy with state-of-the-art binary methods on ImageNet. The accuracy of other binary networks are collected from the original papers, which include BNN [14], PCNN [19], XNOR-Net [51], Bi-RealNet [39], Real-to-Binary Net [42], ReActNet-18 (BN) and ReActNet-A (BN) [38]. “w/o BN” denotes the version without batch normalization; “BN-Free” represents our proposed BN-Free BNNs.

| Binary Network          | Top-1 Acc (%) |
|-------------------------|---------------|
| BNN [14]                | 42.2          |
| PCNN [19]               | 57.3          |
| XNORNet-18 [51]         | 51.2          |
| Bi-RealNet-18 [39]      | 56.4          |
| Real-to-Binary Net [42] | 65.4          |
| ReActNet-18 (BN) [38]   | 65.5          |
| ReActNet-18 (w/o BN)    | 44.6          |
| ReActNet-18 (BN-Free)   | 61.1          |
| ReActNet-A (BN) [38]    | 69.4          |
| ReActNet-A (w/o BN)     | 34.1          |
| ReActNet-A (BN-Free)    | 68.0          |

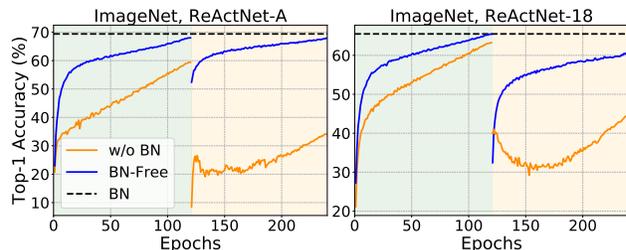


Figure 3. Results of validation accuracy over epochs on ImageNet with ReActNet-18/A. The green background represents the first training step, in which only activations are binarized. And in the orange part, both activations and weights are binary.

As shown in Table 1, compared with the binary neural network without BatchNorm layers, our BN-Free binary neural networks achieve substantial performance improvements. Specifically, we obtain 16.5% and 33.9% accuracy gains for the ReActNet-18 and ReActNet-A on ImageNet, respectively. Note that, our BN-Free ReActNet-A achieves a 68.0% top-1 accuracy, which only has marginal gap

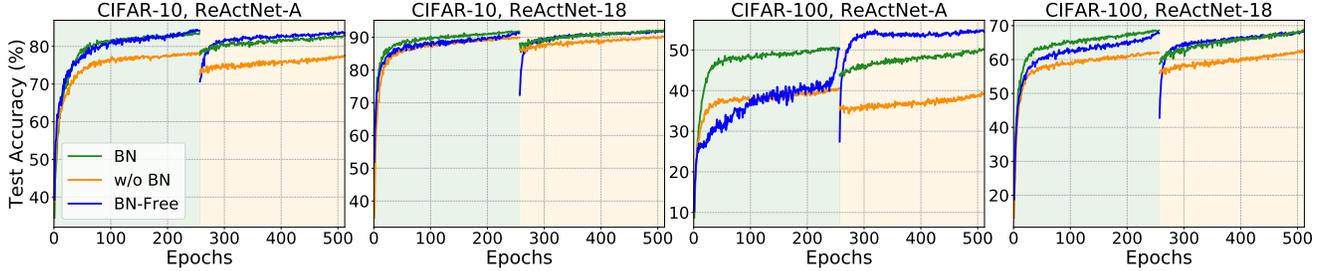


Figure 4. Results of testing accuracy over epochs on CIFAR-10/100 with ReActNet-18/A. The green background represents the first training step, in which only activations are binarized. And in the orange part, both activations and weights are binary.

(i.e., 1.4%) compared to the state-of-the-art ReActNet [38] with batch normalization. Detailed training dynamics are presented in Figure 3. We observe that the proposed BN-Free BNN not only reaches a superior performance, but also leads to more stable training.

**Results on CIFAR-10 and CIFAR-100.** To further evaluate the effectiveness of BN-Free modules in BNN, we implement the three binary networks mentioned in section 4.1, i.e., XNORNet, Bi-RealNet, ReActNet, and compare the performance of their three variants (i.e., BN, w/o BN, BN-Free) on CIFAR-10 and CIFAR-100. With the results in Table 2 and Figure 4, several consistent observations could be drawn as the following:

- The proposed BN-Free approach serves as a remedy for the accuracy degradation caused by the absence of BN layers across all datasets and networks. Specifically, when compared with their counterparts without BN, BN-Free BNNs achieve accuracy improvements of 1.75% ~ 8.29%, 5.74% ~ 15.63% for different binary networks on CIFAR-10 and CIFAR-100.
- Accuracy achieved by BN-Free ReActNet-A surpasses its BN counterpart surprisingly by 0.96% and 4.70% on CIFAR-10 and CIFAR-100, respectively. And BF-ReActNet-18 also achieves comparable performance with its BN version. However, for XNORNet-18 and Bi-RealNet-18, there remains a moderate performance gap between the BN and BN-Free networks.
- Training curves of BN-Free ReActNet on CIFAR-10 in Figure 4, almost overlaps (BN) ReActNet’s curves in both training steps. This indicates our BF networks not only can achieve comparable accuracy but also ensure a stable training process, especially on small datasets.

### 4.3. Ablation Study

In the previous section, we empirically evaluate the effectiveness of BN-Free modules and verify that our BF-ReActNet-A can reach competitive state-of-the-art perfor-

Table 2. Comparison of the top-1 accuracy between the three variants (i.e., BN, w/o BN, BN-Free) of binary networks on CIFAR-10 and CIFAR-100. All networks are modified from ResNet-18 except for ReActNet-A, which is constructed from MobileNetv1.

| Binary Network | CIFAR-10 (%) |        |              | CIFAR-100 (%) |        |              |
|----------------|--------------|--------|--------------|---------------|--------|--------------|
|                | BN           | w/o BN | BN-Free      | BN            | w/o BN | BN-Free      |
| XNORNet-18     | 90.21        | 71.75  | 79.67        | 65.35         | 45.30  | 53.76        |
| Bi-RealNet-18  | 89.12        | 71.30  | 79.59        | 63.51         | 47.72  | 54.34        |
| ReActNet-18    | 92.31        | 90.33  | 92.08        | 68.78         | 62.60  | 68.34        |
| ReActNet-A     | 82.95        | 77.60  | <b>83.91</b> | 50.30         | 39.37  | <b>55.00</b> |

Table 3. Ablation Study of clipping threshold values in AGC on CIFAR-10/100 with ReActNet-18 and ReActNet-A.

| Clipping Value     | ReActNet-18  |              | ReActNet-A   |              |
|--------------------|--------------|--------------|--------------|--------------|
|                    | CIFAR-10     | CIFAR-100    | CIFAR-10     | CIFAR-100    |
| w/o AGC            | 91.08        | 66.15        | 82.39        | 48.12        |
| $1 \times 10^{-2}$ | 91.23        | 65.38        | 82.61        | 49.46        |
| $5 \times 10^{-3}$ | 91.03        | 66.03        | 83.10        | 48.69        |
| $1 \times 10^{-3}$ | <b>92.08</b> | <b>68.34</b> | <b>83.91</b> | 51.32        |
| $8 \times 10^{-4}$ | 91.54        | 67.95        | 83.58        | 52.27        |
| $5 \times 10^{-4}$ | 91.39        | 67.36        | 83.31        | 53.60        |
| $2 \times 10^{-4}$ | 90.43        | 67.07        | 83.03        | <b>55.00</b> |
| $1 \times 10^{-4}$ | 89.62        | 62.45        | 80.81        | 52.28        |

mance. To further investigate the effects of different clipping thresholds in AGC strategy and different components in the proposed BN-Free framework, we provide an ablation study on CIFAR-10 and CIFAR-100 with ReActNet-18 and ReActNet-A as the backbone BNNs.

**Clipping threshold in AGC.** The clipping threshold  $\lambda$  plays an important role in the effectiveness of AGC [7]. In this paragraph, we empirically analyze how does the threshold values affect the training process and final performance. As shown in Table 3, the results indicate that with the growth of the clipping threshold, the final test accuracy first increases to the peak then begins to decline. The results also show that we can get an extra accuracy improvement of 1.00% ~ 6.88% by using an appropriate threshold. In addi-

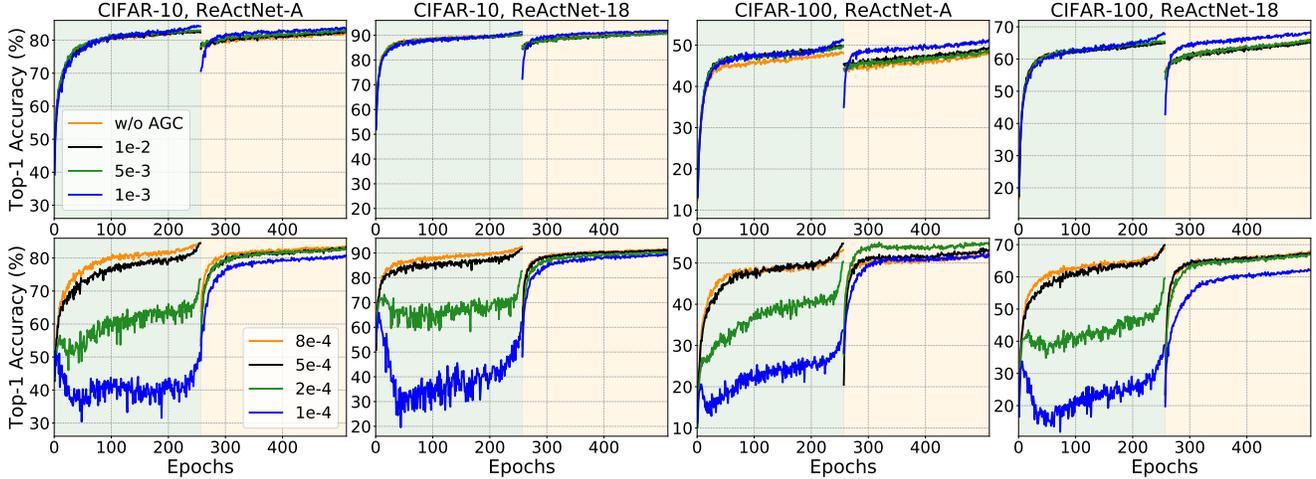


Figure 5. Results of testing accuracy curves of different clipping values in AGC on CIFAR-10/100 with ReActNet-18/A. The green background represents the first training step, where only activations are binarized. In the orange part, both activations and weights are binary.

Table 4. Ablation study of the separate effect of scaled weight standardization and normalizer-free block on CIFAR-10 and CIFAR-100 with two binary networks based on ReActNet. Test accuracies are reported.

| Settings   | ReActNet-18 (%) |              |
|--|-----------------|--------------|
|  | CIFAR-10        | CIFAR-100    |
| BN   | 92.31           | 68.79        |
| w/o BN   | 90.33           | 62.60        |
| WS-Conv  | 91.91           | 68.20        |
| Specialized Block (i.e., $\alpha, \frac{1}{\beta}$ ) | 91.44           | 63.63        |
| BN-Free  | 92.08           | 68.34        |
| Settings   | ReActNet-A (%)  |              |
|  | CIFAR-10        | CIFAR-100    |
| BN   | 82.95           | 50.30        |
| w/o BN   | 77.60           | 39.37        |
| WS-Conv  | 82.34           | 52.37        |
| Specialized Block (i.e., $\alpha, \frac{1}{\beta}$ ) | 80.45           | 54.44        |
| BN-Free  | <b>83.91</b>    | <b>55.00</b> |

tion, the performance of BN-Free networks on CIFAR-10 is less affected by the clipping values. A possible explanation is that the performance on the simple CIFAR-10 classification is saturated and less sensitive. Furthermore, Figure 5 demonstrates that the training process becomes less stable when the threshold in AGC is extremely small. It comes as no surprise that aggressive gradient clipping introduces undesired noise and causes instability.

**Different components in the BN-Free framework.** As described in Section 3.1, the BN-Free module is constructed

with a specialized block for the variance normalization, and a scaled weight standardization technique that is applied to all convolution layers (WS-Conv). To study the effects of different components in the BN-Free framework, we construct five variants on top of baseline BNNs (ReActNet-18 and ReActNet-A): **a)** original baseline (with BN); **b)** baseline (w/o BN); **c)** baseline (w/o BN) + WS-Conv; **d)** baseline (w/o BN) + specialized block; **e)** baseline (w/o BN) + WS-Conv + specialized block **which is equivalent to the complete BN-Free setup**. AGC with the best clipping threshold is adopted. The results are collected in Table 4 and their corresponding training dynamics are presented in Figure 6, from which several observations could be drawn:

- Either specialized block or WS-Conv can improve the performance of binary networks independently, specifically, the separate improvement achieved by WS-Conv ranges from 1.58% to 13.00% and the separate improvement of the specialized block ranges from 1.03% to 15.07%. In addition, the combination of these two approaches can further benefit the BN-Free binary neural networks.
- WS-Conv benefits more than the specialized bottleneck by 0.47%  $\sim$  4.57% performance gains, except the experiment of ReActNet-A on CIFAR-100.

#### 4.4. Visualization

In this section, we provide the visualization of gradient, latent weight, and activation distributions. Three variants of ReActNet-A (i.e., BN, w/o BN, BN-Free) trained on CIFAR-10 are considered.

**Activation distribution.** We visualize the activation distribution in Figure 7. Compared with the network without

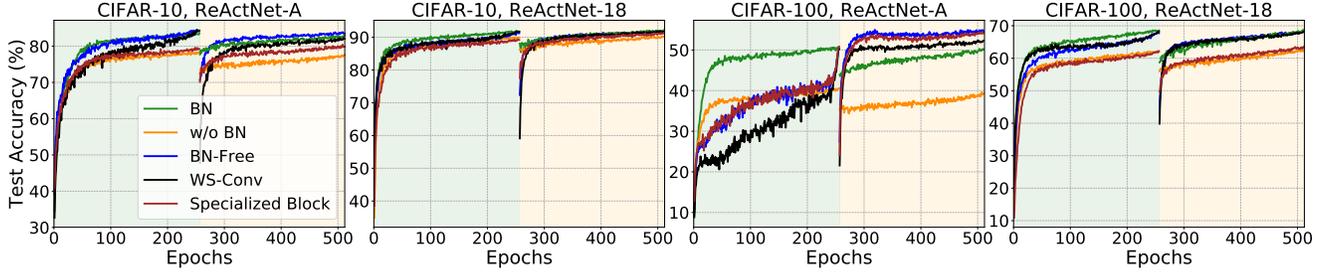


Figure 6. Results of testing accuracy over epochs on CIFAR-10/100 with ReActNet-18/A. The green background represents the first training step, in which only activations are binarized. And in the orange part, both activations and weights are binary.

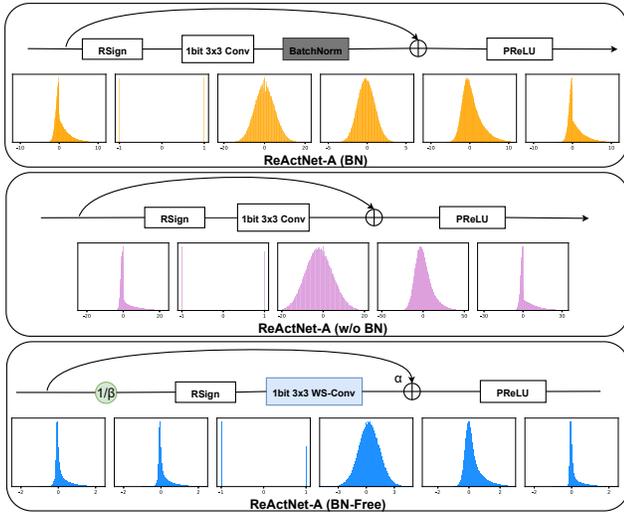


Figure 7. Histogram of the activation distribution inside three variants of ReActNet-A on CIFAR-10: with BN (top), without BN (middle) and BN-Free (bottom).

BN, the values of the activation inside the BN-Free network are consistently concentrated in a smaller region, which provides some insights into the training stability.

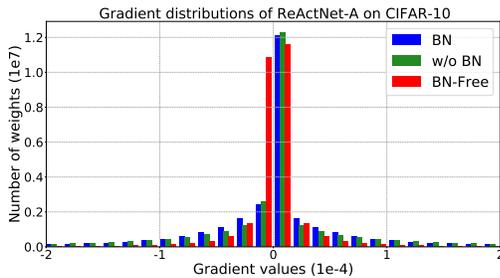


Figure 8. Visualization of gradient distributions of the three variants of ReActNet-A on CIFAR-10: with BN (blue), without BN (green) and BN-Free (red).

**Gradient distribution.** In figure 8, we show histogram visualizations of the gradient distribution. Our proposed BN-Free BNNs (red bars) tend to have a smaller range for gradients, which potentially prevents the emergence of gra-

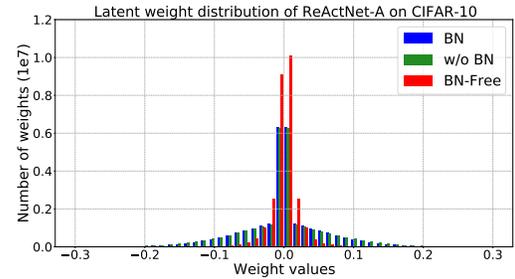


Figure 9. Visualization of latent weight distributions of the three variants of ReActNet-A on CIFAR-10: with BN (blue), without BN (green) and BN-Free (red).

dent exploration caused by training without batch normalization [55].

**Latent weight distribution.** Figure 9 present the latent weight distribution of three variants of ReActNet-A. We observe that BN-Free BNNs have a more zero-centralized weight distribution, which mainly stems from the weight standardization process.

## 5. Conclusions

In this paper, we for the first time propose a framework for training binary neural networks without batch normalization, i.e., BN-Free BNN, which achieves competitive state-of-the-art performance compared to its BN-based counterpart. Specifically, We introduce the scaled weight standardization to deal with the mean-shift in the hidden activation distributions caused by removing BN and apply a specialized bottleneck block for the purpose of variance preserving. Moreover, adaptive gradient clipping is adopted to mitigate the gradient exploration issue and stabilize training, for the BN-Free BNN. With the contributions jointly achieved by these techniques, our BN-Free ReActNet achieves 92.08%, 68.34%, and 68.00% on CIFAR-10, CIFAR-100, and ImageNet, respectively. Note that our BN-Free BNN totally gets rid of batch normalization in both training and inference regimes. In the future, we would be interested to examine the speedup and energy-saving results of the BNN training/inference on a hardware platform.

## References

- [1] Devansh Arpit, Yingbo Zhou, Bhargava Kota, and Venu Govindaraju. Normalization propagation: A parametric technique for removing internal covariate shift in deep networks. In *International Conference on Machine Learning*, pages 1168–1176. PMLR, 2016. 4
- [2] Thomas Bachlechner, Bodhisattwa Prasad Majumder, Huanru Henry Mao, Garrison W Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. *arXiv preprint arXiv:2003.04887*, 2020. 2
- [3] David Balduzzi, Marcus Frean, Lennox Leary, JP Lewis, Kurt Wan-Duo Ma, and Brian McWilliams. The shattered gradients problem: If resnets are the answer, then what is the question? In *International Conference on Machine Learning*, pages 342–350. PMLR, 2017. 3
- [4] Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable methods for 8-bit training of neural networks. *arXiv preprint arXiv:1805.11046*, 2018. 2
- [5] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013. 2
- [6] Andrew Brock, Soham De, and Samuel L Smith. Characterizing signal propagation to close the performance gap in unnormalized resnets. *arXiv preprint arXiv:2101.08692*, 2021. 2, 3, 4
- [7] Andrew Brock, Soham De, Samuel L Smith, and Karen Simonyan. High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*, 2021. 1, 2, 3, 4, 5, 6
- [8] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *arXiv preprint arXiv:1909.13863*, 2019. 2, 4
- [9] Samuel Rota Buló, Lorenzo Porzi, and Peter Kotschieder. In-place activated batchnorm for memory-optimized training of dnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5639–5647, 2018. 2
- [10] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5918–5926, 2017. 2
- [11] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1
- [12] Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*, 2017. 2
- [13] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. *arXiv preprint arXiv:1511.00363*, 2015. 1, 2
- [14] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016. 1, 2, 5
- [15] Soham De and Sam Smith. Batch normalization biases residual blocks towards the identity function in deep networks. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [16] Yonggan Fu, Haoran You, Yang Zhao, Yue Wang, Chaojian Li, Kailash Gopalakrishnan, Zhangyang Wang, and Yingyan Lin. Fractrain: Fractionally squeezing bit savings both temporally and spatially for efficient dnn training. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [17] Igor Gitman and Boris Ginsburg. Comparison of batch normalization and weight normalization algorithms for the large-scale image classification. *arXiv preprint arXiv:1709.08145*, 2017. 2
- [18] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large mini-batch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 3
- [19] Jiaxin Gu, Ce Li, Baochang Zhang, Jungong Han, Xianbin Cao, Jianzhuang Liu, and David Doermann. Projection convolutional neural networks for 1-bit cnns via discrete back propagation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8344–8351, 2019. 2, 5
- [20] Kai Han, Yunhe Wang, Yixing Xu, Chunjing Xu, Enhua Wu, and Chang Xu. Training binary neural networks through learning with noisy supervision. In *International Conference on Machine Learning*, pages 4017–4026. PMLR, 2020. 2
- [21] Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. *arXiv preprint arXiv:1506.02626*, 2015. 2
- [22] Boris Hanin and David Rolnick. How to start training: The effect of initialization and architecture. *arXiv preprint arXiv:1803.01719*, 2018. 2, 3
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. 4
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 4
- [25] Koen Helwegen, James Widdicombe, Lukas Geiger, Zechun Liu, Kwang-Ting Cheng, and Roeland Nusselder. Latent weights do not exist: Rethinking binarized neural network optimization. *arXiv preprint arXiv:1906.02107*, 2019. 2
- [26] Geoffrey Hinton. Neural networks for machine learning. 2012. 2
- [27] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2
- [28] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *arXiv preprint arXiv:1705.08741*, 2017. 3

- [29] Lu Hou, Quanming Yao, and James T Kwok. Loss-aware binarization of deep networks. *arXiv preprint arXiv:1611.01600*, 2016. 2
- [30] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 3, 4, 5
- [31] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 1, 2, 3
- [32] Arthur Jacot, Franck Gabriel, and Clément Hongler. Freeze and chaos for dnns: an ntk view of batch normalization, checkerboard and boundary effects. *CoRR*, abs/1907.05715, 2019. 3
- [33] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V Le, Yunhsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918*, 2021. 1
- [34] Minje Kim and Paris Smaragdis. Bitwise neural networks. *arXiv preprint arXiv:1601.06071*, 2016. 2
- [35] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5
- [36] Zefan Li, Bingbing Ni, Wenjun Zhang, Xiaokang Yang, and Wen Gao. Performance guaranteed network acceleration via high-order residual quantization. In *Proceedings of the IEEE international conference on computer vision*, pages 2584–2592, 2017. 2
- [37] Chunlei Liu, Wenrui Ding, Xin Xia, Baochang Zhang, Jiaxin Gu, Jianzhuang Liu, Rongrong Ji, and David Doermann. Circulant binary convolutional networks: Enhancing the performance of 1-bit dnns with circulant back propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2691–2699, 2019. 2, 5
- [38] Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. Reactnet: Towards precise binary neural network with generalized activation functions. In *European Conference on Computer Vision*, pages 143–159. Springer, 2020. 1, 2, 3, 4, 5, 6
- [39] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European conference on computer vision (ECCV)*, pages 722–737, 2018. 1, 2, 4, 5
- [40] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 1
- [41] Ping Luo, Xinjiang Wang, Wenqi Shao, and Zhanglin Peng. Towards understanding regularization in batch normalization. *arXiv preprint arXiv:1809.00846*, 2018. 3
- [42] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-time binary convolutions. In *International Conference on Learning Representations*, 2020. 2, 5
- [43] Stephen Merity, Nitish Shirish Keskar, and Richard Socher. Regularizing and optimizing lstm language models. *arXiv preprint arXiv:1708.02182*, 2017. 3
- [44] Asit Mishra and Debbie Marr. Apprentice: Using knowledge distillation techniques to improve low-precision network accuracy. *arXiv preprint arXiv:1711.05852*, 2017. 2
- [45] Junting Pan, Elisa Sayrol, Xavier Giro-i Nieto, Kevin McGuinness, and Noel E O’Connor. Shallow and deep convolutional networks for saliency prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 598–606, 2016. 1
- [46] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In *International conference on machine learning*, pages 1310–1318. PMLR, 2013. 3
- [47] Hieu Pham, Qizhe Xie, Zihang Dai, and Quoc V Le. Meta pseudo labels. *arXiv preprint arXiv:2003.10580*, 2020. 1
- [48] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018. 2
- [49] Siyuan Qiao, Huiyu Wang, Chenxi Liu, Wei Shen, and Alan Yuille. Micro-batch training with batch-channel normalization and weight standardization. *arXiv preprint arXiv:1903.10520*, 2019. 2
- [50] Haotong Qin, Ruihao Gong, Xianglong Liu, Xiao Bai, Jingkuan Song, and Nicu Sebe. Binary neural networks: A survey. *Pattern Recognition*, 105:107281, 2020. 1, 2
- [51] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016. 1, 2, 4, 5
- [52] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 1
- [53] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015. 1
- [54] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 5
- [55] Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? *arXiv preprint arXiv:1805.11604*, 2018. 1, 2, 3, 8
- [56] Jianghao Shen, Yue Wang, Pengfei Xu, Yonggan Fu, Zhangyang Wang, and Yingyan Lin. Fractional skipping: Towards finer-grained dynamic cnn inference. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 5700–5708, 2020. 2

- [57] Mingzhu Shen, Kai Han, Chunjing Xu, and Yunhe Wang. Searching for accurate binary neural architectures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
- [58] Cecilia Summers and Michael J Dinneen. Four things everyone should know to improve batch normalization. *arXiv preprint arXiv:1906.03548*, 2019. 2
- [59] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 1
- [60] Yaman Umuroglu, Nicholas J Fraser, Giulio Gambardella, Michaela Blott, Philip Leong, Magnus Jahre, and Kees Visser. Finn: A framework for fast, scalable binarized neural network inference. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 65–74, 2017. 2
- [61] Yue Wang, Ziyu Jiang, Xiaohan Chen, Pengfei Xu, Yang Zhao, Yingyan Lin, and Zhangyang Wang. E2-train: Training state-of-the-art cnns with over 80% energy savings. *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [62] Ziwei Wang, Jiwen Lu, Chenxin Tao, Jie Zhou, and Qi Tian. Learning channel-wise interactions for binary convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 568–577, 2019. 2
- [63] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 2082–2090, 2016. 2
- [64] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016. 2
- [65] Shuang Wu, Guoqi Li, Lei Deng, Liu Liu, Dong Wu, Yuan Xie, and Luping Shi.  $l_1$ -norm batch normalization for efficient training of deep neural networks. *IEEE transactions on neural networks and learning systems*, 30(7):2043–2051, 2018. 2
- [66] Yixing Xu, Kai Han, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Learning frequency domain approximation for binary neural networks. *arXiv preprint arXiv:2103.00841*, 2021. 2
- [67] Zhe Xu and Ray CC Cheung. Accurate and compact convolutional neural networks with trained binarization. *arXiv preprint arXiv:1909.11366*, 2019. 4
- [68] Greg Yang, Jeffrey Pennington, Vinay Rao, Jascha Sohl-Dickstein, and Samuel S Schoenholz. A mean field theory of batch normalization. *arXiv preprint arXiv:1902.08129*, 2019. 3
- [69] Haoran You, Xiaohan Chen, Yonggan Zhang, Chaojian Li, Sicheng Li, Zihao Liu, Zhangyang Wang, and Yingyan Lin. Shiftaddnet: A hardware-inspired deep network. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [70] Hongyi Zhang, Yann N Dauphin, and Tengyu Ma. Fixup initialization: Residual learning without normalization. *arXiv preprint arXiv:1901.09321*, 2019. 2
- [71] Yang Zhao, Xiaohan Chen, Yue Wang, Chaojian Li, Haoran You, Yonggan Fu, Yuan Xie, Zhangyang Wang, and Yingyan Lin. Smartexchange: Trading higher-cost memory storage/access for lower-cost computation. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pages 954–967. IEEE, 2020. 2
- [72] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights. *arXiv preprint arXiv:1702.03044*, 2017. 2
- [73] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016. 1, 2