

# Resilient Binary Neural Network

Sheng Xu<sup>1\*</sup>, Yanjing Li<sup>1\*</sup>, Teli Ma<sup>3\*</sup>, Mingbao Lin<sup>4</sup>,  
Hao Dong<sup>5</sup>, Baochang Zhang<sup>1,2†</sup>, Peng Gao<sup>3</sup>, Jinhua Lu<sup>1,2</sup>

<sup>1</sup> Beihang University, Beijing, P.R.China

<sup>2</sup> Zhongguancun Laboratory, Beijing, P.R.China

<sup>3</sup> Shanghai AI Laboratory, Shanghai, P.R.China

<sup>4</sup> Tencent, P.R.China

<sup>5</sup> Peking University, Beijing, P.R.China

shengxu@buaa.edu.cn, yanjingli@buaa.edu.cn, mateli@pjlab.org.cn, linmb001@outlook.com, hao.dong@pku.edu.cn,  
bczhang@buaa.edu.cn, gaopeng@pjlab.org.cn, lvjinhua@buaa.edu.cn

## Abstract

Binary neural networks (BNNs) have received ever-increasing popularity for their great capability of reducing storage burden as well as quickening inference time. However, there is a severe performance drop compared with real-valued networks, due to its intrinsic frequent weight oscillation during training. In this paper, we introduce a Resilient Binary Neural Network (ReBNN) to mitigate the frequent oscillation for better BNNs' training. We identify that the weight oscillation mainly stems from the non-parametric scaling factor. To address this issue, we propose to parameterize the scaling factor and introduce a weighted reconstruction loss to build an adaptive training objective. For the first time, we show that the weight oscillation is controlled by the balanced parameter attached to the reconstruction loss, which provides a theoretical foundation to parameterize it in back propagation. Based on this, we learn our ReBNN by calculating the balanced parameter based on its maximum magnitude, which can effectively mitigate the weight oscillation with a resilient training process. Extensive experiments are conducted upon various network models, such as ResNet and Faster-RCNN for computer vision, as well as BERT for natural language processing. The results demonstrate the overwhelming performance of our ReBNN over prior arts. For example, our ReBNN achieves 66.9% Top-1 accuracy with ResNet-18 backbone on the ImageNet dataset, surpassing existing state-of-the-arts by a significant margin. Our code is open-sourced at <https://github.com/SteveTsui/ReBNN>.

## Introduction

Deep neural networks (DNNs) have dominated the recent advances of artificial intelligence from computer vision (CV) (Krizhevsky, Sutskever, and Hinton 2012; Russakovsky et al. 2015) to natural language processing (NLP) (Wang et al. 2018; Qin et al. 2019) and many beyond. In particular, large pre-trained models, *e.g.*, ResNet (He et al. 2016) and BERT (Devlin et al. 2018), have continuously broken many records of the state-of-the-art performance. However, the achievement also comes with tremendous demands for

memory and computation resources. These demands pose a huge challenge to the computing ability of many devices, especially resource-limited platforms such as mobile phones and electronic gadgets. In light of this, substantial research efforts are being invested in saving memory usage and computational power for an efficient online inference (He et al. 2018; Rastegari et al. 2016; Qin et al. 2022; Li et al. 2022). Among these studies, network quantization is particularly suitable for model deployment on resource-limited platforms for its great reduction in parameter bit-width and practical speedups supported by general hardware devices.

Binarization, an extreme form of quantization, represents weights and activations of CNNs using a single bit, which well decreases the storage requirements by  $32\times$  and computation cost by up to  $58\times$  (Rastegari et al. 2016). Consequently, binarized neural networks (BNNs) are widely deployed on various tasks such as image classification (Rastegari et al. 2016; Liu et al. 2018; Lin et al. 2022) and object detection (Wang et al. 2020; Xu et al. 2021a, 2022b), and have the potential to be deployed directly on next-generation AI chips. However, the performance of BNNs remains largely unequal to the real-valued counterparts, due primarily to the degraded representation capability and trainability.

Conventional BNNs (Rastegari et al. 2016; Liu et al. 2020) are often sub-optimized, due to their intrinsic frequent weight oscillation during training. We first identify that the weight oscillation mainly stems from the non-parametric scaling factor. Fig. 1(a) shows the epoch-wise oscillation<sup>1</sup> of ReActNet, where weight oscillation exists even when the network is convergent. As shown in Fig. 1(b), the conventional ReActNet (Liu et al. 2020) possesses a channel-wise tri-modal distribution in the 1-bit convolution layers, whose peaks respectively center around the  $\{-1, 0, +1\}$ . Such distribution leads to a magnified scaling factor  $\alpha$ , and thus the quantized weights  $\pm\alpha$  are quite larger than the small weights around 0, which might cause the weight oscillation. As illustrated in Fig. 1(c), in BNNs, the real-valued latent tensor is binarized by the sign function and scaled by the scaling factor

\*These authors contributed equally.

†Corresponding author.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

<sup>1</sup>A toy example of weight oscillation: From iteration  $t$  to  $t+1$ , a misleading weight update occurs causing an oscillation from  $-1$  to  $1$ , and from iteration  $t+1$  to  $t+2$  causes an oscillation from  $1$  to  $-1$ .

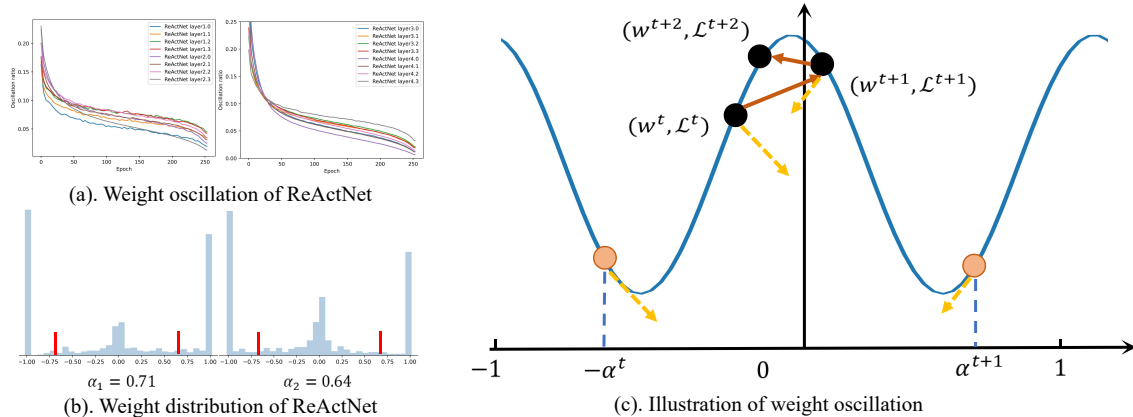


Figure 1: (a) We show the epoch-wise weight oscillation of ReActNet. (b) We randomly select 2 channels of the first 1-bit layer in ReActNet (Liu et al. 2020). Obviously, the distribution is with 3 peaks centering around  $\{-1, 0, +1\}$ , which magnifies the non-parametric scaling factor (red line). (c) We illustrate the weight oscillation caused by such inappropriate scale calculation, where  $w$  and  $\mathcal{L}$  indicate the latent weight and network loss function (blue line), respectively.

(the **orange** dot) in the forward propagation. In the backward propagation, the gradient is computed based on the quantized value  $\pm\alpha$  (indicated by the **yellow** dotted line). However, the gradient of small latent weights is misleading, when the scaling factor is magnified by the weights around  $\pm 1$  as ReActNet (Fig. 1(a)). Then the update is conducted on the latent value (the **black** dot), which leads to the oscillation of the latent weight. With extremely limited representation states, such latent weights with small magnitudes frequently oscillate during the non-convex optimization.

To address the aforementioned problem, we aim to introduce a Resilient Binary Neural Network (ReBNN). The intuition of our work is to re-learn the channel-wise scaling factor as well as the latent weights in a unified framework. Accordingly, we propose to parameterize the scaling factor and introduce a weighted reconstruction loss to build an adaptive training objective. We further show that the oscillation is factually controlled by the balanced parameter attached to the reconstruction loss, which provides a theoretical foundation to parameterize it in back propagation. The oscillation only happens when the gradient possesses a magnitude big enough to change the sign of the latent weight. Consequently, we calculate the balanced parameter based on the maximum magnitude of weight gradient during each iteration, leading to resilient gradients and effectively mitigating the weight oscillation. Our main contributions are summarized as follows:

- We propose a new resilient gradient for learning the binary neural networks (ReBNN), which mitigates the frequent oscillation to better train BNNs.
- We parameterize the scaling factor and introduce a weighted reconstruction loss to build an adaptive learning objective. We prove that the occurrence of oscillation is controlled by the balanced parameter attached to the reconstruction loss. Therefore, we utilize resilient weight gradients to learn our ReBNN and effectively mitigate the weight oscillation.

- Extensive experiments demonstrate the superiority of our ReBNN against other prior state-of-the-arts. For example, our ReBNN achieves 66.9% Top-1 accuracy on the ImageNet dataset, surpassing prior ReActNet by 1.0% with no extra parameters. In particular, our ReBNN also achieves state-of-the-art on fully binarized BERT models, demonstrating the generality of our ReBNN.

## Related Work

BinaryNet, based on BinaryConnect, was proposed to train CNNs with binary weights. The activations are triggered at run-time while the parameters are computed during training. Following this line of research, local binary convolution (LBC) layers are introduced in (Juefei-Xu, Naresh Boddeti, and Savvides 2017) to binarize the non-linearly activations. XNOR-Net (Rastegari et al. 2016) is introduced to improve convolution efficiency by binarizing the weights and inputs of the convolution kernels. More recently, Bi-Real Net (Liu et al. 2018) explores a new variant of residual structure to preserve the information of real activations before the sign function, with a tight approximation to the derivative of the non-differentiable sign function. Real-to-binary (Martinez et al. 2020) re-scales the feature maps on the channels according to the input before binarized operations and adds a SE-Net (Hu, Shen, and Sun 2018) like gating module. ReActNet (Liu et al. 2020) replaces the conventional PReLU (He et al. 2015) and the sign function of the BNNs with RReLU and RSign with a learnable threshold, thus improving the performance of BNNs. RBONN (Xu et al. 2022a) introduces a recurrent bilinear optimization to address the asynchronous convergence problem for BNNs, which further improves the performance of BNNs. However, most of these aforementioned suffer from the weight oscillation mainly stemming from the non-parametric scaling factor.

Unlike prior works, our ReBNN proposes to parameterize the scaling factor and introduces a weighted reconstruction

loss to build an adaptive training objective. We further prove the oscillation is controlled by the balanced parameter. Based on the analysis, we introduce a resilient weight gradient to effectively address the oscillation problem.

## Methodology

### Preliminaries

Given an  $N$ -layer CNN model, we denote its weight set as  $\mathbf{W} = \{\mathbf{w}^n\}_{n=1}^N$  and input feature map set as  $\mathbf{A} = \{\mathbf{a}_{in}^n\}_{n=1}^N$ . The  $\mathbf{w}^n \in \mathbb{R}^{C_{out}^n \times C_{in}^n \times K^n \times K^n}$  and  $\mathbf{a}_{in}^n \in \mathbb{R}^{C_{in}^n \times W_{in}^n \times H_{in}^n}$  are the convolutional weight and the input feature map in the  $n$ -th layer, where  $C_{in}^n$ ,  $C_{out}^n$  and  $K^n$  respectively stand for input channel number, output channel number and the kernel size. Also,  $W_{in}^n$  and  $H_{in}^n$  are the width and height of the feature maps. Then, the convolutional outputs  $\mathbf{a}_{out}^n$  can be technically formulated as:

$$\mathbf{a}_{out}^n = \mathbf{w}^n \otimes \mathbf{a}_{in}^n, \quad (1)$$

where  $\otimes$  represents the convolution operation. Herein, we omit the non-linear function for simplicity. Binary neural network intends to represent  $\mathbf{w}^n$  and  $\mathbf{a}^n$  in a 1-bit format as  $\mathbf{b}^{\mathbf{w}^n} \in \{-1, +1\}^{C_{out}^n \times C_{in}^n \times K^n \times K^n}$  and  $\mathbf{b}^{\mathbf{a}_{in}^n} \in \{-1, +1\}^{C_{in}^n \times W_{in}^n \times H_{in}^n}$  such that the float-point convolutional outputs can be approximated by the efficient XNOR and Bit-count instructions as:

$$\mathbf{a}_{out}^n \approx \alpha^n \circ (\mathbf{b}^{\mathbf{w}^n} \odot \mathbf{b}^{\mathbf{a}_{in}^n}), \quad (2)$$

where  $\circ$  represents the channel-wise multiplication,  $\odot$  denotes the XNOR and Bit-count instructions, and  $\alpha^n = \{\alpha_1^n, \alpha_2^n, \dots, \alpha_{C_{out}^n}^n\} \in \mathbb{R}_+^{C_{out}^n}$  is known as the channel-wise scaling factor vector (Rastegari et al. 2016) to mitigate the output gap between Eq. (1) and its approximation of Eq. (2). We denote  $\mathcal{A} = \{\alpha^n\}_{n=1}^N$ . Most existing implementations simply follow earlier studies (Rastegari et al. 2016; Liu et al. 2018) to optimize  $\mathcal{A}$  and latent weights  $\mathbf{W}$  based on a non-parametric bi-level optimization as:

$$\mathbf{W}^* = \arg \min_{\mathbf{W}} \mathcal{L}(\mathbf{W}; \mathcal{A}^*), \quad (3)$$

$$\text{s.t. } \alpha^{n*} = \arg \min_{\alpha^n} \|\mathbf{w}^n - \alpha^n \circ \mathbf{b}^{\mathbf{w}^n}\|_2^2, \quad (4)$$

where  $\mathcal{L}(\cdot)$  represents the training loss. Consequently, a closed-form solution of  $\alpha^n$  can be derived via the channel-wise absolute mean (CAM) as  $\alpha_i^n = \frac{\|\mathbf{w}_{i,\dots}^n\|_1}{M^n}$  and  $M^n = C_{in}^n \times K^n \times K^n$ . For ease of representation, we use  $\mathbf{w}_i^n$  as an alternative of  $\mathbf{w}_{i,\dots}^n$  in what follows. The latent weight  $\mathbf{w}^n$  is updated via a standard gradient back-propagation algorithm and its gradient is calculated as:

$$\delta_{\mathbf{w}_i^n} = \frac{\partial \mathcal{L}}{\partial \mathbf{w}_i^n} = \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_i^n} \frac{\partial \hat{\mathbf{w}}_i^n}{\partial \mathbf{w}_i^n} = \alpha_i^n \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_i^n} \otimes \mathbf{1}_{|\mathbf{w}_i^n| \leq 1}, \quad (5)$$

where  $\otimes$  denotes the Hadamard product and  $\hat{\mathbf{w}}^n = \alpha^n \circ \mathbf{b}^{\mathbf{w}^n}$ . **Discussion.** Eq. (5) shows weight gradient mainly comes from the non-parametric  $\alpha_i^n$  and the gradient  $\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_i^n}$ .  $\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_i^n}$  is automatically solved in the back propagation and becomes smaller as network convergence, however,  $\alpha_i^n$  is often magnified by the tri-modal distribution (Liu et al. 2020). Therefore,

weight oscillation mainly stems from  $\alpha_i^n$ . Given a single weight  $\mathbf{w}_{i,j}^n$  ( $1 \leq j \leq M^n$ ) centering around zero, the gradient  $\frac{\partial \mathcal{L}}{\partial \mathbf{w}_{i,j}^n}$  is misleading, due to the significant gap between  $\mathbf{w}_{i,j}^n$  and  $\alpha_i^n \mathbf{b}^{\mathbf{w}_{i,j}^n}$ . Consequently, the bi-level optimization leads to frequent weight oscillation. To address this issue, we reformulate traditional bi-level optimization using Lagrange multiplier and show that a learnable scaling factor is a natural training stabilizer.

### Resilient Binary Neural Network

We first give the learning objective in this paper as:

$$\arg \min_{\mathbf{W}, \mathcal{A}} \mathcal{L}(\mathbf{W}, \mathcal{A}) + \mathcal{L}_R(\mathbf{W}, \mathcal{A}), \quad (6)$$

where  $\mathcal{L}_R(\mathbf{W}, \mathcal{A})$  is a weighted reconstruction loss and defined as:

$$\mathcal{L}_R(\mathbf{W}, \mathcal{A}) = \frac{1}{2} \sum_{n=1}^N \sum_{i=1}^{C_{out}^n} \gamma_i^n \|\mathbf{w}_i^n - \alpha_i^n \mathbf{b}^{\mathbf{w}_i^n}\|_2^2, \quad (7)$$

in which  $\gamma_i^n$  is a balanced parameter. Based on the objective, the weight gradient in Eq. (5) becomes:

$$\begin{aligned} \delta_{\mathbf{w}_i^n} &= \frac{\partial \mathcal{L}}{\partial \mathbf{w}_i^n} + \gamma_i^n (\mathbf{w}_i^n - \alpha_i^n \mathbf{b}^{\mathbf{w}_i^n}) \\ &= \alpha_i^n \left( \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_i^n} \otimes \mathbf{1}_{|\mathbf{w}_i^n| \leq 1} - \gamma_i^n \mathbf{b}^{\mathbf{w}_i^n} \right) + \gamma_i^n \mathbf{w}_i^n. \end{aligned} \quad (8)$$

The  $\mathcal{S}_i^n(\alpha_i^n, \mathbf{w}_i^n) = \gamma_i^n (\mathbf{w}_i^n - \alpha_i^n \mathbf{b}^{\mathbf{w}_i^n})$  is an additional term added in the back-propagation process. We add this item given that a too small  $\alpha_i^n$  diminishes the gradient  $\delta_{\mathbf{w}_i^n}$  and causes a constant weight  $\mathbf{w}_i^n$ . In what follows, we state and prove the proposition that  $\delta_{\mathbf{w}_{i,j}^n}$  is a resilient gradient for a single weight  $\mathbf{w}_{i,j}^n$ . We sometimes omit subscript  $i, j$  and superscript  $n$  for an easy representation.

**Proposition 1.** *The additional term  $\mathcal{S}(\alpha, \mathbf{w}) = \gamma(\mathbf{w} - \alpha \mathbf{b}^{\mathbf{w}})$  achieves a resilient training process by suppressing frequent weight oscillation. Its balanced factor  $\gamma$  can be considered as the parameter controlling the occurrence of the weight oscillation.*

**Proof:** We prove the proposition by contradiction. For a single weight  $\mathbf{w}$  centering around zero, the straight-through-estimator  $\mathbf{1}_{|\mathbf{w}| \leq 1} = 1$ . Thus we omit it in the following. Based on Eq. (8), with a learning rate  $\eta$ , the weight updating process is formulated as:

$$\begin{aligned} \mathbf{w}^{t+1} &= \mathbf{w}^t - \eta \delta_{\mathbf{w}^t} \\ &= \mathbf{w}^t - \eta \left[ \alpha^t \left( \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} - \gamma \mathbf{b}^{\mathbf{w}^t} \right) + \gamma \mathbf{w}^t \right] \\ &= (1 - \eta \gamma) \mathbf{w}^t - \eta \alpha^t \left( \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} - \gamma \mathbf{b}^{\mathbf{w}^t} \right) \\ &= (1 - \eta \gamma) \left[ \mathbf{w}^t - \frac{\eta \alpha^t}{(1 - \eta \gamma)} \left( \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} - \gamma \mathbf{b}^{\mathbf{w}^t} \right) \right], \end{aligned} \quad (9)$$

where  $t$  denotes the  $t$ -th training iteration and  $\eta$  represents learning rate. Different weights shares different distances to the quantization level  $\pm 1$ , therefore, their gradients should be modified in compliance with their scaling factors and current

learning rate. We first assume the initial state  $\mathbf{b}^{\mathbf{w}^t} = -1$ , and the analysis process is applicable to the case of initial state  $\mathbf{b}^{\mathbf{w}^t} = 1$ . The oscillation probability from the iteration  $t$  to  $t + 1$  is:

$$P(\mathbf{b}^{\mathbf{w}^t} \neq \mathbf{b}^{\mathbf{w}^{t+1}}) \Big|_{\mathbf{b}^{\mathbf{w}^t} = -1} \leq P\left(\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} \leq -\gamma\right). \quad (10)$$

Similarly, the oscillation probability from the iteration  $t + 1$  to  $t + 2$  is:

$$P(\mathbf{b}^{\mathbf{w}^{t+1}} \neq \mathbf{b}^{\mathbf{w}^{t+2}}) \Big|_{\mathbf{b}^{\mathbf{w}^{t+1}} = 1} \leq P\left(\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^{t+1}} \geq \gamma\right). \quad (11)$$

Thus, the sequential oscillation probability from the iteration  $t$  to  $t + 2$  is:

$$\begin{aligned} & P((\mathbf{b}^{\mathbf{w}^{t+1}} \neq \mathbf{b}^{\mathbf{w}^{t+2}}) \cap (\mathbf{b}^{\mathbf{w}^t} \neq \mathbf{b}^{\mathbf{w}^{t+1}})) \Big|_{\mathbf{b}^{\mathbf{w}^t} = -1} \\ & \leq P\left(\left(\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} \leq -\gamma\right) \cap \left(\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^{t+1}} \geq \gamma\right)\right), \end{aligned} \quad (12)$$

which denotes that the weight oscillation happens only if the magnitudes of  $\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t}$  and  $\frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^{t+1}}$  are both larger than  $\gamma$ . **As a result, its attached factor  $\gamma$  can be considered as a parameter used to control the occurrence of the weight oscillation.**

However, if the conditions in Eq. (12) are met, with Eq. (9) concluded, the gradient of  $\hat{\mathbf{w}}^{t+1}$  is formulated as:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^{t+1}} &= \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} - \eta \frac{\partial^2 \mathcal{L}}{\partial (\hat{\mathbf{w}}^t)^2} \geq \gamma, \\ \eta \frac{\partial^2 \mathcal{L}}{\partial (\hat{\mathbf{w}}^t)^2} &\leq \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}^t} - \gamma \leq -2\gamma. \end{aligned} \quad (13)$$

Note that  $\eta$  and  $\gamma$  are two positive variables, thus the second-order gradient  $\frac{\partial^2 \mathcal{L}}{\partial (\hat{\mathbf{w}}^t)^2} < 0$  holds always. Consequently,  $\mathcal{L}(\hat{\mathbf{w}}^{t+1})$  can only be a local maxima, instead of a minima, which raises a contradiction with convergence in the training process. Such a contradiction indicates that the training algorithm will be convergent until no oscillation occurs, due to the additional term  $\mathcal{S}(\alpha, \mathbf{w})$ . Therefore, we completes our proof.  $\square$

Our proposition and proof reveal that the balanced parameter  $\gamma$  is actually a ‘‘threshold’’. A very small ‘‘threshold’’ fails to mitigate the frequent oscillation effectively while a too large one suppresses the necessary sign inversion and hinders the gradient descent process. To solve this, we devise the learning rule of  $\gamma$  as:

$$\gamma_i^{n,t+1} = \frac{1}{M^n} \|\mathbf{b}^{\mathbf{w}_i^{n,t}} \otimes \mathbf{b}^{\mathbf{w}_i^{n,t+1}} - \mathbf{1}\|_0 \cdot \max_{1 \leq j \leq M^n} \left( \left| \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_{i,j}^{n,t}} \right| \right), \quad (14)$$

where the first item  $\frac{1}{M^n} \|\mathbf{b}^{\mathbf{w}_i^{n,t}} \otimes \mathbf{b}^{\mathbf{w}_i^{n,t+1}} - \mathbf{1}\|_0$  denotes the proportion of weights with sign changed. The second item  $\max_{1 \leq j \leq M^n} \left( \left| \frac{\partial \mathcal{L}}{\partial \hat{\mathbf{w}}_{i,j}^{n,t}} \right| \right)$  is derived from Eq. (12), denoting the gradient with the greatest magnitude of the  $t$ -th iteration. In this way, we suppress the frequent weight oscillation by a resilient gradient.

We further optimize the scaling factor as:

$$\delta \alpha_i^n = \frac{\partial \mathcal{L}}{\partial \alpha_i^n} + \frac{\partial \mathcal{L}_R}{\partial \alpha_i^n}. \quad (15)$$

The gradient derived from softmax loss can be easily calculated according to back propagation. Based on Eq. (7), it is easy to derive:

$$\frac{\partial \mathcal{L}_R}{\partial \alpha_i^n} = -\gamma_i^n (\mathbf{w}_i^n - \alpha_i^n \mathbf{b}^{\mathbf{w}_i^n}) \otimes \mathbf{b}^{\mathbf{w}_i^n}. \quad (16)$$

## Experiments

Our ReBNNs are evaluated first on image classification and object detection tasks for visual recognition. Then, we evaluate ReBNN on the GLUE (Wang et al. 2018) benchmark with diverse NLP tasks. In this section, we first introduce the implementation details of ReBNN. Then we validate the effectiveness of the balanced parameter in the ablation study. Finally, we compare our method with state-of-the-art BNNs on various tasks to demonstrate the superiority of ReBNNs.

### Datasets and Implementation Details

**Datasets:** For its huge scope and diversity, the ImageNet object classification dataset (Russakovsky et al. 2015) is more demanding, which has 1000 classes, 1.2 million training photos, and 50k validation images.

The COCO dataset includes images from 80 different categories. All our experiments on COCO dataset are conducted on the COCO 2014 (Lin et al. 2014) object detection track in the training stage, which contains the combination of 80k images from the COCO `train2014` and 35k images sampled from COCO `val2014`, *i.e.*, COCO `trainval35k`. Then we test our method on the remaining 5k images from the COCO `minival`. We report the average precision (AP) for  $\text{IoUs} \in [0.5: 0.05: 0.95]$ , designated as  $\text{mAP}@[.5,.95]$ , using COCO’s standard evaluation metric. For further analyzing our method, we also report  $\text{AP}_{50}$ ,  $\text{AP}_{75}$ ,  $\text{AP}_s$ ,  $\text{AP}_m$ , and  $\text{AP}_l$ .

The GLUE benchmark contains multiple natural language understanding tasks. We follow (Wang et al. 2018) to evaluate the performance: Matthews correlation for CoLA, Spearman correlation for STS-B, and accuracy for the rest tasks: RTE, MRPC, SST-2, QQP, MNLI-m (matched), and MNLI-mm (mismatched). Also, for machine reading comprehension on SQuAD, we report the EM (exact match) and F1 score.

**Implementation details:** PyTorch (Paszke et al. 2017) is used to implement ReBNN. We run the experiments on 4 NVIDIA Tesla A100 GPUs with 80 GB memory. Following (Liu et al. 2018), we retain weights in the first layer, shortcut, and last layer in the networks as the real-valued.

For the image classification task, ResNets (He et al. 2016) are employed as the backbone networks to build our ReBNNs. We offer two implementation setups for fair comparison. First, we use **one-stage training** on ResNets, with SGD as the optimization algorithm and a momentum of 0.9, and a weight decay of  $1e-4$  following (Xu et al. 2021c).  $\eta$  is set to 0.1. The learning rates are optimized by the annealing cosine learning rate schedule. The number of epochs is set as 200. Then, we employ **two-stage training** following (Liu et al. 2020). Each stage counts 256 epochs. In this implementation, Adam is selected as the optimizer. And the network is supervised by a real-valued ResNet-34 teacher. The weight decay is set as 0 following (Liu et al. 2020). The learning rates  $\eta$  is set as  $1e-3$  and annealed to 0 by linear descent.

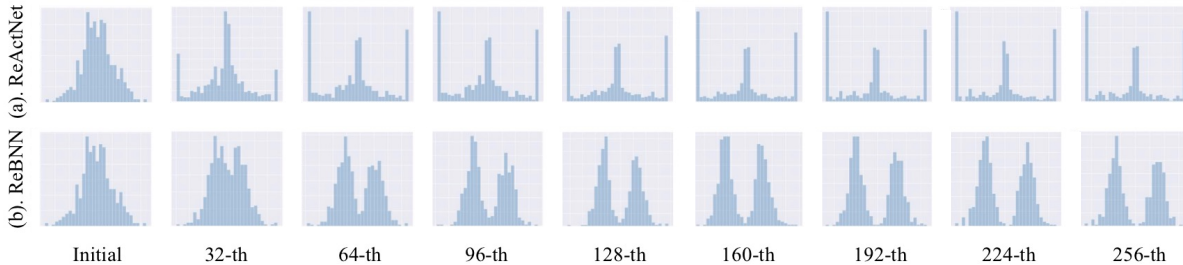


Figure 2: The evolution of latent weight distribution of (a) ReActNet and (b) ReBNN. We select the first channel of the first binary convolution layer to show the evolution. The model is initialized from the first stage training with W32A1 following (Liu et al. 2020). We plot the distribution every 32 epochs.

Value of $\gamma$	Top-1	Top-5
0	65.8	86.3
$1e-5$	66.2	86.7
$1e-4$	66.4	86.7
$1e-3$	66.3	86.8
$1e-2$	65.9	86.5
$\max_{1 \leq j \leq M^n} \left( \left  \frac{\partial \mathcal{L}}{\partial \mathbf{w}_{i,j}^{n,t}} \right  \right)$	66.3	86.2
<b>Eq. (14)</b>	<b>66.9</b>	<b>87.1</b>

Table 1: We compare different calculation method of  $\gamma$ , including constant varying from 0 to  $1e-2$  and gradient-based calculation.

For objection detection, we use the Faster-RCNN (Ren et al. 2016) and SSD (Liu et al. 2016), which are based on ResNet-18 (He et al. 2016) and VGG-16 (Simonyan and Zisserman 2015) backbone, respectively. We fine-tune the detector on the dataset for object detection. For SSD and Faster-RCNN, the batch size is set to 16 and 8, respectively, with applying SGD optimizer.  $\eta$  is equal to 0.008. We use the same structure and training settings as BiDet (Wang et al. 2020) on the SSD framework. The input resolution is  $1000 \times 600$  for Faster-RCNN and  $300 \times 300$  for SSD, respectively.

For the natural language processing task, we conduct experiments based on BERT<sub>BASE</sub> (Devlin et al. 2018) (with 12 hidden layers) architecture following BiBERT (Qin et al. 2022), respectively. The detailed training setups are the same as BiBERT. We extend the ReBNN to multi-layer perceptrons (MLPs) and use the Bi-Attention following BiBERT (Qin et al. 2022).

### Ablation Study

Since no extra hyper-parameter is introduced, we first evaluate the different calculation of  $\gamma$ . Then we show how our ReBNN achieves a resilient training process. In the ablation study, we use the ResNet-18 backbone initialized from the first stage training with W32A1 following (Liu et al. 2020).

**Calculation of  $\gamma$ :** We compare the different calculations of  $\gamma$  in this part. As shown in Tab. 1, the performances increase first and then decrease when increasing the value of constant  $\gamma$ . Considering that the gradient magnitude varies layer-wise and channel-wise, a subtle  $\gamma$  can hardly be manually set as

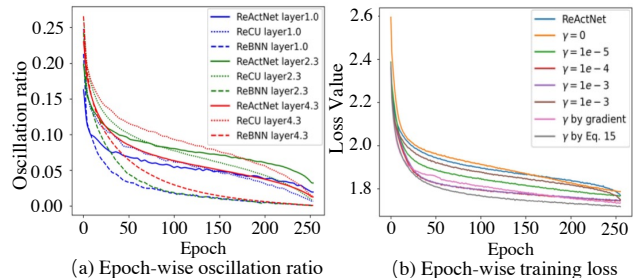


Figure 3: (a) The epoch-wise weight oscillation ratio of ReActNet (solid), ReCU (dotted) and ReBNN (dashed). (b) Comparing the loss curves of ReActNet and our ReBNN with different calculation of  $\gamma$ .

a global value. We further compare the gradient-based calculation. To avoid extreme values, we set the upper bound as  $2e-4$  and lower bound as  $1e-5$ . As shown in the bottom lines, we first use  $\max_{1 \leq j \leq M^n} \left( \left| \frac{\partial \mathcal{L}}{\partial \mathbf{w}_{i,j}^{n,t}} \right| \right)$ , the maximum intra-channel gradient of last iteration, which shows a similar performance compared with the constant  $1e-4$ . This indicates that only using the maximum intra-channel gradient may also suppress necessary sign flip, thus hindering the training. Inspired by this, we use Eq. (14) to calculate  $\gamma$  and improve the performance by 0.6%, showing that considering the weight oscillation proportion allows the necessary sign flip and leads a more effective training. We also show the training loss curves in Fig. 3(b). As plotted, the curves of  $\mathcal{L}$  almost demonstrate the degrees of training sufficiency. Thus we draw the conclusion that ReBNN with  $\gamma$  calculated by Eq. (14) achieves the lowest training loss as well as an efficient training process.

Note that the loss may not be minimal at each training iteration, but our method is just a reasonable version of gradient descent algorithms by nature, which can be used to solve the optimization problem as general. We empirically prove ReBNN’s capability of mitigating the weight oscillation, leading to a better convergence.

**Resilient training process:** We first show the evolution of the latent weight distribution is this section. We plot the

Network	Method	#Bits	Size <sub>(MB)</sub>	OPs <sub>(10<sup>8</sup>)</sub>	Top-1	Top-5	
ResNet-18	Real-valued	32-32	46.76	18.21	69.6	89.2	
	BNN				42.2	67.1	
	XNOR-Net				51.2	73.2	
	Bi-Real Net				56.4	79.5	
	RBNN	1-1	4.15	1.63	59.6	81.6	
	ReCU				61.0	82.6	
	<b>ReBNN<sub>1</sub></b>				<b>61.6</b>	<b>83.4</b>	
	ReActNet				65.9	86.2	
	FDA-BNN	1-1	4.15	1.63	65.8	86.4	
	ReCU				66.4	86.5	
	<b>ReBNN<sub>2</sub></b>				<b>66.9</b>	<b>87.1</b>	
ResNet-34	Real-valued	32-32	87.19	36.74	73.3	91.3	
	Bi-Real Net				62.2	83.9	
	RBNN	1-1	5.41	1.93	63.1	84.4	
	ReCU				65.1	85.8	
	<b>ReBNN<sub>1</sub></b>				<b>65.8</b>	<b>86.2</b>	
	ReActNet				69.3	88.6	
		<b>ReBNN<sub>2</sub></b>	1-1	5.41	1.93	<b>69.9</b>	<b>88.9</b>

Table 2: A performance comparison with SOTAs on ImageNet with different training strategies. #Bits denotes the bit width of weights and activations. We report the Top-1 (%) and Top-5 (%) accuracy performances. ReBNN<sub>1</sub> and ReBNN<sub>2</sub> denote our ReBNN learned with one-stage and two-stage training.

Framework	Backbone	Method	#Bits	Size <sub>(MB)</sub>	OPs <sub>(G)</sub>	mAP @ [.5, .95]	AP <sub>50</sub>	AP <sub>75</sub>	AP <sub>s</sub>	AP <sub>m</sub>	AP <sub>l</sub>
Faster-RCNN	ResNet-18	Real-valued	32-32	47.48	434.39	26.0	44.8	27.2	10.0	28.9	39.7
		DoReFa-Net	4-4	6.73	55.90	22.9	38.6	23.7	8.0	24.9	36.3
		XNOR-Net				10.4	21.6	8.8	2.7	11.8	15.9
		Bi-Real Net				14.4	29.0	13.4	3.7	15.4	24.1
		BiDet	1-1	2.39	8.58	15.7	31.0	14.4	4.9	16.7	25.4
		<b>ReBNN</b>				<b>19.6</b>	<b>37.6</b>	<b>20.4</b>	<b>7.0</b>	<b>20.1</b>	<b>33.1</b>
SSD	VGG-16	Real-valued	32-32	105.16	31.44	23.2	41.2	23.4	5.3	23.2	39.6
		DoReFa-Net	4-4	29.58	6.67	19.5	35.0	19.6	5.1	20.5	32.8
		XNOR-Net				8.1	19.5	5.6	2.6	8.3	13.3
		Bi-Real Net				11.2	26.0	8.3	3.1	12.0	18.3
		BiDet	1-1	21.88	2.13	13.2	28.3	10.5	5.1	14.3	20.5
		<b>ReBNN</b>				<b>18.1</b>	<b>33.9</b>	<b>17.5</b>	4.2	<b>17.9</b>	<b>25.9</b>

Table 3: Comparison of mAP@[.5, .95](%), AP (%) with different IoU threshold and AP for objects in various sizes with state-of-the-art binarized object detectors on COCO minival. #Bits denotes the bit width of weights and activations

distribution of the first channel of the first binary convolution layer per 32 epochs in Fig. 2. As seen, our ReBNN can efficiently redistribute the BNNs towards resilience. Conventional ReActNet (Liu et al. 2020) possesses a tri-modal distribution, which is unstable due to the scaling factor with large magnitudes. In contrast, our ReBNN is constrained by the balanced parameter  $\gamma$  during training, thus leading to a resilient bi-modal distribution with fewer weights centering around zero. We also plot the ratios of sequential weight oscillation of ReBNN and ReActNet for the 1-st, 8-th, and 16-th binary convolution layers of ResNet-18. As shown in Fig. 3(a), the dashed lines gain much lower magnitudes than the solid (ReActNet) and dotted (ReCU (Xu et al. 2021c)) lines with the same color, validating the effectiveness of our ReBNN in suppressing the consecutive weight oscillation. Besides, the sequential weight oscillation ratios of ReBNN are gradually decreased to 0 as the training converges.

## Image Classification

We first show the experimental results on ImageNet with ResNet-18 and ResNet-34 (He et al. 2016) backbones in

Tab. 2. We compare ReBNN with BNN (Courbariaux, Bengio, and David 2015), XNOR-Net (Rastegari et al. 2016), Bi-Real Net (Liu et al. 2018), RBNN (Lin et al. 2020), and ReCU (Xu et al. 2021c) for the one-stage training strategy. For the two-stage training strategy, we compare with ReActNet (Liu et al. 2020), FDA-BNN (Xu et al. 2021b), and ReCU (Xu et al. 2021c).

We evaluate our ReBNN with one-stage training following (Xu et al. 2021c). ReBNN outperforms all of the compared binary models in both Top-1 and Top-5 accuracy, as shown in Tab. 2. ReBNN-based ResNet-18 respectively achieves 61.6% and 83.4% in Top-1 and Top-5 accuracy, with 0.6% increases over state-of-the-art ReCU. ReBNN further outperforms all compared methods with ResNet-34 backbone, achieving 65.8% Top-1 accuracy. We further evaluate our ReBNN with two-stage training following (Liu et al. 2020), where ReBNN surpasses ReActNet by 1.0%/0.6% Top-1 accuracy with ResNet-18/34 backbones, respectively.

In this paper, we use memory usage and OPs following (Liu et al. 2018) in comparison to other tasks for further reference. As shown in Tab. 2, ReBNN theoretically

Method	#Bits	Size <sub>(MB)</sub>	OPs <sub>(G)</sub>	MNLI-m/mm	QQP	QNLI	SST-2	CoLA	STS-B	MRPC	RTE	Avg.
Real-valued	32-32-32	418	22.5	84.9/85.5	91.4	92.1	93.2	59.7	90.1	86.3	72.2	83.9
Q-BERT	2-8-8	43.0	6.5	76.6/77.0	-	-	84.6	-	-	68.3	52.7	-
Q2BERT	2-8-8	43.0	6.5	47.2/47.3	67.0	61.3	80.6	0	4.4	68.4	52.7	47.7
TernaryBERT	2-2-2	28.0	1.5	40.3/40.0	63.1	50.0	80.7	0	12.4	68.3	54.5	45.5
BinaryBERT	-	-	-	35.6/35.3	66.2	51.5	53.2	0	6.1	68.3	52.7	41.0
BiBERT	1-1-1	16.5	0.4	66.1/67.5	84.8	72.6	88.7	25.4	33.6	72.5	57.4	63.2
<b>ReBNN</b>	-	-	-	<b>69.9/71.3</b>	<b>85.2</b>	<b>79.2</b>	<b>89.3</b>	<b>28.8</b>	<b>38.7</b>	<b>72.6</b>	56.9	<b>65.8</b>

Table 4: Comparison of BERT quantization methods without data augmentation. #Bits denotes the bit width of weights, word embedding, and activations. “Avg.” denotes the average results.

Network	Method	#Bits	Size <sub>(MB)</sub>	Memory Saving	Latency <sub>(ms)</sub>	Acceleration
ResNet-18	Real-valued	32-32	46.76	-	583.1	-
	ReBNN	1-1	4.15	11.26×	67.5	8.64×
ResNet-34	Real-valued	32-32	87.19	-	1025.6	-
	ReBNN	1-1	5.41	16.12×	113.6	9.03×

Table 5: Comparing ReBNN with real-valued models on hardware (single thread).

accelerates ResNet-18/34 by 11.17× and 19.04×, which is significant for real-time applications.

## Object Detection

On COCO, the proposed ReBNN is compared against state-of-the-art 1-bit neural networks such as XNOR-Net (Rastegari et al. 2016), Bi-Real Net (Liu et al. 2018), and BiDet (Wang et al. 2020). We present the performance of the 4-bit DoReFa-Net (Zhou et al. 2016) for reference.

As shown in Tab. 3, compared with state-of-the-art XNOR-Net, Bi-Real Net, and BiDet, our method improves the mAP@[.5,.95] by 9.2%, 5.2%, and 3.9% using the Faster-RCNN framework with the ResNet-18 backbone. Moreover, on other APs with different IoU thresholds, our ReBNN clearly beats others. Compared to DoReFa-Net, a quantized neural network with 4-bit weights and activations, our ReBNN obtains only 3.3% lower mAP. Our method yields a 1-bit detector with a performance of only 6.1% mAP lower than the best-performing real-valued counterpart (19.6% vs. 26.0%). Similarly, using the SSD300 framework with the VGG-16 backbone, our method achieves 18.1% mAP@[.5,.95], outperforming XNOR-Net, Bi-Real Net, and BiDet by 10.0%, 6.9%, and 4.9% mAP, respectively. Our ReBNN also achieves highly efficient models by theoretically accelerating Faster-RCNN and SSD by 50.62× and 14.76×.

## Natural Language Processing

In Tab. 4, we show experiments on the BERT<sub>BASE</sub> architecture and the GLUE benchmark without data augmentation following BiBERT (Qin et al. 2022). Experiments show that outperforms other methods on the development set of GLUE benchmark, including TernaryBERT (Zhang et al. 2020), BinaryBERT (Bai et al. 2020), Q-BERT (Shen et al. 2020), Q2BERT (Shen et al. 2020), and BiBERT (Qin et al. 2022). Our ReBNN surpasses existing methods on BERT<sub>BASE</sub> architecture by a clear margin in the average accuracy. For example, our ReBNN surpasses BiBERT by 6.6% accuracy on QNLI dataset, which is significant for the natural language

processing task. We observe our ReBNN brings improvements on 7 out of total 8 datasets, thus leading to a 2.6% average accuracy improvement. Our ReBNN also achieves highly efficient models by theoretically accelerating the BERT<sub>BASE</sub> architecture by 56.25×.

## Deployment Efficiency

We implement the 1-bit models achieved by our ReBNN on ODROID C4, which has a 2.016 GHz 64-bit quad-core ARM Cortex-A55. With evaluating its real speed in practice, the efficiency of our ReBNN is proved when deployed into real-world mobile devices. We leverage the SIMD instruction SSL on ARM NEON to make the inference framework BOLT (Feng 2021) compatible with ReBNN. We compare ReBNN to the real-valued backbones in Tab. 5. We can see that ReBNN’s inference speed is substantially faster with the highly efficient BOLT framework. For example, the acceleration rate achieves about 8.64× on ResNet-18, which is slightly lower than the theoretical acceleration rate. For ResNet-34 backbone, ReBNN can achieve 9.03× acceleration rate with BOLT framework on hardware, which is significant for the computer vision on real-world edge devices.

## Conclusion

In this paper, we analyze the influence of frequent weight oscillation in binary neural networks and proposed a Resilient Binary Neural Network (ReBNN) to provide resilient gradients for latent weights updating. Our method specifically proposes to parameterize the scaling factor and introduces a weighted reconstruction loss to build an adaptive training objective. We further manifest that the balanced parameter can serve as an indicator to reflect the frequency of the weight oscillation during back propagation. Our ReBNN reaches resilience by learning the balanced parameter, leading to a great reduction of weight oscillation. ReBNN shows strong generalization to gain impressive performance on various tasks such as image classification, object detection, and natural language processing tasks, demonstrating the superiority of the proposed method over state-of-the-art BNNs.



## Acknowledgments

This work was supported by National Natural Science Foundation of China under Grant 62076016, 62206272, 62141604, 61972016, 62032016, Beijing Natural Science Foundation L223024, National Key R&D Program of China (NO.2022ZD0160100), and in part by Shanghai Committee of Science and Technology (Grant No. 21DZ1100100).

## References

- Bai, H.; Zhang, W.; Hou, L.; Shang, L.; Jin, J.; Jiang, X.; Liu, Q.; Lyu, M.; and King, I. 2020. Binarybert: Pushing the limit of bert quantization. In *Proc. of ACL*, 4334–4348.
- Courbariaux, M.; Bengio, Y.; and David, J.-P. 2015. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Proc. of NeurIPS*, 3123–3131.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Feng, J. 2021. Bolt. <https://github.com/huawei-noah/bolt>. Accessed: 2020-03-10.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proc. of ICCV*, 1026–1034.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. of CVPR*, 770–778.
- He, Y.; Kang, G.; Dong, X.; Fu, Y.; and Yang, Y. 2018. Soft filter pruning for accelerating deep convolutional neural networks. In *Proc. of IJCAI*, 2234–2240.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *Proc. of CVPR*, 7132–7141.
- Juefei-Xu, F.; Naresh Boddeti, V.; and Savvides, M. 2017. Local binary convolutional neural networks. In *Proc. of CVPR*, 19–28.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Proc. of NeurIPS*, 1097–1105.
- Li, Y.; Xu, S.; Zhang, B.; Cao, X.; Gao, P.; and Guo, G. 2022. Q-ViT: Accurate and Fully Quantized Low-bit Vision Transformer. *arXiv preprint arXiv:2210.06707*.
- Lin, M.; Ji, R.; Xu, Z.; Zhang, B.; Chao, F.; Lin, C.-W.; and Shao, L. 2022. Siman: Sign-to-magnitude network binarization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–12.
- Lin, M.; Ji, R.; Xu, Z.; Zhang, B.; Wang, Y.; Wu, Y.; Huang, F.; and Lin, C.-W. 2020. Rotated Binary Neural Network. In *Proc. of NeurIPS*, 1–9.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Proc. of ECCV*, 740–755.
- Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. Ssd: Single shot multibox detector. In *Proc. of ECCV*, 21–37.
- Liu, Z.; Shen, Z.; Savvides, M.; and Cheng, K.-T. 2020. Re-ActNet: Towards Precise Binary Neural Network with Generalized Activation Functions. In *Proc. of ECCV*, 143–159.
- Liu, Z.; Wu, B.; Luo, W.; Yang, X.; Liu, W.; and Cheng, K.-T. 2018. Bi-Real Net: Enhancing the Performance of 1-Bit CNNs with Improved Representational Capability and Advanced Training Algorithm. In *Proc. of ECCV*, 722–737.
- Martinez, B.; Yang, J.; Bulat, A.; and Tzimiropoulos, G. 2020. Training binary neural networks with real-to-binary convolutions. In *Proc. of ICLR*, 1–11.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch. In *Proc. of NeurIPS Workshop*, 1–4.
- Qin, H.; Ding, Y.; Zhang, M.; Yan, Q.; Liu, A.; Dang, Q.; Liu, Z.; and Liu, X. 2022. BiBERT: Accurate Fully Binarized BERT. In *Proc. of ICLR*, 1–24.
- Qin, L.; Che, W.; Li, Y.; Wen, H.; and Liu, T. 2019. A Stack-Propagation Framework with Token-Level Intent Detection for Spoken Language Understanding. In *Proc. of EMNLP*, 2078–2087.
- Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Proc. of ECCV*, 525–542.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2016. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(6): 1137–1149.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision*, 115(3): 211–252.
- Shen, S.; Dong, Z.; Ye, J.; Ma, L.; Yao, Z.; Gholami, A.; Mahoney, M. W.; and Keutzer, K. 2020. Q-bert: Hessian based ultra low precision quantization of bert. In *Proc. of AAAI*, 8815–8821.
- Simonyan, K.; and Zisserman, A. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *Proc. of ICLR*, 1–13.
- Wang, A.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. R. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proc. of ICLR*, 1–20.
- Wang, Z.; Wu, Z.; Lu, J.; and Zhou, J. 2020. BiDet: An Efficient Binarized Object Detector. In *Proc. of CVPR*, 2049–2058.
- Xu, S.; Li, Y.; Wang, T.; Ma, T.; Zhang, B.; Gao, P.; Qiao, Y.; Lü, J.; and Guo, G. 2022a. Recurrent bilinear optimization for binary neural networks. In *Proc. of ECCV*, 19–35.
- Xu, S.; Li, Y.; Zeng, B.; Ma, T.; Zhang, B.; Cao, X.; Gao, P.; and Lü, J. 2022b. IDa-Det: An Information Discrepancy-Aware Distillation for 1-Bit Detectors. In *Proc. of ECCV*, 346–361.



Xu, S.; Zhao, J.; Lu, J.; Zhang, B.; Han, S.; and Doermann, D. 2021a. Layer-Wise Searching for 1-Bit Detectors. In *Proc. of CVPR*, 5682–5691.

Xu, Y.; Han, K.; Xu, C.; Tang, Y.; Xu, C.; and Wang, Y. 2021b. Learning frequency domain approximation for binary neural networks. In *Proc. of NeurIPS*, 25553–25565.

Xu, Z.; Lin, M.; Liu, J.; Chen, J.; Shao, L.; Gao, Y.; Tian, Y.; and Ji, R. 2021c. Recu: Reviving the dead weights in binary neural networks. In *Proc. of ICCV*, 5198–5208.

Zhang, W.; Hou, L.; Yin, Y.; Shang, L.; Chen, X.; Jiang, X.; and Liu, Q. 2020. Ternarybert: Distillation-aware ultra-low bit bert. In *Proc. of EMNLP*, 509–521.

Zhou, S.; Wu, Y.; Ni, Z.; Zhou, X.; Wen, H.; and Zou, Y. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*.