

# BiPer: Binary Neural Networks using a Periodic Function

Edwin Vargas<sup>1,†</sup>, Claudia V. Correa<sup>1</sup>, Carlos Hinojosa<sup>2</sup>, Henry Arguello<sup>1</sup>

<sup>1</sup>Universidad Industrial de Santander (UIS); <sup>2</sup>King Abdullah University of Science and Technology

<sup>†</sup>edwin.vargas4@correo.uis.edu.co

## Abstract

*Quantized neural networks employ reduced precision representations for both weights and activations. This quantization process significantly reduces the memory requirements and computational complexity of the network. Binary Neural Networks (BNNs) are the extreme quantization case, representing values with just one bit. Since the sign function is typically used to map real values to binary values, smooth approximations are introduced to mimic the gradients during error backpropagation. Thus, the mismatch between the forward and backward models corrupts the direction of the gradient causing training inconsistency problems and performance degradation. In contrast to current BNN approaches, we propose to employ a binary periodic (BiPer) function during binarization. Specifically, we use a square wave for the forward pass to obtain the binary values and employ the trigonometric sine function with the same period of the square wave as a differentiable surrogate during the backward pass. We demonstrate that this approach can control the quantization error by using the frequency of the periodic function and improves network performance. Extensive experiments validate the effectiveness of BiPer in benchmark datasets and network architectures, with improvements of up to 1% and 0.69% with respect to state-of-the-art methods in the classification task over CIFAR-10 and ImageNet, respectively. Our code is publicly available at <https://github.com/edmav4/BiPer>.*

## 1. Introduction

Deep Neural Networks (DNN) have achieved unprecedented results in many high-level tasks, such as classification, segmentation, and detection, with a tremendous concurrent impact in computer vision, natural language processing, information retrieval, and many others [42]. Typically, DNNs rely on full-precision (32 bit) weights and activation functions. Accurate and precise models, however, become expensive in terms of computation, storage and number of parameters. For this reason, DNN deployment is usually prohibited for devices with limited re-

sources, such as mobile, hand-held or wearables. Different approaches to reduce computation requirements include efficient neural network architecture design [20, 22, 40], network pruning [19], knowledge distillation [15], low rank tensor decomposition [31], hashing [12], and network quantization [33, 37, 47]. Among them, network quantization has become one of the most promising techniques, aiming at compressing large models usually stored as floating-point weights with low bitwidth numbers. Binary Neural Networks (BNNs) are the extreme quantization case, where weights and activation functions are constrained to just one bit, i.e., binary values, typically +1 or -1. In contrast to DNNs, BNNs replace heavy matrix computations by bit-wise operations, yielding to 32× memory compression, and 58× speed-up on CPUs [36]. Thus, this approach drastically reduces the computational requirements and accelerates inference, making BNNs particularly appealing for resource-constrained environments such as edge devices and mobile applications.

Despite significant advantages for efficient BNN deployment in hardware with limited capabilities, the binarization of full-precision models severely degrades the accuracy performance in high-level tasks such as object detection, classification, segmentation, and others [37]. For instance, in large datasets such as ImageNet, one of the earliest BNN models, the XNOR-Net [36], achieved an accuracy degradation of around 18% compared to the full precision ResNet-18 architecture. Recent efforts have been devoted to close the performance gap of BNN with respect to their real-valued counterparts. Nonetheless, state-of-the-art approaches still exhibit accuracy degradations of approximately 8% [45].

Binarization of real-valued weights and activations is generally performed using the sign function during the feed-forward procedure. A relevant limitation of the sign function is that its gradient is null everywhere except in zero, which makes it incompatible with error back-propagation methods, due to the non-differentiability of binary operations. To overcome this issue, various techniques like the straight-through estimator (STE) and relaxed training approaches have been adopted [21]. STE essentially substi-

tutes the sign function for the identity function to calculate the gradients during the backwards process. Since there exists a mismatch between the forward and backward pass caused by the STE approximation, research efforts have focused on designing better smooth and differentiable functions to estimate the gradient of the sign function [14, 25]. Although these approaches have improved the accuracy of BNNs, gradient instability persists when the quantization error is minimized, i.e., when the approximation functions are close to the sign function. Therefore, there is still a significant performance gap between real and binary models, and the pursuit of refined quantization functions for BNN, with lower accuracy degradation and higher gradient stability, remains an open challenge.

Instead of using the Sign function, in this work, we propose to address the aforementioned issues of extreme 1-bit quantization by using a binary periodic (BiPer) function or square wave function to promote binary weight values. Thus, opposite to the sign function which is always negative for negative values or positive for positive values, the proposed periodic function can reach positive and negative values in the whole domain of the latent weights. To the best of our knowledge, this is the first time a binary periodic function is employed for the binarization problem in quantized neural networks. Since the gradient of the periodic function still faces the problem of being zero almost everywhere, it cannot be directly integrated within a back-propagation algorithm based on gradient descent. We solved this problem by employing a sinusoidal function with the same fundamental frequency of the periodic function as a differentiable surrogate during the backward pass. The continuity and differentiable characteristics of the sine function, make it suitable for stochastic gradient methods. In contrast to existing BNN methods that smoothly and progressively approximate the sign function to reduce the quantization error (QE), we will show that in the proposed BiPer approach the QE can be controlled by the frequency of the periodic function. We further leverage this property to provide an initialization of the weights that better balances the trade-off between the estimation error and performance accuracy. Experimental results demonstrate that BiPer provides the best network performance for the classification task, with respect to state-of-the-art BNN approaches on the CIFAR-10 and ImageNet data sets.

**Paper Contribution.** The contributions of our work are summarized as follows:

- We propose a simple yet powerful and effective modification in the binarization process, by including a binary periodic function.
- We introduce a continuous, periodic sinusoidal function as a differentiable surrogate of the binary periodic function during the back-propagation process, suitable for stochastic gradient methods.

- We mathematically analyze the quantization error of BiPer and show that it can be controlled by the frequency of the periodic function. This leverages a flexible initialization scheme for the binary weights, that balances the QE and network performance accuracy.
- Experiments on benchmark data sets demonstrate the advantages of BiPer for the classification task with respect to state-of-the-art BNN approaches. BiPer outperforms prior works by up to 1% on CIFAR-10, and 0.4% on ImageNet.

## 2. Related Work

### 2.1. Binary Neural Networks

Among pioneering work on binary neural networks, BinaryConnect introduced a training scheme involving binary weights during training, and real-valued weights during the forward pass for weight updates [8]. Subsequent approaches have focused on designing differentiable surrogates for the sign function to reduce the effect of the quantization error. For instance, a differentiable soft quantization (DSQ) approach was proposed in [14], which introduced a tanh-alike differentiable asymptotic function to estimate the forward and backward of the sign function. An improved BNN (BNN+) [10] employs a SignSwish activation function to modify the back-propagation of the sign function and includes a regularization that encourages the weights around binary values. A more recent approach proposed an approximation of the sign function that varies along the training process [25]. The authors in [11] introduced an approach to training BNN by focusing on the regularization of activation distributions. Such regularization term encourages the binary activations to follow a balanced distribution during training. On the other hand, the rotated BNN [26] incorporates rotation operations into the training process, through a rotation-invariant loss function that mitigates the sensitivity of BNNs to input rotations. IR-Net [34] employs a mechanism that preserves information in the form of real numbers during forward propagation and utilizes it in the backward pass. ReCU [45] addresses the issue of dead weights in binary neural networks, by introducing a reconfiguration mechanism that revives inactive weights during training. ResTE [43] proposed an approach for training binary BNNs with a Rectified Straight Through Estimator, which addresses the limitations of the traditional STE. A rectification was introduced into the gradient estimation process during back-propagation, to enhance the accuracy of gradient signals, leading to improved training stability and convergence for BNNs. MST [41] proposed an approach for compressing and accelerating BNNs through the use of a Minimum Spanning Tree (MST). The MST-compression method leverages the structural characteristics of the network to identify an optimal subset of weights, so

that redundant connections are pruned, leading to a more compact BNN.

## 2.2. Periodic Functions in Deep Neural Networks

Periodic functions were successfully introduced as activations within neural network architectures in [39], where a sinusoidal function is used in a multilayer perceptron (MLP) to encode a spatial field into the network weights. In contrast to typical non-linear activation functions such as ReLU or tanh, periodic functions are capable of representing complex natural signals and their derivatives. For this reason, they have been since employed for 3D-aware image synthesis to represent scenes as view-consistent radiance fields [7, 30], as well as encoding model equations, like Partial Differential Equations (PDE), as a component of physics informed neural networks (PINN) [9]. It should be noted that all the aforementioned approaches use full-precision models, while this work employs a binary periodic function to promote binary weights in the forward model, and a sinusoidal function to better approximate the gradients of the binary weights in the backpropagation process. To the best of our knowledge, this is the first work that exploits the characteristics of periodic functions to improve the performance of BNNs.

## 3. BNN Preliminaries

In this section, we introduce the essential mathematical formulations of neural network binarization. The main feature of BNN is to constrain the weights of a DNN as well as its activations to have binary values. The most common approach to obtain binary weights and activations consists of applying the sign function to a real-valued variable as

$$w^q = \phi(w) = \text{Sign}(w) = \begin{cases} 1, & \text{if } w \geq 0, \\ -1, & \text{otherwise,} \end{cases} \quad (1)$$

where  $w^q$  is the binarized variable and,  $w$  is the real-valued variable. Here, we consider the binarization of a CNN model, however, these ideas can be extended to any neural model. Considering the binary weights  $\mathbf{w}^q$  and the activations  $\mathbf{a}^q$ , the convolution operation can be formulated as

$$\mathbf{y} = \mathbf{w}^q \oplus \mathbf{a}^q, \quad (2)$$

where the symbol  $\oplus$  denotes bit-wise operations including XNOR and BITCOUNT. It has been previously shown that the direct binarization of the convolution result in (2) introduces large quantization errors. To alleviate this problem [36] proposed to introduce real-valued scaling factors  $\alpha$  and  $\beta$  to the weights and activations, respectively. Thus, the binary convolution can be expressed as

$$\mathbf{y} = \alpha\beta \odot (\mathbf{w}^q \oplus \mathbf{a}^q), \quad (3)$$

where  $\odot$  denotes element-wise multiplication. Further, the work in [4], showed that combining  $\alpha$  and  $\beta$  in a single scaling factor  $\gamma$  results in better performance. Since the gradient of the sign function is zero almost everywhere, the optimization of the weight parameters is incompatible with back-propagation algorithms based on gradient descent. To solve this problem, the work in [8], introduced the straight-through estimator (STE) to approximate the gradient with respect to weights as

$$\frac{\partial \mathcal{L}}{\partial w} = \frac{\partial \mathcal{L}}{\partial w^q} \frac{\partial w^q}{\partial w} \approx \frac{\partial \mathcal{L}}{\partial w^q}. \quad (4)$$

More precisely STE uses the identity function as a differentiable surrogate of the sign function. It means that the gradient with the latent full-precision weights straightly equals to the gradient of the binarized outputs, which is also the origin of the name straight through estimator. To reduce the gradient error the clip function is also commonly employed, yielding to the approximation

$$\frac{\partial \mathcal{L}}{\partial w} = \frac{\partial \mathcal{L}}{\partial w^q} \frac{\partial w^q}{\partial w} \approx \frac{\partial \mathcal{L}}{\partial w^q} \mathbf{1}_{\mathcal{A}}(w), \quad (5)$$

where  $\mathbf{1}_{\mathcal{A}}$  is the indicator function of the clipping set  $\mathcal{A} = \{w \in \mathbb{R} : |w| \leq 1\}$ . While, in principle, the binarization of activations can be performed in the same manner as for the weights, several approaches have shown that using a different approximation for binarizing activations leads to improved performance [14, 28, 34]. In this work, following [28], we use a piece-wise polynomial function to approximate the gradient with respect to the activations as follows

$$\frac{\partial \mathcal{L}}{\partial a} = \frac{\partial \mathcal{L}}{\partial a^q} \cdot \frac{\partial a^q}{\partial a} \approx \frac{\partial \mathcal{L}}{\partial a^q} \cdot \frac{\partial F(a)}{\partial a} \quad (6)$$

with

$$\frac{\partial F(a)}{\partial a} = \begin{cases} 2 + 2a, & \text{if } -1 \leq a < 0 \\ 2 - 2a, & \text{if } 0 \leq a < 1 \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

where  $a$  represents the real activation values, while  $a^q$  represents the binary activations. The main limitation of the STE is that the gradient approximation introduces a considerable inconsistency between the forward and backward passes. To reduce the degree of inconsistency, several previous studies have explored alternative gradient estimators. Nevertheless, it has been shown that reducing the estimation error often results in highly divergent gradients, that harm the model training and increase the risk of gradient vanishing and gradient exploding [43].

## 4. BiPer

To overcome the gradient and quantization error challenges from existing binarization methods and their gradient approximation functions, we propose to use a binary periodic

function or square wave function (see Fig. 1) instead of just the sign function to model the binary weights. In contrast to the sign function described in section 3 and depicted in Fig. 1(a), which is always negative for negative values of  $w$ , the proposed periodic function (Fig. 1(b)) can reach positive and negative values in the whole domain of the latent weights.

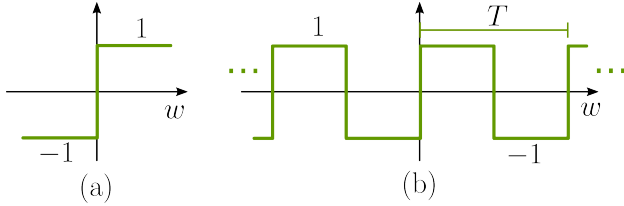


Figure 1. (a) Sign function. (b) Binary periodic function.

It should be pointed out that the gradient of the periodic function still faces the problem of being zero almost everywhere, therefore, it cannot be directly integrated within a back-propagation algorithm based on gradient descent. To solve this problem, we first rewrite the square wave function as

$$w^q = \text{Sign}(\sin(\omega_0 w)), \quad (8)$$

where  $\omega_0 = \frac{2\pi}{T}$  is the angular frequency. We note that this corresponds to applying the sign function to the first harmonic of the periodic function. Based on (8), we can approximate the gradient with respect to the weights as

$$\frac{\partial \mathcal{L}}{\partial w} = \frac{\partial \mathcal{L}}{\partial w^q} \frac{\partial w^q}{\partial w} \approx \frac{\partial \mathcal{L}}{\partial w^q} \frac{\partial \hat{w}}{\partial w} \quad (9)$$

where

$$\hat{w} = \sin(\omega_0 w). \quad (10)$$

Note that the last differential term in (9) corresponds to the gradient of a continuous differentiable sinusoidal function, which is also a smooth periodic function, and proportional to the frequency  $\omega_0$ . Comparing Eq. (9) to the STE approximation in (4), it can be seen that the proposed approximation is the product of the STE estimator of a periodic function and a proportional constant given by the frequency. In the following section, we explore this property to control the quantization error with the frequency  $\omega_0$ .

#### 4.1. Quantization Error Analysis

This section shows that an additional advantage of using the periodic function is its flexibility to control the quantization error. In particular, we mathematically demonstrate how a lower quantization error can be achieved by setting the fundamental period of the wave function. To this end, let us first assume that the latent weights roughly follow the zero-mean Laplace distribution, i.e.,  $\mathcal{W} \sim La(0, b)$  [1, 2, 48].

Since the weights  $\hat{w}$  in (10) before quantization are a function of a random variable, they are also a random variable  $\hat{\mathcal{W}} \in [-1, 1]$ . Computing the probability density function (pdf) of a random variable  $\mathcal{Y} = g(\mathcal{X})$  from the pdf of  $\mathcal{X}$  ( $f_{\mathcal{X}}(x)$ ) can be easily done employing the method of transformation [32], if the function  $g$  is differentiable and strictly increasing or decreasing, i.e., strictly monotonic. Thus, the pdf of  $\mathcal{Y}$  can be computed as

$$f_{\mathcal{Y}}(y) = \begin{cases} \frac{f_{\mathcal{X}}(x_1)}{|g'(x_1)|} = f_{\mathcal{X}}(x_1) \cdot \left| \frac{dx_1}{dy} \right| & \text{where } g(x_1) = y \\ 0 & \text{if } g(x) \neq y. \end{cases}$$

The more general case in which  $g$  is not monotonic requires splitting the domain into  $n$  intervals, so that  $g$  is strictly monotonic and differentiable on each partition. Then, the pdf can be obtained as

$$f_{\mathcal{Y}}(y) = \sum_{k=1}^n \frac{f_{\mathcal{X}}(x_k)}{|g'(x_k)|} = \sum_{k=1}^n f_{\mathcal{X}}(x_k) \cdot \left| \frac{dx_k}{dy} \right|, \quad (11)$$

where  $x_1, \dots, x_n$  are real solutions to  $g(x) = y$ . For BiPer, since the periodic function from Eq. (10) is not monotonic, we can use (11) to compute the pdf of  $\hat{\mathcal{W}}$  using the pdf of  $\mathcal{W}$ . Letting  $f_{\mathcal{W}}(w) = \frac{1}{b} \exp(-|w|/b)$  denote the pdf of  $\mathcal{W}$ , and setting  $g$  as the sine function, we can divide the sinusoidal function into subsequent intervals of  $T/2$  where it is strictly increasing or decreasing, alternately. For simplicity, let us assume  $\omega_0 = 1$ , so that  $T = 2\pi$ . Then, the solutions  $w_k$  to  $\sin(w) = \hat{w}$  are given by

$$w_k = (-1)^k \arcsin(\hat{w}) + \pi k, \quad (12)$$

for some  $k \in \mathbb{Z}$ . Considering that the latent weights follow a Laplace distribution, using  $w_k$  and Eq. (11), we obtain the pdf of the weights before binarization as

$$f_{\hat{\mathcal{W}}}(\hat{w}) = \sum_{k=-\infty}^{\infty} \frac{1}{2b} \frac{1}{\sqrt{1-\hat{w}^2}} \exp\left(\frac{-|(-1)^k \arcsin(\hat{w}) + \pi k|}{b}\right). \quad (13)$$

As it is shown in the supplementary material, the summation in (13) converges to the probability density function of the latent weights before binarization  $\hat{w}$  for an arbitrary frequency  $\omega_0$  given by

$$f_{\hat{\mathcal{W}}}(\hat{w}) = \frac{1}{2b\omega_0} \frac{1}{\sqrt{1-\hat{w}^2}} \exp\left(\frac{-|\arcsin(\hat{w})|}{b\omega_0}\right) + \frac{1}{2b\omega_0} \frac{1}{\sqrt{1-\hat{w}^2}} \cosh\left(\frac{\arcsin(\hat{w})}{b\omega_0}\right) \frac{1}{e^{\pi/b\omega_0} - 1}. \quad (14)$$

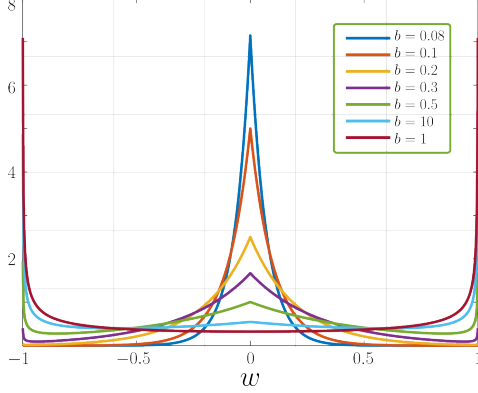


Figure 2. Probability density function of  $\hat{w} = \sin(\omega_0 w)$  assuming that the random variable  $w$  follows a Laplace distribution with parameter  $b$  and a fixed value of  $\omega_0 = 1$ .

Figure 2 depicts the distribution of the weights for different values of the Laplace distribution parameter  $b$  and a fixed frequency  $\omega_0 = 1$ . Note that different from the random variable  $\mathcal{W}$  which can take any real value, the codomain of the random variable  $\hat{\mathcal{W}}$  is  $[-1, 1]$ . From Fig. 2 we can observe that when the value of  $b$  increases the pdf of  $\hat{w}$  behaves as an arcsin distribution with values concentrated around -1 and 1. This reduces the quantization error in comparison to the Laplacian distribution. Also, a similar behavior occurs when the frequency value increases for a fixed  $b$ . To further analyze these observations consider the QE defined as

$$QE = \int_{-\infty}^{+\infty} f_{\mathcal{W}}(w) (\sin(\omega_0 w) - \gamma \text{sign}(\sin(\omega_0 w)))^2 dw, \quad (15)$$

where  $f_{\mathcal{W}}$  is the density distribution function of the latent weights. Using the fact that  $|x| = x \text{Sign}(x)$  along with the properties of the absolute value, we can rewrite Eq. (15) as

$$QE = \int_0^{+\infty} \frac{1}{b} \exp\left(\frac{-w}{b}\right) (|\sin(\omega_0 w)| - \gamma)^2 dw. \quad (16)$$

The solution to this integral is given by (see supplementary material for more details on this calculation)

$$QE = \frac{2(\omega_0 b)^2}{4(\omega_0 b)^2 + 1} - \frac{2\gamma\omega_0 b (e^{\pi/\omega_0 b} + 1)}{(\omega_0 b)^2 + 1 (e^{\pi/\omega_0 b} - 1)} + \gamma^2. \quad (17)$$

On the other hand, the optimal solution of the scaling factor  $\gamma$  in (15) can be computed as

$$\gamma = \mathbb{E}\{|\sin(\omega_0 w)|\} = \frac{\omega_0 b (e^{\pi/\omega_0 b} + 1)}{(\omega_0 b)^2 + 1 (e^{\pi/\omega_0 b} - 1)}. \quad (18)$$

Replacing  $\gamma$  from (18) into Eq. (17), we can rewrite the QE as a function of the frequency  $\omega_0$  and the parameter  $b$ . Figure 3 illustrates the QE as a function of the frequency  $\omega_0$  for different values of  $b$ . It can be seen that the maximum QE is 0.102835, which occurs when the product  $b\omega_0 \approx 0.954882$ . This means that for a fixed value of  $b$  the maximum QE occurs at  $\omega_0 \approx 0.954882/b$ . Furthermore, from this value, we note that by increasing the frequency to  $\infty$ , the QE converges to  $0.5 - \frac{4}{\pi^2}$ . Nonetheless, increasing  $\omega_0$  poses an additional issue since the gradient is proportional to the frequency value, thus it will diverge. This result is consistent with recent literature, which states that minimizing the QE yields to divergent gradients [43]. On the other hand, from Fig. 3, we can also see that the quantization error can be reduced by decreasing the frequency value. However, this corresponds to the trivial case when  $\gamma \text{Sign}(\sin(\omega_0 w))$  goes to zero as well as the weights. Clearly, this is not a practical approach either since there would not be information available to update the gradients.

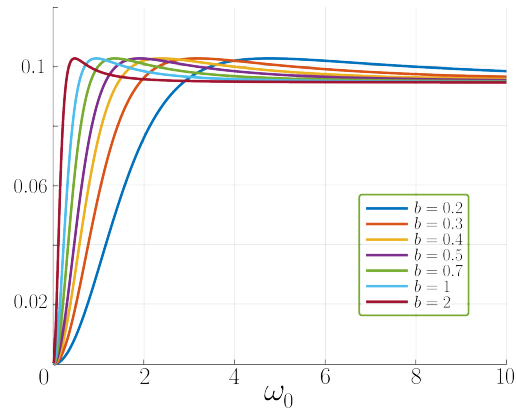


Figure 3. Quantization error as a function of the frequency  $\omega_0$  for different values of  $b$ . The proposed BiPer approach is able to control QE with the frequency of the periodic function.

It is worth noting that in contrast to current approaches that progressively reduce the QE to zero, BiPer does not meet this QE value. Nonetheless, further explorations can adapt state-of-the-art surrogate estimators to smoothly converge from the sine function to the square wave. These studies however, fall beyond the scope of this paper, which explores an alternative binarization technique.

## 5. Experiments

We comprehensively evaluated BiPer for image classification, with widely used neural network architectures, i.e., ResNet and VGG-Small, trained on benchmark datasets (CIFAR-10 and ImageNet). In the following, we first describe the experiments setup. Then, we present ablation experiments conducted on CIFAR-10. Finally, we compare

BiPer with state-of-the-art BNN approaches in terms of performance and complexity on both data sets.

## 5.1. Experiments Setup

### 5.1.1 Data Sets

**CIFAR-10 [23]:** Consists of 60,000 images of size  $32 \times 32$ , divided into 10 categories. Training set contains 50,000 images and the remaining 10,000 are used for testing.

**ImageNet [24]:** A challenging data set because of its larger size and more diverse image categories. Among its multiple versions, we adopted the widely used ILSVRC12 version, which is divided into 1,000 categories, from which 1.2 million are training images and, 50,000 test images. ImageNet is the most widely used data set to report results on binary networks and, it allows us to show for the first time that binary networks can perform competitively on a large-scale data set.

### 5.1.2 Network Architectures

On CIFAR-10, we evaluated BiPer using ResNet-18/20 and VGG-Small. For ResNet-18/20, we adopted the double skip connections as in [28], to provide fair comparisons. On ImageNet (ILSVRC-2012) we chose to binarize ResNet-18/34. Following [3], the downsampling layers are not quantized, and the double skip connections [28] were included. Following standard procedures of the comparison methods, we binarized all layers but the first and last.

### 5.1.3 Training Details and Procedures

BiPer approach was implemented in PyTorch because of its flexibility and powerful automatic differentiation mechanisms. Our models were executed on a NVIDIA RTX 3090 GPU. All experiments used stochastic gradient descent (SGD) optimization with 0.9 momentum. We followed the data augmentation strategies in [18], which include random crops and horizontal flips.

**Two stage training:** Recent works have shown that an appropriate initialization is often required to improve network performance. BNN initialization strategies employ an adaptation of high precision Neural Networks (NNs) [13, 18], or start from a pre-trained real-valued NN. To alleviate feature quantization adverse effects, two-stage training strategies are generally employed [5, 29]. Particularly, in the first stage the network is trained with real weights and binary features. Then, in the second stage, a warm weight initialization is employed based on the binary representation of the output weights from the first stage, and the model is fully trained to binarize the weights. Thus, the problem splits into two sub-problems: weight and feature binarization. In BiPer, we propose a two-stage training where, the

first stage uses real-valued weights  $\hat{w}$  as in Eq. (10), and the second stage uses the weight binarization from Eq.(8). If not otherwise specified, the weight-decay was fixed at  $5e - 4$  for the first stage and  $5e - 5$  for the second stage. The learning rate was set to 0.1 in the first stage and, 0.01 in the second stage. In both stages the learning rate was adjusted by the cosine scheduler.

## 5.2. Ablation Studies

To investigate the actual contributions of BiPer training stages, we conducted ablation studies on CIFAR-10. In all these experiments, ResNet-18/20 was used as the backbone.

### 5.2.1 Impact of Frequency $\omega_0$

BiPer introduces the frequency of the periodic function ( $\omega_0$ ) as a relevant hyperparameter. Section 4.1 showed that  $\omega_0$  can control QE, and particularly, the weight QE depends on the product of  $\omega_0$  and the parameter  $b$ . Thus, for a given value of  $b$ , increasing the frequency implies lower QE. To verify these theoretical results, we trained BiPer Stage 1 for 200 epochs, with a lr= 0.05, weight decay  $5e - 4$ , and varying  $\omega_0$ . Figure 4 illustrates the obtained stage 1 results for (a) Top-1 validation precision, (b) QE and (c) maximum likelihood estimated parameter  $b$ . These results validate our theoretical findings. However, consistently with the state of the art, reducing the QE degrades network performance. It is worth noting that during the first stage, there is no QE for the weights and, the values reported in Table 1 correspond to the QE obtained from weight binarization at the end of stage 1 training. This same value corresponds to the QE of the initialization for stage 2.

We further verified the effect of  $\omega_0$  on the full BiPer model. To this end, we trained BiPer Stage 2 for different learning rates, using the outputs of Stage 1 as warm start. The results are also presented in Figure 4. Interestingly, Fig. 4 (a) shows that the best Stage 1 model ( $\omega_0 = 10$ ) does not result in the best performance of the full binary model. Furthermore, ensuring the lowest QE in Stage 1, does not result in the best network performance either. For the frequency hyperparameter, we found that an intermediate frequency  $\omega_0 = 20$  balances the initial QE and precision of the full binary model. Moreover, in some cases the final BiPer model achieves higher precision than the initial model trained with real weights, e.g.  $\omega_0 = 20, 30$ . To the best of our knowledge, this is the first empirical analysis of the impact of the quantization error in the initialization.

A third experiment was conducted in order to verify the contribution of the periodic function in the binarization process. In particular, we are interested on evaluating whether BiPer improvements on classification precision are due to the use of the periodic function instead of the two-stage training methodology. Therefore, we trained BiPer follow-

ing the proposed two-stage training, but we just employ the sign function to binarize the weights and activations as in (1). Table 1 summarizes the results for this experiment. We can see that stage 1 attained higher precision and QE compared to stage 2. Specifically, precision degradation reached 0.9% and QE decreased approximately 0.2. This behavior is similar to that from the case when  $\omega_0 = 5$  in Fig. 4, where we obtained high precision and high QE on stage 1, but in the second stage, the performance drops. This occurs because lower frequencies promote small argument values for the sinusoidal function, and  $\sin(x) \approx x$  for small values of  $x$ .

Table 1. BiPer two-stage training without using a periodic function in the binarization.

Stage 1		Stage 2	
Precision	QE	Precision	QE
93.93	0.2351	93.01	0.0584

### 5.2.2 Binary Weights - Real Activations

This experiment considers only weight binarization using the proposed BiPer approach. Thus, we trained a resnet-20 for the CIFAR-10 data set, keeping real activation values. Quantitative results are reported in Table 2, where it can be noted that the proposed approach has competitive performance with the state-of-the-art approach LCR-BNN [38]. Furthermore, it is important to highlight that the degradation accuracy compared to a full precision model is less than 1%

Table 2. Performance comparison for binary weights and real activations in BiPer, compared to state-of-the-art methods on CIFAR-10. W/A: bit length of weights and activations. FP: Full precision model.

Network	Method	W/A	Top-1
ResNet-20	FP	32/32	91.7%
	DoReFa	1/32	90.0%
	LQ-Net	1/32	90.1%
	DSQ	1/32	90.2%
	IR-Net	1/32	90.8%
	<b>LCR-BNN</b>	1/32	<b>91.20%</b>
	<b>BiPer (Ours)</b>	1/32	<b>91.20%</b>

### 5.3. Comparison with SOTA methods

#### 5.3.1 Experiments on CIFAR-10

For ResNet-18, we compared the proposed BiPer approach with existing methods such as RAD[11], IR-Net [34], LCR-BNN [38], RBNN [27], ReCU [45], ReSTE [43], MST [41], and DIR-Net [35]. For VGG-Small, we compared BiPer with the XNOR-Net [36], BNN [21], DoReFa [49], IR-Net [34], RBNN [27], DSQ, SLB [46], ReCU [45] and ReSTE

Table 3. BiPer performance comparison with state-of-the-art BNN on CIFAR-10. W/A: bit length of weights and activations. FP: full precision model.

Network	Method	W/A	Top-1
ResNet-18	FP	32/32	94.8%
	RAD	1/1	90.5%
	IR-Net	1/1	91.5%
	LCR-BNN	1/1	91.80%
	RBNN	1/1	92.2%
	ReCU	1/1	92.8%
	ReSTE	1/1	92.63%
	MST	1/1	93.2%
	DIR-Net	1/1	92.8%
<b>BiPer (Ours)</b>	1/1	<b>93.75%</b>	
ResNet-20	FP	32/32	92.1%
	DoReFa	1/1	79.3%
	DSQ	1/1	84.1%
	SLB	1/1	85.5%
	IR-Net	1/1	86.5%
	ReCU	1/1	87.4%
	<b>BiPer (Ours)</b>	1/1	<b>87.5%</b>
VGG-small	FP	32/32	94.1%
	XNOR-Net	1/1	89.8%
	BNN	1/1	89.9%
	DoReFa	1/1	90.2%
	IR-Net	1/1	90.4%
	RBNN	1/1	91.3%
	DSQ	1/1	91.7%
	SLB	1/1	92.0%
	<b>ReCU</b>	1/1	<b>92.2%</b>
	<b>BiPer (Ours)</b>	1/1	<b>92.11%</b>

[43]. The classification precision results are shown in Table 3, and it can be seen that BiPer provides the best performance for ResNet-18 and ResNet-20, while it provides competitive results for the VGG-small network. Specifically, the proposed approach obtains a Top-1 accuracy of 93.75% when binarizing ResNet-18, improving classification precision in almost 1% against ReCU and 0.55% against the recent Minimum Spanning Tree (MST) strategy. Furthermore, the accuracy loss against the full precision model is reduced to just 1%. Recall that BiPer is capable of providing these results without requiring additional network modules or regularization functions.

#### 5.3.2 Experiments on ImageNet

We also evaluate the proposed BiPer approach using ResNet-18 and ResNet-34, and training on the large-scale ImageNet dataset. Table 4 shows a number of SOTA quantization methods over ResNet-18 and ResNet-34, including XNOR-Net [36], Bi-Real Net [28], PCNN [16], IR-Net [34], BONN [17], LCR-BNN [38], HWGQ [6], RBNN

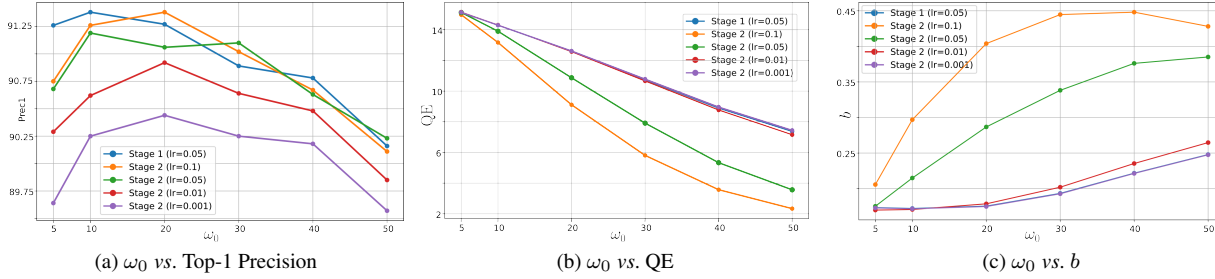


Figure 4. Impact of the frequency of the periodic function  $\omega_0$  on BiPer (a) Top-1 classification precision, (b) Quantization error and, (c) Weight distribution parameter  $b$ , for the CIFAR-10 data set.

Table 4. BiPer performance comparison with state-of-the-art BNN on ImageNet. W/A: bit length of weights and activations. FP: full precision model.

Network	Method	W/A	Top-1	Top-5
ResNet-18	FP	32/32	69.6%	89.2%
	XNOR-Net	1/1	51.2%	73.2%
	Bi-Real Net	1/1	56.4%	79.5%
	PCNN	1/1	57.3%	80.0%
	IR-Net	1/1	58.1%	80.0%
	BONN	1/1	59.3%	81.6%
	LCR-BNN	1/1	59.6%	81.6%
	HWGQ	1/1	59.6%	82.2%
	RBNN	1/1	59.9%	81.9%
	FDA	1/1	60.2%	82.3%
	ReSTE	1/1	60.88%	82.59%
	ReCU	1/1	61.0%	82.6%
	DIR-Net	1/1	60.4%	81.9%
	<b>BiPer (Ours)</b>	1/1	<b>61.4%</b>	<b>83.14%</b>
ResNet-34	FP	32/32	73.3%	91.3%
	Bi-Real Net	1/1	62.2%	83.9%
	IR-Net	1/1	62.9%	84.1%
	RBNN	1/1	63.1%	84.4%
	ReSTE	1/1	65.05%	85.78%
	ReCU	1/1	65.1%	85.8%
	DIR-Net	1/1	64.1%	85.3%
	<b>BiPer (Ours)</b>	1/1	<b>65.73%</b>	<b>86.39%</b>

[27], FDA [44], ReSTE [43], ReCU [45], and DIR-Net [35]. We can observe that the proposed BiPer approach in the 1W/1A setting achieves the best Top-1 and top-5 accuracy for both network architectures. Specifically, for ResNet-18, we attained a top-1 validation accuracy of 61.4%, outperforming the second-best result of 61.0% achieved by ReCu. Furthermore, our top-5 performance reached 83.14%, surpassing the second-best result of 82.6%, also achieved by ReCu. Likewise, for ResNet-34, we achieved the highest top-1 and top-5 accuracies, namely 65.73% and 86.39%, respectively. These results improve the second-best method (ReCu) by 0.63% and 0.59% in top-1 and Top-5 accuracies,

respectively. The extensive comparison results presented in Table 3 and Table 4 demonstrate the effectiveness of BiPer on classification tasks.

## 6. Conclusions

An approach for neural network binarization using a binary periodic function or square wave, dubbed BiPer, has been proposed. To improve gradient stability we employed a sinusoidal function with the same period of the square wave as a differentiable surrogate during the backward pass. This simple, yet powerful modification tackles the problem of standard gradient mismatch between forward and backward steps during network training, providing a suitable alternative that can be incorporated within back-propagation algorithms based on stochastic gradient descent. Mathematical analysis of BiPer quantization error demonstrated that it can be controlled by the frequency of the periodic function. This leverages a flexible initialization scheme for the binary weights, that balances the QE and network performance accuracy. The advantages of BiPer for the classification task were demonstrated through experimental results for two benchmark data sets, i.e., CIFAR-10 and ImageNet. Comparisons with respect to state-of-the-art BNN approaches showed that BiPer outperforms prior works by up to 1% on CIFAR-10, and up to 0.63% on Imagenet, respectively. Although this work tested the BiPer approach for classification, it can be easily extended to other high-level tasks without increasing the complexity. Thus, BiPer opens new horizons for neural network quantization by analyzing the frequency of the periodic function and its relation to the network quantization error. In this regard, we have shown that the QE can be controlled by the frequency. These advancements collectively contribute to the maturation of BNNs as viable alternatives for real-world applications, offering a compelling trade-off between model size, computational efficiency, and accuracy.

**Acknowledgment.** This work was supported by the Vicerrectoría de Investigación y Extensión of UIS, Colombia under the research project 3735.



## References

- [1] Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. *Advances in Neural Information Processing Systems*, 32, 2019. 4
- [2] Chaim Baskin, Natan Liss, Eli Schwartz, Evgenii Zheltonozhskii, Raja Giryes, Alex M Bronstein, and Avi Mendelson. Uniq: Uniform noise injection for non-uniform quantization of neural networks. *ACM Transactions on Computer Systems (TOCS)*, 37(1-4):1–15, 2021. 4
- [3] Joseph Bethge, Haojin Yang, and Christoph Meinel. Training accurate binary neural networks from scratch. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 899–903. IEEE, 2019. 6
- [4] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *arXiv preprint arXiv:1909.13863*, 2019. 3
- [5] Adrian Bulat, Georgios Tzimiropoulos, Jean Kossaiji, and Maja Pantic. Improved training of binary networks for human pose estimation and image recognition. *arXiv preprint arXiv:1904.05868*, 2019. 6
- [6] 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. 7
- [7] Eric R. Chan, Marco Monteiro, Petr Kellnhofer, Jiajun Wu, and Gordon Wetzstein. Pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5799–5809, 2021. 3
- [8] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2*, page 3123–3131, Cambridge, MA, USA, 2015. MIT Press. 2, 3
- [9] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(88), 2022. 3
- [10] Sajad Darabi, Mouloud Belbahri, Matthieu Courbariaux, and Vahid Partovi Nia. Bnn+: Improved binary network training. In *ICLR*, 2019. 2
- [11] Ruizhou Ding, Ting-Wu Chin, Zeye Liu, and Diana Marculescu. Regularizing activation distribution for training binarized deep networks. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11400–11409, 2019. 2, 7
- [12] Elad Eban, Yair Movshovitz-Attias, Hao Wu, Mark Sandler, Andrew Poon, Yerlan Idelbayev, and Miguel Á. Carreira-Perpiñán. Structured multi-hashing for model compression. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11900–11909, 2020. 1
- [13] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010. 6
- [14] R. Gong, X. Liu, S. Jiang, T. Li, P. Hu, J. Lin, F. Yu, and J. Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4851–4860, Los Alamitos, CA, USA, 2019. IEEE Computer Society. 2, 3
- [15] Jianping Gou, Baosheng Yu, Stephen J. Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129:1789–1819, 2021. 1
- [16] 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*, pages 8344–8351, 2019. 7
- [17] Jiaxin Gu, Junhe Zhao, Xiaolong Jiang, Baochang Zhang, Jianzhuang Liu, Guodong Guo, and Rongrong Ji. Bayesian optimized 1-bit cnns. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4909–4917, 2019. 7
- [18] 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. 6
- [19] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 1
- [20] 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. *CoRR*, abs/1704.04861, 2017. 1
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. In *NeurIPS*, 2016. 1, 7
- [22] Forrest N. Iandola, Matthew W. Moskewicz, Khalid Ashraf, Song Han, William J. Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 1mb model size. *ArXiv*, abs/1602.07360, 2016. 1
- [23] Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009. 6
- [24] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, page 1097–1105, Red Hook, NY, USA, 2012. Curran Associates Inc. 6
- [25] Huu Le, Rasmus Kjør Høier, Che-Tsung Lin, and Christopher Zach. Adaste: An adaptive straight-through estimator to train binary neural networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 460–469, Los Alamitos, CA, USA, 2022. IEEE Computer Society. 2

- [26] Mingbao Lin, Rongrong Ji, Zihan Xu, Baochang Zhang, Yan Wang, Yongjian Wu, Feiyue Huang, and Chia-Wen Lin. Rotated binary neural network. In *Advances in Neural Information Processing Systems*, pages 7474–7485. Curran Associates, Inc., 2020. 2
- [27] Mingbao Lin, Rongrong Ji, Zihan Xu, Baochang Zhang, Yan Wang, Yongjian Wu, Feiyue Huang, and Chia-Wen Lin. Rotated binary neural network. *Advances in neural information processing systems*, 33:7474–7485, 2020. 7, 8
- [28] Zechun Liu, Wenhan Luo, Baoyuan Wu, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Binarizing deep network towards real-network performance. *International Journal of Computer Vision*, 128:202–219, 2020. 3, 6, 7
- [29] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-to-binary convolutions. In *International Conference on Learning Representations*, 2019. 6
- [30] Roy Or-El, Xuan Luo, Mengyi Shan, Eli Shechtman, Jeong Joon Park, and Ira Kemelmacher-Shlizerman. Stylesdf: High-resolution 3d-consistent image and geometry generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13503–13513, 2022. 3
- [31] Anh-Huy Phan, Konstantin Sobolev, Konstantin Sozykin, Dmitry Ermilov, Julia Gusak, Petr Tichavský, Valeriy Glukhov, Ivan Oseledets, and Andrzej Cichocki. Stable low-rank tensor decomposition for compression of convolutional neural network. In *Computer Vision – ECCV 2020*, pages 522–539, Cham, 2020. Springer International Publishing. 1
- [32] Hossein Pishro-Nik. *Introduction to probability, statistics, and random processes*. Kappa Research, LLC Blue Bell, PA, USA, 2014. 4
- [33] Haotong Qin, Ruihao Gong, Xianglong Liu, Xiao Bai, Jingkuan Song, and Nicu Sebe. Binary neural networks: A survey. *Pattern Recognition*, 105:107281, 2020. 1
- [34] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. Forward and backward information retention for accurate binary neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2, 3, 7
- [35] Haotong Qin, Xiangguo Zhang, Ruihao Gong, Yifu Ding, Yi Xu, and Xianglong Liu. Distribution-sensitive information retention for accurate binary neural network. *International Journal of Computer Vision*, 131(1):26–47, 2023. 7, 8
- [36] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *Computer Vision – ECCV 2016*, pages 525–542. Springer International Publishing, 2016. 1, 3, 7
- [37] Ratshih Sayed, Haytham Azmi, Heba Shawkey, A. H. Khalil, and Mohamed Refky. A systematic literature review on binary neural networks. *IEEE Access*, 11:27546–27578, 2023. 1
- [38] Yuzhang Shang, Dan Xu, Bin Duan, Ziliang Zong, Liqiang Nie, and Yan Yan. Lipschitz continuity retained binary neural network. In *European conference on computer vision*, pages 603–619. Springer, 2022. 7
- [39] Vincent Sitzmann, Julien N.P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. In *Proc. NeurIPS*, 2020. 3
- [40] Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 1
- [41] Quang Hieu Vo, Linh-Tam Tran, Sung-Ho Bae, Lok-Won Kim, and Choong Seon Hong. Mst-compression: Compressing and accelerating binary neural networks with minimum spanning tree. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023. 2, 7
- [42] C. Wang, A. Bochkovskiy, and H. Liao. Yolov7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7464–7475, Los Alamitos, CA, USA, 2023. IEEE Computer Society. 1
- [43] Xiao-Ming Wu, Dian Zheng, Zuhao Liu, and Wei-Shi Zheng. Estimator meets equilibrium perspective: A rectified straight through estimator for binary neural networks training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 17055–17064, 2023. 2, 3, 5, 7, 8
- [44] Yixing Xu, Kai Han, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Learning frequency domain approximation for binary neural networks. *Advances in Neural Information Processing Systems*, 34:25553–25565, 2021. 8
- [45] Zihan Xu, Mingbao Lin, Jianzhuang Liu, Jie Chen, Ling Shao, Yue Gao, Yonghong Tian, and Rongrong Ji. Recu: Reviving the dead weights in binary neural networks. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5178–5188, 2021. 1, 2, 7, 8
- [46] Zhaohui Yang, Yunhe Wang, Kai Han, Chunjing Xu, Chao Xu, Dacheng Tao, and Chang Xu. Searching for low-bit weights in quantized neural networks. *Advances in neural information processing systems*, 33:4091–4102, 2020. 7
- [47] Chunyu Yuan and Sos S. Agaian. A comprehensive review of binary neural network. *Artificial Intelligence Review*, 56: 12949–13013, 2023. 1
- [48] Zhaoyang Zhang, Wenqi Shao, Jinwei Gu, Xiaogang Wang, and Ping Luo. Differentiable dynamic quantization with mixed precision and adaptive resolution. In *International Conference on Machine Learning*, pages 12546–12556. PMLR, 2021. 4
- [49] 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. 7