

Soft Threshold Ternary Networks

Weixiang Xu^{1*}, Xiangyu He¹, Tianli Zhao¹, Qinghao Hu^{1,2}, Peisong Wang^{1,2} and Jian Cheng^{1,2}

¹Institute of Automation, Chinese Academy of Sciences

²CAS Center for Excellence in Brain Science and Intelligence Technology

{xuweixiang2018,hexiangyu2017,zhaotianli2019}@ia.ac.cn,

{qinghao.hu,peisong.wang,jcheng}@nlpr.ia.ac.cn,

Abstract

Large neural networks are difficult to deploy on mobile devices because of intensive computation and storage. To alleviate it, we study ternarization, a balance between efficiency and accuracy that quantizes both weights and activations into ternary values. In previous ternarized neural networks, a hard threshold Δ is introduced to determine quantization intervals. Although the selection of Δ greatly affects the training results, previous works estimate Δ via an approximation or treat it as a hyper-parameter, which is suboptimal. In this paper, we present the Soft Threshold Ternary Networks (STTN), which enables the model to automatically determine quantization intervals instead of depending on a hard threshold. Concretely, we replace the original ternary kernel with the addition of two binary kernels at training time, where ternary values are determined by the combination of two corresponding binary values. At inference time, we add up the two binary kernels to obtain a single ternary kernel. Our method dramatically outperforms current state-of-the-arts, lowering the performance gap between full-precision networks and extreme low bit networks. Experiments on ImageNet with AlexNet (Top-1 55.6%), ResNet-18 (Top-1 66.2%) achieves new state-of-the-art.

1 Introduction

Deploying deep neural networks to resource limited devices is still a challenging problem due to the requirements of abundant computing and memory resources. A variety of methods have been proposed to reduce the parameter size and accelerate the inference phase, such as compact model architecture design (in a handcrafted way or automatic search way) [Howard *et al.*, 2017; Zoph and Le, 2016], network pruning [Han *et al.*, 2015; He *et al.*, 2019], knowledge distilling [Hinton *et al.*, 2015], low-rank approximation [Jaderberg *et al.*, 2014; Jaderberg *et al.*, 2014], etc. Besides, recent works [Courbariaux *et al.*, 2014; Gupta *et al.*, 2015] show that full-precision weights and activations are not necessary

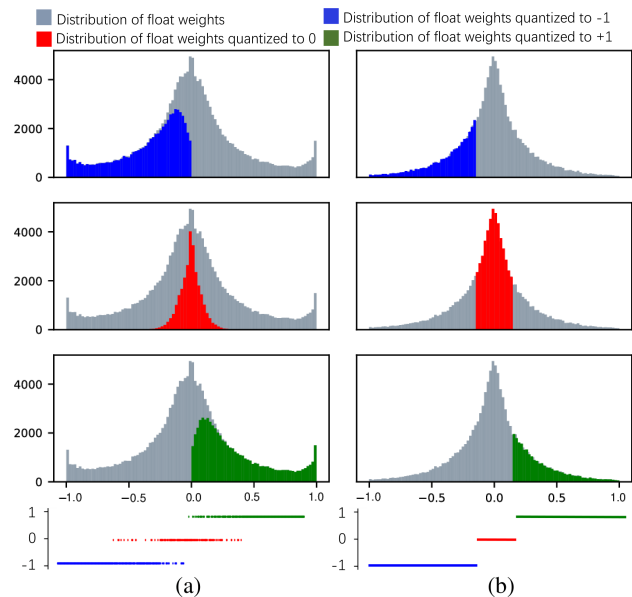


Figure 1: Comparison between *soft* and *hard* threshold. The weight distributions are drawn from ResNet-18 on ImageNet. We take Layer1.0.conv1 as an example. (a) Our soft threshold ternarization. (b) Hard threshold ternarization which splits intervals with Δ . The blue/red/green distribution comes from floating point weights quantized to $-1/0/1$, respectively. The blue/red/green lines below distributions denote positions of floating point weights that will be quantized to $-1/0/1$ when sorting weights from small to large.

for networks to achieve high performances. This discovery indicates that both weights and activations in neural networks can be quantized to low-bit formats. In this way, both storage and computation resources can be saved.

The extreme cases of network quantization are binary and ternary quantization ($\{-1, 1\}$ and $\{-1, 0, 1\}$). In computationally heavy convolutions, multiply-accumulate consumes most of the operation time. Through binarization or ternarization, multiply-accumulate can be replaced by cheap bitwise *xnor* and *popcount* operations [Courbariaux *et al.*, 2016]. Although binary networks can achieve high compression rate and computing efficiency, they inevitably suffer from a long training procedure and noticeable accuracy drops owing to poor representation ability.

In order to make a balance between efficiency and accura-

*Contact Author

cy, ternary CNNs convert both weights and activations into ternary values. Theoretically, ternary CNNs have stronger representation ability than binary CNNs, and should have a better performance. Although many works [Li *et al.*, 2016; Zhu *et al.*, 2016; Wang *et al.*, 2018; Wan *et al.*, 2018] provide tremendous efforts on ternary neural networks, we argue that there are two main issues in existing works: 1) All previous ternary networks use a *hard* Δ to divide floating point weights into three intervals (see Eq.3 for more details), which introduces an additional constraint into ternary networks. 2) Although the theoretical optimal Δ can be calculated through optimization, calculating it at each training iteration is too time-consuming. Therefore, previous methods either use a hyper-parameter to estimate Δ or only calculate the exact Δ at the first training iteration then keep it for later iterations, which further limits ternary network accuracy. A brief summary: they introduce an additional constraint based on Δ . Further more, a suboptimal Δ is used instead of the exact one.

In this paper, we propose Soft Threshold Ternary Networks (STTN), an innovative method to train ternary networks. STTN avoids the *hard* threshold and enables the model to automatically determine which weights to be $-1/0/1$ in a *soft* manner. We convert training ternary networks into training equivalent binary networks (more details in section 3.2). STTN enjoys several benefits: 1) The constraint based on Δ is removed. 2) STTN is free of calculation of Δ . 3) As shown in Figure 1 (a), weights are ternarized in a soft manner, which is why we name it *soft* threshold.

Our contributions are summarized as follows.

- We divide previous ternary networks into two catalogues: Optimization-based method and Learning-based method. Analysis is given about issues in the existing ternary networks that prevent them from reaching high performance.
- We propose STTN, an innovative way to train ternary networks by decomposing one ternary kernel into two binary kernels during training. No extra parameters or computations are needed during inference. Quantizing to $-1/0/1$ is determined by two binary variables rather than decided by a hard threshold.
- We show that our proposed ternary training method provides competitive results on the image classification, i.e. CIFAR-10, CIFAR-100 and ImageNet datasets. Qualitative and quantitative experiments show that it outperforms previous ternary works.

2 Related Work

Low-bit quantization of deep neural networks has recently received increasing interest of deep learning communities. By utilizing fixed-point weights and feature representations, not only the model size can be dramatically reduced, but also inference time can be saved.

The extreme cases of network quantization are binary neural networks (BNN) and ternary neural networks (TNN).

BNN [Courbariaux *et al.*, 2016] constrains both the weights and activations to either $+1$ or -1 , which produces

reasonable results on small datasets, such as MNIST and CIFAR-10. However, there is a significant accuracy drop on large scale classification datasets, such as ImageNet. Some improvements based on BNN have been investigated. For example, [Darabi *et al.*, 2018] introduces a regularization function that encourages training weights around binary values. [Tang *et al.*, 2017] proposes to use low initial learning rate. Nevertheless, there is still non-negligible accuracy drop. To improve the quality of the binary feature representations, XNOR-Net [Rastegari *et al.*, 2016] introduces scale factors for both weights and activations during binarization process. DoReFa-Net [Zhou *et al.*, 2016] further improves XNOR-Net by approximating the activations with more bits. Since ABC-Net [Lin *et al.*, 2017], several works propose to decompose a single convolution layer into K binary convolution operations [Liu *et al.*, 2019; Zhu *et al.*, 2019]. Although higher accuracy can be achieved, $K \times$ extra parameters and computations are needed for both training and inference time, which defeats the original purpose of binary networks.

To make a balance between the efficiency and accuracy, ternary CNNs convert both weights and activations into ternary values. We will review previous ternary works in detail in section 3.1.

3 Methodology

We first revisit the formulation of previous ternary networks. We divide them into two catalogues and show their common issues in calculating the appropriate Δ . We then present our novel Soft Threshold Ternary Networks in detail, including scaling coefficient constraint and backward approximation. We show that our method can avoid the previous issue in a simple but effective way.

3.1 Review of Previous Ternary Networks

Problem Formulation

Full precision convolution or inner-product can be formed as: $Y = \sigma(W^T X)$. Here $\sigma(\cdot)$ represents the nonlinear activation function, such as ReLU. $W \in \mathbb{R}^{n \times chw}$ and $X \in \mathbb{R}^{chw \times HW}$ are float point weights and inputs respectively, where (n, c, h, w) are filter number, input channels, kernel height and kernel width, and (H, W) are height and width of output feature maps. Ternary networks convert both weights and inputs into ternary values: $T \in \{+1, 0, -1\}^{n \times chw}$ and $X^t \in \{+1, 0, -1\}^{chw \times HW}$.

As for weights ternarization, α is used to estimate the floating weight W alone with ternary weight T .

$$W \approx \alpha T \quad (1)$$

Previous works formulate ternarization as a weight approximation optimization problem:

$$\begin{cases} \alpha^*, T^* = \arg \min_{\alpha, T} J(\alpha, T) = \|W - \alpha T\|_2^2 \\ s.t. \quad \alpha \geq 0, T_i \in \{-1, 0, 1\}, i = 1, 2, \dots, n. \end{cases} \quad (2)$$

A *hard* threshold Δ is then introduced by previous works to divide quantization intervals, which sets an additional con-

straint on ternary networks:

$$T_i = \begin{cases} -1, & \text{if } W_i > \Delta \\ 0, & \text{if } |W_i| \leq \Delta \\ +1, & \text{if } W_i < -\Delta \end{cases} \quad (3)$$

As illustrated in Figure 1(b), Δ plays a role as hard threshold here, splitting the floating point weights into $\{-1, 0, 1\}$. We will show that the above constraint can be avoided in our method. According to methods of calculating α and Δ , we divide previous ternary networks into 1) optimization-based ternary networks, and 2) learning-based ternary networks.

Optimization-based Ternary Networks

TWN [Li *et al.*, 2016] and TSQ [Wang *et al.*, 2018] formulate the ternarization as Eq.(2).

In TWN, T is obtained according to threshold-based constraint Eq.(3). They give exact optimal α^* and Δ^* as follows:

$$\alpha_{\Delta}^* = \frac{1}{|I_{\Delta}|} \sum_{i \in I_{\Delta}} |W_i|; \quad \Delta^* = \arg \max_{\Delta > 0} \frac{1}{|I_{\Delta}|} \left(\sum_{i \in I_{\Delta}} |W_i| \right)^2 \quad (4)$$

Note that Δ^* in Eq.(4) has no straightforward solution. It can be time and computing consuming if using discrete optimization to solve it. In TWN, they adopt an estimation $\Delta' \approx 0.7 \cdot E(|W|)$ to approximate the optimal Δ^* . Obviously there is a gap between Δ^* and Δ' . What's more, the optimal scaling factor α^* in Eq.(4) will also be affected because of dependence on Δ .

In TSQ, they try to obtain T^* by directly solving Eq.(2). They give exact optimal α^* and T^* as follows¹:

$$\alpha^* = \frac{W^T T}{T^T T}; \quad T^* = \arg \max_T \frac{(W^T T)^2}{T^T T} \quad (5)$$

Note that there is still no straightforward solution for T^* in Eq.(5). TSQ proposes OTWA(Optimal Ternary Weights Approximation) to obtain T^* . However, sorting for $|W_i|$ is needed in their algorithm. Time complexity is higher than $O(N \log(N))$, where N is the number of elements in kernel weights. To avoid time consuming sorting, they only calculate Δ at the first training iteration and then keep it for later iterations, which is suboptimal.

Learning-based Ternary Networks

TTQ [Zhu *et al.*, 2016] tries another way to obtain scaling factor α and threshold Δ . They set α as a learnable parameter and set Δ as a hyper-parameter. Based on TTQ, RTN [Li *et al.*, 2019] introduces extra float point scaling and offset parameters to reparameterize ternary weights.

Because the search space of hyper-parameter Δ is too large, both of TTQ and RTN set Δ the same across all layers, e.g. $\Delta = 0.5$ in RTN. We argue that ternary methods with hyper-parameter are not reasonable as following:

In Figure 2, we visualize the weight distribution of ResNet-18 on ImageNet from different layers and from different kernels in the same layer, respectively. From the figure, we can observe that the distributions are different from layer to layer. Even for kernels in the same layer, their weight distributions

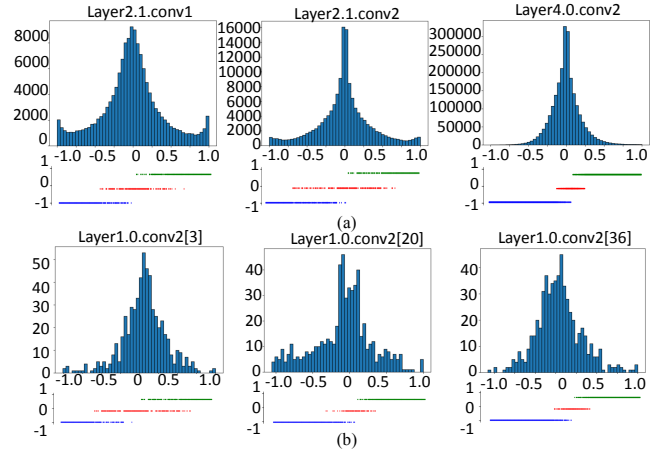


Figure 2: Weight distribution of ResNet-18 on ImageNet. (a) Distributions from different layers. (b) Distributions from different kernels in the same layer. "Layer1.0.conv2[36]" denotes the 37th kernel in Layer1.0.conv2. The blue/red/green lines below each distribution denote positions of floating point weights that will be quantized to $-1/0/1$, which are obtained through our STTN.

can be diverse. Therefore, 1) it is not reasonable for previous works to set a fixed threshold for all kernels (e.g. $\Delta = 0.5$ in TTQ and RTN). 2) What's more, their ternarization also use hard threshold: once Δ is set, their quantized intervals are determined according to Eq.(3).

So can we remove the hard Δ constraint Eq.(3), solving ternary networks in another way?

3.2 Soft Threshold Ternary Networks

Due to the issues analyzed above, we propose our novel STTN. Our motivation is to enable the model to automatically determine which weights to be $-1/0/1$, avoiding the hard threshold Δ .

We introduce our methods via convolution layers. The inner-product layer has a similar form. Concretely, at *training* time, we replace the ternary convolution filter T with t -two parallel binary convolution filters B_1 and B_2 . They are both binary-valued and have the same shape with ternary filter: $n \times chw$. Due to the additivity of convolutions with the same kernel sizes, a new kernel can be obtained by:

$$T = B_1 + B_2 \quad (6)$$

A key requirement for T to have ternary values is that those two binary filters should have the same scaling factors: α . With $\alpha B_1, \alpha B_2 \in \{+\alpha, -\alpha\}^{n \times chw}$, the sum of αB_1 and αB_2 is ternary-valued, i.e. $\alpha T \in \{+2\alpha, 0, -2\alpha\}^{n \times chw}$. Note that we only decompose the filters at training time. After the training, the two trained parallel binary filters are added up to obtain the ternary filter. Thus there is no extra computation when deploying trained models to devices. An example is illustrated in Figure 3.

Zeroes are introduced in ternary filters (white squares in Figure 3 at positions where two parallel filters have opposite values). And $-1/1$ is obtained at positions where two parallel filters have the same signs. In this way, ternary values are determined by the combination of two corresponding binary values.

¹ W^T denotes transpose of matrix W . T denotes ternary weights.

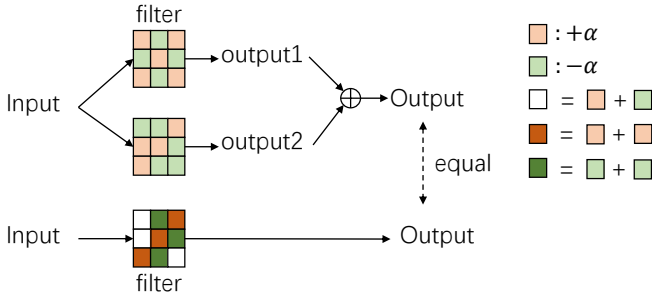


Figure 3: We illustrate our method with a 2D convolution for simplicity. The top of the figure are two training-time binary convolution kernels. Their weights are binary-valued: $\{+\alpha, -\alpha\}$. The bottom of the figure is an inference-time ternary convolution kernel who takes the same input. The inference-time kernel can be obtained by easily adding the corresponding two binary kernels. The weights are ternary-valued: $\{+2\alpha, 0, -2\alpha\}$. The additivity of convolutions tells us: the inference-time model can produce the same outputs as the training-time.

It is obvious that the outputs of training-time model are equal to the outputs of inference-time model:

$$Y = \sigma(\alpha B_1^T X + \alpha B_2^T X) = \sigma(\alpha T^T X) \quad (7)$$

In our proposed STTN, we convert training ternary networks into training equivalent binary networks. We abandon the constraint in Eq.(3). Not only Δ can be avoided, but also quantization intervals can be divided in a soft manner.

Scaling Coefficient

As mentioned above, a key requirement is that the corresponding two parallel binary filters should have the same scaling factors to guarantee the sum of them are ternary. Taking the two parallel binary filters into consideration, we obtain the appropriate $\{\alpha_1, \alpha_2\}$ by minimizing the following weight approximation problem:

$$J(\alpha_1, \alpha_2, B_1, B_2) = \|W_1 - \alpha_1 B_1\|_2^2 + \|W_2 - \alpha_2 B_2\|_2^2 \quad (8)$$

s.t. $\alpha_1 = \alpha_2; \alpha_1, \alpha_2 \geq 0$

Here B_1 and B_2 are the two parallel binary filters. And W_1 and W_2 are the corresponding float point filters. With constraint $\alpha_1 = \alpha_2$, we use α to denote them. Through expanding Eq.(8), we have

$$J(\alpha_1, \alpha_2, B_1, B_2) = \alpha^2 (B_1^T B_1 + B_2^T B_2) - 2\alpha (B_1^T W_1 + B_2^T W_2) + C \quad (9)$$

$C = W_1^T W_1 + W_2^T W_2$ is a constant because W_1 and W_2 are known variables. In order to get α^* , the optimal B_1^* and B_2^* should be determined. Since $B_1, B_2 \in \{+1, -1\}^{n \times chw}$, $B_1^T B_1 + B_2^T B_2 = 2nchw$ is also a constant. From Eq.(9), B_1^* and B_2^* can be achieved by maximizing $B_1^T W_1 + B_2^T W_2$ with constraint condition that $B_1, B_2 \in \{+1, -1\}^{n \times chw}$. Obviously the optimal solution can be obtained when binary kernel has the same sign with the corresponding float point kernel at the same positions, i.e. $B_1^* = \text{sign}(W_1)$, $B_2^* = \text{sign}(W_2)$.

Based on the optimal B_1^* and B_2^* , α^* can be easily calculated as:

$$\alpha^* = \frac{B_1^T W_1 + B_2^T W_2}{B_1^T B_1 + B_2^T B_2} = \frac{1}{2N} \left(\sum_{i=1}^N |W_{1i}| + \sum_{i=1}^N |W_{2i}| \right) \quad (10)$$

where $N = nchw$, is the number of elements in each weight. W_{1i} and W_{2i} are elements of W_1 and W_2 , respectively.

Backward Approximation

Since we decompose one ternary filter into two parallel binary filters at training time, binary weights approximation is needed in both forward and backward processes. During the forward propagation, the two related weights can be binarized through sign function along with the same scaling factor calculated by Eq.(10). However, during the backward propagation, the gradient of sign function is almost everywhere zero.

Assume ℓ as the loss function and $\widetilde{W} = \alpha B = \alpha \cdot \text{sign}(W)$ as the approximated weights, XNOR-Net [Rastegari *et al.*, 2016] alleviates this problem through Straight-Through Estimator (STE):

$$\frac{\partial \ell}{\partial W_i} = \frac{\partial \ell}{\partial \widetilde{W}_i} \frac{\partial \widetilde{W}_i}{\partial W_i} = \frac{\partial \ell}{\partial \widetilde{W}_i} \left(\frac{1}{N} + \alpha \frac{\partial \text{sign}(W_i)}{\partial W_i} \right) \quad (11)$$

Here, $\text{sign}(W_i)$ is approximated with $W_i \mathbf{1}_{|W_i| \leq 1}$. N is the number of elements in each weight.

However, note that an important requirement in our STTN is that the related two parallel binary filters should have the same scaling factors. The exact approximated weights should be:

$$\widetilde{W} = \alpha B = \frac{1}{2N} \left(\sum_{i=1}^N |W_{1i}| + \sum_{i=1}^N |W_{2i}| \right) \cdot \text{sign}(W) \quad (12)$$

Because Eq.(10) indicates that α is dependent on both W_1 and W_2 . When calculating the derivatives of W_{1i} , the effect of other kernels W_{1j} and W_{2j} should be considered. But Eq.(11) ignores the effect of W_{1j} and W_{2j} , which is not suitable for our backward approximation.

Taking above analysis into consideration, we propose to calculate derivatives of W in a more precise way:

$$\begin{aligned} \frac{\partial \ell}{\partial W_{1i}} &= \sum_{k=1}^2 \sum_{j=1}^N \frac{\partial \ell}{\partial \widetilde{W}_{kj}} \frac{\partial \widetilde{W}_{kj}}{\partial W_{1i}} \\ &= \sum_{k=1}^2 \sum_{j=1}^N \frac{\partial \ell}{\partial \widetilde{W}_{kj}} \left[\frac{1}{2N} \frac{\partial |W_{1i}|}{\partial W_{1i}} \text{sign}(W_{kj}) + \alpha \frac{\partial \text{sign}(W_{1i})}{\partial W_{1i}} \right] \\ &= \frac{1}{2N} \text{sign}(W_{1i}) \sum_{k=1}^2 \sum_{j=1}^N \left[\frac{\partial \ell}{\partial \widetilde{W}_{kj}} \text{sign}(W_{kj}) \right] \\ &\quad + \alpha \frac{\partial \text{sign}(W_{1i})}{\partial W_{1i}} \sum_{k=1}^2 \sum_{j=1}^N \frac{\partial \ell}{\partial \widetilde{W}_{kj}} \end{aligned} \quad (13)$$

Here $W_k = [W_{k1}, W_{k2}, \dots, W_{kN}]$ ($k \in \{1, 2\}$) are the two parallel kernels respectively. $\partial \ell / \partial W_{2i}$ can be calculated in the same way as Eq.(13).

Activation

In this paper, we also convert activations into ternary values. We use the same ternarization function as RTN [Li *et al.*, 2019]. Given floating point activation X , the ternary activation is calculated by the following equation. The difference between ours and RTN is that we do not introduce extra parameters or calculations.

$$X_i^t = \text{Ternarize}(X_i) = \begin{cases} \text{sign}(X_i), & \text{if } |X_i| > 0.5 \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

Method	TWN	TTQ	Ours
$\ W - \alpha T\ _2^2$	537.76	439.46	379.25

Table 1: L_2 distance between approximated ternary weights and float point weights. (We add up distances of all convolution layers together.)

During the backward process, as previous binary/ternary works [Courbariaux *et al.*, 2016; Rastegari *et al.*, 2016; Liu *et al.*, 2018], gradients propagated through ternarization function are estimated by Straight-Through Estimator (STE).

4 Experiments

In this section, we evaluate the proposed STTN in terms of qualitative and quantitative studies. Our experiments are conducted on three popular image classification datasets: CIFAR-10, CIFAR-100 and ImageNet (ILSVRC12). We test on several representative CNNs including: AlexNet, VGG-Net, and ResNet.

4.1 Implementation Details

We adopt the standard data augmentation scheme. In all CIFAR experiments, we pad 2 pixels in each side of images and randomly crop 32×32 size from padded images during training. As for ImageNet experiments, we first proportionally resize images to $256 \times N$ ($N \times 256$) with the short edge to 256. Then we randomly sub-crop them to 224×224 patches with mean subtraction and randomly flipping. No other data augmentation tricks are used during training.

Following RTN [Li *et al.*, 2019], we modify the block structure as BatchNorm \rightarrow Ternarization \rightarrow TernaryConv \rightarrow Activation. Following XNOR-Net [Rastegari *et al.*, 2016], we place a dropout layer with $p = 0.5$ before the last layer for AlexNet. For VGG-Net, we use the same architecture VGG-7 as TWN [Li *et al.*, 2016] and TBN [Wan *et al.*, 2018]. We do not quantize the first and the last layer as previous binary/ternary works. We replace all 1×1 downsampling layers with max-pooling in ResNet.

We use Adam with default settings in all our experiments. The batch size for ImageNet is 256. We set weight decay as $1e^{-6}$ and momentum as 0.9. All networks on ImageNet are trained for 110 epochs. The initial learning rate is 0.005, and we use cosine learning rate decay policy. All our models are trained from scratch.

4.2 Weight Approximation Evaluation

In this section, we explore the effect of the proposed STTN from qualitative view.

Previous works quantize weights into $\{-1, 0, 1\}$ by setting a hard threshold Δ . Different from them, the proposed STTN generates soft threshold, quantizing weights more flexibly. We illustrate the impact of threshold calculation on the performance of TNN based on ResNet-18. We first calculate the distance between trained floating weights W and trained ternary weights T . L_2 norm is used as the criterion for measurement as Eq.(2). We compare our method with TWN and TTQ. The results are shown in Table 1.

We can see that STTN obtains the smallest gap between trained floating weights and trained ternary weights, which

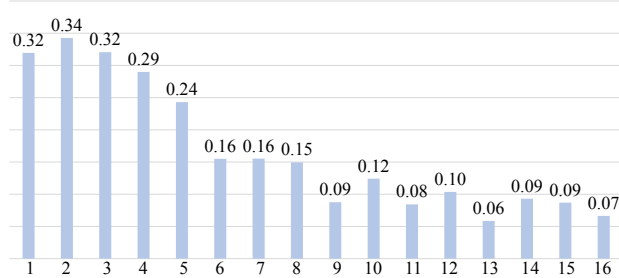


Figure 4: Weight sparsity rate of different layer in our STTN on ImageNet with ResNet-18. Here we illustrate it with 16 convolution layers in building blocks of ResNet.

indicates that our methods can realize a good weight approximation. Intuitively, the smaller the approximated weight error we get, the higher precision the model can obtain. The results show the effect of STTN qualitatively and further quantitative analyses are given in section 4.3.

We also analyze the reason why our STTN can obtain a smaller approximation error than previous ternary method. An essential effect of STTN is to quantize weights into $\{+1, 0, -1\}$ in a soft manner, just as shown in Figure 1(a) and Figure 2. The effect comes from that STTN throws the hard threshold away. That is, Eq.(3) is no more a constraint to the weight approximation optimization problem. By comparing Figure 1(a) with (b): this provides more flexible ternarization intervals.

Besides, we find that STTN adjust the weight sparsity regularly. Figure 4 shows the weight sparsity rates (the percentage of zeros) of different layers in our STTN on ResNet-18. From the figure we can see that the sparsity rates gradually decrease from the first layer to the last layer. This probably because high-level semantics need dense kernels to encode.

4.3 Network Ternarization Results

In this section, we evaluate the STTN from quantitative view by comparing with the state-of-the-art low-bit networks on various architectures. We ternarize both weights and activations. Experiments on only quantizing weights are also given.

Results on CIFAR-10

We first conduct experiments on CIFAR-10 dataset. We use the same network architecture as TWN, denoted as VGG-7. Compared with the architecture VGG-9 adopted in BNN and XNOR, the last two fully connection layers are removed. Table 2 shows the STTN results. Note that for VGG-7, STTN with ternary weights and activations can even obtain better performance than the full-precision model.

Results on CIFAR-100

In addition, we also evaluate STTN on CIFAR-100 dataset. We compare our STTN with a strong multi-bit baseline CBCN [Liu *et al.*, 2019]. CBCN replaces each convolution layer with several parallel binarized convolution layers. For fair comparisons, we use the same architecture as CBCN (ResNet-18 with 32-64-128-256 kernel stage). Note that in CBCN, the number of channels in one layer is $4 \times$. Table 3 shows our results that although CBCN uses $4 \times$ channels than ours, we obtain higher accuracy with fewer computations.

Bit-width	Method	Error(%)
32+32	Floating point [Li <i>et al.</i> , 2016]	7.12
1+32 [†]	BWN [Courbariaux <i>et al.</i> , 2015]	8.27
1+1 [†]	BNN [Courbariaux <i>et al.</i> , 2016]	10.15
1+1 [†]	XNOR [Rastegari <i>et al.</i> , 2016]	9.98
2+32	TWN [Li <i>et al.</i> , 2016]	7.44
1+2	TBN [Wan <i>et al.</i> , 2018]	9.15
2+2	Ours	7.07

Table 2: The error rates on CIFAR-10 with VGG-7. The number before and after “+” in the first column denotes the weight and activation bit-width respectively. † denotes the architecture is VGG-9, which adds two more FC layers at last.

Model	Kernel Stage	Accuracy(%)
Float(32+32)	32-64-128-256	73.62
CBCN(1+1)	(32-64-128-256)×4	70.07
Ours(2+2)	32-64-128-256	72.10

Table 3: Accuracy on CIFAR-100 with ResNet-18 (32-64-128-256). Our ternary networks can outperform multi-bit method significantly.

From this experiment, we argue that ternary networks can be considered before resorting to multi-bit methods.

Results on ImageNet

For the large-scale dataset, we evaluate our STTN over AlexNet and ResNet-18 on ImageNet. We compare our method with several exiting state-of-the-art low-bit quantization methods: 1) only quantizing weights: BWN [Courbariaux *et al.*, 2015], TWN [Li *et al.*, 2016] and TTQ [Zhu *et al.*, 2016]; 2) quantizing both weights and activations: XNOR [Rastegari *et al.*, 2016], Bi-Real [Liu *et al.*, 2018], ABC [Lin *et al.*, 2017], TBN [Wan *et al.*, 2018], HWGQ [Cai *et al.*, 2017], PACT [Choi, 2018] and RTN [Li *et al.*, 2019].

The overall results based on AlexNet and ResNet-18 are shown in Table 4 and 5. We highlight our accuracy improvement (up to **15%** absolute improvement compared with XNOR-Net and up to **1.7%** compared with state-of-the-art ternary models, **without pre-training**). These results show that the STTN outperforms the best previous ternary methods. Such improvement indicates that our soft threshold significantly benefits extreme low-bit networks.

What’s more, compared with PACT and RTN, we highlight additional improvements apart from accuracy: 1) Both PACT and RTN introduce extra floating point parameters (the activation clipping level parameter in PACT and reparameterized scale/offset in RTN) into the networks, which needs extra storage space and computation. 2) The extra introduced learnable parameters in PACT and RTN need careful manual adjustments, such as learning rate, weight decay and so on. Extensive manual tuning have to be tested for different networks on different datasets. However, our method is free of extra hyper-parameters to be tuned. Further more, we argue that our method can combine with those methods for further accuracy improvement. 3) For RTN, they argue that initial-

Model	Bit-width	Top-1(%)	Top-5(%)
XNOR	1+1	44.2	69.2
TBN	1+2	49.7	74.2
HWGQ	2+2	52.7	76.3
PACT*	2+2	55.0	—
RTN	2+2	53.9	—
Ours	2+2	55.6	78.9

Table 4: Comparison with the state-of-the-art methods on ImageNet with AlexNet. “—” means the accuracy is not reported. “*” indicates the networks use quaternary values instead of ternary values for 2 bits representation.

Model	Bit-width	Top-1(%)	Top-5(%)
Floating	32+32	69.3	89.2
TWN	2+32	61.8	84.2
TWN**	2+32	65.3	86.2
TTQ**	2+32	66.6	87.2
RTN***	2+32	68.5	—
Ours	2+32	68.8	88.3
XNOR	1+1	51.2	73.2
Bi-Real***	1+1	56.4	79.5
ABC-5	(1×5)+(1×5)	65.0	85.9
TBN	1+2	55.6	79.0
HWGQ	1+2	56.1	79.7
PACT*	2+2	64.4	—
RTN***	2+2	64.5	—
Ours	2+2	66.2	86.4

Table 5: Comparison with the state-of-the-art methods on ImageNet with ResNet-18. “*” indicates the model uses quaternary values instead of ternary values for 2 bits representation. “**” indicates the filter number of the network is 1.5×. “***” indicates the model needs full-precision pre-trained models to initialize. “×” in “ABC-5” denotes multi-bit networks with multi-branch.

ization from pre-trained full-precision models is vitally important for their methods (since the small architecture modification such as changing the order of BN and Conv layer in low-bit quantization, full-precision models released by open model zoo can not be used directly). However, our method shows that training from scratch can still obtain state-of-the-art results.

5 Conclusion

In this paper, we propose a simple yet effective ternarization method, Soft Threshold Ternary Networks. We divide previous ternary works into two catalogues and show that their hard threshold is suboptimal. By simply replacing the original ternary kernel with two parallel binary kernels at training, our model can automatically determine which weights to be -1/0/1 instead of depending on a hard threshold. Experiments on various datasets show that STTN dramatically outperforms current state-of-the-arts, lowering the performance gap between full-precision networks and extreme low bit networks.

References

- [Cai *et al.*, 2017] Zhaowei Cai, Xiaodong He, Jian Sun, and Nuno Vasconcelos. Deep learning with low precision by half-wave gaussian quantization. pages 5406–5414, 2017.
- [Choi, 2018] Jungwook Choi. Pact: Parameterized clipping activation for quantized neural networks. *arXiv: Computer Vision and Pattern Recognition*, 2018.
- [Courbariaux *et al.*, 2014] Matthieu Courbariaux, Jean-Pierre David, and Yoshua Bengio. Low precision storage for deep learning. *arXiv preprint arXiv:1412.7024*, 2014.
- [Courbariaux *et al.*, 2015] Matthieu Courbariaux, Yoshua Bengio, and Jeanpierre David. Binaryconnect: training deep neural networks with binary weights during propagations. pages 3123–3131, 2015.
- [Courbariaux *et al.*, 2016] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- [Darabi *et al.*, 2018] Sajad Darabi, Mouloud Belbahri, Matthieu Courbariaux, and Vahid Partovi Nia. Bn-n+: Improved binary network training. *arXiv preprint arXiv:1812.11800*, 2018.
- [Gupta *et al.*, 2015] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, pages 1737–1746, 2015.
- [Han *et al.*, 2015] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [He *et al.*, 2019] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019.
- [Hinton *et al.*, 2015] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- [Howard *et al.*, 2017] Andrew Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, M Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv: Computer Vision and Pattern Recognition*, 2017.
- [Jaderberg *et al.*, 2014] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866*, 2014.
- [Li *et al.*, 2016] Fengfu Li, Bo Zhang, and Bin Liu. Ternary weight networks. *arXiv preprint arXiv:1605.04711*, 2016.
- [Li *et al.*, 2019] Yuhang Li, Xin Dong, Sai Qian Zhang, Haoli Bai, Yuanpeng Chen, and Wei Wang. Rtn: Reparameterized ternary network. *arXiv preprint arXiv:1912.02057*, 2019.
- [Lin *et al.*, 2017] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *Advances in Neural Information Processing Systems*, pages 345–353, 2017.
- [Liu *et al.*, 2018] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 722–737, 2018.
- [Liu *et al.*, 2019] ChunLei Liu, Wenrui Ding, Xin Xia, Baochang Zhang, Jiaxin Gu, Jianzhuang Liu, Rongrong Ji, and Doermann David. Circulant binary convolutional networks: Enhancing the performance of 1-bit dcnns with circulant back propagation. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2019.
- [Rastegari *et al.*, 2016] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: ImageNet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pages 525–542. Springer, 2016.
- [Tang *et al.*, 2017] Wei Tang, Gang Hua, and Liang Wang. How to train a compact binary neural network with high accuracy? In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [Wan *et al.*, 2018] Diwen Wan, Fumin Shen, Li Liu, Fan Zhu, Jie Qin, Ling Shao, and Heng Tao Shen. Tbn: Convolutional neural network with ternary inputs and binary weights. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 315–332, 2018.
- [Wang *et al.*, 2018] Peisong Wang, Qinghao Hu, Yifan Zhang, Chunjie Zhang, Yang Liu, and Jian Cheng. Two-step quantization for low-bit neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4376–4384, 2018.
- [Zhou *et al.*, 2016] 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.
- [Zhu *et al.*, 2016] Chenzhuo Zhu, Song Han, Huizi Mao, and William J Dally. Trained ternary quantization. *arXiv preprint arXiv:1612.01064*, 2016.
- [Zhu *et al.*, 2019] Shilin Zhu, Xin Dong, and Hao Su. Binary ensemble neural network: More bits per network or more networks per bit? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4923–4932, 2019.
- [Zoph and Le, 2016] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016.