Towards Accurate Binary Neural Networks via Modeling Contextual Dependencies

Xingrun Xing¹, Yangguang Li², Wei Li³, Wenrui Ding¹, Yalong Jiang^{(⊠)1}, Yufeng Wang^{(⊠)1}, Jing Shao², Chunlei Liu¹, and Xianglong Liu¹

¹ Beihang University

{sy2002215, ding, all enyljiang, wyfeng, liuchunlei, xlliu }@buaa.edu.cn $^2~{\rm SenseTime~Group}$

liyangguang@sensetime.com, shaojing@senseauto.com ³ Nanyang Technological University wei.l@ntu.edu.sg

Abstract. Existing Binary Neural Networks (BNNs) mainly operate on local convolutions with binarization function. However, such simple bit operations lack the ability of modeling contextual dependencies, which is critical for learning discriminative deep representations in vision models. In this work, we tackle this issue by presenting new designs of binary neural modules, which enables BNNs to learn effective contextual dependencies. First, we propose a binary multi-layer perceptron (MLP) block as an alternative to binary convolution blocks to directly model contextual dependencies. Both short-range and long-range feature dependencies are modeled by binary MLPs, where the former provides local inductive bias and the latter breaks limited receptive field in binary convolutions. Second, to improve the robustness of binary models with contextual dependencies, we compute the contextual dynamic embeddings to determine the binarization thresholds in general binary convolutional blocks. Armed with our binary MLP blocks and improved binary convolution, we build the BNNs with explicit Contextual Dependency modeling, termed as BCDNet. On the standard ImageNet-1K classification benchmark, the BCDNet achieves 72.3% Top-1 accuracy and outperforms leading binary methods by a large margin. In particular, the proposed BCDNet exceeds the state-of-the-art ReActNet-A by 2.9% Top-1 accuracy with similar operations. Our code is available at https://github.com/Sense-GVT/BCDNet.

Keywords: binary neural network, contextual dependency, binary MLP

1 Introduction

Over the last decade, deep learning methods have shown impressive results for a multitude of computer vision tasks. However, these models require massive parameters and computation to achieve strong performance, hindering their application in practical embedded devices with limited storage and computing resources. Aiming at reduce computational cost in deep models, model compression



Fig. 1: Comparisons of binary operations: (a) The binary convolutions have inductive bias but limited local perceptions; (b) The binary token-mixing MLPs are sharing contextual perceptions but difficult to optimize; (c) Our proposed binary MLPs achieve inductive bias in short-range and explore long-range dependencies concurrently.

has drawn growing interests and developed into various methods, such as quantization [20,15,48], pruning [11,43], knowledge distillation [9,35], etc. Binary neural networks (BNNs) [23,10] can be viewed as an extreme case of low-bit quantization and become one of the most prevailing approaches to save computational cost at an extremely high rate. BNNs binarize the weights and activations to save at most $64 \times$ operations and $32 \times$ memory. However, current BNNs for vision tasks are built upon simple local convolutions, contributing to limited perception field in each layer. As such, it is of significant interest to explore more efficient contextual binary operations for BNNs.

Despite the advantage of inductive bias (e.g. locality) of convolutions for vision tasks, recent studies have shown improving interests in contextual operations, including the vision transformers (ViTs) and MLPs, that achieves stateof-the-art performance. Different from local convolutions, ViTs and MLPs treat images as tokens and model the dependencies across all tokens at once. To explore the binary contextual interaction in BNNs, we start by implementing a vanilla binary MLP-Mixer model, which we find is difficult to optimize compared with traditional binary convolutions. The model also fails to improve BNN accuracy with the help of contextual dependencies. This suggests local inductive bias is more crucial for BNNs compared with real-valued models.

Motivated by this, we aim to take advantage of contextual information while maintaining local inductive bias concurrently. A long-short range binary MLP operation is first proposed as an alternative to binary convolutions, where the short-range and long-range dependencies are introduced upon the standard binary convolutions, as shown in Fig. 1(c). We indicate the original location modeling as pointwise branch. The pointwise and short-range dependencies provide inductive bias similar to convolutions, meanwhile we break the local perceptions with long-range modeling. In summary, the three modeling patterns jointly perform a contextual interaction at each location. In this way, the proposed long-



Fig. 2: The performance of binary sparse MLPs express state-of-the-art binary CNNs; class activation maps (CAM) suggest that contextual binary MLPs are better to detect salient features, for example the 'beak' for the 'bird' class.

short range binary MLPs take the essence of convolutions and MLP-Mixer while alleviate their problems (Fig. 1 (a) and (b)).

Furthermore, we leverage the proposed long-short-range binary MLP blocks to replace some of the last convolutional blocks to improve the contextual perception field of BNNs while maintaining almost the same calculation. In addition, to access global perceptions in general binary convolutions, we further propose to compute the contextual dynamic embeddings to determine binarization thresholds. In binary convolutional blocks, we first aggregate global features with global average pooling, and then transform pooled features as global embeddings, which adapts binarization at inference. With the binary contextual MLPs and dynamic convolutions, we present a new binary networks named BCDNet to enhance the Binary Contextual Dependencies that is deficient in previous binary CNNs. As shown in Fig. 2, with similar architecture and computational cost, BCDNets achieves 2.9% accuracy gain against the state-of-the-art ReActNet-A model [23]. It also detects more discriminative features in the contextual range.

Our contributions are summarized as follows:

• We make the first attempt to break local binary convolutions and explore more efficient contextual binary operations.

• We design a long-short range binary MLP module to enhance modeling contextual dependencies in BNNs.

• To our best knowledge, the proposed BCDNet reports the best BNN performance on the large-scale ImageNet dataset currently.

2 Related Work

2.1 Binary Neural Networks

1-bit neural networks are first introduced by BNN [4] but encounters severe quantity error, especially in large scale datasets. To improve performance, some works focus to optimize binary weights. For example, XNOR-Net [32] reduce binarization error with scaling factors; RBNN [17] learns rotation matrixes to

reduce angular bias. SA-BNN [19] reduces flips in training; AdamBNN [22] explores influence of optimizers. Others try to modify architectures to improve expression ability. For instance, Bi-Real Net [24] adds more skip connections and achieves better performance; ReActNet [23] employs channel shift and the MobileNet-V1 [8] architecture; CP-NAS [25] adopts architecture search to explore connections. However, previous works only focus on CNN-based BNNs and ignore the performance upper-bound of binary convolutions. This work provides MLP based BNNs and exceed state of the arts.

2.2Vision Transformers and MLPs

Recently, transformers [49,16,5] are widely used to explore contextual dependencies. ViTs [6] first define to represent an image as 16×16 tokens using an embedding layer, and then stack self-attention and MLP layers. To improve performance of ViTs, DeiTs [38] propose a distillation token which benefits transformers a lot. More recently, Swin transformers [21], Twins [3] and NesT [47] explore various self-attentions to save computational cost, which makes transformers become the first choice in vision tasks. At the same time, works such as MLP-Mixers [36], ResMLP [37] simplify transformers using spatial and channel MLPs, but encounter overfeating. To this end, works such as S^2 MLPs [45], CycleMLPs [2] adopt surrounding shift operations to replace token-mixing MLPs. However, these works only focus on shifts and full precision circumstances. For BNNs, there still lacks in efficient contextual operations.

3 Background

In this section, we first review the binary convolution in recent binary convolutional neural networks [4,32,23]. Then, we follow the design in the binary convolution and make an initial attempt to binarize the MLP layer in vision MLP models [36,37,45,2]. We further conduct pilot studies to evaluate the performance of a binarized MLP-Mixer [36] and a ResNet+MLP architecture on ImageNet.

Binary Operations 3.1

Binary Convolution. By replacing arithmetic operations with more efficient bitwise operations, binary convolutions can drastically reduce memory allocation and accesses. Typically, full precision weights W and activations A are first binarized as W_b and A_b using the quantization function $Q_b(\cdot)$. As such, convolutions can be implemented as XNOR-Bitcounting operations [4,32]:

$$\texttt{BiConv}(\boldsymbol{A}) = \alpha Q_b(\boldsymbol{W})^T Q_b(\boldsymbol{A}) = \frac{\|\boldsymbol{W}\|_{\ell_1}}{k \times k \times c_{in}} \texttt{bitcount}(\boldsymbol{W}_b^T \oplus \boldsymbol{A}_b), \quad (1)$$

where k and c_{in} indicate the kernel size and input channels respectively, and α serves as the scale factor $\frac{\|\mathbf{W}\|_{\ell_1}}{k \times k \times c_{in}}$. Following [24], the quantization function

4

 $Q_b(\cdot)$ can be implemented in the form of a differentiable polynomial function:

$$Q_b(x) = \begin{cases} -1 & \text{if } x < -1\\ 2x + x^2 & \text{if } -1 \le x < 0\\ 2x - x^2 & \text{if } 0 \le x < 1\\ 1 & \text{otherwise} \end{cases}, \quad \frac{\partial Q_b(x)}{\partial x} = \begin{cases} 2 + 2x & \text{if } -1 \le x < 0\\ 2 - 2x & \text{if } 0 \le x < 1\\ 0 & \text{otherwise} \end{cases}$$
(2)

where gradients are approximated during back-propagation to optimize model weights W. Compared with real-valued convolutions with the 32-bit weight parameters, binary convolutions output 1-bit representations and can obtain up to $32 \times$ memory and $64 \times$ operation savings, leading to much faster test-time inference and lower power consumption.

Binary MLP. The MLP layer conducts regular matrix multiplications, which can be treated as 1×1 convolution that are repeatedly applied across either spatial locations or feature channels. For spatial dimensions, A single token-mixing MLP layer can be binarized in the form of:

$$\texttt{TM-BiMLP}(\boldsymbol{A}) = \frac{\|\boldsymbol{W}\|_{\ell_1}}{h \times w} \texttt{bitcount}(\boldsymbol{W}_b^T \oplus \boldsymbol{A}_b), \tag{3}$$

where h and w indicate the height and width of input features. Next, we binarize MLP-Mixer-S as shown in Fig. 3 (left). For binary MLP-Mixer, one MLP block is formulated by 2 binarized token-mixing MLP layers and 2 binarized token-wise MLP layers, and the expending ratio is 4. Later, we further explore a unified architecture composed of binary convolutions and MLPs. As shown in Fig. 3 (middle), our initial attempt is to replace the last 3 binary convolution blocks as binarized MLP-Mixer blocks according to similar operations, which aims at modeling contextual dependencies in high level. Because MLP-Mixer blocks have large expending ratio and multiple feed-forward layers, the binarized MLP-Mixer block presents similar or larger model size than the binary convolution.

3.2 Analysis

In order to evaluate the performance of token-mixing binary MLP layers introduced above, we binarize operations in MLP-Mixer-S [36] except the first and last layers. For comparison, we also evaluate binary ResNet-18 implemented by ReActNet and a binary ResNet+MLP architecture mentioned above, where the convolutional and MLP blocks are implemented by the ReActNet and binarized MLP-Mixer ones respectively. We train binary models on the public ILSVRC2021 ImageNet dataset [33] for 100 epochs. Complete training configurations appear in the Supplementary A.

Summary. As shown in Fig. 3, although MLP-Mixers achieves larger perception field than CNNs, the binary version cannot even achieve 40% top1 accuracy, which is worse than the first BNN model [4]. Compared with the binary MLPs that directly omit inductive bias, binary convolutions achieve best accuracy with

6 X. Xing et al.



Fig. 3: Initial implementations of binary MLP-Mixer (left) and a binary ResNet adding MLP architecture (middle). The comparison of performance with binary MLP-Mixer-S, ResNet-18 and ResNet+MLP is shown on the right.

the help of local perceptions and transformation invariance. Moreover, in the unified architecture of binary convolutions and MLPs, binary token-mixing layers still can not benefit from contextual perceptions, but much easier to train than stand-alone binary MLPs. Here we draw the following analysis:

• In the design of contextual binary operations, one need consider the inductive bias sufficiently. However, only preventing local dependencies in binary convolutions also barriers BNNs to capture long-range interactions. To this end, comprehensive design of local and contextual modeling is desired.

• In the design of BNN architectures, binary convolution stage is necessary, as stand-alone binary MLP-Mixers are difficult to optimize.

4 Method

Our proposed binary model is built on modeling contextual dependencies. An overview of the BCDNet architecture is illustrate in Fig. 4. We first describe the fundamentals of binary MLP block. Then, we introduce the contextual design of binary convolutions with dynamic embeddings. Finally, we present the complete BCDNet architecture with modeling contextual dependencies.

4.1 Binary MLP Block with Long-Short Range Dependencies

We first introduce three basic components in our binary MLP blocks, including a long-short range sampling scheme, a token reconstruction function and binary MLP operations, and then indicate the overall binary MLP block.

Long-Short Range Sampling. As shown in Fig. 1(c), we design a comprehensive sampling function which consists of three sampling patterns: the pointwise, short-range and long-range samplings. First, pointwise and short-range tokens



Fig. 4: Overall architecture of BCDNet, which consists of M binary convolution blocks and N binary MLP blocks. In practice, M = 11 and N = 9.

provide enough local inductive bias like convolutions. Second, long-range tokens are responding to explore contextual perceptions like self-attentions and token-mixing MLPs. In every block, one pointwise, four short-range, and four long-range tokens are sampled at each location, which compose a sparse contextual sampling. Given a sampling distance (r_1, r_2) in spatial, we use a sampling function to indicate the index of a token shifted from a location (x_1, x_2) :

$$S(r_1, r_2) = \{ y : y = (x_1 + r_1, x_2 + r_2) \}.$$
(4)

In each location, we donate sampled long-range, short-range and the pointwise tokens as $\{S(0, \pm h/2), S(\pm w/2, 0)\}$, $\{S(0, \pm 1), S(\pm 1, 0)\}$ and S(0, 0) respectively. Note that, when an index comes out of range, we just pad features from the other side and regard the index as recurrent (shown in Fig. 1 (c)).

Token Reconstruction across Samplings. Based on sampled tokens above, one simple way to aggregate them is concatenation. However, this will increase the channel numbers and calculations. To save calculations, we select and concatenate parts of channels of sampled tokens similar to [40], which keeps channels unchanged in reconstructed tokens. We choose the spatial shift operation [40] to reconstruct four sampled long-range tokens or four short-range tokens respectively, which is zero parameters and FLOPs. In each location, we can reconstruct a new short-range token A_b^S from parts of four sampled short-range tokens:

$$\mathbf{A}_{b}^{S} = Cat\{\mathbf{A}_{b}[0:c/4]_{S(-1,0)}, \mathbf{A}_{b}[c/4:c/2]_{S(1,0)}, \\
\mathbf{A}_{b}[c/2:3c/4]_{S(0,-1)}, \mathbf{A}_{b}[3c/4:c]_{S(0,1)}\},$$
(5)

where A_b is a binarized input token, $S(r_1, r_2)$ indexes its location, and [:] indexes its channels. Similarly, a new long-range token A_b^L is reconstructed from parts of four long-range tokens:

$$\boldsymbol{A}_{b}^{L} = Cat\{\boldsymbol{A}_{b}[0:c/4]_{S(-h/2,0)}, \boldsymbol{A}_{b}[c/4:c/2]_{S(h/2,0)}, \\ \boldsymbol{A}_{b}[c/2:3c/4]_{S(0,-w/2)}, \boldsymbol{A}_{b}[3c/4,c]_{S(0,w/2)}\}.$$
(6)

Here, we set the short-range interaction distance as one token and long-range as h/2 or w/2 tokens in space. In summary, we can obtain an pointwise token, a reconstructed short-range token and a long-range token $(\mathbf{A}_b, \mathbf{A}_b^S, \mathbf{A}_b^L)$ for each location, so that a sparse contextual token sampling is formulated.

Token-wise Binary MLP. In the following, we model three kinds of sampled tokens $(\mathbf{A}_b, \mathbf{A}_b^S, \mathbf{A}_b^L)$ independently using three token-wise binary MLPs. In MLP blocks, We follow ReActNets [23] and apply Sign(.) and RSign(.) to binarize weights and activations respectively. Given binarized weights \mathbf{W}_b and activations \mathbf{A}_b , the token-wise binary MLP is defined as a binarized fully connected layer across the channel dimension:

$$\texttt{TW-BiMLP}(\boldsymbol{A}) = \frac{\|\boldsymbol{W}\|_{\ell_1}}{c_{in}}\texttt{popcount}(\boldsymbol{W}_b^T \oplus \boldsymbol{A}_b), \tag{7}$$

where $\|\boldsymbol{W}\|_{\ell_1}$ indicates the L1 norm of real valued weights across the channel dimension. Similar to XNOR-Net [32], we compute a scaling factor $\frac{\|\boldsymbol{W}_b\|_{\ell_1}}{c_{in}}$ to minimize binarization error in weights. Next, the binary matrix product can be implemented by XNOR and bit-counting. The difference between Eq. 3 and 7 is fully connected dimensions.

Binary MLP Block. To explore block-wise contextual interactions, we propose the long-short range binary MLP block as an alternative to binary convolutional residual blocks in a network. Different from real-valued MLP-Mixers that aggregate all tokens indistinctively, our primary design is independently sampling pointwise, short-range, and long-range tokens according to a long-short sparse sampling rule. After a token reconstruction function to downsample channels, we use three token-wise binary MLPs to model three kinds of dependencies comprehensively. As shown in Fig. 4, we first binarize activations using $RSign(\cdot)$. In every location, three binary MLPs $\{P, S, L\}$ independently model three kinds of dependencies and are then added together:

$$\mathbf{A}^{i} = P\left(\mathbf{A}_{b}^{i-1}\right) + S\left(\mathbf{A}_{b}^{S}\right) + L\left(\mathbf{A}_{b}^{L}\right),\tag{8}$$

where $\{A_b, A_b^S, A_b^L\}$ indicates pointwise binary tokens, reconstructed short-range tokens and long-range tokens respectively. Like residual blocks, batch norm, skip connection and activation functions are also attached as in Fig. 4.

When the channel size is large (e.g., 512 and 1024), a binary MLP block consumes about 1/3 operations compared with a typical 3×3 binary convolution with the same input and output channels. Our binary MLPs can be easily adopted as stand-alone or plug-in building blocks in BNNs With similar configurations of regular layers.

4.2 Binary Convolutions with Dynamic Contextual Embeddings

To enhance model robustness with contextual dependencies, we improve the binary convolutions by introducing dynamic embedding to determine binarization thresholds $\alpha \in \mathbb{R}^c$ before the sign functions. Unlike ReActNets [23] that learns overall thresholds from the whole dataset without representing image-specific characteristics, we propose to infer binarization thresholds according to the input images dynamically. Specifically, in each residual block, we first obtain the mean value of each channel using the global average pooling (GAP), followed by the matrix product operation:

$$\alpha(\boldsymbol{A}) = \mathsf{GAP}(\boldsymbol{A})\boldsymbol{W}_1 + \boldsymbol{b}_\alpha,\tag{9}$$

where $\boldsymbol{A} \in \mathbb{R}^{c \times h \times w}$ is an input tensor, and $\{\boldsymbol{W}_1 \in \mathbb{R}^{c \times \frac{c}{4}}, \boldsymbol{b}_\alpha \in \mathbb{R}^{\frac{c}{4}}\}$ are learnable parameters. Note that, this matrix product increases negligible calculation due to the GAP function. Based on the contextual feature $\alpha(\boldsymbol{A})$, the channel-wise thresholds $\beta(\boldsymbol{A}) \in \mathbb{R}^c$ is obtained by a matrix product operation:

$$\beta(\boldsymbol{A}) = \alpha(\boldsymbol{A})\boldsymbol{W}_2 + \boldsymbol{b}_\beta,\tag{10}$$

with learnable parameters $\{W_2 \in \mathbb{R}^{\frac{c}{4} \times c}, b_\beta \in \mathbb{R}^c\}$. After merging $\beta(A)$ with the following sign function, we quantize inputs using instance specific thresholds:

$$Q_b(x) = \operatorname{sign}(x - \beta(\mathbf{A})) = \begin{cases} +1, x > \beta(\mathbf{A}) \\ -1, x \le \beta(\mathbf{A}). \end{cases}$$
(11)

To compensate the activation distribution variance caused by subtracting binarization thresholds, we learn additional embedding $\gamma(\mathbf{A})$ after binary convolutions:

$$\gamma(\boldsymbol{A}) = \alpha(\boldsymbol{A})\boldsymbol{W}_3 + \boldsymbol{b}_{\gamma},\tag{12}$$

where dynamic embeddings $\gamma(\mathbf{A}) \in \mathbb{R}^{c'}$ are also transformed by a matrix product with parameters $\mathbf{W}_3 \in \mathbb{R}^{\frac{c}{4} \times c'}$, $\mathbf{b}_{\gamma} \in \mathbb{R}^{c'}$. As such, the output tensor \mathbf{A}' is given by:

$$\mathbf{A}' = \mathbf{A}' + \gamma(\mathbf{A}),\tag{13}$$

In summary, two dynamic embeddings that model contextual dependencies are attached before and after the binary convolution respectively, formulating a dynamic binarization strategy for given inputs.

4.3 Overall Network

In Fig. 4, we outline the overall architecture of BCDNet. BCDNet is composed of two stages: the binary CNN-embedding stage with M binary convolutional blocks, and the binary MLP stage with N proposed binary MLP blocks. In a standard ViT or MLP architecture, the first embedding layer are implemented by a non-overlapping stride convolution. However, as suggested by [41], this patchify stem layer leads to optimization problems and can be sensitive to training strategies. In section 3.2, we find stand-alone contextual binary models are difficult to optimize. Instead, we adopt our binary MLP modules to enhance high level features. First, we replace several binary convolution blocks with our binary MLPs based on ReActNet-A, which we indicate as BCDNet-A. To exploit the

trade-off between convolution and MLP blocks, we perform a grid search under the same computational cost, and determine 11 convolutional blocks and 9 MLP blocks in Fig. 5. Second, we apply both improved binary convolution and MLP blocks and introduce a BCDNet-B model to further improve performance.

5 Experiments

5.1 Implementation Details

We train and evaluate the proposed BNNs on the large-scale ImageNet-1K dataset [33], which contains 1.28M images for training and 50K images for validation. We use 224×224 resolution images with standard data augmentations similar to Real-To-Binary [28]. Following ReActNet [23] and Real-to-Binary [28], we conduct a two-step training strategy with knowledge distillation and weight decay. We first train binary activation and real-valed weight networks for 300K iterations. Then, we initialize model weights from the first step, and train binary weight and activation models for 600K iterations. We use the AdamW [26] optimizer with a cosine learning rate schedule and batch size of 256. In two steps, we set the learning rate 1×10^{-4} and 5×10^{-5} and the weight decay 1×10^{-5} and 0, accordingly. For supervision, we use real labels with smooth rate 0.1 and distill from a full precision ResNet50. Detailed evaluation of distillation affect is shown in Supplementary B. For dynamic embeddings, we initialize W_2, W_3 and all biases as zeros and finetune another 300K iterations in Table 1.

Considering one binary MLP block has almost 1/3 operations of a 3×3 binary convolution block, BCDNet-A simply replace the last three 3×3 convolution blocks in ReActNet-A with 9 binary MLP blocks and keeps almost the same computational cost (Table 1); 1×1 binary convolutions in ReActNet-A remain unchanged. The replacing number is determined by a grid search in Fig. 5. Based on BCDNet-A, we further attach contextual embeddings in the CNN stage to improve binary convolutions and formulate BCDNet-B. We follow previous works [1] and use BOPs, FLOPs and OPs to measure the computational cost, where OPs=1/64BOPs + FLOPs. Note that the FLOPs indicate full precision MACs in BNNs as [24].

5.2 Comparisons with State of the Arts

In Table 1, we present comparisons of BCDNets with state-of-the-art binary or low-bit neural networks. To our best of knowledge, all BNNs are based on convolutional networks. BCDNets achieve the highest accuracy, which indicates the necessity to improve binary convolutions with contextual interactions. With similar architecture and computational cost, we observe that BCDNet-A is able to exceed ReActNet-A by a large margin (+2.4% top1 accuracy). Moreover, with additional dynamic embeddings, we find BCDNet-B further improves 0.5% top-1 accuracy with a little computational overhead, which validates the effectiveness of contextual dynamic embeddings. Also, BCDNet-B outperforms the full-precision MobileNet-V2 for the first time with a binary model. Table 1: Comparisons with state of the arts. "W/A" is bit-width of weights and activations. We underline ReActNet-A with consideration of our models share the similar architecture and operations. "†" indicates operations reported by ReActNets which may omit some small operations. We also report operations of BCDNets only considering conv. and fc. in "(.)".

Methods	W/A	BOPs $(\times 10^9)$	FLOPs $(\times 10^8)$	OPs $(\times 10^8)$	Top1 (%)	Top5 (%)
Mobile-V1 [8]	32/32	0	5.69	5.69	70.6	-
Mobile-V2 [34]	32/32	0	3.00	3.00	72.0	-
ResNet-18 [7]	32/32	0	18.14	18.14	69.3	89.2
BWN [32]	1/32	_	_	_	60.8	83.0
LQ-Net [46]	1/2	—	—	_	62.6	84.3
DoReFa [50]	2/2	—	—	—	62.6	84.4
SLB [44]	1/8	—	—	_	66.2	86.5
ABC-Net [18]	$(1/1) \times 5$	_	_	-	65.0	85.9
Bi-Real-34 [24]	1/1	3.53	1.39	1.93	62.2	83.9
Real-to-Bin [28]	1/1	1.68	1.56	1.83	65.4	86.2
FDA-BNN* [42]	1/1	—	—	_	66.0	86.4
SA-BNN-50 [19]	1/1	—	—	—	68.7	87.4
MeliusNet-22 [1]	1/1	4.62	1.35	2.08	63.6	84.7
MeliusNet-42 [1]	1/1	9.69	1.74	3.25	69.2	88.3
MeliusNet-59 [1]	1/1	18.3	2.45	5.25	71.0	89.7
ReActNet-A [23]	1/1	4.82	$0.31 \ (0.12\dagger)$	$1.06(0.87\dagger)$	69.4	_
ReActNet-B [23]	1/1	4.69	$0.61 (0.44\dagger)$	$\overline{1.34(1.17\dagger)}$	70.1	-
ReActNet-C [23]	1/1	4.69	$1.57(1.40\dagger)$	$2.30(2.14\dagger)$	71.4	-
BCDNet-A	1/1	4.82	$0.32 \ (0.12)$	1.08(0.87)	71.8 (+2.4)	90.3
BCDNet-B	1/1	4.82	$0.34 \ (0.14)$	$1.09 \ (0.89)$	$\overline{72.3\ (+2.9)}$	90.5

Table 2: Comparison with ResNet-18 architectures on the ImageNet dataset. W/A indicates bit-width of weights and activations respectively.

Methods	W/A	Top1 Acc.	Methods	W/A	Top1 Acc.
ResNet-18	32/32	69.3	RBNN [17]	1/1	59.9
XNOR-Net [32]	1/1	51.2	SA-BNN [19]	1/1	61.7
Bi-Real [24]	1/1	56.4	ReActNet (R18) $[23]$	1/1	65.5
Real-to-Bin [28]	1/1	65.4	FDA-BNN [42]	1/1	60.2
IR-Net [31]	1/1	58.1	FDA-BNN* [42]	1/1	66.0
SLB [44]	1/1	61.3	BCDNet-A (R18)	1/1	66.9
SLB [44]	1/8	66.2	BCDNet-B (R18)	1/1	67.9

Generalization to ResNet-18. The proposed binary MLP block can be easily used as plug-in blocks and improve contextual modeling ability in many BNNs. In Table 2, we compare results in the ResNet-18 (R18) architecture. We replace the last three convolutional layers by 9 binary MLP blocks so as to keep the same operations. Compared with ReActNet (R18), without changing the first three

Table 3: Evaluation on fine-grained classification benchmarks. "†" indicates the comparison with ImageNet pretraining.

Method	W/A	CUB-200	Flowers	Aircraft	Cars	Dogs	Avg.
ResNet18	32/32	52.3	44.2	66.8	39.2	47.1	49.9
BiRealNet18	1/1	19.5	27.0	18.6	14.8	17.9	19.6
ReActNet-A	1/1	29.9	29.8	18.4	19.4	18.7	23.2
BCDNet-A	1/1	34.3	25.0	18.1	25.5	20.7	24.7
ReActNet-A [†]	1/1	83.7	80.8	74.2	85.4	70.3	78.9
BCDNet-A†	1/1	90.6	83.0	81.1	91.1	75.2	84.2

stages, BCDNet-A (R18) improves 1.4% top-1 accuracy with binary MLPs. For BCDNet-B (R18), significant improvement (+1.0%) is obtained with the help of dynamic embeddings.

Generalization to Fine-grained Datasets. As shown in Table 3, we also report results on 5 fine-grained datasets: CUB-200-2011 [39], Oxford-flowers102[30], Aircraft[27], Stanford-cars[13], Stanford-dogs[12]. Detailed training strategies are reported in Supplementary C. These datasets have much less images than ImageNet but more concentrate on detailed differences of attributions. As shown in Table 3, BCDNet-A exceeds ReActNet-A 1.5% average accuracy, and are able to work well in most cases. We also report results with ImageNet pretraining. Pre-training helps BCDNet-A more than pure binary convolution based ReActNet-A. Modeling contextual dependencies helps BCDNet-A to exceed ReActNet-A by 5.3% on average, which shows BCDNet is more friendly to pretraining.

5.3 Ablation Studies

Method	Stages	Top1 Acc.	Setting	Block OPs $(\times 10^6)$	Top1 Acc.
ReActNet-18	_	61.89%	BiConv	$3.813 (\times 1.00)$	61.89%
BCDNet-18	4	63.58%	$BiMLP \times 1$	$1.505~(\times 0.39)$	61.32%
BCDNet-18	3, 4	63.47%	$BiMLP \times 2$	$3.011 (\times 0.79)$	62.64%
BCDNet-18	1, 2, 3, 4	63.10%	$BiMLP \times 3$	$4.516 (\times 1.18)$	63.58%

Table 4: Replacement stage settings Table 5: Replacement number settings

Efficiency of Binary MLPs. We first study where to replace binary convolutional blocks as proposed binary MLP blocks (BiMLPs) in a ReActNet (RN18), and then report the influence of different configuration of replacements. Training settings in Table 4, 5 are reported in Supplementary A. In Table 4, we gradually replace all 3×3 binary conv stages of a ReActNet RN18 (except downsampling layers). In practice, we use the same MLP sampling range settings as the last stage MLPs. Results exceed original ReActNet in each drop-in replacement. We



Fig. 5: Performance of replacing binary convolutional blocks with binary MLP blocks. We start from full convolutional ReActNet-A.

Table 6: Evaluation of different range modeling. We report checkpoint results

for different combinations of long- and short-range branches.

71.98/90.56

73.13/91.05

P-S-S

P-S-L

Traing time Stage1 Stage1+200K Stage1+400K Stage1+600K Bitwidth W/A 32/11/11/11/1Convolution 71.58/90.02 66.30/86.61 68.27/87.84 70.31/89.05 P-L-L 72.85/90.82 69.14/88.42 70.79/89.39 67.06/87.32

68.24/88.10

68.63/88.45

69.87/89.07

70.46/89.37

71.72/89.97

71.76/90.32

find only replacing the last stage is the most efficient. In Table 5, we replace every 3×3 binary convolutions with 1, 2 or 3 BiMLPs in the last stage of a ReActNet RN18 (except the downsampling layer). One BiMLP has 0.39 times single block OPs, while drops accuracy slightly; three BiMLPs have 1.18 times block OPs while significantly improve accuracy. In the MobileNet-V1 backbone, we simply replace one convolutional block by three binary MLP blocks once at a time in the ReActNet-A. Fig. 5 presents comparisons between MLPs and convolutions, where overall OPs are basically the same $(1.06 \times 10^8 \sim 1.08 \times 10^8)$. Due to the influence of randomness, $9 \sim 12$ MLP blocks make out the best for a ReActNet-A architecture. For simplicity, we just choose the setting of 11 (1,1,2,2,4,1) blocks

for each resolution) convolutional and 9 MLP blocks in other experiments.

Evaluation of Contextual Dependencies. Contextual dependencies come from long-short range sampling and the MLP architecture. In Table 6, we keep three branches in every binary MLP block and evaluate different sampling range combinations. 'P', 'S' and 'L' indicate the pointwise (original location) branch, the short-range branch and the long-range branch, with token sampling ranges of 0, 1 and half of the resolution respectively. We also report results of 200k, 400k and 600k training iterations. In Table 6, combination of long-short ranges ("P-S-L") with different receptive field achieves better performance. Compared with only long ranges ("P-L-L"), "P-S-L" has inductive bias and achieve better performance. Compared with only short ranges ("P-S-S"), "P-S-L" places long-

13

14 X. Xing et al.



Fig. 6: Binarization error of different branches in MLP layers.

range modeling and can be easier to converge in early training times (e.g., 200k, 400k iterations). In the MLP architecture, even though without long-ranges, "P-S-S" still exceeds binary CNNs because it decouples a 3×3 binary convolution to three BiMLPs and expands reception field in each layer (from 3×3 to 7×7).

Binarization Error Analysis. We define the average binarization error as:

$$error = \frac{1}{n} \sum \left| \frac{\|sign(\boldsymbol{W})\|_{\ell_1}}{c_{in}} sign(\boldsymbol{W}) - \boldsymbol{W} \right|$$
(14)

where $\frac{\|sign(\mathbf{W})\|_{\ell_1}}{c_{in}}$ follows Eq. 14; \mathbf{W} is the weight matrix in a binary MLP layer with n numbers. For BCDNet-A, we calculate the average binarization error of each MLP layer in Fig. 6. Note that, when binarization error increases, gradient estimation is inaccurate when backpropagation. We find in almost every block, errors in 3 branches follows a order of: pointwise>short-range>long-range. Binarization error is improved by decreasing of exchange distance, which indicates the long-range branch is easier to optimize than the short-range branch.

6 Conclusion

This work improves binary neural networks by modeling contextual dependencies. A long-short range binary MLP module is proposed to explore comprehensive contextual perceptions, which can be an efficient alternative to traditional binary convolutions. Equipped with our binary MLPs and improved binary convolutions, BCDNets exceed state of the arts significantly.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (U20B2042), the National Natural Science Foundation of China (62076019) and Aeronautical Science Fund (ASF) of China (2020Z071051001).

References

- Bethge, J., Bartz, C., Yang, H., Chen, Y., Meinel, C.: Meliusnet: Can binary neural networks achieve mobilenet-level accuracy? arXiv preprint arXiv:2001.05936 (2020)
- Chen, S., Xie, E., Ge, C., Liang, D., Luo, P.: Cyclemlp: A mlp-like architecture for dense prediction. arXiv preprint arXiv:2107.10224 (2021)
- 3. Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C.: Twins: Revisiting spatial attention design in vision transformers (2021)
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., Bengio, Y.: Binarized neural networks: Training deep neural networks with weights and activations constrained to+ 1 or-1. arXiv preprint arXiv:1602.02830 (2016)
- Dai, Z., Cai, B., Lin, Y., Chen, J.: Up-detr: Unsupervised pre-training for object detection with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1601–1610 (2021)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2020)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. arXiv preprint arXiv:1704.04861 (2017)
- Ji, M., Shin, S., Hwang, S., Park, G., Moon, I.C.: Refine myself by teaching myself: Feature refinement via self-knowledge distillation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10664– 10673 (2021)
- Jiang, X., Wang, N., Xin, J., Li, K., Yang, X., Gao, X.: Training binary neural network without batch normalization for image super-resolution. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1700–1707 (2021)
- Joo, D., Yi, E., Baek, S., Kim, J.: Linearly replaceable filters for deep network channel pruning. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8021–8029 (2021)
- Khosla, A., Jayadevaprakash, N., Yao, B., Li, F.F.: Novel dataset for fine-grained image categorization: Stanford dogs. In: Proc. CVPR workshop on fine-grained visual categorization (FGVC). vol. 2. Citeseer (2011)
- Krause, J., Stark, M., Deng, J., Fei-Fei, L.: 3d object representations for finegrained categorization. In: Proceedings of the IEEE international conference on computer vision workshops. pp. 554–561 (2013)
- 14. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
- Lee, J., Kim, D., Ham, B.: Network quantization with element-wise gradient scaling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6448–6457 (2021)
- Lin, K., Wang, L., Liu, Z.: End-to-end human pose and mesh reconstruction with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1954–1963 (2021)
- Lin, M., Ji, R., Xu, Z., Zhang, B., Wang, Y., Wu, Y., Huang, F., Lin, C.W.: Rotated binary neural network. Advances in Neural Information Processing Systems 33 (2020)

- 16 X. Xing et al.
- Lin, X., Zhao, C., Pan, W.: Towards accurate binary convolutional neural network. Advances in neural information processing systems **30** (2017)
- Liu, C., Chen, P., Zhuang, B., Shen, C., Zhang, B., Ding, W.: Sa-bnn: Stateaware binary neural network. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2091–2099 (2021)
- Liu, X., Ye, M., Zhou, D., Liu, Q.: Post-training quantization with multiple points: Mixed precision without mixed precision. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8697–8705 (2021)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
- 22. Liu, Z., Shen, Z., Li, S., Helwegen, K., Huang, D., Cheng, K.T.: How do adam and training strategies help bnns optimization? In: International Conference on Machine Learning. PMLR (2021)
- Liu, Z., Shen, Z., Savvides, M., Cheng, K.T.: Reactnet: Towards precise binary neural network with generalized activation functions. In: European Conference on Computer Vision (ECCV) (2020)
- Liu, Z., Wu, B., Luo, W., Yang, X., Liu, W., Cheng, K.T.: 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). pp. 722–737 (2018)
- Li'an Zhuo, B.Z., Chen, H., Yang, L., Chen, C., Zhu, Y., Doermann, D.: Cp-nas: Child-parent neural architecture search for 1-bit cnns. IJCAI (2020)
- Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
- Maji, S., Rahtu, E., Kannala, J., Blaschko, M., Vedaldi, A.: Fine-grained visual classification of aircraft. arXiv preprint arXiv:1306.5151 (2013)
- 28. Martinez, B., Yang, J., Bulat, A., Tzimiropoulos, G.: Training binary neural networks with real-to-binary convolutions. arXiv preprint arXiv:2003.11535 (2020)
- 29. Netzer, Y., Wang, T., Coates, A., Bissacco, A., Wu, B., Ng, A.Y.: Reading digits in natural images with unsupervised feature learning (2011)
- Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: 2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing. pp. 722–729. IEEE (2008)
- Qin, H., Gong, R., Liu, X., Shen, M., Wei, Z., Yu, F., Song, J.: Forward and backward information retention for accurate binary neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2250–2259 (2020)
- Rastegari, M., Ordonez, V., Redmon, J., Farhadi, A.: Xnor-net: Imagenet classification using binary convolutional neural networks. In: European conference on computer vision. pp. 525–542. Springer (2016)
- 33. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International journal of computer vision 115(3), 211–252 (2015)
- Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018)
- 35. Song, L., Wu, J., Yang, M., Zhang, Q., Li, Y., Yuan, J.: Robust knowledge transfer via hybrid forward on the teacher-student model. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 2558–2566 (2021)

- Tolstikhin, I., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Keysers, D., Uszkoreit, J., Lucic, M., et al.: Mlp-mixer: An all-mlp architecture for vision. arXiv preprint arXiv:2105.01601 (2021)
- Touvron, H., Bojanowski, P., Caron, M., Cord, M., El-Nouby, A., Grave, E., Joulin, A., Synnaeve, G., Verbeek, J., J'egou, H.: Resmlp: Feedforward networks for image classification with data-efficient training. ArXiv abs/2105.03404 (2021)
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & amp; distillation through attention. In: International Conference on Machine Learning. vol. 139, pp. 10347–10357 (July 2021)
- Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
- Wu, B., Wan, A., Yue, X., Jin, P., Zhao, S., Golmant, N., Gholaminejad, A., Gonzalez, J., Keutzer, K.: Shift: A zero flop, zero parameter alternative to spatial convolutions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 9127–9135 (2018)
- Xiao, T., Dollar, P., Singh, M., Mintun, E., Darrell, T., Girshick, R.: Early convolutions help transformers see better. Advances in Neural Information Processing Systems 34 (2021)
- Xu, Y., Han, K., Xu, C., Tang, Y., Xu, C., Wang, Y.: Learning frequency domain approximation for binary neural networks. Advances in Neural Information Processing Systems 34, 25553–25565 (2021)
- 43. Yamamoto, K.: Learnable companding quantization for accurate low-bit neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5029–5038 (2021)
- 44. Yang, Z., Wang, Y., Han, K., Xu, C., Xu, C., Tao, D., Xu, C.: Searching for low-bit weights in quantized neural networks. Advances in neural information processing systems 33, 4091–4102 (2020)
- Yu, T., Li, X., Cai, Y., Sun, M., Li, P.: S2-mlp: Spatial-shift mlp architecture for vision. arXiv preprint arXiv:2106.07477 (2021)
- 46. Zhang, D., Yang, J., Ye, D., Hua, G.: Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In: Proceedings of the European conference on computer vision (ECCV). pp. 365–382 (2018)
- 47. Zhang, Z., Zhang, H., Zhao, L., Chen, T., Pfister, T.: Aggregating nested transformers. arXiv preprint arXiv:2105.12723 (2021)
- Zhao, K., Huang, S., Pan, P., Li, Y., Zhang, Y., Gu, Z., Xu, Y.: Distribution adaptive int8 quantization for training cnns. In: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (2021)
- Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6881–6890 (2021)
- 50. Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., Zou, Y.: Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. arXiv preprint arXiv:1606.06160 (2016)