

Estimator Meets Equilibrium Perspective: A Rectified Straight Through Estimator for Binary Neural Networks Training

Xiao-Ming Wu¹, Dian Zheng¹, Zuhao Liu¹, Wei-Shi Zheng^{1,2,3,4*}

¹School of Computer Science and Engineering, Sun Yat-sen University, China, ²Pengcheng Lab, China,

³Guangdong Province Key Laboratory of Information Security Technology, China,

⁴Key Laboratory of Machine Intelligence and Advanced Computing, Ministry of Education, China

{wuxm65, zhengd35, liuzh327}@mail2.sysu.edu.cn, wszheng@ieee.org

Abstract

*Binarization of neural networks is a dominant paradigm in neural networks compression. The pioneering work BinaryConnect uses Straight Through Estimator (STE) to mimic the gradients of the sign function, but it also causes the crucial inconsistency problem. Most of the previous methods design different estimators instead of STE to mitigate it. However, they ignore the fact that when reducing the estimating error, the gradient stability will decrease concomitantly. These highly divergent gradients will harm the model training and increase the risk of gradient vanishing and gradient exploding. To fully take the gradient stability into consideration, we present a new perspective to the BNNs training, regarding it as the equilibrium between the estimating error and the gradient stability. In this view, we firstly design two indicators to quantitatively demonstrate the equilibrium phenomenon. In addition, in order to balance the estimating error and the gradient stability well, we revise the original straight through estimator and propose a power function based estimator, **Rectified Straight Through Estimator (ReSTE for short)**. Comparing to other estimators, ReSTE is rational and capable of flexibly balancing the estimating error with the gradient stability. Extensive experiments on CIFAR-10 and ImageNet datasets show that ReSTE has excellent performance and surpasses the state-of-the-art methods without any auxiliary modules or losses.*

1. Introduction

Deep neural networks have revolutionary development in recent years since its admirable ability to learn discriminate features [31, 20, 15, 38, 45]. But they tend to require massive computational cost and memory cost, which is unsuitable to deploy at some resource-limited devices. To this

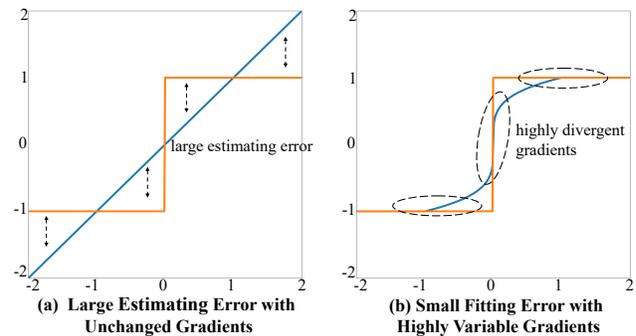


Figure 1: The intuitive illustrations of the equilibrium perspective of BNNs training, i.e., the equilibrium between the estimating error and the gradient stability. When reducing the estimating error, the gradients will become highly divergent, which harms the model training and increases the risk of gradient vanishing and gradient exploding. Blue, orange lines represent the estimators and sign function respectively.

end, many network compression methods have been proposed [25, 19, 23], such as pruning [39, 59, 58, 13], tiny model design [55, 25, 49, 29], distillation [40, 50, 56] and tensor decomposition[43]. Among them, network quantization [8, 44, 32, 21, 1] is a kind of excellent method with high compression ratio and little performance degradation. Binary Neural Networks (BNNs) [8, 9, 27], an extreme case of network quantization which aims to quantize 32-bit inputs into 1-bit, attract great research enthusiasm in recent years due to its extremely high compression ratio and great performance in neural networks compression.

In BNNs research, the pioneering work BinaryConnect [8] proposes to apply sign function to binary the full-precision inputs in forward process, and use the straight through estimator (STE) to mimic the gradients of the sign function when backpropagation, which achieves great performance. However, the difference between the forward

* denotes the corresponding author.

and the backward processes causes the crucial inconsistency problem in BNNs training. To reduce the degree of inconsistency, many previous works design different estimators instead of STE, attempting to narrow the estimating error. Nevertheless, they neglect the fact that when reducing the estimating error, the gradient stability will decrease concomitantly. This will make the gradients highly divergent, harming the model training and increasing the risk of gradient vanishing and gradient exploding.

To fully take the gradient stability into consideration, we present a new perspective to the BNNs training, regarding it as the equilibrium between the estimating error and the gradient stability, as shown in Fig. 1. In this view, we firstly design two indicators to measure the degree of the equilibrium between estimating error and the gradient instability. With these indicators, we can quantitatively demonstrate the equilibrium phenomenon. In addition, to balance the estimating error with the gradient stability well, we revise the original straight through estimator (STE) and propose a power function based estimator, **Rectified Straight Through Estimator**, **ReSTE** for short. The insight is from the fact that STE is a special case of the power function. With this design, ReSTE is always rational, i.e., having less estimating error than STE, and capable of flexibly balancing the estimating error and the gradient stability, which are the two main advantages of ReSTE comparing to other estimators.

Sufficient experiments on CIFAR-10 [30] and large-scale ImageNet ILSVRC-2012 [11] datasets show that our method has good performance and surpasses the state-of-the-art methods without any auxiliaries, e.g., additional modules or losses. Moreover, by two carefully-designed indicators, we demonstrate the equilibrium phenomenon and further show that ReSTE can flexibly balance the estimating error and the gradient stability. Our source code is available at <https://github.com/DravenALG/ReSTE>, hoping to help the development of the BNNs community.

2. Revisiting Binary Neural Networks

Binary Neural Networks (BNNs) [8, 9, 44, 6] aim to binarize full-precision inputs, weights or features (also called activations in BNNs literature) in each layers into 1-bits, which is an extreme case of network quantization. Essentially, the optimization of BNNs is a constraint optimization problem. Naively using brute-force algorithms to solve this problem is intractable due to the huge combinatorial probabilities when the dimensions of input are large.

The exploration of tractable solutions to binary neural networks training can be traced back to many pioneering works [28, 7, 48]. Among them, BinaryConnect [8] forms the main optimization paradigm in this domain due to its great performance. BinaryConnect connects a sign function between the full-precision inputs and the following calculation modules in forward process. Since the gradients of the

sign function are zero almost everywhere, BinaryConnect uses an identity function to substitute for the sign function when calculating the gradients in backward process, which is also known as straight through estimator (STE) [22, 4]. For convenience, we respectively denote \mathbf{z} and \mathbf{z}_b as the full-precision inputs and the binarized outputs. The forward and backward processes of the binary procedure in BinaryConnect are as follows:

$$\begin{aligned} \text{Forward: } \mathbf{z}_b &= \text{sign}(\mathbf{z}), \\ \text{Backward: } \frac{\partial \mathcal{L}}{\partial \mathbf{z}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{z}_b}, \end{aligned} \quad (1)$$

where \mathcal{L} is the loss function and **sign** represents the element-wise sign function. It means that the gradients with respect to the full-precision inputs straightly equals to the gradients of the binarized outputs, which is also the origin of the name straight through estimator.

To improve the performance of binary neural networks, many different improvement strategies have been proposed. Some works try to modify the model architectures of the backbone, which heightens the expressive ability of the binary neural networks [37, 36]. In spite of the performance improvement, these works revise the architectures of the backbone, which is not universal to all architectures and adds additional computational and memory cost in inference. In addition, some other works focus on improving the forward process with some additional assistance, e.g., modules [54, 33, 53, 51], losses [3, 17, 18, 35, 46, 52] and even distillation [50]. This type of methods significantly increase parameters and the computational cost when training.

Besides, many works mainly focus on the essential and vital component of binary neural networks training, i.e., the estimator to mimic the gradients of the sign function. BNN+ [10] designs a SignSwish function, Bi-Real-Net [37] models a piece-wise polynomial function, DSQ [16] proposes a tanh-based function, IR-Net [42] gives the EDE function and FDA [53] applies Fourier series to simulate the gradients. Although they have achieved excellent performance, they ignore the fact that when reducing the estimating error, the gradient stability will decrease concomitantly, which means that the gradient will become highly divergent, harming the model training and increasing the risk of gradient vanishing and gradient exploding. To fully consider the gradient stability in BNNs training, we present a new perspective, viewing it as the equilibrium between the estimating error and the gradient stability. From this perspective, we revise the original STE and propose a power function based estimator, Rectified Straight Through Estimator (ReSTE for short). Comparing to the estimators above, ReSTE is rational, i.e., having less or equal estimating error than STE and capable of flexibly balancing the estimating error and gradient stability. Sufficient experiments show that our method surpasses the state-of-the-art methods

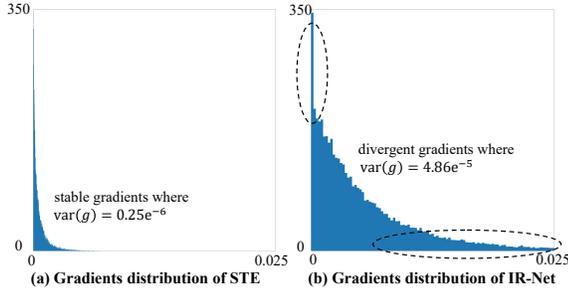


Figure 2: Illustrations of the gradient distributions of STE (left) and IR-Net (right). X-axes represent the values of the gradients, y-axes are the frequency.

without any auxiliaries, e.g., additional modules or losses.

3. Estimator Meets Equilibrium Perspective

3.1. Equilibrium Perspective

The **inconsistency problem** is inevitable but crucial in BNNs training since we use estimators to mimic the gradients of sign function in backpropagation. To mitigate the degree of the inconsistency, lots of follow-up works design different estimators instead of STE, aiming to reduce the estimating error. Although they improve the performance of BNNs, they only care about reducing the estimating error and ignore the concomitant gradient instability. The gradients will become highly divergent, which increases the risk of gradients vanishing and gradients exploding, as shown in Fig. 1. For more persuasive, we visualize the gradient distributions of STE [44] and the influential work IR-Net [42] in Fig. 2. Although IR-Net attempts to reduce the estimating error by approximating the sign function as it claims, it suffers from the problem of highly divergent gradients, which will harm the model training.

To fully take the gradients stability into consideration, we present a new perspective, considering BNNs training as the equilibrium between the estimating error and the gradient stability. For clear description, we first give the definition of the estimating error and the gradient stability. We define that the **estimating error** is the difference between the sign function and the estimator, which reflects how close between the estimator and sign function. We define the **gradient stability** as the divergence of the gradients of all parameters in each iteration. The insight is that when we use estimator to close to sign function, the gradients of all parameters in one iteration are inevitably divergent, which is intuitively shown in Fig. 1. This may lead to a wrong updating direction and harm the model training.

With the definitions, we now formally discuss our **equilibrium perspective**. Since the BNNs training is the equilibrium between the estimating error and the gradient stability, we should not reduce estimating error without limits.

Instead, we should design an estimator which can easily adjust the degree of equilibrium to obtain better performance.

3.2. Indicators of Estimating Error and Gradient Instability

To quantitatively and clearly demonstrate the equilibrium phenomenon, we firstly design two indicators to quantitatively analyze the degree of the estimating error and the gradient instability.

Since the estimating error is the difference between the sign function and the estimator, we stipulate that the estimating error can be evaluated by the distance between the results through the element-wise sign function and the results through the estimator in each iteration. We define $f(\cdot)$ as the estimator and D as the distance metric. The degree of estimating error can be formally described as:

$$e = D(\text{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z})), \quad (2)$$

where $D(\cdot)$ is L2-norm in our method. We call e as the **estimating error indicator**.

In addition, to measure the degree of the gradient stability, we design a **gradient instability indicator**. Since the gradient stability is the divergence of the gradients of all parameters in each iteration, we use the variance of gradients of all the parameters in each iteration to evaluate it. We design the indicator as follows:

$$s = \text{var}(|\mathbf{g}|), \quad (3)$$

where \mathbf{g} donates the gradients, $|\cdot|$ is the element-wise absolute operation and $\text{var}(\cdot)$ stands for the variance operator. Here we use absolute operation since we only care about the gradients magnitude (the updating directions are not relevant to the gradient stability). Note that s is the gradient instability indicator that the magnitude of s reflects the degree of the instability.

3.3. Rectified Straight Through Estimator

To balance the estimating error and the gradient stability, we should design an estimator that can easily adjust the degree of equilibrium well. Before that, we firstly claim that sign function and STE are two extremes in gradient stability. The sign function has zero gradients almost everywhere and has infinite gradients at the origin of the coordinate, whose gradients are completely vanishing or exploding. Therefore sign function has the highest gradient instability. In contract, STE uses linear function to estimate the gradients of sign function, which not at all changes the gradients backward in the estimating process. So STE is with the lowest instability. We want to design an estimator close to sign function to get less estimating error, but not too much unstable to train. Based on this intuition, we design two properties that an estimator should

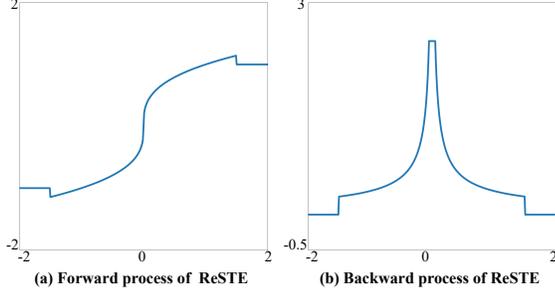


Figure 3: Illustrations of the forward (left) and backward (right) processes of ReSTE.

satisfy: **1) Rational property:** It should always have less or equal estimating error than straight through estimator (the identity function), which can be formally described as $D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z})) - D(\mathbf{sign}(\mathbf{z}), \mathbf{z}) \leq 0$. The rational property is rational since the fact that if an estimator has more estimating error than STE in some ranges, directly applying STE to mimic the gradients in these ranges is more reasonable, which is more stable and has less estimating error. **2) Flexible property:** It should be capable of flexibly adjusting the degree of the estimating error and the gradient stability to adjust the degree of equilibrium. The flexible property consists of two aspects. First, the estimator can change from STE to sign function. Second, the changing should be gradually. Gradually changing means that each point can move a small step closer to sign function when we adjust the estimator, which is the key to find a suitable degree of the equilibrium.

To achieve these goals, we revise the STE and propose a power function based estimator, **Rectified Straight Through Estimator, ReSTE** for short. The inspiration of ReSTE is that the STE strategy (identity function) is a special case of the power function, when the power is one for specific. When the power function is close to STE, the gradient is stable, but the estimating error is large. When the power increases, the power function will close to sign function and have less estimating error, while increasing the instability of the gradients. In a word, power function can easily change from STE to sign function, demonstrating its ability of adjusting the degree of equilibrium.

Under such observation, we propose to use power function as the estimator in backward process to balance the estimating error and the gradient stability. Our ReSTE function has the following form:

$$\mathbf{f}(\mathbf{z}) = \mathbf{sign}(\mathbf{z})|\mathbf{z}|^{\frac{1}{o}}, \quad (4)$$

s.t. $o \geq 1$,

where o are hyper-parameters controlling the power, which is also the degree of the equilibrium. In detail, o decides the ratified degree of ReSTE to balance the estimating error

and gradient stability. Note that when $o = 1$, the ReSTE function is the basic STE. With o increasing, the ReSTE function closes to sign function, which has less estimating error gradually. With simple derivation, the gradients of the ReSTE function is:

$$\mathbf{f}'(\mathbf{z}) = \frac{1}{o} |\mathbf{z}|^{\frac{1-o}{o}}. \quad (5)$$

Comparing to other estimators, ReSTE satisfies the properties proposed above, i.e., rational and capable of flexibly balancing the estimating error and the gradient stability, which are the two main advantages of our method. To prove that, we firstly give the following lemma.

Lemma 3.1. *If $o_1 \geq o_2$, $D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z}, o_1)) - D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z}, o_2)) \leq 0$ holds. The proof is as follows:*

$$\begin{aligned} D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z}, o_1)) &= \sum_{i=1}^d (\mathbf{sign}(z_i) - \mathbf{f}(z_i, o_1))^2 \\ &= \sum_{i=1}^d (\mathbf{sign}(z_i) - \mathbf{sign}(z_i)|z_i|^{\frac{1}{o_1}})^2 \\ &= \sum_{i=1}^d |1 - |z_i|^{\frac{1}{o_1}}|^2, \end{aligned} \quad (6)$$

where $|\cdot|$ is the absolute operation. Since $o_1 \geq o_2$, with the nature of the power function, we can achieve that when $|z_i| \leq 1$, $|1 - |z_i|^{\frac{1}{o_1}}| = 1 - |z_i|^{\frac{1}{o_1}} \leq 1 - |z_i|^{\frac{1}{o_2}} = |1 - |z_i|^{\frac{1}{o_2}}|$, and when $|z_i| \geq 1$, $|1 - |z_i|^{\frac{1}{o_1}}| = |z_i|^{\frac{1}{o_1}} - 1 \leq |z_i|^{\frac{1}{o_2}} - 1 = |1 - |z_i|^{\frac{1}{o_2}}|$. Thus, $|1 - |z_i|^{\frac{1}{o_1}}| \leq |1 - |z_i|^{\frac{1}{o_2}}|$ always holds. Then, we can write:

$$\begin{aligned} D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z})) &= \sum_{i=1}^d |1 - |z_i|^{\frac{1}{o_1}}|^2 \\ &\leq \sum_{i=1}^d |1 - |z_i|^{\frac{1}{o_2}}|^2 \\ &= D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z}, o_2)). \end{aligned} \quad (7)$$

Under the lemma, we give the proof of the two properties. As for the rational property, since STE equals to $\mathbf{f}(\mathbf{z}, 1)$ and ReSTE has the condition $o \geq 1$, we can easily get that $D(\mathbf{sign}(\mathbf{z}), \mathbf{f}(\mathbf{z})) - D(\mathbf{sign}(\mathbf{z}), \mathbf{z}) \leq 0$ always holds by lemma 3.1. About the flexible property, we know that STE equals to $\mathbf{f}(\mathbf{z}, 1)$ and when $o \rightarrow \infty$, $\mathbf{f}(\mathbf{z}) \rightarrow \mathbf{sign}(\mathbf{z})$, thus ReSTE can change from STE to sign function. Moreover, from the proof of lemma 3.1 we can observe that if $o_1 \geq o_2$, $|1 - |z_i|^{\frac{1}{o_1}}| \leq |1 - |z_i|^{\frac{1}{o_2}}|$ always holds for any z_i , thus $|\mathbf{sign}(z_i) - \mathbf{f}(z_i, o_1)| \leq |\mathbf{sign}(z_i) - \mathbf{f}(z_i, o_2)|$ always holds for any z_i . So the changing of ReSTE is gradually, where

Backbone	Method	W/A	Auxiliary	Acc(%)
ResNet-18	FP	32/32	-	94.84
	RAD [12]	1/1	Loss	90.50
	IR-Net [42]	1/1	Module	91.50
	LCR-BNN [46]	1/1	Loss	91.80
	RBNN [33]	1/1	Module	92.20
	ReSTE (ours)	1/1	-	92.63
ResNet-20	FP	32/32	-	91.70
	DSQ [16]	1/1	-	84.11
	DoReFa-Net [57]	1/1	-	84.44
	IR-Net [42]	1/1	Module	85.40
	LCR-BNN [46]	1/1	Loss	86.00
	FDA	1/1	Module	86.20
	RBNN [33]	1/1	Module	86.50
	ReSTE (ours)	1/1	-	86.75
	IR-Net * [42]	1/1	Module	86.50
	LCR-BNN * [46]	1/1	Loss	87.20
	RBNN * [33]	1/1	Module	87.50
ReSTE * (ours)	1/1	-	87.92	
ResNet-34	FP	32/32	-	91.70
	DoReFa-Net [57]	1/32	-	90.00
	LQ-Net [54]	1/32	-	90.10
	DSQ [16]	1/32	-	90.20
	IR-Net [42]	1/32	Module	90.80
	LCR-BNN [46]	1/32	Loss	91.20
	ReSTE (ours)	1/32	-	91.32
	VGG-small	FP	32/32	-
LBA [24]		1/1	-	87.70
Xnor-Net [44]		1/1	-	89.80
BNN [9]		1/1	-	89.90
RAD [12]		1/1	Loss	90.00
IR-Net [42]		1/1	Module	90.40
RBNN [33]		1/1	Module	91.30
ReSTE (ours)		1/1	-	92.55

Table 1: Performance comparison with SOTA methods in CIFAR-10 dataset. Auxiliary refers to whether some additional assistance is used (module or loss). FP is the full-precision version of the backbone. * donates the method with Bi-Real structure. W/A is the bit width of weights or activations. Best results are shown in black bold font.

any z_i moves a small step closer to sign function when increasing o . Therefore, ReSTE satisfies the flexible property. The rational and flexible properties are designed based on the equilibrium perspective and form the main advantages between ReSTE and other estimators in previous methods.

In addition, for more stable gradients, we use some gradients truncation tricks to our estimator. First, we clip the gradients where the corresponding full-precision inputs with the absolute value larger than a threshold t to zero, which considers the saturation in BNNs training [9]. Next, since the gradients of ReSTE may be large when the input is sufficiently small, we set a threshold m and the gradients within the threshold $(0, m)$, $(-m, 0)$ use numerical method $(f(m) - f(0))/m$, $(f(0) - f(-m))/m$ to simulate.

For clear illustration, we demonstrate the forward and backward processes of ReSTE in Fig. 3.

3.4. Overall Binary Method

We summarize the overall Binary procedure of our method. As for the forward process of binarization, we

Backbone	Method	W/A	Auxiliary	Top-1(%)	Top-5(%)	
ResNet-18	FP	32/32	-	69.60	89.20	
	ABC-Net [34]	1/1	-	42.70	67.60	
	Xnor-Net [44]	1/1	-	51.20	73.20	
	BNN+ [5]	1/1	Loss	53.00	72.60	
	DoReFa-Net [57]	1/2	-	53.40	-	
	Bi-Real [37]	1/1	-	56.40	79.50	
	Xnor-Net++ [5]	1/1	-	57.10	79.90	
	IR-Net [42]	1/1	Module	58.10	80.00	
	LCR-BNN [46]	1/1	Loss	59.60	81.60	
	RBNN [33]	1/1	Module	59.90	81.90	
	FDA [53]	1/1	Module	60.20	82.30	
	ReSTE (ours)	1/1	-	60.88	82.59	
	ResNet-34	FP	32/32	-	69.60	89.20
		SQ-BWN [14]	1/32	-	58.40	81.60
BWN [44]		1/32	-	60.80	83.00	
HWGQ [6]		1/32	-	61.30	83.20	
TWN [2]		2/32	-	61.80	84.20	
SQ-TWN [14]		2/32	-	63.80	85.70	
BWHN [26]		1/32	-	64.30	85.90	
IR-Net [42]		1/32	Module	66.50	86.60	
LCR-BNN [46]		1/32	Loss	66.90	86.40	
ReSTE (ours)		1/32	-	67.40	87.20	
ResNet-50		FP	32/32	-	73.30	91.30
		ABC-Net [34]	1/1	-	52.40	76.50
		Bi-Real [37]	1/1	-	62.20	83.90
		IR-Net [42]	1/1	Module	62.90	84.10
	RBNN [33]	1/1	Module	63.10	84.40	
	LCR-BNN [46]	1/1	Loss	63.50	84.60	
	ReSTE(ours)	1/1	-	65.05	85.78	
	ResNet-101	FP	32/32	-	73.30	91.30
IR-Net [42]		1/32	Module	70.40	89.50	
ReSTE(ours)		1/32	-	70.74	89.50	

Table 2: Performance comparison with SOTA methods in ImageNet dataset. Auxiliary refers to whether some additional assistance is used (module or loss). FP is the full-precision version of the backbone. W/A is the bit width of weights or activations. Best results are in black bold font.

employ DoReFa-Net [57] as most of the previous methods do[42, 33, 53, 46], which uses sign function to binarize the inputs and endows a layer-level scalar $\beta = \|\mathbf{z}\|_{l1}/n$ (n is the dimensions of \mathbf{z}) for binarization to enhance the representative ability. In backpropagation, we apply ReSTE as the estimator to simulate the gradients of the sign function. About the hyper-parameter o to adjust the degree of equilibrium, we use the progressive adjusting strategy, which is proposed in [42] and widely used in recent works[33, 53]. We change o from 1 to o_{end} when training, which we use $o_{\text{end}} = 3$ in our experiments. Comparing to the fixed strategy, the progressive adjusting strategy ensures sufficient updating at the beginning and accurate gradients at the end of the training. Experiments about the design for the tuning strategies of o are shown in supplementary materials.

In BNNs literature, there have two types of options to binarize a neural network. The first type is that only the weights are binarized and the second type is weights and activations are both to be binarized, which significantly improves the inference speed via XNOR and Bitcount operations [9, 42]. After binarization, the model size decreases 32x comparing to the original full-precision model and the inference process is accelerated.

Estimators	Formula	Type	Rational	Flexible	Acc(%)
DSQ [16]	$\mathbf{f}(\mathbf{z}) = l + \Delta (\mathbf{i} + (\mathbf{s}\mathbf{t}\mathbf{a}\mathbf{n}\mathbf{h}(k(\mathbf{z} - \mathbf{m})) + 1)/2)$	Tanh-alike	Not rational	Little flexible	84.11
STE [57]	$\mathbf{f}(\mathbf{z}) = \mathbf{z}$	Identity function	Rational	Not flexible	84.44
EDE [42]	$\mathbf{f}(\mathbf{z}) = k\mathbf{t}\mathbf{a}\mathbf{n}\mathbf{h}(t\mathbf{z})$	Tanh-alike	Not rational	Little flexible	85.20
FDA † [53]	$\mathbf{f}(\mathbf{z}) = \frac{4}{\pi} \sum_{i=0}^k \mathbf{s}\mathbf{i}\mathbf{n}((2i+1)\omega\mathbf{z})/(2i+1)$	Fourier series	Not rational	Little flexible	85.80
RBNN † [33]	$\mathbf{f}(\mathbf{z}) = k \cdot \left(-\mathbf{s}\mathbf{i}\mathbf{g}\mathbf{n}(\mathbf{z}) \frac{t^2\mathbf{z}^2}{2} + \sqrt{2}t\mathbf{z} \right)$	Polynomial function	Not rational	Little flexible	85.87
ReSTE (ours)	$\mathbf{f}(\mathbf{z}) = \mathbf{s}\mathbf{i}\mathbf{g}\mathbf{n}(\mathbf{z}) \mathbf{z} ^{\frac{1}{\alpha}}$	Power function	Rational	Flexible	86.75

Table 3: Results of the estimators comparison. † means we only use the estimators for fair comparison (without some additional modules, the overall comparison can be found in Sec. 4.2). "Rational" means that the estimator satisfies the rational property proposed in Sec. 3.3 while "Not rational" indicates dissatisfaction. "Flexible" means that the estimator satisfies the flexible property proposed in Sec. 3.3 while "Not flexible" and "Little flexible" means dissatisfaction. "Not flexible" implies that the estimator can not reduce the estimator. "Little flexible" indicates that the estimator can reduce the estimating error in some kind but not fully satisfy the flexible property. The best results are shown in black bold font.

4. Experiments

4.1. Datasets and Settings

Datasets. In this work, we choose two datasets, i.e. CIFAR-10 [30] and ImageNet ILSVRC-2012 [11], which are widely-used in binary neural networks literature [42, 33, 53]. CIFAR-10 is a common datasets for image classification, which contains 50k training images and 10k testing images with 10 different categories. Each image is of size 32x32 with RGB color channels. ImageNet ILSVRC-2012 is a large-scale dataset with over 120k training images and 50k verification images. Each image contains 224x224 resolutions with RGB color channels. It has 1000 different categories.

Implementation Details. we follow the same setting as other binary methods [42, 33] used for fair comparison. For specific, we apply RandomCrop, RandomHorizontalFlip and Normalize for both CIFAR-10 and ImageNet pre-processing. We use SGD and set learning rate beginning from 0.1. Cosine learning rate descent schedule is adopted when training. What's more, we only use cross entropy as the loss function for classification. As for the hyper-parameter α_{end} , we set $\alpha_{\text{end}} = 3$ in all the experiments. We find that this value is suitable and robust to balance the estimating error and the gradient stability. Regarding the hyper-parameters t and i for gradient truncation, we simply set $t = 1.5$ and $i = 0.1$. All the models are implemented with PyTorch [41] on NVIDIA RTX3090 GPUs or NVIDIA RTX A6000 GPUs. For more details about the experiments parameters, please refer to our published codes and the README file in GitHub.

4.2. Performance Study

To prove the performance of our method, we conduct performance study in comparison with other binary methods. Note that our method only modify the estimators in

backward process without other auxiliaries, e.g., additional modules or losses. To highlight the superiority of our approach, we add a column to note the auxiliaries used in other methods in the result tables.

We first test the performance of ReSTE on CIFAR-10 [30] with the SOTA methods. In detail, we binarize three backbone models, ResNet-18, ResNet-20 [20] and VGG-small [47]. We compare a list of SOTA methods to validate our performance, including LBA [24], RAD [12], DSQ [16], Xnor-Net [44], DoReFa-Net [57], LQ-Net [54], IR-Net [42], LCR-BNN [46], RBNN [33], FDA [53]. For ResNet-20, we both evaluate the performance of our method in the basic ResNet architecture and the Bi-Real architecture [37]. Experiments results are exhibited in Table 1. From the table we can find that our ReSTE shows excellent performance, outperforming all the SOTA methods both at the setting of 1W/1A and 1W/32A without any assistance, e.g., modules or losses. For example, with ResNet-20 as the backbone, ReSTE respectively obtains 0.25% and 0.45% enhancement over the SOTA method RBNN [33] in the basic ResNet architecture and in the Bi-Real architecture [37], at the setting of 1W/1A, even that RBNN additionally adds a rotation module into the training. As for the setting of 1W/32A, ReSTE has 0.12% improvement over the SOTA method LQ-Net [54], which additional uses a Lipschitz loss to improve the training.

Moreover, we employ ReSTE on ResNet-18, ResNet-34 [20] and validate the performance on large-scale ImageNet ILSVRC-2012 [11]. In this setting, we compare ReSTE with ABC-Net [34], BWN [44], TWN [2], SQ-BWN and SQ-TWN [14], Xnor-Net [44], HWGO [6], BWHN [26], BNN+ [10], DoReFa-Net [57], Bi-Real [37], Xnor-Net++ [5], IR-Net [42], LCR-BNN [46], RBNN [33], FDA [53]. At the setting of 1W/1A, we use the Bi-Real architecture as most previous methods [42, 33, 53, 50] do for fair comparison. The results are shown in Table 2. Similar

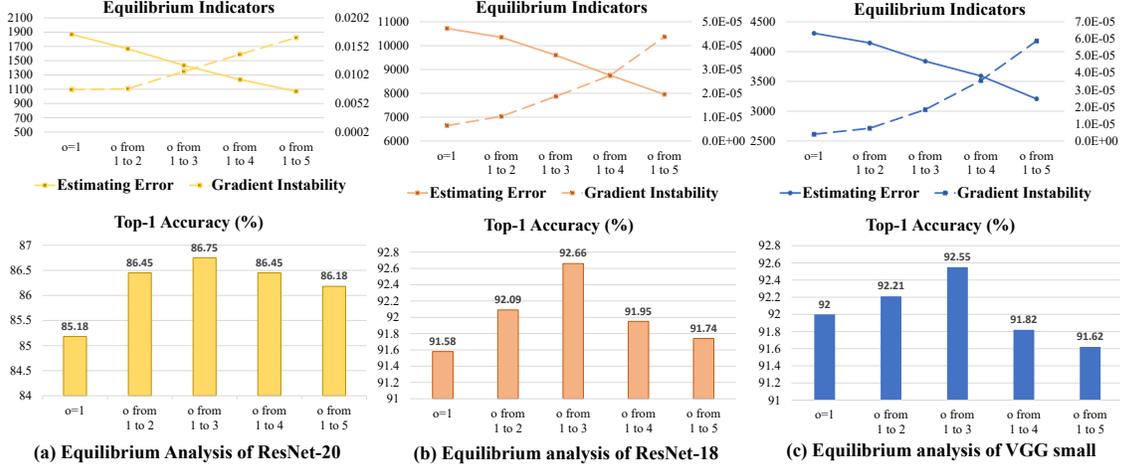


Figure 4: Illustrations of the estimating error indicators (above), gradient instability indicators (above) and the Top-1 accuracy (below) at different scales of o_{end} on CIFAR-10 dataset.

as the analysis on CIFAR-10 dataset, ReSET also displays excellent performance and outperforms all the SOTA methods without any assistance, e.g., modules or losses. For example, with ResNet-18 as backbone, ReSTE has 0.68% over the SOTA method FDA [53], at the setting of 1W/1A, even that FDA [53] adds a noise adaptation module to help the training. About the 1W/32A setting, ReSTE also has 0.50 improvement over the SOTA method LQ-Net [54], which has an additional loss to assist the training.

To sum up, we can conclude that ReSTE has excellent performance and outperforms the SOTA methods in both CIFAR-10 and large-scale ImageNet ILSVRC-2012 datasets. The reason is that our ReSTE is always rational, with less estimating error than STE, as well as that we obtain the desirable degree of the equilibrium by the ReSTE, which is capable of flexibly balancing the estimating error and the gradient stability. Moreover, ReSTE surpasses other binary methods without any assistance of additional modules or losses, showing the importance of fully considering the gradient stability and finding the suitable degree of equilibrium to BNNs training.

4.3. Estimators Comparison

To further evaluate the effectiveness of our approach, we compare ReSTE with other estimators in the same and fair setting without other auxiliaries, e.g., modules or additional losses.

Specifically, we use ResNet-20 as our backbone, comparing ReSTE with STE [9], DSQ [16], EDE [42], FDA [53], RBNN [33] on CIFAR-10 [30] at the setting of 1W/1A. Note that FDA here doesn't contain the noise adaptation module [53] and RBNN doesn't use the rotation procedure since we only use the sign function with scalar in forward process for fair comparison. Experiments results

are shown in Table 3. From the table we can observe that although ReSTE is concise, it significantly surpasses all the estimators in SOTA binary methods at the fully fair setting, with about 0.88% and 0.95% improvement over the estimators in RBNN and FDA. There are two facets of reasons. First is that our ReSTE always guarantees the rational property, with less estimating error than STE. Second is that we find out the desirable degree of the equilibrium with the assistance of the excellent ReSTE, which is capable of flexibly balancing the estimating error and the gradient stability.

4.4. Analysis of the Equilibrium Perspective

To quantitatively and clearly demonstrate the equilibrium phenomenon and show the balancing ability of ReSTE, we adjust o_{end} at different scales and meanwhile test the estimating error, gradient stability and the model performance. To make the results more convincing, we conduct the experiments with three widely-used backbones, ResNet-20, ResNet-18 and VGG-small. All the experiments are conducted on CIFAR-10 dataset at the setting of 1W/1A. We evaluate the estimating error and gradient stability layers by layers with the indicators proposed in Sec. 3.1 and use the average results of all the binarized layers. Meanwhile, we will collect the results from different training epochs to obtain the final indicators for an overall training, as shown in Fig. 4.

From the figures we can observe that with o_{end} increasing, the estimating error becomes smaller and smaller, while the gradient instability becomes bigger and bigger. This observation shows that although the estimating error can be reduced by adjusting the estimator close to the sign function, the gradient stability will decline along with. In addition, the model performance increases first and then decreases with the change of o_{end} , which implies that the large gra-

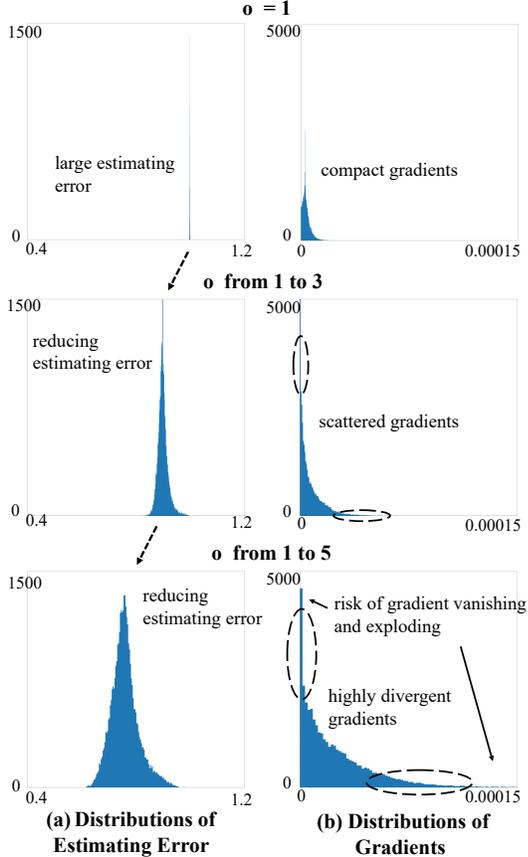


Figure 5: Illustrations of distributions of the estimating error (left) and the gradients (right) at different scales of o_{end} . X-axes represent the values of the estimating error and the gradients, y-axes are the frequency.

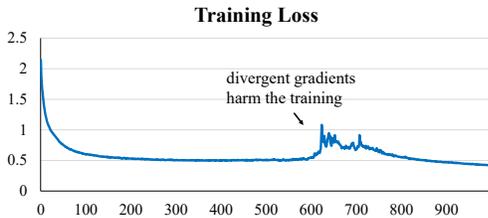


Figure 6: Illustrations of an example that divergent gradients ($o_{\text{end}} = 10$) will harm the BNNs training.

gradient instability will harm the model performance. Such changes clearly reflect the equilibrium phenomenon and validate our claim that highly divergent gradients will harm the BNNs training.

In addition, it can also be seen from the figures that ReSTE can adjust the degree of equilibrium by easily changing the hyper-parameter o_{end} . Moreover, the desirable degrees of equilibrium, i.e., the desirable o_{end} to produce high performance, are same in all the backbones, showing

the robustness and universality of ReSTE. When applying ReSTE at different backbones for different applications, we can simply adjust o_{end} to find out the suitable degree of the equilibrium and obtain good performance. More experiments about equilibrium analysis are shown in supplementary materials.

To obtain intuitive visualizations of the equilibrium phenomenon, we additionally visualize the distributions of the estimating error and the distributions of gradient at different scales of o . We use ResNet-18 as backbone and conduct the experiment on CIFAR10 dataset at the setting of 1W/1A. The results are shown in Fig. 5. From the figure we can observe that with o_{end} increasing, the peak values of the estimating error distribution become smaller, but the gradients become more divergent, which harms the model training and increases the risk of gradient vanishing or exploding. This visualization further demonstrates the equilibrium phenomenon and highlights the importance of finding the suitable degree of it.

To further validate our claim that highly divergent gradients will harm the model training, we demonstrate an example in Fig. 6. In this example, we use $o_{\text{end}} = 10$ with ResNet-20 as backbone and test on CIFAR-10 dataset at the setting of 1W/1A. We can observe that the training loss has huge fluctuations at about 600 to 700 epochs due to the divergent gradients, causing the final accuracy to decrease from 86.75 to 82.86. When o_{end} further increases, the training will fail irreversibly. This phenomenon verifies the harm of highly divergent gradients to model training and further demonstrates the importance of the equilibrium perspective.

5. Conclusion

In this work, we consider BNNs training as the equilibrium between the estimating error and the gradient stability. In this view, we firstly design two indicators to quantitatively and clearly demonstrate the equilibrium phenomenon. In addition, to balance the estimating error and the gradient stability well, we look back to the original STE and revise it into a new power function based estimator, rectified straight through estimator (ReSTE). Comparing to other estimators, ReSTE is rational and is capable of flexibly balancing the estimating error and the gradient stability. Extensive performance study on two datasets has demonstrated the effectiveness of ReSTE, surpassing state-of-the-art methods. By two carefully-designed indicators, we demonstrate the equilibrium phenomenon and show the ability of ReSTE to adjust the degree of equilibrium.

6. Acknowledgments

This work was supported partially by the NSFC (U21A20471, U1911401, U1811461), Guangdong NSF Project (No. 2023B1515040025, 2020B1515120085).

References

- [1] Thalaiyasingam Ajanthan, Kartik Gupta, Philip Torr, Richard Hartley, and Puneet Dokania. Mirror descent view for neural network quantization. In *AISTATS*, 2021.
- [2] Hande Alemdar, Vincent Leroy, Adrien Prost-Boucle, and Frédéric Pétrot. Ternary neural networks for resource-efficient ai applications. In *IJCNN*, 2017.
- [3] Yu Bai, Yu-Xiang Wang, and Edo Liberty. Proxquant: Quantized neural networks via proximal operators. In *ICLR*, 2019.
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv*, 2013.
- [5] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *arXiv*, 2019.
- [6] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. In *CVPR*, 2017.
- [7] Zhiyong Cheng, Daniel Soudry, Zexi Mao, and Zhenzhong Lan. Training binary multilayer neural networks for image classification using expectation backpropagation. *arXiv*, 2015.
- [8] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *NeurIPS*, 2015.
- [9] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv*, 2016.
- [10] Sajad Darabi, Mouloud Belbahri, Matthieu Courbariaux, and Vahid Partovi Nia. Bnn+: Improved binary network training. *arXiv*, 2018.
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [12] Ruizhou Ding, Ting-Wu Chin, Zeye Liu, and Diana Marculescu. Regularizing activation distribution for training binarized deep networks. In *CVPR*, 2019.
- [13] Xiaohan Ding, Tianxiang Hao, Jianchao Tan, Ji Liu, Jungong Han, Yuchen Guo, and Guiguang Ding. Resrep: Lossless cnn pruning via decoupling remembering and forgetting. In *ICCV*, 2021.
- [14] Yinpeng Dong, Renkun Ni, Jianguo Li, Yurong Chen, Jun Zhu, and Hang Su. Learning accurate low-bit deep neural networks with stochastic quantization. In *BMVC*, 2017.
- [15] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- [16] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. In *ICCV*, 2019.
- [17] 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 backpropagation. In *AAAI*, 2019.
- [18] Jiaxin Gu, Junhe Zhao, Xiaolong Jiang, Baochang Zhang, Jianzhuang Liu, Guodong Guo, and Rongrong Ji. Bayesian optimized 1-bit cnns. In *ICCV*, 2019.
- [19] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv*, 2015.
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [21] Koen Helwegen, James Widdicombe, Lukas Geiger, Zechun Liu, Kwang-Ting Cheng, and Roeland Nusselder. Latent weights do not exist: Rethinking binarized neural network optimization. In *NeurIPS*, 2019.
- [22] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. Neural networks for machine learning. *Coursera, video lectures*, 2012.
- [23] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv*, 2015.
- [24] Lu Hou, Quanming Yao, and James T Kwok. Loss-aware binarization of deep networks. In *ICLR*, 2017.
- [25] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv*, 2017.
- [26] Qinghao Hu, Peisong Wang, and Jian Cheng. From hashing to cnns: Training binary weight networks via hashing. In *AAAI*, 2018.
- [27] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *JMLR*, 2017.
- [28] Kyuyeon Hwang and Wonyong Sung. Fixed-point feedforward deep neural network design using weights+ 1, 0, and -1. In *SiPS*, 2014.
- [29] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv*, 2016.
- [30] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Technical report, University of Toronto*, 2009.
- [31] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 2015.
- [32] Cong Leng, Zesheng Dou, Hao Li, Shenghuo Zhu, and Rong Jin. Extremely low bit neural network: Squeeze the last bit out with admm. In *AAAI*, 2018.
- [33] Mingbao Lin, Rongrong Ji, Zihan Xu, Baochang Zhang, Yan Wang, Yongjian Wu, Feiyue Huang, and Chia-Wen Lin. Rotated binary neural network. In *NeurIPS*, 2020.
- [34] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *NeurIPS*, 2017.
- [35] Chunlei Liu, Wenrui Ding, Yuan Hu, Baochang Zhang, Jianzhuang Liu, Guodong Guo, and David Doermann. Rectified binary convolutional networks with generative adversarial learning. *IJCV*, 2021.

- [36] Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. Reactnet: Towards precise binary neural network with generalized activation functions. In *ECCV*, 2020.
- [37] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *ECCV*, 2018.
- [38] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.
- [39] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *ICCV*, 2017.
- [40] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *CVPR*, 2019.
- [41] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [42] 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 *CVPR*, 2020.
- [43] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. Introduction to tensor decompositions and their applications in machine learning. *arXiv*, 2017.
- [44] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, 2016.
- [45] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016.
- [46] Yuzhang Shang, Dan Xu, Bin Duan, Ziliang Zong, Liqiang Nie, and Yan Yan. Lipschitz continuity retained binary neural network. In *ECCV*, 2022.
- [47] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [48] Daniel Soudry, Itay Hubara, and Ron Meir. Expectation backpropagation: Parameter-free training of multilayer neural networks with continuous or discrete weights. In *NeurIPS*, 2014.
- [49] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019.
- [50] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. *arXiv*, 2019.
- [51] Zhijun Tu, Xinghao Chen, Pengju Ren, and Yunhe Wang. Adabin: Improving binary neural networks with adaptive binary sets. In *ECCV*, 2022.
- [52] Sheng Xu, Yanjing Li, Tiancheng Wang, Teli Ma, Baochang Zhang, Peng Gao, Yu Qiao, Jinhu Lv, and Guodong Guo. Recurrent bilinear optimization for binary neural networks. In *ECCV*, 2022.
- [53] Yixing Xu, Kai Han, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Learning frequency domain approximation for binary neural networks. In *NeurIPS*, 2021.
- [54] Dongqing Zhang, Jiaolong Yang, Dongqiangzi Ye, and Gang Hua. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *ECCV*, 2018.
- [55] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *CVPR*, 2018.
- [56] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *CVPR*, 2022.
- [57] 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*, 2016.
- [58] Mingjian Zhu, Yehui Tang, and Kai Han. Vision transformer pruning. *arXiv*, 2021.
- [59] Zhuangwei Zhuang, Minghui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. Discrimination-aware channel pruning for deep neural networks. In *NeurIPS*, 2018.