

BCNN: A Binary CNN With All Matrix Ops Quantized To 1 Bit Precision

Arthur J. Redfern Lijun Zhu
Texas Instruments, Inc
12500 TI Boulevard, Dallas TX 75243
{redfern, lijun}@ti.com

Molly K. Newquist
Georgia Institute of Technology
North Avenue NW, Atlanta, GA 30332
mnewquist3@gatech.edu

Abstract

This paper describes a CNN where all CNN style 2D convolution operations that lower to matrix matrix multiplication are fully binary. The network is derived from a common building block structure that is consistent with a constructive proof outline showing that binary neural networks are universal function approximators. 71.24% top 1 accuracy on the 2012 ImageNet validation set was achieved with a 2 step training procedure and implementation strategies optimized for binary operands are provided.

1. Introduction

CNNs have achieved wide spread success in vision applications. The feature encoder portion of an image classification network typically forms the backbone of networks designed for semantic segmentation, object detection, object segmentation, depth estimation and motion estimation. Minimizing the complexity of the backbone while achieving a high level of accuracy is a key to enabling the use of CNNs in resource constrained environments.

The complexity of CNNs is dominated by the CNN style 2D convolution operator. This operator takes a 3D input tensor and a 4D weight tensor and produces a 3D output tensor. CNN style 2D convolution can be lowered to matrix matrix multiplication via appropriate arrangement of tensors into matrices [3].

Remaining network operators such as average and max pooling, depth wise 2D convolution, unbatched linear layers, bias, scale and various nonlinearities reduce to either pointwise operations, vector vector operations or vector matrix operations. These operations have a lower arithmetic intensity and / or fewer parameters and tend to contribute substantially less to the complexity of common network designs.

Binarizing both the inputs and weights is a method for drastically reducing the implementation complexity of CNN style 2D convolution relative to variants with 32 bit float (common training precision) or 8 bit fixed point (com-

mon inference precision) inputs and weights. While binarization benefits the implementation complexity, it typically results in a model with lower accuracy. As such, there is an interest in the design and training of high accuracy CNNs which use fully binary CNN style 2D convolution.

This paper:

- Describes a network structure built from a common building block where all CNN style 2D convolution operations are implemented with binary inputs and binary weights; this includes the network stem and down sampling portions of the network
- Provides the outline of a constructive universal function approximator proof for binary neural networks which is consistent with the building block design
- Uses an existing 2 step training method to achieve 71.24% top 1 accuracy on the 2012 ImageNet validation set; to the best of the authors' knowledge, this is the highest accuracy on the 2012 ImageNet validation set for a network with all matrix ops fully binary

2. Related work

There are a variety of approaches to minimizing CNN complexity while achieving a high level of accuracy. In practice, implementation complexity is a joint function of network design, software mapping and hardware architecture. As such, consider the network designs mentioned in this section as a starting point that need to be further considered in the context of a specific software and hardware architecture.

2.1. Filter sizes

Initial CNN designs used different filter sizes in the CNN style 2D convolutional layers. Early attempts at complexity reduction replaced a single layer with a larger filter with multiple layers with 3x3 filters, e.g., replacing a single 5x5 layer with two 3x3 layers. While not mathematically equivalent, the receptive field size is maintained while the number of parameters and MACs is reduced [17]. A somewhat hidden tradeoff, though, is the potential latency and feature

map data movement increase, the consequences of which are dependent on the specific hardware architecture.

2.2. Sparsity

A logical next step in complexity reduction was the realization that the combined spatial and channel mixing of 3x3 CNN style 2D convolution can be achieved with a depth wise 3x3 2D convolution layer for spatial mixing followed by a separable 1x1 CNN style 2D convolution layer for channel mixing. As before, while not mathematically equivalent, spatial and channel mixing are still enabled and the parameters and MACs are reduced [10].

Depth wise 2D convolution is an extreme form of feature map grouping which can be viewed as structured sparsity where only a subset of input channels are connected to each output channel [21]. Other forms of sparsity exist beyond grouped convolution and can be used to reduce complexity. The more random the sparsity, the more opportunity there is for accuracy vs complexity optimization, but there are also more challenges from an implementation perspective.

2.3. Building block design

The above approaches are typically part of the design of building blocks used to construct networks. The optimal specification of building blocks in terms of input size, channel width and network depth is a key component of efficient network design. Historically, channels doubled at each down sampling level and building blocks at level 4 were repeated the most as they provided a good mix of receptive field size increase while not increasing the number of parameters as much as level 5. Recently, large experiments have been used to determine optimal channel width increases and building block repeat patterns for a variety of network complexity design points [15, 18].

2.4. Quantization

Quantization can be applied to all of the afore mentioned methods to reduce the precision at which computation is performed. While quantization does not change the theoretical complexity of computations, it has a large effect on the practical implementation complexity.

Memory and data movement, assuming that size changes do not change memory hierarchy locations, scale with the number of bits: if the number of bits per element increases by 2x the memory size increases by 2x. Integer addition and comparison operations scale with the number of bits used to represent both inputs: if the number of bits in both operands increases by 2x the implementation complexity increases by $\sim 2x$.

Integer multiplication, however, scales with the number of bits used to represent each input: if the number of bits in both operands increases by 2x the implementation complexity increases by $\sim 4x$. As multiplication is more complex

from an implementation perspective than addition or comparison, the importance of the precision in layers with lots of multiplications is critical. Hence the focus on the quantization of the CNN style 2D convolution operation.

Historically, CNN training is done at 32 bit float precision with 24 bits for the sign and mantissa (resolution) and 8 bits for the exponent (range). Networks can be relatively easily quantized from 32 bit float to bfloat16, with 8 bits for the sign and mantissa and 8 bits for the exponent, as networks tend to be less sensitive to resolution above some threshold and the range stays the same. The current sweet spot for practical CNN inference is 8 bit fixed point, ideally with range optimization and fixed point quantization built into some portion of the training. Other combinations of precision for the inputs and weights have been used beyond those mentioned above.

2.5. Binary CNN design

The most efficient potential quantization is the extreme end point: 1 bit for the inputs and 1 bit for the weights. An abbreviated arch of papers using binary CNN style 2D convolution includes:

- XNOR-Net used binary CNN style 2D convolution with output scaling as an approximation of real CNN style 2D convolution [16]; extensions such as XNOR-Net++ looked at different options for output scaling to improve accuracy [1]
- ABC-Net improved the binary CNN style 2D convolution approximation of real CNN style 2D convolution via a 2 level parallel structure [11]; different branches in the parallel structure used different biases before the sign function to create different binarizations
- Bi-Real Net, consistent with the data processing inequality, introduced real identity connections around binary CNN style 2D convolution to minimize information loss [13]; most all high performing binary CNNs have since adopted this strategy as it provides a good tradeoff of accuracy vs complexity for a range of cases
- Group-Net considered different serial and parallel combinations of binary CNN style 2D convolution with identity connections [25]; philosophically, similar to how ABC-Net generalized XNOR-Net, Group-Net generalized Bi-Real Net
- MeliusNet focused on building a better building block via a DenseNet style addition of channels followed by an improvement of those new channels [2]; however, both the network stem and down sampling stages remained real valued
- ReActNet modified the MobileNetV1 structure via making the depth wise separable convolution operations binary 3x3 CNN style 2D convolution (vs depth

wise 2D convolution) followed by binary 1x1 CNN style 2D convolution [12]; it also introduced trainable biases before the sign and activation function to allow for learnable binarizations and included a binary down sampling stage with channel replication; however, the network stem remained real valued

2.6. Binary CNN training

The not so hidden challenge of working with binary CNNs is training. Specifically, almost everywhere there's no gradient propagation through sign functions during automatic differentiation with reverse mode accumulation. Approaches for addressing this include using real teacher networks and feature map and / or output distribution matching, different functions in the forward and backward paths and / or gradually converging from a real network to a binary network.

3. Theory

A 3 layer real neural network can approximate arbitrarily closely any continuous function on a compact subset of \mathbb{R}^K [5]. The universal function approximator property of neural networks underlies their successful application to a wide variety of problems. As such, it's useful to understand if binary neural networks maintain this same property.

3.1. Proof outline

The following is the outline of a constructive proof, reminiscent of the Reimann integral, showing that a 3 layer binary neural network with a particular building block structure maintains the universal function approximator property. The proof outline will initially consider the simpler case of showing that a binary neural network can approximate arbitrarily closely a compact function that maps a scalar input $x \in [0, 1]$ to a continuous scalar output $0 \leq f(x) < \infty$, then remove restrictions.

Consider the 3 layer binary neural network structure shown in Figure 1 with input x swept from 0 to 1.

Layer 1 uses $2(1/d - 1)$ input branches to create $1/d - 1$ branches, each generating a rectangle of height 1 and width d centered at $d, 2d, \dots, 1 - d$, respectively. For convenience this omits a $1/2$ width rectangle at the beginning and end of the unit interval which can be added if more formality is desired.

Layer 2 expands the branch centered at d to $\text{round}(Qf(d))$ sub branches, the branch centered at $2d$ to $\text{round}(Qf(2d))$ sub branches, \dots , and the branch centered at $1 - d$ to $\text{round}(Qf(1 - d))$ sub branches. Each sub branch generates a rectangle of height 1 and width d centered at the same location as the branch.

Layer 3 combines all of the sub branches of all of the branches and scales the result to generate the final output.

Let $d \rightarrow 0$ to approximate all $x \in [0, 1]$ and $Q \rightarrow \infty$ to approximate the corresponding value $f(x)$ to complete the initial portion of the proof outline.

Restrictions in the initial proof outline can be removed as follows:

- For $x \in [x_{\min}, x_{\max}]$ where x_{\min} and x_{\max} are finite, choose branch locations $x_{\min} + d, x_{\min} + 2d, \dots, x_{\max} - d$ to uniformly tile the full interval
- For $f(x)$ with a countable number of discontinuities, center an additional branch at each discontinuity to handle the jump there exactly
- For $f_{\min} \leq f(x) \leq f_{\max}$ where x_{\min} and x_{\max} are finite and each can be negative, 0, or positive, approximate the non negative function $f(x) + f_{\min}$ then include a bias term at the end from an additional layer as $-f_{\min}$
- For a vector input $x \in \mathbb{R}^K$ replicate layer 1 for each element of the vector, in layer 2 generate $\text{round}(Q(f(d)))$ K dimensional rectangles for each combination of layer 1 branches and add the K dimensional rectangles together in layer 3
- For a vector output $f(x) \in \mathbb{R}^M$ replicate layer 3 for each of the M outputs

This completes the outline of the constructive proof that the 3 layer binary neural network structure in Figure 1 is a universal function approximator.

3.2. Implications

In the corresponding real valued constructive proof, $d \rightarrow 0$ implies more branches. From a practical perspective, more branches are needed in regions of higher function variation and fewer branches are needed in regions of less function variation. This matches what would be expected by traditional sampling theory.

The binary valued constructive proof adds the additional constraint $Q \rightarrow \infty$, which implies even more branches are needed in the binary neural network case (being somewhat loose with the concept of infinity). From a practical perspective, to maintain a given accuracy in the function approximation more sub branches are needed for functions with large range, fewer sub branches are needed by functions with small range. This matches what would be expected by traditional quantization theory, and also implies that a binary neural network will tend to have more channels than a similar accuracy real valued neural network.

4. Design

This section describes the design of an ImageNet image classification CNN where all operations that can be lowered to matrix matrix multiplication are fully binary.

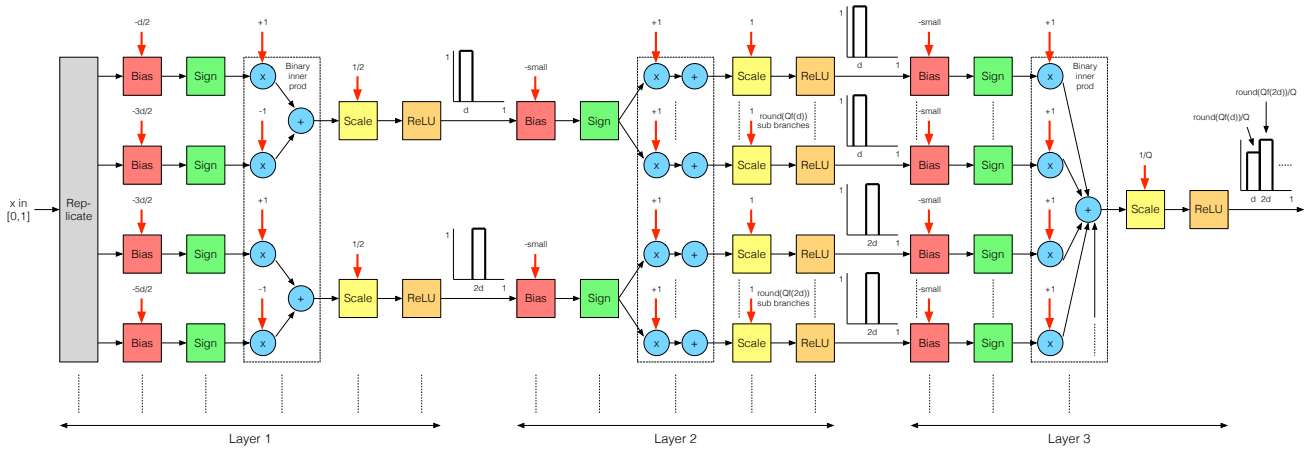


Figure 1. A 3 layer binary neural network that is a universal function approximator.

4.1. Network structure

The ImageNet image classification CNN used in this paper is shown in Figure 2. While it's a relatively typical parameterized structure with 5 levels of down sampling and multiple repeated building blocks, there are a few features to note.

In the data loader a 4th intensity channel is created from the initial 3 RGB channels. Additionally, places in the network that increase the number of channels by an integer factor were configured to always choose a power of 2 integer factor. The result is that the number of channels for each level is a power of 2, a convenient value for hardware implementations to work with. Note that this is not a requirement.

The stem uses the same building block structure as levels 1 - 5. While this not unheard of, it's a bit less common than a more traditional option like using real 3x3 CNN style 2D convolution. However, it was done to keep all CNN style 2D convolutions in the network fully binary.

4.2. Building block

The building block includes an identity path and a residual path composed of 1x1, 3x3 and 1x1 convolution modules.

4.2.1 Convolution module

The 1x1 convolution modules include their own identity path and residual path with learnable bias and sign for the binarization of activations, fully binary 1x1 CNN style 2D convolution for mixing across channels, batch norm for per channel scale and bias, convolution module identity path addition and PReLU. This follows a Bi-Real Net style approach to the identity connection and a ReActNet style approach to the learnable bias before sign.

Binary 1x1 CNN style 2D convolution implies that each input feature is equally present (+1) or not present (-1) and equally positively (+1) or negatively weighted (-1). Larger numbers of input and output channels allow more nuance in the values produced by this operation.

The 3x3 convolution module includes real 3x3 depth wise 2D convolution (a vector op) for spatial mixing, batch norm and PReLU. It's reasonable to ask: Why not use binary 3x3 CNN style 2D convolution for this module? Using a binary filter for spatial mixing to create a new center value in a 3x3 region as a +-1 combination of the center and neighboring values is not consistent with the typical shape of real filters. However, successful binary CNNs have made use of binary 3x3 CNN style 2D convolution. But for the above reason, it was not used here. On the upside, this operation reduces to a vector op and the number of parameters and MACs is low, so it is in keeping with the theme of implementation complexity reduction.

Inside the binary convolution modules it's possible to use P parallel residual branches, each with a different learned bias, binary weights and batch norm parameters as shown in Figure 4. This is in the same spirit as ABC-Net and Group-Net.

This parallel approach points to a possible alternative for binary $F \times F$ CNN style 2D convolution. Using the parallel building block structure in Figure 4, let there be $P = F^2$ parallel branches, with an appropriate feature map shift and 0 pad applied before the bias operation of each. This effectively turns binary $F \times F$ CNN style 2D convolution into F^2 binary 1x1 CNN style 2D convolution operations, where each parallel branch benefits from a real scale in the batch norm before combination to enable learning filter weights closer to what would be found in a real network.

4.2.2 Channel increase

Because of the identity connections around binary CNN style 2D convolution and the fully grouped nature of the real 3x3 depth wise 2D convolution, it's not convenient to increase the number of channels with the convolution operations in the building block. Instead, when it's necessary to modify the number of channels, channels are simply replicated by an integer factor $R > 1$ at the start of the building block. In practice, no extra memory is needed for this as the replication can be done by bookkeeping.

Theoretically, R could be chosen as a non integer value larger than 1 such that not all the channels are replicated. Likewise, R could be chose as a non integer value less than 1 with the behavior defined as averaging in the channel dimension. Channel replication was also used as part of the

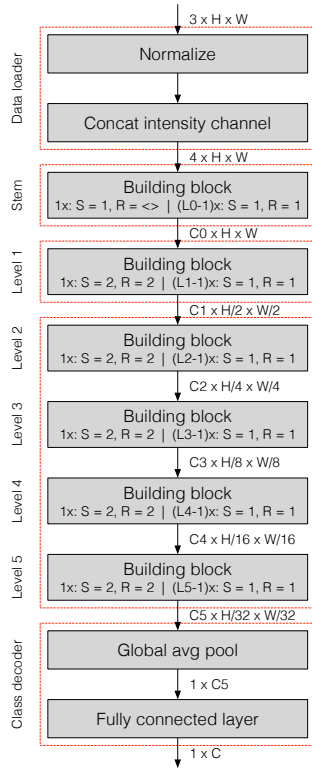


Figure 2. An ImageNet network structure based on a common building block. Each of the 5 levels has 1 common building block configured for channel replication and down sampling followed by $L_x - 1$ common building blocks with no channel replication or down sampling, where L_x is the total number of common building blocks in level x . The stem uses the common building block configured for channel replication only (no down sampling), the selection of the replication factor here determines the number of channels for the remainder of the network. The decoder uses global average pooling to spatially aggregate features and abstract away the input image size followed by a fully connected layer to classify.

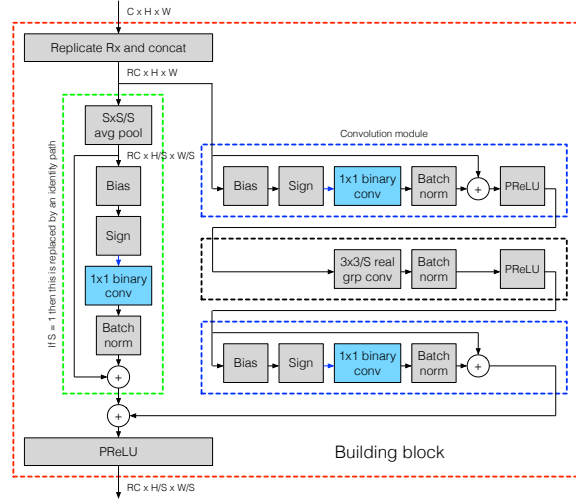


Figure 3. The stem and level 1 - 5 building block structure capable of integer factor channel increase and integer spatial down sampling; note that all CNN style 2D convolution operations that can be lowered to matrix matrix multiplication are fully binary, the depth wise 3x3 2D convolution that can be lowered to vector matrix multiplication is real.

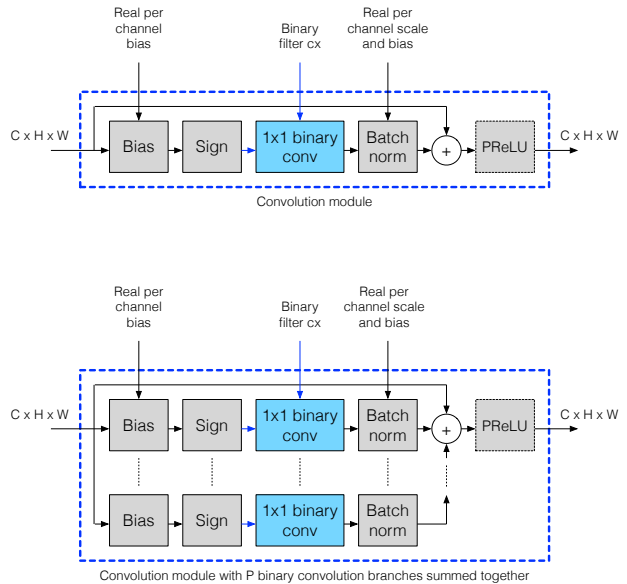


Figure 4. Convolution module (top) and alternative convolution module with P parallel binary convolution residuals (bottom).

ReActNet channel increase strategy and is in line with the universal function approximation proof outline.

4.2.3 Spatial down sampling

When it's necessary to spatially down sample by a factor of S , in a spirit similar to common modifications to

ResNet [9], the identity path is replaced by $S \times S / S$ average pooling for spatial mixing followed by a binary 1×1 convolution module for channel mixing. Likewise, in the residual path, $3 \times 3 / S$ depth wise convolution is used in the real 3×3 convolution module for spatial mixing (followed by channel mixing in the subsequent binary convolution module). A positive aspect of this strategy is that it allows the identity connections inside the binary 1×1 convolution modules to remain identity connections.

4.2.4 Connections to real networks

Connecting with recent real network building block designs, if the binary 1×1 convolution modules were replaced by real 1×1 convolution modules and larger group sizes were used in the real 3×3 convolution module, then this structure would effectively be a RegNetX building block [15]. Many other network building blocks also use an identity path with variants of a 1-3-1 residual structure initially made famous by ResNet [8].

5. Training

Both 1 step (direct training of binary weights) and 2 step (step 1 training of real weights followed by step 2 binarization and training of binary weights) methods of training were tried. 2 step training provided a slightly higher accuracy, though it's not clear if this is a fundamental to real pre training or just a consequence of the specific hyper parameters used in the training. Both 1 and 2 step methods borrowed from strategies used by [12, 13, 14], the 2 step method is described below.

Note that none of the recent training methods commonly employed to boost ImageNet image classification accuracy for real networks were used [4, 9, 19, 20, 22, 23, 24]. Some combination of these would likely improve the reported accuracy results.

5.1. Step 1: Binary activations real weights

In step 1 of training, the input to binary convolution is binary but the weights were allowed to take on real values (they will be binarized in step 2). Step 1 hyper parameters included:

- Batch size 128 (4 GPUs, 32 images per GPU)
- Random resized crops to $3 \times 224 \times 224$, random horizontal flip and 0 mean unit variance normalization
- Distributional loss with output pmf targets provided by ResNet34
- Adam optimizer
- 5 epochs of linear warm up followed by 20 epochs of half wave cosine decay
- A max learning rate of 0.0005, initial learning rate scale of 0.01x and final learning rate scale of 0.001x

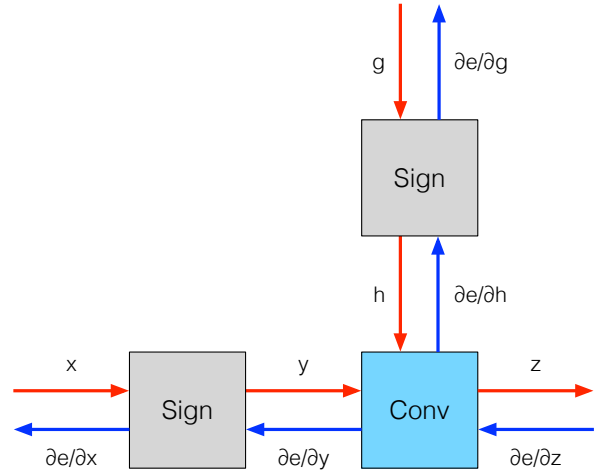


Figure 5. Zooming in on the (binary) convolution operation to look at the forward flow of weights and activations (red) along with the backward flow of associated error gradients (blue); as both y and h are binary, the convolution operator is a binary convolution operator.

- L2 weight decay applied to the convolution parameters with a weighting of $1e-5$

Step 1 was used to get weights in the ballpark, not to maximize accuracy of binary activations and real weights (only $5 + 20 = 25$ epochs were used). An alternative weight decay that may degrade the accuracy of step 1 but provide a better starting point for step 2 is described in [6].

5.2. Step 2: Binary activations binary weights

Weights were transferred from the last checkpoint in step 1 to initialize step 2. The real weights for the binary convolution module were used for gradient accumulation in the backward pass, but passed through a sign function to enable fully binary convolution in the forward pass.

Step 2 used the same hyper parameters as step 1 with the following exceptions:

- 5 epochs of linear warm up followed by 100 epochs of half wave cosine decay
- No L2 weight decay

Typically, weight decay is needed during real network training to improve generalization and achieve high accuracies on ImageNet. Binarization has a similar regularizing effect and prevents any of the CNN style 2D convolution weights from becoming large. Clipping of the associated real accumulated weights was not used.

5.3. The sign function

Figure 5 zooms in on the binary convolution operation showing the sign operation in the data path and implicit sign

operation applied to real weights to create binary weights. A key for training binary CNNs is propagating the sensitivity of the error with respect to feature maps and weights through the sign operation as the derivative of the output of the sign operation with respect to the input is 0 almost everywhere.

The training strategy used here followed the strategy in [12] with the backward function decoupled from the forward function through the sign operator. The backward function used to propagate the sensitivity of the error with respect to the activations is

$$\begin{aligned} \partial y / \partial x &= -2|x| + 2, & x \in [-1, 1] \\ &= 0, & \text{otherwise} \end{aligned} \quad (1)$$

which corresponds to (unused) piecewise forward function:

$$\begin{aligned} y &= -1, & x < -1 \\ &= x^2 + 2x, & -1 \leq x < 0 \\ &= -x^2 + 2x, & 0 \leq x < 1 \\ &= 1, & 1 \leq x \end{aligned} \quad (2)$$

The backward function used to propagate the sensitivity of the error with respect to the weights is:

$$\begin{aligned} \partial y / \partial x &= 1, & x \in [-1, 1] \\ &= 0, & \text{otherwise} \end{aligned} \quad (3)$$

which corresponds to (unused) forward forward function:

$$\begin{aligned} y &= -1, & x < -1 \\ &= x, & -1 \leq x < 1 \\ &= 1, & 1 \leq x \end{aligned} \quad (4)$$

A complementary method of pushing error information deeper into the network with a teacher network (e.g., step 1 as a teacher for step 2 or a different variant of the network) and forming additional errors via intermediate feature map matching was not used [7].

6. Results

6.1. Accuracy

For the network defined in Section 4 configured as shown in Table 1 and using the 2 step training defined in Section 5, the top 1 accuracy on the 2012 ImageNet validation set after 5 + 20 epochs in step 1 was 66.63% and after 5 + 100 epochs in step 2 was 68.46%. After an additional cycle of 5 + 100 epochs in step 2 with the max learning rate decreased to 0.00025 the top 1 accuracy improved to 68.96%. Training hyper parameter optimization could likely improve this further.

For the same network with $P = 2$ parallel branches in the convolution module, using the same 2 step 2 cycle training procedure, top 1 accuracy on the 2012 ImageNet validation set improved to 71.24% after the final step.

	Blocks	S	R	RC	H/S	W/S
Stem replicate	1	1	8	32	224	224
Stem normal	0	1	1	32	224	224
Level 1 down	1	2	2	64	112	112
Level 1 normal	0	1	1	64	112	112
Level 2 down	1	2	2	128	56	56
Level 2 normal	0	1	1	128	56	56
Level 3 down	1	2	2	256	28	28
Level 3 normal	1	1	1	256	28	28
Level 4 down	1	2	2	512	14	14
Level 4 normal	5	1	1	512	14	14
Level 5 down	1	2	2	1024	7	7
Level 5 normal	1	1	1	1024	7	7
Global avg pool	1			1024	1	1
Fully connected	1			1000	1	1

Table 1. Network design configuration. The feature map size at the level output is RC x H/S x W/S.

6.2. Operations

Per the discussion in Section 2, it’s not fully appropriate to convert binary operations to equivalent real operations by dividing by a single factor. The approach taken in this sub section is to simply list the specific operations, which while slightly more cumbersome, is more appropriate for evaluating the complexity on a given implementation (though still not fully acceptable).

The number of parameters and number of operations is listed in Table 2 for the network specified in Table 1. Note that $\sim 90\%$ of the real parameters are in the class decoder fully connected layer. Reducing the memory requirements in this layer, potentially via binarization, is a direction for future work.

Item	$P = 1$	$P = 2$
Binary parameters	9.04 e6	18.09 e6
Binary MACs	2.41 e9	4.83 e9
Real parameters	1.15 e6	1.19 e6
Real MACs	0.04 e9	0.04 e9
Real adds	62.67 e6	109.64 e6
Real mults	19.57 e6	35.22 e6
Sign	15.65 e6	31.31 e6
PReLU	16.41 e6	16.41 e6

Table 2. Parameter and operation counts in the proposed network.

In Table 3 the real MACs, real adds, real mults, sign and PReLU are aggregated into real ops / 2 (to match the units

of MACs) with powers chosen to better enable binary to real comparisons. As the implementation complexity of binary CNN style 2D convolution is small, the implementation complexity of the remaining operations and associated data movement plays a larger role in the ultimate system performance.

Item	$P = 1$	$P = 2$
Binary parameters	9.04 e6	18.09 e6
Binary MACs	2.41 e9	4.83 e9
Real parameters	1.15 e6	1.19 e6
Real OPs/2	0.09 e9	0.13 e9

Table 3. Parameters and aggregated number of operations.

6.3. Comparisons

Table 4 includes comparisons to 2 recent top performing networks which use mostly binary CNN style 2D convolution: MeliusNet and ReActNet. The comparisons are grouped to an ImageNet 2012 top 1 validation accuracy of $69.2\% \pm 0.2\%$ and $71.2\% \pm 0.2\%$ where Mb is million binary parameters, BGM is binary GMACs, MB is million real parameters and RGO is real GOPs/2. For reference, the top 1 accuracy of MobileNetV1 is 70.6%, so the accuracy of these binary networks are in the ballpark of a real network optimized for smaller devices.

Network	Acc%	Mb	BGM	MB	RGO
MeliusNet42	69.2		9.69		0.17
ReActNetA	69.4	28.25	4.82	1.08	0.03
BCNN P=1	69.0	9.04	2.41	1.15	0.09
MeliusNet59	71.0		18.3-		0.25
ReActNetC	71.4	27.55	4.69	1.77	0.16
BCNN P=2	71.2	18.09	4.83	1.19	0.13

Table 4. Comparisons to other binary networks.

Combining real operations with binary operations via the often used but as previously mentioned lacking rule of thumb normalized OPs = binary OPs / 64 + real OPs and combining real parameters (assuming int8 format) with binary parameters allows for the normalized results in Table 5.

Network	Acc%	Param MB	Norm MOPs/2
MeliusNet42	69.2	10.1-	325.—
ReActNetA	69.4	4.61	108.40
BCNN P=1	69.0	2.28	131.28
MeliusNet59	71.0	17.4-	532.—
ReActNetC	71.4	5.22	232.51
BCNN P=2	71.2	3.45	208.15

Table 5. Normalized comparisons to other binary networks.

Around 69.2% accuracy, BCNN P=1 requires only $\sim 50\%$ of the parameter storage but $\sim 20\%$ more normalized operations than ReActNetA. Around 71.2% accuracy, BCNN P=2 requires only $\sim 66\%$ of the parameter storage and $\sim 10\%$ less operations than ReActNetC. These results were obtained with no efficient frontier of accuracy vs complexity optimization for BCNN.

Note that MeliusNet and ReActNet use real non fully grouped convolutions for the stem and / or transitions. Doing the same for BCNN could likely be used to improve accuracy, but our goal was binarizing all matrix matrix multiplication operations.

For reference, other well known networks with various parts binarized (and associated top 1 accuracy on the 2012 ImageNet validation set) include BNNs (42.2%), XNOR-Net (51.2%), Bi-RealNet-152 (64.5%) and Real-To-Bin Net (65.4%). These networks were omitted from the complexity comparison, however, as they are at lower accuracies.

7. Conclusion

This paper presented a CNN designed for image classification where all operations that can be lowered to matrix matrix multiplication are fully binary. This was enabled via a common building block design motivated by a constructive universal function approximator proof. A realization of the proposed network achieved 71.24% top 1 accuracy on the 2012 ImageNet validation data set.

Logical extensions and next steps include:

- Attempting to improve training via feature map matching with a teacher network, L2 weight regularization during step 1 to $\{-1, 1\}$, alternative real to binary transformations and improved hyper parameter tuning
- Finding an efficient frontier of accuracy vs complexity for different network configuration parameters
- Experimenting with binary $F \times F$ CNN style 2D convolution using the parallel branch structure and $P = F^2$ binary 1×1 CNN style 2D convolution operations with additional feature maps shift and 0 pad operations before the bias
- Designing an inverted residual building block using the same strategy to reduce real feature map memory in the identity path
- Using the binary CNN feature encoder in more complex vision problems with binarized decoders

References

- [1] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *arXiv preprint arXiv:1909.13863*, 2019. 2
- [2] Adrian Bulat, Georgios Tzimiropoulos, Jean Kossaifi, and Maja Pantic. Improved training of binary networks for human pose estimation and image recognition. *arXiv preprint arXiv:1904.05868*, 2019. 2
- [3] Kumar Chellappilla, Sidd Puri, and Patrice Simard. High performance convolutional neural networks for document processing. 2006. 1
- [4] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. *arXiv preprint arXiv:1805.09501*, 2018. 6
- [5] G Cybenko. Approximations by superpositions of a sigmoidal function. *Mathematics of Control, Signals, and Systems*, pages 303–314, 1989. 3
- [6] Sajad Darabi, Mouloud Belbahri, Matthieu Courbariaux, and Vahid Partovi Nia. Regularized binary network training. *arXiv: 1812.11800*, 2018. 6
- [7] Ruizhou Ding, Ting-Wu Chin, Zeye Liu, and Diana Marculescu. Regularizing activation distribution for training binarized deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11408–11417, 2019. 7
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [9] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2019. 6
- [10] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2
- [11] Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *Advances in Neural Information Processing Systems*, pages 345–353, 2017. 2
- [12] Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. Reactnet: Towards precise binary neural network with generalized activation functions. *arXiv preprint arXiv:2003.03488*, 2020. 3, 6, 7
- [13] 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. 2, 6
- [14] Brais Martinez, Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Training binary neural networks with real-to-binary convolutions. *arXiv preprint arXiv:2003.11535*, 2020. 6
- [15] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10428–10436, 2020. 2, 6
- [16] 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. 2
- [17] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [18] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 2
- [19] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In *Advances in Neural Information Processing Systems*, pages 8252–8262, 2019. 6
- [20] Cihang Xie, Mingxing Tan, Boqing Gong, Jiang Wang, Alan L Yuille, and Quoc V Le. Adversarial examples improve image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 819–828, 2020. 6
- [21] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 2
- [22] Zhe Xu and Ray CC Cheung. Accurate and compact convolutional neural networks with trained binarization. *arXiv preprint arXiv:1909.11366*, 2019. 6
- [23] I Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *arXiv preprint arXiv:1905.00546*, 2019. 6
- [24] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 6
- [25] Bohan Zhuang, Chunhua Shen, Minghui Tan, Lingqiao Liu, and Ian Reid. Structured binary neural networks for accurate image classification and semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 413–422, 2019. 2