and highly compressed ternary weights.

- Our method significantly outperforms existing work in terms of bit-width, model size, and accuracy. For example, on the ImageNet, our method can compress a ResNet-18 model from 46 MB to 939 KB (49× compressed) while the accuracy is only 4% lower than the original accuracy. It is the best result among the existing results (41×, 6.5% accuracy drop). Other structures, such as ResNet-50, MobileNetV2 and Mask R-CNN, are also highly compressed to demonstrate the generalizability of our method.

- We combine 2-bit activation quantization with our method. Compared with the leading 2-bit or 1-bit quantization methods, ours maintains the similar accuracy, but the model size is 2× smaller.

2. Related Work

A notable amount of research efforts has been devoted to model compression methods to reduce model size and accelerate the inference, including quantization [8, 41, 55] and pruning [17, 26]. A recent comprehensive overview of model compression can be found in [38].

2.1. Quantization

Low-bit quantization methods convert float values of weights and activations to lower bit values [8, 9, 35, 41]. These methods make it possible to substantially reduce the computational cost during inference on CPU [49].

In weight sharing quantization, a finite set of vectors (codebook) are constructed to represent the blocks of weight values to achieve a higher compression rate [5, 15, 48]. Common approaches cluster weights through k-means [15, 47, 48, 52] and further finetune the clustering center by minimizing the reconstruction error in the network [14, 48, 52].

Mixed-precision quantization methods include reinforcement learning [50], integer programming [4], and differentiable neural architecture search [51]. Such methods assign different bit-widths to layer weights based on various importance measures, including hardware feedback [50, 54] and second-order information [12, 46].

2.2. Pruning

Pruning consists of structured and unstructured methods. It can greatly compress redundancy and maintain a high accuracy. However, only structured pruning methods can reduce the inference latency. Unstructured pruning uses criteria such as gradient [24, 37], and magnitude [17, 36] information to remove individual weights whereas structured pruning [1, 23, 26] aims to remove unimportant channels of the convolutional layer based on similar criteria. Unstructured pruning results in sparse model matrices that are difficult to accelerate inference speed without customized hardware [16]. In contrast, structured pruning produces models that are easier to accelerate [20, 26] because the original weight structures of the model are preserved.

Different from previous work, we show that minimizing the weight decay term to its minimum with L2 normalization and pruning can convert weight values to near discrete. Our findings show that near discrete weights can be used to produce more accurate and highly compressed ternary weights, which surpasses the current model compression methods and has a wide range of application scenarios.

3. Preliminaries

3.1. Ternary Quantization

Ternary quantization constrains the weights to ternary values: +1, 0 and -1, which helps to compress neural networks and reduce the inference time [57]. To convert a matrix $\mathbf{W}$ into ternary representation, we conduct the ternary operation over each element $w \in \mathbf{W}$ by:

$$\hat{w} = Quant(w, \Delta) = \begin{cases} 1 : w > \Delta, \\ 0 : |w| \leq \Delta, \\ -1 : w < -\Delta, \end{cases} \tag{1}$$

where Δ is a positive learnable threshold.

3.2. Unstructured Pruning

Pruning removes redundant weights from a trained model for compression purpose. During unstructured pruning, the less salient weight values are removed element-by-element. The weight values are sorted by their magnitude first, and only the top $k\%$ of them are kept:

$$t = Top(\mathbf{W}, k), \tag{2}$$

where t is a threshold at the top-k percent position of $\mathbf{W}$. The unstructured pruning is:

$$\textit{Prune}(w, t) = \begin{cases} 0 & |w| \leq t \\ w & |w| > t \end{cases}, \tag{3}$$

for each element $w \in \mathbf{W}$ where $|w| \leq t$ will be set to zero. Subsequently, we use pruning to represent unstructured pruning for simplicity.

3.3. The Minimum of the Weight Decay Term

It is well known that the weight decay penalty term favors many small weights rather than a few large ones [3, 39]. We further show that this feature will help to generate ternary as the magnitude of such smaller weights tends to

become identical when optimizing the weight decay term to its minimum.

Considering a neural network layer with a single output, the weight vector is:

$$\mathbf{w} = [w_1, w_2, ..., w_n], \tag{4}$$

$$\sum_{i=1}^{n} |w_i| = \omega, \tag{5}$$

where $w_i \in \mathbb{R}$, and $\omega \neq 0$. A standard error function E with weight decay (Chapter 9.5.4 in [3]):

$$\tilde{E} = E + \frac{\nu}{2} \sum_{i=1}^{n} w_i^2, \tag{6}$$

where $\sum_{i=1}^{n} w_i^2$ is the weight decay term and ν controls the extent to which the weight decay term influences values of $\mathbf{w}$'s elements. We would ideally like to estimate how a change in weight $\mathbf{w}$ would impact its weight decay term $\sum_{i=1}^{n} w_i^2$, so as to be able to find the connection between the minimum of weight decay term and ω in Equation 5.

Definition 1. Let $w_i = \frac{\omega}{n}$, then the weight decay term becomes:

$$A = \sum_{i=1}^{n} w_i^2 = \sum_{i=1}^{n} (\frac{\omega}{n})^2 = \frac{\omega^2}{n^2} n = \frac{\omega^2}{n}. \tag{7}$$

When $w_i = \frac{a_i \omega}{n}$, $a_i \in \mathbb{R}$, $a_i \neq 0$, and $\sum_{i=1}^{n} a_i = n$, then:

$$\sum_{i=1}^{n} w_i = \sum_{i=1}^{n} \frac{a_i \omega}{n} = \omega, \tag{8}$$

where $w_i = \frac{a_i \omega}{n}$ can represent arbitrary non-zero real number in $\mathbf{w}$ and the weight decay term becomes:

$$B = \sum_{i=1}^{n} w_i^2 = \omega^2 (\frac{a_1^2}{n^2} + \frac{a_2^2}{n^2} + ... + \frac{a_n^2}{n^2}) = \frac{\omega^2}{n^2} \sum_{i=1}^{n} a_i^2. \tag{9}$$

Theorem 1. For any non-zero real number $w_i = \frac{a_i \omega}{n}$, $a_i \in \mathbb{R}$ subject to $\sum_{i=1}^{n} a_i = n$, the minimum of $\mathbf{w}$'s weight decay term is $A = \frac{\omega^2}{n}$. When minimizing the weight decay term $\sum_{i=1}^{n} w_i^2$ to the minimum $A = \frac{\omega^2}{n}$ is equivalent to pushing $w_i = \frac{\omega}{n}$.

Proof. Since :

$$B - A = \frac{\omega^2}{n^2} (-n + \sum_{i=1}^{n} a_i^2), \tag{10}$$

where $\frac{\omega^2}{n^2} > 0$. we need to find the connection between n and $\sum_{i=1}^{n} a_i^2$. Let $f(a_1, a_2, .., a_n) = \sum_{i=1}^{n} a_i^2$ subject to the constraint $\sum_{i=1}^{n} a_i = n$, form the Lagrangian:

$$L(a_1, a_2, .., a_n, \lambda) = \sum_{i=1}^{n} a_i^2 + \lambda(-n + \sum_{i=1}^{n} a_i) \tag{11}$$

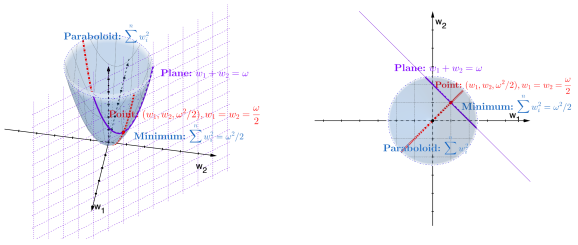

(a) Side view. (b) Top view.

Figure 1. The minimum of the weight decay term. The red dashed line is the intersecting line between the paraboloid $\sum_{i=1}^{n} w_i^2$ and the plane $w_1 = w_2$. **Point**$(w_1, w_2, \frac{\omega^2}{2})$ shows the conclusion, i.e., the minimum of the weight decay term $\sum_{i=1}^{n} w_i^2$ subject to the constraint $\sum_{i=1}^{n} |w_i| = \omega$ is $\frac{\omega^2}{2}$, where $w_1 = w_2 = \frac{\omega}{2}$ (and $n = 2$ in this case). For simplicity, we only demonstrate the result with $w_i > 0$

The first-order partial derivative of a_i and λ are:

$$\frac{\partial L}{\partial a_i} = \lambda + 2a_i, i \in [1, 2, .., n] \tag{12}$$

$$\frac{\partial L}{\partial \lambda} = -n + \sum_{i=1}^{n} a_i \tag{13}$$

Solving the system below based on Equation(12) and Equation(13):

$$\sum_{i=1}^{n} \frac{\partial L}{\partial a_i} = n\lambda + 2 \sum_{i=1}^{n} a_i = n\lambda + 2n = 0. \tag{14}$$

Then $\lambda = -2$ and $a_1 = a_2 = ... = a_n = 1$. The minimum of $f(a_1, a_2, .., a_n) = \sum_{i=1}^{n} a_i^2$ subject to the constraint of $\sum_{i=1}^{n} a_i = n$ is:

$$f(a_1, a_2, .., a_n) = \sum_{i=1}^{n} a_i^2 = n. \tag{15}$$

Then:

$$-n + \sum_{i=1}^{n} a_i^2 \geq 0, \tag{16}$$

thus $B - A \geq 0$ and A is the minimum of the weight decay term $\sum_{i=1}^{n} w_i^2$ (Figure 1).

4. Method

We find that minimizing the weight decay term to its minimum will not only push the magnitudes of weight values close to zeros [21, 39], but also near identical (Theorem 1). Having near identical weight magnitudes means the weight values are close to discrete. However, the existence of gradient update makes weight values far away from zeros [21]. Such conflict prevents weight magnitudes from being identical and discrete.

In this section, based on the minimum of the weight decay term, we show that using L2 normalization and pruning can make continuous weight values near discrete and having identical magnitudes, which can reduce the weight discrepancy between the forward pass and backward pass in STE. We also propose resetting trick and refined gradient methods to further relieve such discrepancy. Finally, the discrete weight values are quantized into highly compact and accurate ternary values through the discrepancy-reduced STE.

4.1. Make Weight Values Near Discrete

During training, the weight values tend to decrease due to the weight decay (Chapter 9.5.4 in [3]) [22, 39]. The minimum of $\frac{w^2}{n}$ (Equation 7) will also keep decreasing. On the contrary, the gradient update will offset such decay and maintain the magnitude of important weight values [21]. The final magnitude of weight values is determined by its usefulness in reducing error. The less important ones will be close to zero [21,39]. Then, the minimum inside the weight decay term is unreachable due to the conflict between gradient update and weight decay. Thus it is impossible to make weight magnitudes near identical (Theorem 1) only by minimizing weight decay term. We show that the L2 normalization and pruning can work with weight decay to make model weights have near identical and discrete magnitudes.

4.1.1 Connect Ternary with Weight Decay Term

Suppose that it is possible to have the weight decay term to reach its minimum $\frac{\omega^2}{n}$ (Equation 7), then:

$$\frac{\mathbf{w}}{\|\mathbf{w}\|} = \frac{\mathbf{w}}{\sqrt{\frac{\omega^2}{n}}} \Leftrightarrow w_i = \frac{\frac{\omega}{n}}{\sqrt{\frac{\omega^2}{n}}} = \pm\frac{\sqrt{n}}{n}. \tag{17}$$

And for ternary weight values:

$$\frac{\hat{\mathbf{w}}}{\|\hat{\mathbf{w}}\|} = \frac{\hat{\mathbf{w}}}{\sqrt{n}} \Leftrightarrow \hat{w}_i = \pm\frac{\sqrt{n}}{n}, \tag{18}$$

where $\hat{\mathbf{w}} = Quant(\mathbf{w}, 0)$. Equation 17 and 18 indicate that when minimizing the weight decay term to its minimum, the L2 normalized weight values is equivalent to the L2 normalized ternary weight values. To satisfy the L2 normalization condition, the model must be converted to L2 normalized version first. Fortunately, Sphere-Face [32], Weight-Normalization [43], and other studies [10,31] show that applying L2 normalization during training can promote convergence and generalization capabilities. Therefore, the ternary weight values and the weight decay term are connected by simply applying L2 normalization during training.

4.1.2 Eliminate Smaller Weight Values

In practice, it is impossible to have the weight decay term to reach its minimum (Equation 7) because of the above-mentioned conflict between the gradient update and weight decay (Section 4, 4.1). And we cannot obtain ternary from minimizing the weight-decay term with L2 normalization (Equation 17, 18) directly.

However, it is worth noting that, in Equation 17, the L2 normalized magnitude of w_i solely depends on n, i.e., the amount of larger weight values in $\mathbf{w}$ rather than the amount of non-zero values. Because removing smaller values in $\mathbf{w}$ will not affect ω significantly.

That is to say, when minimizing weight decay term with pruning (reducing n in Equation 17), as the less important weight values are eliminated, the L2 normalized magnitude of remaining weight values are reaching $\frac{\sqrt{n}}{n}$ (Equation 17, Figure 2), i.e., the normalized ternary values. In other words, pruning solves the conflict between gradient update and weight decay under the circumstance of L2 normalization.

Figure 2 shows that, after pruning, the mean magnitude of remaining weight values (L2 normalized) are very close to $\frac{\sqrt{n}}{n}$ (Equation 17). As the pruning ratio increases, the mean of the L2 normalized magnitude of remaining weight values will also increase (Figure 2).

4.1.3 The Impact of L2 Normalization

Suppose the L2 normalized $\mathbf{w}$ is used only in forward pass during training (note: the original weight $\mathbf{w}$ is unchanged, so the weight decay term is not normalized) [32, 43], then the error function becomes:

$$E = (\frac{\mathbf{w}^T}{\|\mathbf{w}\|}\mathbf{x} - \mathbf{y})^2. \tag{19}$$

The gradient update rule is:

$$\frac{\partial E}{\partial \mathbf{w}} = \frac{\partial E}{\partial \tilde{\mathbf{w}}}\frac{\partial \tilde{\mathbf{w}}}{\partial \mathbf{w}} = M_w\frac{\partial E}{\partial \tilde{\mathbf{w}}}, \tag{20}$$

where $\tilde{\mathbf{w}} = \frac{\mathbf{w}}{\|\mathbf{w}\|}$, and M_w is defined as (See [43]):

$$M_w = \frac{1}{\|\mathbf{w}\|}(\mathbf{I} - \frac{\mathbf{w}\mathbf{w}^T}{\|\mathbf{w}\|^2}), \tag{21}$$

4.2. Weight Quantization

Having had the L2 normalized magnitudes of weight values close to identical (Equation 17, 18), we still need to perform weight quantization together with pruning to further remove the less important weight values. We introduce a learnable quantization threshold Δ (Equation 1) to continue prune the less important values in order to facilitate weight quantization. Because of the use of L2 normalization, we refine the gradient of STE to reduce the discrepancy:

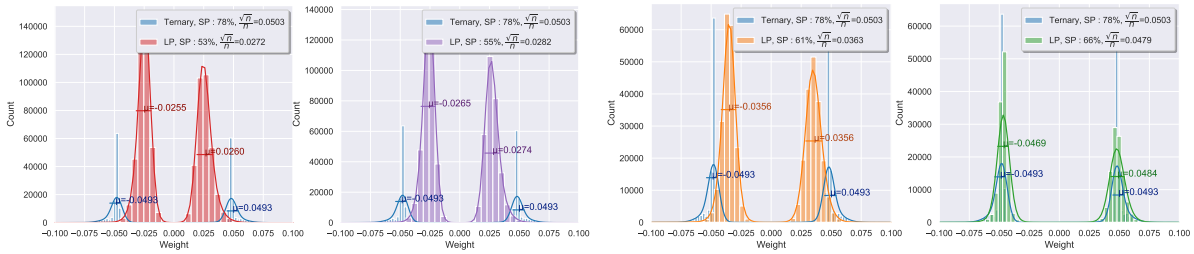

Figure 2. The relation between pruning ratio and the weight decay minimum of a middle layer of a ResNet-50 model. "Ternary" represents the ternary weight values (L2 normalized) from a ternary model. "LP" represents weight values (L2 normalized) from a pruned model with L2 normalization. "SP" represents the pruning ratio. "μ" is the mean value of the weight magnitudes. "$\frac{\sqrt{n}}{n}$" of "LP" is the magnitude of weight values (L2 normalized) corresponding to the minimum of the weight decay term (Equation 17). The charts show that the mean of the normalized weight magnitude, which is very close to the theoretical value $\frac{\sqrt{n}}{n}$, is increasing as the pruning ratio increases; and it is getting close to the optimized ternary weight values (blue colored).

Forward:

$$\hat{\mathbf{w}} = Quant(\mathbf{w}, \Delta) \tag{22}$$

$$\hat{\mathbf{w}} = \frac{\hat{\mathbf{w}}}{\|\hat{\mathbf{w}}\|}, \tag{23}$$

Backward:

$$\frac{\partial E}{\partial \mathbf{w}} = \frac{\partial E}{\partial \hat{\mathbf{w}}} \frac{\partial \hat{\mathbf{w}}}{\partial \mathbf{w}} M_w \underset{STE}{\approx} \frac{\partial E}{\partial \hat{\mathbf{w}}} M_w, \tag{24}$$

where M_w is defined in Equation 21. $\frac{\partial \hat{\mathbf{w}}}{\partial \mathbf{w}} = 1$, which is the same as the original STE [2]. We apply L2 normalization after quantization (Equation 23) in forward pass and scale the gradient by M_w in backward pass according to Equation 21. We adopt similar method from Trained Ternary Quantization (TTQ) [57] to optimize the quantization threshold Δ. The derivative of Δ is defined as [57]:

$$\frac{\partial L}{\partial \mathbf{\Delta}} = \sum_{i=1}^{n} \frac{\partial L}{\partial w_i}, w_i \in \mathbf{w} \text{ and } w_i \neq 0. \tag{25}$$

The gradient of $\frac{\partial L}{\partial w_i}$ where $w_i = 0$ is ignored (Equation 25).

4.2.1 Perform resetting Trick

Based on our observation, the weight values are gradually reduced during training even if L2 normalization and pruning are applied. The gap between important and less important values are reducing. Therefore, it is hard to distinguish the important and less important values by using the learned threshold Δ during quantization.

In our method, we reset the weight values into ternary after pruning ($\mathbf{w} = Quant(\mathbf{w}, 0)$), then perform regular training. The idea is to reduce the weight discrepancy between the forward pass and the backward pass. Resetting can put the important weight values closer to the normalized discrete ternary values $\frac{\sqrt{n}}{n}$ (Equation 17, 18), thus resisting weight decaying. After resetting, the less important values decays faster than the important weights during training. At the end of the training, the weight values are closer to the values being used in the forward pass in STE (Equation 22, Figure 2). We further investigate the impact of weight resetting in the Ablation Study (5.4).

4.3. Implementation Details

Our proposed method is shown in Algorithm 1. The overall process can be summerized as:

1. Apply L2 normalization to the pretrained model weights and train the model until it is converged.

2. Prune the model weights and reset the weight values to $\{-1, 0, 1\}$ based on their signs.

3. Perform ternary quantization and update the weights through STE.

4. Encode the ternary weights with a fixed codebook.

4.4. Activation Quantization

We combine our method with activation quantization methods to examine its performance in convolution layers. The LSQ [13] activation quantization is defined as:

$$y_q = round(\text{clip}(y/s, 0, 2^{k-1} - 1)), \tag{26}$$

where s is a learned quantizer step size, k is the quantization bit width, and y_q denotes a set of quantization levels. After clipping, each element of y is projected onto the quantization levels y_q.

Algorithm 1 PTQ training approach

Input: Input image $\mathbf{x}$; pretrained full precision weights $\mathbf{w}$; pruning ratio $k = 0.7$; Δ is a learned quantization threshold. $\phi(.)$ denotes a nonlinear activation function such as the ReLU.

Result: Quantized ternary networks for inference

L2 Normalization:
while *not converged* **do**
  $output = \phi(\frac{\mathbf{w}^T}{\|\mathbf{w}\|}\mathbf{x})$
  Perform SGD, and calculate $\frac{\partial L}{\partial \mathbf{w}}$ and update $\mathbf{w}$
end while

Pruning:
$t = Top(\mathbf{w}, k)$ *// pruning threshold*
$\mathbf{w} = Prune(\mathbf{w}, t)$
$\mathbf{w} = Quant(\mathbf{w}, 0)$ *// resetting trick*
while *not converged* **do**
  $output = \phi(\frac{\mathbf{w}^T}{\|\mathbf{w}\|}\mathbf{x})$
  Perform SGD, and calculate $\frac{\partial L}{\partial \mathbf{w}}$ and update $\mathbf{w}$
end while

Ternary Quantization:
while *not converged* **do**
  $\hat{\mathbf{w}} = Quant(\mathbf{w}, \Delta)$ *//pruning and quantization*
  $output = \phi(\frac{\hat{\mathbf{w}}^T}{\|\hat{\mathbf{w}}\|}\mathbf{x})$
  Perform SGD, and calculate $\frac{\partial L}{\partial \mathbf{w}}$, and $\frac{\partial L}{\partial \Delta}$ via refined STE, Equation 24-25
  Update $\mathbf{w}$, Δ
end while

4.5. The Fixed Codebook

Ternary values $\{-1, 0, 1\}$ are the ideal components of a codebook for model compression. In our work, we take every 3 continuous weight values as the codewords. For the linear layer, we pad zeros to match the length of the codewords. The total length of the codebook is 27 (5-bit). Our work only needs 15-bit (three codewords) to store each convolution kernel. However, previous work [34, 48] uses float number as codewords, and the memory footprint is not reduced during convolutional operations. As the minimum bit width of a number is 8-bit in a common operation system, we use external compression method [11, 40] to reduce the model's disk footprint.

5. Experiments

Our experiments involve two tasks, image classification and object detection. We evaluate our method on the ImageNet dataset [42] with MobileNetV2 [44] and ResNet-18/50 [19] architectures for image classification. For object detection, we use MS COCO [29] dataset and Mask R-CNN [18] with ResNet-50 FPN [53] as backbone. The pretrained weights are provided by the PyTorch zoo and Detectron2 [53]. We further quantize the model activations to 2-bit and compare against representative 2-bit activation quantization approaches. In the Ablation Study section, we study the impact of weight resetting during training.

5.1. Experimental Setup

For image classification, the batch size is 128. The weight decay is 0.0001, and the momentum of stochastic gradient descent (SGD) is 0.9. We use the cosine annealing schedule with restarts every 10 epochs [33] to adjust the learning rates. The initial learning rate is 0.001. We empirically prune the model with the ratio of 70%.

All of the experiments use 16-bit Automatic Mixed Precision (AMP) training provided by Nvidia and PyTorch to accelerate the training process. Thus, all parameters (except codebook and codewords) are stored as 16-bit floating point values. For further compression, we fuse the parameters of the batch normalization layers into two vectors [34, 48] and use 7-zip [11, 40] to compress the codewords.

5.2. Image Classification

Following the best practices of mainstream model compression work [12, 34, 48], we quantize all of the weights of the convolution and the fully-connected layers to 2-bit (except the first convolutional layer). We compare our results with leading model compression results from PQF [34], ABGD [48], BRECQ [28], HAWQ [12], and TQNE [14]. We also compare our work with existing milestone approaches, such as ABC-Net [30], Deep Compression (DC) [17], Hardware-Aware Automated Quantization (HAQ) [50], Hessian AWare Quantization (HAWQ) [12], LR-Net [45], and BWN [41].

The image classification results (Table 1) show that our proposed method significantly outperforms leading model compression methods in terms of bit-width, compression ratio and accuracy (Figure 3). For ResNet structures, our method reduces the storage footprint by more than 32× with a slight drop of 2% in accuracy compared with the original accuracy. Specifically, we compress ResNet18 from 46 MB to 1.28MB (36× compressed) while maintaining high accuracy (67.03% vs. 69.7% of the original model), and compress ResNet50 from 99 MB to 3.1MB (32× compressed) with an accuracy of 74.7% (vs. 76.15% of the original model). In the extreme cases, our method produces much smaller models with less than 4% accuracy drop, for example, 935KB out of 46MB (49× compressed) for ResNet-18 and 2.6MB out of 99MB (38× compressed) for ResNet-50. For MobileNetV2, our method has the best compression-accuracy result at above 15× compression level.

Models	Comp. Level	Methods	Bits (W/A)	Accuracy	Model Size	Ratio
ResNet-18 Acc.: 69.76 Size: 45 MB	∼30×	**PTQ (OURS)**	**2/16**	**67.03**	**1.28 MB**	**36×**
		PQF (2021) [34]	32/32	66.74	1.54 MB	29×
		ABGD (2020) [48]	32/32	65.81	1.6 MB	28×
	∼40×	**PTQ (OURS)**	**2/16**	**65.43**	**939 KB**	**49×**
		PQF (2021) [34]	32/32	63.33	1.03 MB	43×
		ABGD (2020) [48]	32/32	61.18	1.03 MB	43×
ResNet-50 Acc.: 76.15 Size: 99 MB	15−30×	**PTQ (OURS)**	2/16	74.7	**3.1MB**	**32×**
		HAWQ (2019) [12]	**4/2**	**75.4**	7.9 MB	12×
		PQF (2021) [34]	32/32	75.04	5.09 MB	19×
		TQNE (2020) [14]	32/32	74.3	-	19×
		ABGD (2020) [48]	32/32	73.79	5 MB	20×
		HAQ (2019) [50]	2∼8/2∼8	70.63	6.3MB	16×
		DC (2015) [17]	2/32	68.95	6.32MB	16×
	∼30×	**PTQ (OURS)**	**2/16**	**73.87**	**2.67 MB**	**38×**
		PQF (2021) [34]	32/32	72.18	3.19 MB	31×
		TQNE (2020) [14]	32/32	68.8	-	32×
		ABGD (2020) [48]	32/32	68.21	3.1 MB	32×
MobileNetV2 Acc.: 71.88 Size: 14 MB	15−20×	**PTQ (OURS)**	2/16	58.736	**0.74 MB**	**19×**
		HAQ (2019) [50]	**2∼6/2∼6**	**66.75**	0.95 MB	15×
		BRECQ (2021) [28]	2/8	56.29	0.83 MB	17×

Table 1. Model compression results on the ImageNet. Top-1 accuracy, bit-width, and model size are compared. "Bits (W/A)" denotes the bit-width of weight and activation. "Comp. Level" means compression level.

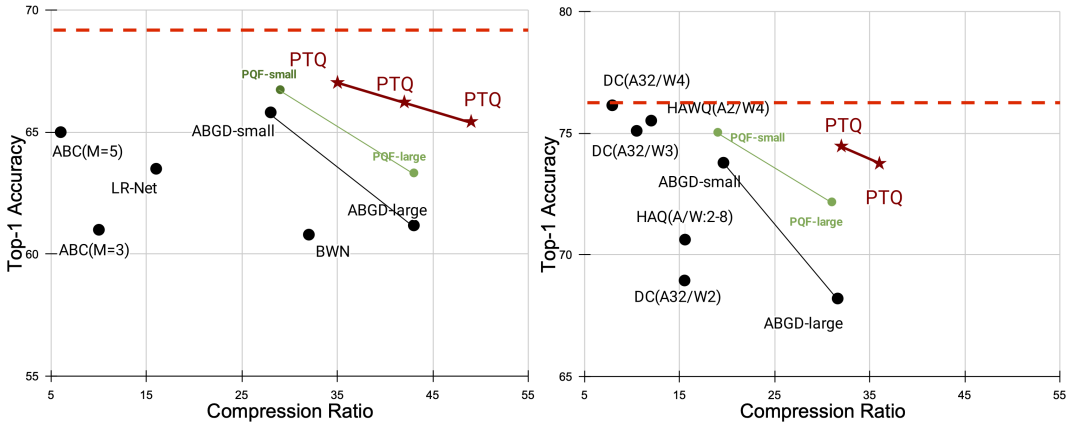

(a) ResNet-18 on the ImageNet.

(b) ResNet-50 on the ImageNet.

Figure 3. The compression ratio and accuracy of ResNet-18/50 on the ImageNet dataset. Our method achieves a much higher accuracy and compression ratio compared to existing work. The dash-line indicates the full-precision accuracy of models from PyTorch zoo.

We integrate 2-bit activation quantization, i.e., LSQ [13], to our work. We compare our results with publicly available 2-bit activation quantization models: APoT [27],LQ-Net [55], and DoReFa [56].

The results with 2-bit activation (Table 2) show that our method achieves similar top-1 accuracy with about 2× smaller model size.

5.3. Object Detection and Segmentation

Similar as previous work [28,34,48], we test our method on the Mask R-CNN [18] architecture with ResNet-50 [19] backbone to verify its generalizability. The original source code and the pretrained model are provided by Detectron2 [53]. We apply our method to the entire model except the first convolutional layer. We compare against recent base-

Models	Methods	Bits (W/A)	Acc.	Size
ResNet-18 Acc.: 69.76	**PTQ (OURS)**	2/2	**65.8**	**3 MB**
	LQ-NET (2018)	2/2	64.9	6 MB
	APOT (2019)	2/2	64.7	5 MB
	DOREFA (2016)	2/2	62.6	5.2 MB
ResNet-50 Acc.: 76.15	**PTQ (OURS)**	2/2	**71.6**	**4.8 MB**
	LQ-NET (2018)	2/2	71.5	13.8 MB

Table 2. Top-1 accuracy on the ImageNet with 2-bit activation quantization.

Methods	Bits (W/A)	AP^{bb}	AP^{mk}	Size	Ratio
BASELINE	32/32	38.6	35.2	170MB	1×
PTQ (OURS)	2/16	35.8	32.4	**5MB**	**34**×
ABGD (2020)	32/32	33.9	30.8	6.6MB	26×
PQF (2021)	32/32	36.3	33.5	6.6MB	26×
BRECQ (2021)	2/8	34.23	-	-	-

Table 3. For the object detection task, we compare the model size and Average Precision (AP) with bounding box (bb) and mask (mk). The "BASELINE" is the pretrained Mask-RCNN with ResNet-50 FPN [53].

lines, such as the ABGD [48], PQF [34], and BRECQ [28]. As shown in the Table(3), compared with ABGD and PQF, our method brings a higher compression ratio and similar or better recognition result.

We directly apply our method into existing Mask R-CNN source code, i.e., apply L2 normalization to the layer weights. During quantization training, we modify the forward pass and backward pass as described in Section 4.2. All of the hyperparameter settings and training strategies are the same as Detectron2's. We first convert the pretrained model into L2 normalized one, then prune the model by 80%, and quantize the model according to Section 4.3.

5.4. Ablation Study

Unlike other quantization work that uses STE to optimize the low-bit models directly, we apply resetting trick to reset weight values into ternary right after pruning, so as to reduce the discrepancy in STE. During training, we compare two experiment settings. One is referred to as STE-Only and the other is STE-Reset. The STE-Only experiment only uses STE to optimize the model, whereas the STE-Reset experiment applies extra resetting after pruning, then perform regular training. Both experiments prune the model to 70%. The results in Table 4 show that model performance improves as we reset the weight to ternary. Note that without resetting and pruning, our method becomes a variant of TWN [25].

Model	Methods	Accuracy
ResNet-18	STE-Only	65.2
	STE-Reset	68.0
ResNet-50	STE-Only	72.9
	STE-Reset	75.19

Table 4. The classification results on the ImageNet with/without applying the resetting after pruning. The results validates our strategy of reducing the discrepancy in STE. We skip the first and last layer during quantization as others [7, 13, 57]

6. Limitations

Our work has three limitations. Firstly, our method does not fully leverage activation quantization, therefore the inference efficiency is yet to be maximized. Fortunately, our method is compatible with mainstream activation methods, e.g., LSQ. However, we still need more experiments on ternary or binary activation to maximize the power of ternary quantization. Secondly, the existing computational architecture makes it difficult to efficiently store the 5-bit codebook. We have to use an external compression algorithm to reduce the disk footprint like other work does. The actual bit-width of the memory footprint depends on the type of hardware devices. We need more experiments to test our method across different devices. Thirdly, we have to apply L2 normalization before performing pruning and quantization. Therefore, our training time triples.

As for limitations in our assumptions, we assume that the proper pruning ratios are 70% for the image classification task and 80% for the object detection task. These ratios are based on our practical experience. However, in real-world application, it is hard to estimate the ratio without multiple experiments. A higher pruning ratio may reduce the model capacity and cause underfitting, whereas a lower pruning ratio may slows down convergence. Therefore, it is important to have an estimation of the pruning ratio.

7. Conclusions

We propose a novel method, Pruning Ternary Quantization, to construct sparse ternary weights by unifying L2 normalization, pruning, and quantization. The proposed method outperforms state-of-the-art model compression results. A major advantage of our method is the use of L2 normalization and pruning to enhance the model sparsity, performance, and compression capability at the same time. Pruning provides a range of size-accuracy trade-offs. The use of the balanced ternary weights $\{-1, 0, 1\}$ reduces a 3×3 convolutional kernel memory footprint from 288 bits to 15 bits. The 2-bit ResNet-18 ternary model with an accuracy of 65.43% and a size of 0.93 MB marks a major step forward in real-world applications.

Future work may combine our method with ternary activation quantization [6], [27] to further speed up the infer-

ence.

References

[1] Jose M Alvarez and Mathieu Salzmann. Compression-aware training of deep networks. In *Advances in Neural Information Processing Systems*, pages 856–867, 2017.

[2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

[3] Christopher M Bishop et al. *Neural networks for pattern recognition*. Oxford university press, 1995.

[4] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13169–13178, 2020.

[5] Miguel A Carreira-Perpinán and Yerlan Idelbayev. Model compression as constrained optimization, with application to neural nets. part ii: Quantization. *arXiv preprint arXiv:1707.04319*, 2017.

[6] Peng Chen, Bohan Zhuang, and Chunhua Shen. Fatnn: Fast and accurate ternary neural networks, 2020.

[7] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018.

[8] Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pages 3123–3131, 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 preprint arXiv:1602.02830*, 2016.

[10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019.

[11] L. Peter Deutsch, Jean-Loup Gailly, Mark Adler, L. Peter Deutsch, and Glenn Randers-Pehrson. Gzip file format specification version 4.3. RFC 1952, RFC Editor, May 1996. http://www.rfc-editor.org/rfc/rfc1952.txt.

[12] Zhen Dong, Zhewei Yao, Amir Gholami, Michael Mahoney, and Kurt Keutzer. Hawq: Hessian aware quantization of neural networks with mixed-precision, 2019.

[13] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019.

[14] Angela Fan, Pierre Stock, Benjamin Graham, Edouard Grave, Rémi Gribonval, Herve Jegou, and Armand Joulin. Training with quantization noise for extreme model compression. *arXiv preprint arXiv:2004.07320*, 2021.

[15] Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014.

[16] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: Efficient inference engine on compressed deep neural network. *ACM SIGARCH Computer Architecture News*, 44(3):243–254, 2016.

[17] 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.

[18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.

[19] 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.

[20] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1389–1397, 2017.

[21] Geoffrey E Hinton et al. Learning distributed representations of concepts. In *Proceedings of the eighth annual conference of the cognitive science society*, volume 1, page 12. Amherst, MA, 1986.

[22] Geoffrey E Hinton and Drew Van Camp. Keeping the neural networks simple by minimizing the description length of the weights. In *Proceedings of the sixth annual conference on Computational learning theory*, pages 5–13, 1993.

[23] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *arXiv preprint arXiv:1607.03250*, 2016.

[24] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990.

[25] Fengfu Li, Bo Zhang, and Bin Liu. Ternary weight networks. *arXiv preprint arXiv:1605.04711*, 2016.

[26] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.

[27] Yuhang Li, Xin Dong, and Wei Wang. Additive powers-of-two quantization: An efficient non-uniform discretization for neural networks. In *International Conference on Learning Representations*, 2019.

[28] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brecq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426*, 2021.

[29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.

[30] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *Advances in neural information processing systems*, pages 345–353, 2017.

[31] Hao Liu, Xiangyu Zhu, Zhen Lei, and Stan Z Li. Adaptiveface: Adaptive margin and sampling for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11947–11956, 2019.

[32] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 212–220, 2017.

[33] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.

[34] Julieta Martinez, Jashan Shewakramani, Ting Wei Liu, Ioan Andrei Bârsan, Wenyuan Zeng, and Raquel Urtasun. Permute, quantize, and fine-tune: Efficient compression of neural networks. *arXiv preprint arXiv:2010.15703*, 2021.

[35] Mark D McDonnell. Training wide residual networks for deployment using a single bit for each weight. *arXiv preprint arXiv:1802.08530*, 2018.

[36] Dmitry Molchanov, Arsenii Ashukha, and Dmitry Vetrov. Variational dropout sparsifies deep neural networks. *arXiv preprint arXiv:1701.05369*, 2017.

[37] Michael C Mozer and Paul Smolensky. Using relevance to reduce network size automatically. *Connection Science*, 1(1):3–16, 1989.

[38] James O' Neill. An overview of neural network compression. *arXiv preprint arXiv:2006.03669*, 2020.

[39] Steven J Nowlan and Geoffrey E Hinton. Simplifying neural networks by soft weight-sharing. *Neural computation*, 4(4):473–493, 1992.

[40] Igor[Pavlov. Lempel–ziv–markov chain algorithm, Oct 1998.

[41] 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.

[42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.

[43] Tim Salimans and Durk P Kingma. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in neural information processing systems*, pages 901–909, 2016.

[44] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.

[45] Oran Shayer, Dan Levi, and Ethan Fetaya. Learning discrete weights using the local reparameterization trick. *arXiv preprint arXiv:1710.07739*, 2017.

[46] Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Q-bert: Hessian based ultra low precision quantization of bert. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8815–8821, 2020.

[47] Sanghyun Son, Seungjun Nah, and Kyoung Mu Lee. Clustering convolutional kernels to compress deep neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 216–232, 2018.

[48] Pierre Stock, Armand Joulin, Rémi Gribonval, Benjamin Graham, and Hervé Jégou. And the bit goes down: Revisiting the quantization of neural networks. In *International Conference on Learning Representations (ICLR)*, 2020.

[49] Vincent Vanhoucke, Andrew Senior, and Mark Z Mao. Improving the speed of neural networks on cpus. 2011.

[50] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8612–8620, 2019.

[51] Bichen Wu, Yanghan Wang, Peizhao Zhang, Yuandong Tian, Peter Vajda, and Kurt Keutzer. Mixed precision quantization of convnets via differentiable neural architecture search. *arXiv preprint arXiv:1812.00090*, 2018.

[52] Jiaxiang Wu, Cong Leng, Yuhang Wang, Qinghao Hu, and Jian Cheng. Quantized convolutional neural networks for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4820–4828, 2016.

[53] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. https://github.com/facebookresearch/detectron2, 2019.

[54] Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael Mahoney, et al. Hawq-v3: Dyadic neural network quantization. In *International Conference on Machine Learning*, pages 11875–11886. PMLR, 2021.

[55] Dongqing Zhang, Jiaolong Yang, Dongqiangzi Ye, and Gang Hua. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *Proceedings of the European conference on computer vision (ECCV)*, pages 365–382, 2018.

[56] 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. *CoRR*, abs/1606.06160, 2016.

[57] Chenzhuo Zhu, Song Han, Huizi Mao, and William J Dally. Trained ternary quantization. *arXiv preprint arXiv:1612.01064*, 2016.