

Single-bit-per-weight deep convolutional neural networks without batch-normalization layers for embedded systems

Mark D. McDonnell

Computational Learning Systems Laboratory,
University of South Australia
Mawson Lakes SA, Australia
mark.mcdonnell@unisa.edu.au

Hesham Mostafa

Institute for Neural Computation,
University of California, San Diego
La Jolla, USA

Runchun Wang and André van Schaik

International Centre for
Neuromorphic Systems,
Western Sydney University
Penrith NSW, Australia

Abstract—Batch-normalization (BN) layers are thought to be an integrally important layer type in today’s state-of-the-art deep convolutional neural networks for computer vision tasks such as classification and detection. However, BN layers introduce complexity and computational overheads that are highly undesirable for training and/or inference on low-power custom hardware implementations of real-time embedded vision systems such as UAVs, robots and Internet of Things (IoT) devices. They are also problematic when batch sizes need to be very small during training, and innovations such as residual connections introduced more recently than BN layers could potentially have lessened their impact. In this paper we aim to quantify the benefits BN layers offer in image classification networks, in comparison with alternative choices. In particular, we study networks that use shifted-ReLU layers instead of BN layers. We found, following experiments with wide residual networks applied to the ImageNet, CIFAR 10 and CIFAR 100 image classification datasets, that BN layers do not consistently offer a significant advantage. We found that the accuracy margin offered by BN layers depends on the data set, the network size, and the bit-depth of weights. We conclude that in situations where BN layers are undesirable due to speed, memory or complexity costs, that using shifted-ReLU layers instead should be considered; we found they can offer advantages in all these areas, and often do not impose a significant accuracy cost.

I. INTRODUCTION

Following its introduction in 2015 [1], Batch Normalization (BN) layers rapidly became a default layer type in state-of-the-art deep convolutional neural networks (CNNs). In particular for CNNs designed as image classifiers, batch-normalization is essential for state of the art accuracy on difficult datasets like Imagenet [2]. However, BN layers have several disadvantages such as:

- BN layers typically increase the GPU memory requirements and computational load during training, due to by-default storage of feature maps for every layer, for use in gradient calculations¹; additional multiplications needed for BN layers also slows down inference;
- BN layers and the computation of their parameters are implemented slightly differently in different popular deep

¹This extra storage can be avoided by recomputing BN layer outputs from its input feature map, but this slows down training slightly.

learning libraries, resulting in difficulty replicating results and in model portability;

- BN layers are not well-suited to custom-hardware implementations of training, such as those that use a pipeline approach where all processing needs to be done on one sample independently of all other samples.
- BN layers create challenges for efficient training when batches are split across multiple GPUs, and as a result sometimes non-exact approximations are used [3];
- BN layers can be less effective for models that have to be trained using small batches [4], such as when input images or model sizes are large enough to fill up GPU RAM;
- BN layers can be problematic for datasets with high class-imbalance or when samples in a minibatch are not independent [5].

The last four disadvantages are due to the fact that BN layers require calculation of batch-wise statistics, namely the mean and variance of every channel, calculated over all locations in the channel’s feature map and all samples in a minibatch, at each point in a network where a BN layer is used. These statistics are less robust for both smaller numbers of samples in a minibatch, and if a batch includes rarely chosen or non-independent samples [5].

Some of the potential challenges of BN layers are not as frequently of importance for inference, such as minibatch size. However in the near future, it has been predicted that there will be many applications where it will be desirable to *train* deep neural networks on low-power custom hardware, such as when networks need to be updated frequently using new data, and communication to a data center is too slow or unavailable [6].

For all these reasons, it is desirable to know whether BN layers really are generally necessary for best performance, or whether for particular datasets state-of-the-art accuracy can be achieved without them. However, our main motivation is drawn from seeking to design deep CNN classifiers implemented in custom hardware, that run as fast and efficiently as possible in an embedded system. It is already well-established that convolutional layers can be implemented without use of

multipliers, by reducing representation precision of weights or feature maps to 1 bit, hence potentially saving large amounts of chip space and power usage [7], [8], [6], [9], [10]. However, existing investigations in this area often still use BN layers, which if implemented exactly introduce significant complexity, including mandating the use of minibatches, and the need for multipliers, and it is desirable to know whether they can be removed or replaced.

In this paper, we train deep CNNs with and without BN layers, and analyze the resulting accuracy changes.

The paper is structured as follows. In Section II we review the BN layer definition, and discuss relevant prior research related to the goals of this paper. Then, in Section III, we describe the CNN architectures we chose to use for this study, how we vary the use of BN layers in this architecture, and adaptations we need to introduce to enable them to work effectively. Next, Section IV contains our results. Finally, we discuss the implications and significance of our results in Section V.

II. RELEVANT PREVIOUS RESEARCH

Depth is crucial for effective learning [11] in modern neural networks as it allows a network to learn a hierarchical representation of the data by successively composing simpler features into more complex features. Depth, however, poses a number of challenges when training neural networks. The large number of parameters in deep networks typically restricts the optimization methods that could be feasibly used to first order methods like stochastic gradient descent. The loss surface of a deep neural network, however, is highly non-convex and there is no guarantee that starting from an initial parameter point and traversing this loss surface using gradient descent will land in a good local minimum that allows the network to generalize well on unseen data [12], [13]. A considerable number of techniques and tricks have been developed to allow gradient descent to practically succeed in training deep networks, ranging from random parameter initialization strategies [14], [15], unsupervised pre-training [16], [17], [18], or augmenting gradient descent with gradient history information to more effectively traverse the loss surface [19], [20].

The performance of gradient descent is highly dependent on the way in which the optimization problem is parameterized. For example, by re-parameterizing a deep neural network, i.e., by changing the scale and shift of the network’s parameters, we can drastically change the curvature of the loss surface [21] and the behavior of gradient descent. One of the famous deleterious effects of bad parameterization is the vanishing and exploding gradients problem that plagued early neural networks [22]. To combat such problems, and to yield a more favorable parameterization for gradient descent in general, a broad class of techniques attempt to re-parameterize deep neural networks through various forms of normalization. Such normalization techniques operate on various quantities in the network, such as activations or weights.

Perhaps the most popular normalization technique is batch-normalization [1], which is described in more detail in the next

subsection. Batch-normalization not only accelerates gradient descent learning, it also has a beneficial regularization effect [23]. Batch-normalization operates on pre-activations (the inputs to neurons before the activation function is applied); it rescales and shifts the pre-activations of each single neuron so that they have zero mean and unity variance across the mini-batch samples, in each channel. During training, the network’s response to an example thus depends on the other examples that accompany it in a mini-batch. This dependence, however, could prove undesirable in some situations [24]. Batch-normalization requires the use of a mini-batch that is large enough to yield reliable pre-activation statistics. Increasing mini-batch size, however, often leads to a decrease in the network’s generalization performance [25].

A related normalization technique that does not depend on mini-batch statistics is layer normalization [26]. Layer normalization normalizes the pre-activations of neurons in a layer so that they have zero mean and unity variance. These statistics are calculated across all pre-activations in a layer, unlike batch-normalization which calculates them for each neuron individually across the mini-batch. One downside, however, is that different neurons, especially in convolutional layers, can have widely different input statistics, so normalizing all of them using the same coefficients is poorly motivated. Instead of normalizing the activations, ref [27] normalizes by the norm of the input weight vector of each neuron. An extra parameter is introduced for each weight vector to explicitly control its length. The performance of weight normalization closely matches that of batch-normalization while avoiding its major downside: the dependence on mini-batch statistics. Another innovation introduced by [27] is the “mean-only BN” layer, in which layer inputs are centered according to the mean of a batch, in conjunction with weight normalization.

Recently, yet another method was introduced and shown to outperform BN: group normalization [4] while avoiding some of its downsides.

The choice of activation function has a large impact on the performance of gradient descent or backpropagation. When errors are backpropagated through a layer, they are scaled by the derivative of the activation function of the neurons in that layer. If these derivatives are small, then errors are effectively blocked from propagating backwards. Saturating activation functions such as the logistic sigmoid, or hyperbolic tan, are particularly vulnerable to this effect as their derivatives are very small when their input is far from zero [28]. The use of non-saturating activation functions such as Rectified Linear Units (ReLUs) has partially alleviated this problem. A ReLU activation is zero for negative inputs and the identity for positive inputs. ReLUs, however, cannot have an output that is zero mean across the training examples as they do not produce negative outputs. Exponential Linear Units (ELUs) [29] address this issue by having a saturating negative output if the input is less than zero. ELUs have been shown to perform extremely well without any form of explicit normalization suggesting that they intrinsically normalize the levels of activation in the network.

However, shifted Rectified Linear Units (sReLU) offer nearly all the same benefits as ELUs, but without the need to calculate exponentials, and given our motivation of minimizing computational load, they are our focus here instead of ELUs. Indeed, our results in Figures 5 and 6 show that networks using sReLU and ELU activations do not have significant differences in accuracy.

For similar reasons, we aim to establish if we can avoid alternatives to BN like group normalization [4], where complexity and additional computation is introduced for computing statistics and carrying out normalizations by non-constant factors.

A. Review of BN layers

Batch-normalization (BN) layers have been shown to help deep neural networks produce better accuracy following training. The usual explanation for this is that they help “reduce internal covariate shift” [1]. This view has been recently challenged [30]. However it is clear that the use of BN layers enables higher learning rates, and hence faster convergence during training, and diminishes the importance of good initial conditions for convolutional layers [1], [5].

BN layers are applied to a minibatch of feature maps, which typically can be represented as a 4-axis tensor, $\mathbf{x} \in \mathcal{R}^{K,H,W,C}$, where K is the size of the minibatch, H and W are the height and width of the feature maps, and C is the number of channels. BN layers operate on a per-channel basis, with channel-specific gain and shift parameters, and make use of channel-specific calculations of the mean and variance of the inputs to the channel in a minibatch. Mathematically, if $x_{k,h,w,c}$ is the value of the feature map at location h, w in the c -th channel of the k -th sample, then BN layers transform that value according to

$$\hat{x}_{k,h,w,c} = g_c \left(\frac{x_{k,h,w,c} - \mu_c}{\sqrt{\sigma_c^2 + \epsilon}} \right) + o_c, \quad (1)$$

where μ_c and σ_c^2 are calculated as

$$\mu_c = \frac{1}{KHW} \sum_k \sum_h \sum_w x_{k,h,w,c} \quad (2)$$

and

$$\sigma_c^2 = \frac{1}{KHW} \sum_k \sum_h \sum_w (x_{k,h,w,c} - \mu_c)^2. \quad (3)$$

The remaining parameters are:

- gain, g_c , and shift o_c , which are usually learned in the same manner as weights in convolutional layers;
- ϵ , which ensures division by zero cannot happen in the event of zero variance, and can also act like a regularizer; it is not widely appreciated that accuracy can depend on the exact choice of ϵ and that inference performance is sensitive to inadvertent changes in ϵ compared to training. It is also noteworthy that some popular deep learning libraries define ϵ differently, by taking it outside the square root, which changes subtly but sometimes significantly the performance of otherwise identical models.

There are two major differences between BN layers and other layers:

- 1) calculation of layer outputs depend on all samples in a minibatch and cannot be computed independently for each sample;
- 2) During training, minibatch means and variances are calculated. For inference, each mean and variance is an additional parameter that forms part of the trained model, but unlike most parameters derived from training data, these are not learned, but rather are computed. Most popular libraries compute these values during training by calculating exponential moving averages over batches as training progresses. We have found that this can sometimes give misleading and sub-par performance during monitoring on a validation set as training progresses, because parameters change over the window in which averages are created. Instead, we favour calculation of batch means and variances using multiple training batches while training is frozen, as described in Algorithm 2 in the original BN paper [1].

B. BN layer scale and shift can be detrimental

Recent work demonstrated the surprising result that not learning BN scale and shift parameters can actually be beneficial, leading to reduced error rates for CIFAR 10 and CIFAR 100 in a wide residual network [10]. This result was found to be the case for both full-precision networks, and for “1-bit-per-weight” versions of the same networks.

C. Shifted Rectified Linear Unit and Exponential Linear Unit

It has already been shown that removal of BN layers can somewhat be compensated for by using exponential linear units (ELUs) [29] instead of the combination of BN and ReLU. The mathematical definition of these two activation functions are

$$\text{ELU}(x) = x\mathcal{I}(x) + (\exp(x) - 1)(1 - \mathcal{I}(x)), \quad (4)$$

where $\mathcal{I}(x)$ is the Heaviside step function, and

$$\text{sReLU}(x) = \max(-1, x). \quad (5)$$

Both can be generalized to a different minimum output value, but we consider only the case of -1 . See Figure 1 for plots of the sReLU and ELU characteristics, in comparison with ReLU.

When ELUs were introduced, it was shown that they outperform shifted rectified linear units (sReLU) [29]. However, in our experiments with more recently advanced networks and training methods not available at the time of [29], we did not observe any significant difference between ELUs and sReLU. Moreover, ELUs are computationally expensive compared with sReLU, both because sReLU do not require calculations of exponentials in the forward pass, nor storage of their values for use in the backward pass during training. We show here that shifted ReLU are an effective and efficient alternative to ELUs, and seek to quantify how much accuracy penalty is incurred by replacing BN layers.

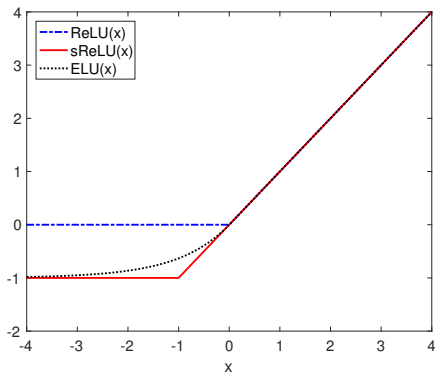


Fig. 1. **Shifted Rectified Linear Unit (sReLU) activation function.** The sReLU activation function lets negative inputs pass through, between 0 and some negative constant, in this case equal to -1 . While the Exponential Linear Unit (ELU) is more popular, we have found sReLU to be equally effective, and less computationally demanding, due to avoiding calculation of an exponential.

III. METHODS

A. Baseline network architecture and training

Our experiments are based on wide residual networks [31] (see Figure 2) that use the post-activation architecture [32]. Using the nomenclature of [31], for CIFAR 10 and CIFAR 100 we used depth 20 and widths of $4\times$ and $10\times$ and for ImageNet depth 18 and widths $1\times$ and $2.5\times$, meaning in both cases that the first layer of convolutional weights had either 64 or 160 output channels. The design we used has a few differences to [31], such as that the final weights layer is a 1×1 convolutional layer applied before the global average pooling layer—see Figure 2 which illustrates this aspect, and the overall design. More details are described in [10]. Note also that in the design of [10] a BN layer was applied to the RGB input channels prior to the first convolution layer. Here we do the same for all variations, including the sReLU one. Provided that the shift and scale factors are not learned, this BN layer is equivalent to preprocessing raw data, and does not need to be considered to form a network layer.

Training was carried out similarly to [10], i.e. we used backpropagation and stochastic gradient descent, with mini-batches of size 125 samples, momentum of 0.9 and weight decay with a value of 0.0005. No biases are used. Weights were initialized using the method of [15], while BN gains were initialized to 1 and shifts to 0. No convolutional layer biases were used. Unlike [10], we did not use a warm restart learning rate schedule, as we found this could sometimes lead to divergence following a restart with the sReLU networks. However, we did use a cosine learning-rate decay schedule, starting at an initial value of 0.1 and finishing at 10^{-5} after 300 epochs (CIFAR) or 60 epochs (ImageNet). For networks where BN layers were used, we computed the mean and variance statistics following the end of training, by calculating averages over all minibatches in 1 epoch, with learning turned off.

During training, each image selected for a minibatch was augmented using standard methods as in [10]. In addition,

cutout augmentation [33] was used for CIFAR 10 and CIFAR 100, with the same design as in [10], with a patch size of 18 pixels.

B. 1-bit-per-weight networks

As well as training networks with the usual 32 bit floating point precision for all variables, including learned weights, we also trained networks using a method for enabling storage of learned weights and inference to take place using 1-bit values. We followed the method of [10]; differences in training are summarised for the case of sReLU in Figure 3.

We emphasize that the main motivation of this paper is to examine whether methods for enabling reduced-precision representations in deep neural networks, such as [10], still work effectively when batch-normalization layers are removed.

C. Shifted ReLUs

We conducted experiments using the same architecture as in Fig. 2, but where (i) all BN layers except the final one closest to the output are replaced by shifted ReLUs and (ii) where all BN layers are replaced by shifted ReLUs. The architecture for the latter case is shown in Figure 4.

D. Training innovations for networks with BN layers removed

For the model where all BN layers but the final one were removed, we found we did not need to change any aspect of training relative to our baseline models.

However, without any BN layers, we found that training diverged, or converged very slowly. A partial remedy to this problem is to simply reduce the learning rate, a fact consistent with one of the advantages of using BN layers, i.e. that the initial learning rate can be set higher [1] than in networks without BNs. However, this leads to slow convergence and increased error rates.

We investigated why this was the case. We found, surprisingly, that using just a single BN layer between the final convolutional layer and the global average pooling layer was sufficient to enable the network to be trained exactly like the baseline models.

Based on this observation, we hypothesized that the main reason that networks with BN layers enable a larger learning rate is due to the standardized scaling the final BN layer imparts on gradients from the output prior to backpropagation to weight layers. Indeed, we observed that without a final BN layer, the distribution of gradients at the output early in training have a longer tail than when BN is included.

We therefore introduced in our shifted ReLU networks a constant scaling layer to replace the BN layer between the final convolutional layer and global average pooling layer. For simplicity we have set the constant scaling identically for all channels (which in our design at this point in the network is equal to the number of classes).

Due to the linearity of the global average pooling layer, changing this scaling is equivalent to changing the temperature in the softmax layer, from its default value of 1. That is, our

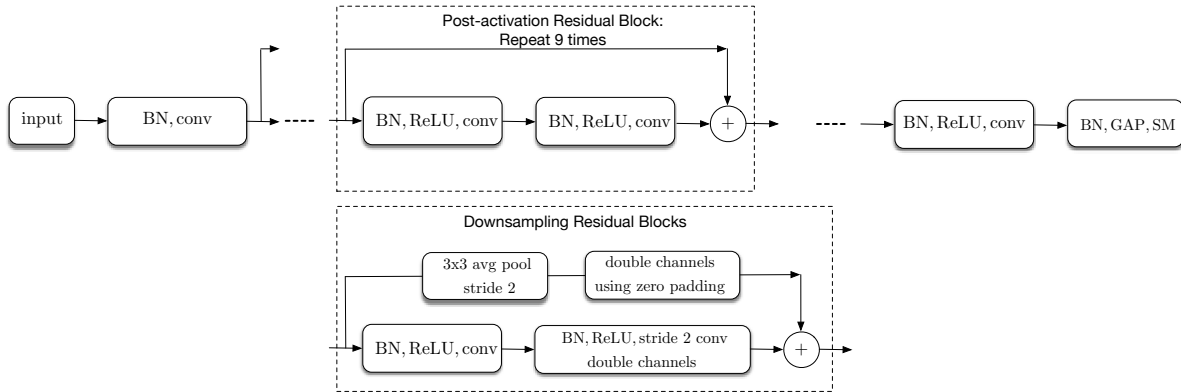


Fig. 2. **Wide ResNet architecture for Baseline CIFAR models where BN layers are used.** This architecture is nearly identical to that of [10], except here there is no optional ReLU applied to the input. Note the ordering of the final layers, where global average pooling (GAP) is used after a final 1×1 convolutional layer, that reduces the number of channels to equal the number of classes, and then feeds directly to the softmax output (SM).




Fig. 3. **Changes when training for 1-bit-per-weight.** When we train 1-bit-per-weight networks following the method of [10], we apply the sign operator to full-precision copies of weights during training, and then scale by a constant equal to the initial standard deviation of the weights according to the method of [15].

approach corresponds to employing a final softmax layer of the form

$$SM_i(x) := \frac{\exp\left(\frac{x_i}{T}\right)}{\sum_{j=1}^N \exp\left(\frac{x_j}{T}\right)}, \quad (6)$$

where T is the temperature—see, e.g. [34], [35].

We found a value for the temperature between approximately 30 and 100 to enable training to take place identically to the baseline models, including the same high initial learning rate. Lower values of T tended to result in failure to converge shortly after training commenced.

Note that although the gradient propagated back is linearly scaled by $1/T$, changing from the default of $T = 1$ is not equivalent to simply changing the learning rate by a factor of T . This is due to the nonlinearity in the softmax layer. Increasing T has the effect of moving softmax outputs away from 0 or 1, thereby increasing the entropy of the output vector. In turn, this means gradients propagating backwards early in training have a distribution with lower standard deviation.

That this temperature scaling is all that is needed to ensure a high learning rate suggests that the main problem with

larger learning rate for models without BN layers is simply that early in training, the gradients calculated at the output of the network are too high for high learning rates, and that the normalization of the final BN layer compensates for this.

IV. EXPERIMENTS AND RESULTS

We trained the following variations of our wide residual networks on both CIFAR 10 and CIFAR 100, all for both width 4 and width 10.

- 1) Baseline 1: networks as in Figure 2, with conventional BN layers where scales and shifts are learned.
- 2) Baseline 2: the same as Baseline 1, but without any scales and shifts learned, as in [10].
- 3) A single BN layer at the end of the networks only (no scale and shift learned), with all other BN-ReLU combinations replaced by sReLU.
- 4) No BN layers – all replaced by sReLU, as in Figure 4.
- 5) No BN layers – all replaced by ELU [29].
- 6) No BN layers – all replaced by sReLU, except for a ‘mean-only BN’ layer at the end of the network [27], without a learned bias.

Our results for width-4 and width-10 ResNets are summarized in Tables I and II, while Figures 5 and 6 show the mean and spread of accuracies for width-4 ResNets for CIFAR 10 and CIFAR 100 over 10 runs.

For ImageNet, we compared width-1 and width-2.5 networks, and an ensemble of 3 width-1 networks, for the case of Baseline 1, and case 3 in the above list. The results are summarised in Table III.

Our conclusions drawn from the results are left for Section V.

V. DISCUSSION

A. The impact of removing BN layers

Our results indicate the following:

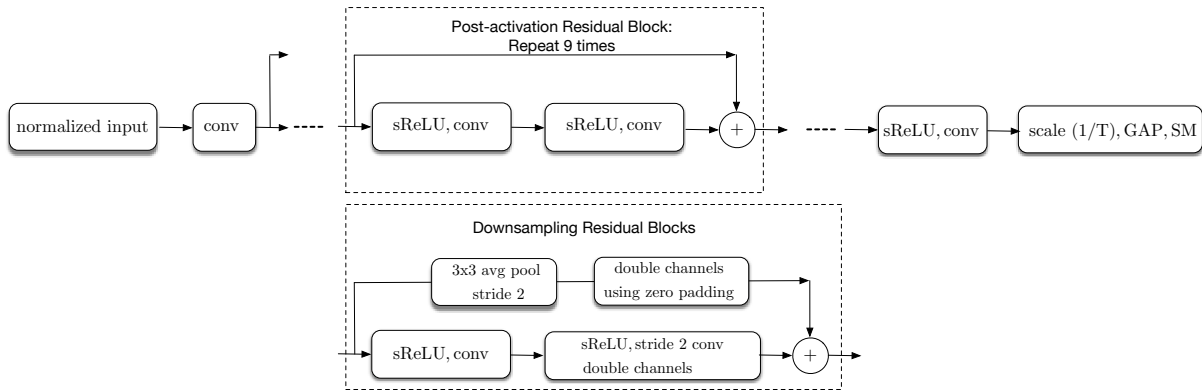


Fig. 4. **Wide ResNet architecture for CIFAR when all BN layers are replaced by sReLUs.** The architecture is identical to that of Figure 2 except that (i) all BNs have been removed and ReLUs have been replaced by shifted ReLUs (sReLU); (ii) a scale layer (multiplies all inputs by a constant) has been inserted before the global average pooling (GAP) layer; and (iii) the input is now normalized using pre-preprocessing applied to each of the RGB channels.

TABLE I

CIFAR 10: Test-set error-rates. THE ENTRIES FOR RESNET 20-4 NETWORKS ARE THE MEAN VALUE FROM 10 REPEATS. THE ENTRIES FOR RESNET 20-10 NETWORKS ARE FOR SINGLE RUNS. THE FINAL COLUMN SHOWS THE DIFFERENCE BETWEEN THE ERRORS FOR sReLU ONLY NETWORKS COMPARED WITH THE BEST RESULT IN EACH ROW, AS INDICATED IN BOLD FONT.

Model	Baseline 1	Baseline 2	Final BN only	sReLU only	sReLU gap
ResNet 20-4, 32 bit weights	3.97%	3.80%	4.42 %	4.67%	0.87%
ResNet 20-10, 32 bit weights	3.29%	3.57%	3.79%	4.36%	1.07%
ResNet 20-4, 1 bit weights	4.51%	4.65%	4.47%	4.66%	0.19%
ResNet 20-10, 1 bit weights	3.83%	3.65%	4.00%	3.74%	0.09%

TABLE II

CIFAR 100: Test-set error-rates. THE ENTRIES FOR RESNET 20-4 NETWORKS ARE THE MEAN VALUE FROM 10 REPEATS. THE ENTRIES FOR RESNET 20-10 NETWORKS ARE FOR SINGLE RUNS. THE FINAL COLUMN SHOWS THE DIFFERENCE BETWEEN THE ERRORS FOR sReLU ONLY NETWORKS COMPARED WITH THE BEST RESULT IN EACH ROW, AS INDICATED IN BOLD FONT.

Model	Baseline 1	Baseline 2	Final BN only	sReLU only	sReLU gap
ResNet 20-4, 32 bit weights	22.10%	19.53%	20.18%	22.24%	2.69%
ResNet 20-10, 32 bit weights	20.99%	17.05%	18.25%	20.97%	3.92%
ResNet 20-4, 1 bit weights	22.82%	22.07%	22.31%	23.69%	1.60%
ResNet 20-10, 1 bit weights	20.61%	18.22%	19.42%	20.74%	2.52%

- The importance of BN layers is data-set and model-size dependent. Our findings show that the accuracy loss when replacing all BN layers with shifted ReLUs is larger on CIFAR 100 than on CIFAR 10. For all three datasets, and especially ImageNet, the accuracy loss is very large for width-1 (non-wide), but not for wide variants, indicating BN is much more important for smaller models.
- For CIFAR, there is no clear advantage in using ELU layers rather than sReLU.
- The impact of removing BN is larger for full-precision weights compared with 1-bit-per weight.
- Baseline 1, where gains and shifts are learned is about as good for CIFAR 10 as Baseline 2 where they are not, but Baseline 2 is clearly better for CIFAR 100.
- Using a single BN-layer at the end of networks was mostly markedly better than using all sReLUs, and in some cases nearly as good as the best baseline.
- For width 4 networks, the variations within models across runs for CIFAR 10 and 1-bit-per-weight indicated that any model could produce best results. For CIFAR 100, the first three models exhibited this effect, but the sReLU network clearly had a small gap in performance compared to other models. For 32-bit-per-weight models, gaps were larger, but the top two models for CIFAR 10 were the two baselines, while for CIFAR 100 they were Baseline 2 and the final BN only model.
- The mean-only-BN model outperforms the All ReLU and All ELU networks for CIFAR 100 and 1-bit per weight, halving the error rate gap to Baseline 2.
- Consistent with [10], for CIFAR 100, Baseline 1 where scales and shifts are learned causes a dramatic drop in accuracy relative to Baseline 2. Unlike [10], this is not

TABLE III

ImageNet: Validation-set top-5 error-rates. CENTRE CROP MEANS A SINGLE 224×224 CROP FROM THE CENTRE OF THE VALIDATION IMAGE WAS USED. MULTICROP MEANS 25 DIFFERENT CROPS WERE RUN USED AND THEIR PREDICTIONS AVERAGED BEFORE CLASSIFYING, SIMILAR TO [32].

Bits per weight	Model	# Learned parameters	Test mode	Baseline 1	Final BN only	Gap
32	ResNet 18-1	11.5M	centre crop	12.41%	15.50%	3.01%
32	ResNet 18-1	11.5M	multi crop	9.03%	14.70%	5.67%
32	Ensemble of 3 ResNet 18-1	34.5M	centre crop	10.70%	14.00%	3.30%
32	Ensemble of 3 ResNet 18-1	34.5M	multi crop	7.95%	12.30%	4.35%
32	ResNet 18-2.5	70M	centre crop	9.20%	9.51%	0.3%
32	ResNet 18-2.5	70M	multi crop	6.91%	8.83%	1.92%
1	ResNet 18-1	11.5M	centre crop	17.55%	23.66%	6.11%
1	ResNet 18-1	11.5M	multi crop	12.80%	19.94%	7.14%
1	Ensemble of 3 ResNet 18-1	34.5M	centre crop	15.48%	22.18%	6.70%
1	Ensemble of 3 ResNet 18-1	34.5M	multi crop	11.38%	18.85%	7.47%
1	ResNet 18-2.5	70M	centre crop	11.51%	11.81%	0.3%
1	ResNet 18-2.5	70M	multi crop	8.48%	10.03%	1.58%




Fig. 5. **Spread of results: CIFAR 10, Width 4.** The circle markers show the mean from 10 repeated runs for each of the 4 model types, using different random seeds for each repeat, but the same seed for each model. The error bars indicate the maximum and minimum errors over the 10 repeated runs.




Fig. 6. **Spread of results: CIFAR 100, Width 4.** The circle markers show the mean from 10 repeated runs for each of the 4 model types, using different random seeds for each repeat, but the same seed for each model. The error bars indicate the maximum and minimum errors over the 10 repeated runs.

longer the case for CIFAR 10, which might be because here we do not use a ReLU applied to the input.

From these observations, we propose the following:

- **Conclusion 1:** Removing all BN layers can be expected to cause an accuracy penalty, but this penalty is po-

tentially small, depending on the dataset and network architecture.

- **Conclusion 2:** Removing all but the final BN layer is a potentially viable option for getting the benefits of removing most BN layers, but without as much accuracy loss as removing all.
- **Conclusion 3:** There is no consistently best way to design networks with BNs; in some cases learning scales and shifts is beneficial, but in other cases it causes a big drop in performance.
- **Conclusion 4:** There is less impact on accuracy in the case of 1-bit-per-weight compared with full precision weights. Indeed, for CIFAR 10, sReLU networks and 1-bit-per-weight had an accuracy gap no larger than 0.2%, and only 0.3% for width-2.5 ImageNet with centre cropping, suggesting sReLU as a viable option when 1-bit-per-weight is used.

B. Significance for custom hardware implementations

Hardware acceleration of CNNs is necessary to provide real-time embedded vision systems, e.g., UAVs and Internet of Things (IoT) devices, because existing systems using CPUs are too slow. To increase speed, most software-based CNNs use GPUs. However, it is difficult to deploy GPUs in embedded systems, since they consume a significant amount of power. Thus, custom rather than general-purpose hardware-based CNNs are desired for low-power and real-time embedded vision systems.

In hardware-based acceleration systems, most power consumption comes from the computation modules, e.g., the multipliers and adders, and from accessing of the data, particularly the weights. By using binary weights, the multiplier can be replaced by a multiplexer, which consumes orders of magnitude less power. In modern CMOS technologies, e.g., 28nm, a binary convolutional operation with multiplexers achieves a power efficiency up to 230 1b-TOPS/W [36]. More importantly, using binary weights enables use of only on-chip memories, e.g., SRAMs, to store these weights. Accessing data stored in external memories would consume an order of magnitude more power ($10\times$) than the computation itself [37].

The use of batch-normalization layers might maintain maximum classification accuracy but at the cost of extra silicon area and computation and thus more power consumption. Particularly, it will require significant amount of silicon area to implement the nonlinear square and square-root operations in the conventional batch-normalization. What makes it worse is that these operations impede low bit-width quantization techniques [38]. One can easily implement the shifted ReLU activation function with a tiny silicon area, while achieving comparable accuracies to networks with BN layers.

C. Future work

In future, it will be valuable to try to devise new low-complexity methods that narrow the accuracy gap between networks that use BN layers, and the ones outlined here.

ACKNOWLEDGMENT

This work was supported by a Discovery Project funded by the Australian Research Council (project number DP170104600).

REFERENCES

- [1] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, vol. abs/1502.03167, 2015. [Online]. Available: <http://arxiv.org/abs/1502.03167>
- [2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," Microsoft Research, Tech. Rep., 2015, arxiv.1512.03385.
- [3] P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, A. Tulloch, Y. Jia, and K. He, "Accurate, Large Minibatch SGD: Training ImageNet in 1 Hour," *ArXiv e-prints*, June 2017.
- [4] Y. Wu and K. He, "Group Normalization," *ArXiv e-prints*, Mar. 2018.
- [5] S. Ioffe, "Batch renormalization: Towards reducing minibatch dependence in batch-normalized models," *CoRR*, vol. abs/1702.03275, 2017. [Online]. Available: <http://arxiv.org/abs/1702.03275>
- [6] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer, "Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <1MB model size," *CoRR*, vol. abs/1602.07360, 2016. [Online]. Available: <http://arxiv.org/abs/1602.07360>
- [7] M. Courbariaux, Y. Bengio, and J. David, "BinaryConnect: Training Deep Neural Networks with binary weights during propagations," *CoRR*, vol. abs/1511.00363, 2015. [Online]. Available: <http://arxiv.org/abs/1511.00363>
- [8] P. Merolla, R. Appuswamy, J. V. Arthur, S. K. Esser, and D. S. Modha, "Deep neural networks are robust to weight binarization and other non-linear distortions," *CoRR*, vol. abs/1606.01981, 2016. [Online]. Available: <http://arxiv.org/abs/1606.01981>
- [9] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: Imagenet classification using binary convolutional neural networks," *CoRR*, vol. abs/1603.05279, 2016. [Online]. Available: <http://arxiv.org/abs/1603.05279>
- [10] M. D. McDonnell, "Training wide residual networks for deployment using a single bit for each weight," 2018, in Proc. ICLR 2018; arxiv: 1802.08530.
- [11] J. Ba and R. Caruana, "Do deep nets really need to be deep?" in *Advances in neural information processing systems*, 2014, pp. 2654–2662.
- [12] A. Choromanska, M. Henaff, M. Mathieu, G. Arous, and Y. LeCun, "The loss surfaces of multilayer networks," in *Artificial Intelligence and Statistics*, 2015, pp. 192–204.
- [13] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," *arXiv preprint arXiv:1611.03530*, 2016.
- [14] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks." in *AISTATS*, vol. 9, 2010, pp. 249–256.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification," in *Proc. IEEE International Conference on Computer Vision (ICCV) (see arXiv:1502.01852)*, 2015.
- [16] D. Erhan, Y. Bengio, A. Courville, P. Manzagol, P. Vincent, and S. Bengio, "Why does unsupervised pre-training help deep learning?" *The Journal of Machine Learning Research*, vol. 11, pp. 625–660, 2010.
- [17] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng, "Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations," in *Proceedings of the 26th Annual International Conference on Machine Learning*, ser. ICML '09. New York, NY, USA: ACM, 2009, pp. 609–616.
- [18] M. Ranzato, F. Huang, Y.-L. Boureau, and Y. LeCun, "Unsupervised learning of invariant feature hierarchies with applications to object recognition," in *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*. IEEE, 2007, pp. 1–8.
- [19] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [20] M. Zeiler, "Adadelta: an adaptive learning rate method," *arXiv preprint arXiv:1212.5701*, 2012.
- [21] L. Dinh, R. Pascanu, S. Bengio, and Y. Bengio, "Sharp minima can generalize for deep nets," *arXiv preprint arXiv:1703.04933*, 2017.
- [22] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [23] A. Morcos, D. Barrett, N. Rabinowitz, and M. Botvinick, "On the importance of single directions for generalization," *arXiv preprint arXiv:1803.06959*, 2018.
- [24] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training GANs," in *Advances in Neural Information Processing Systems*, 2016, pp. 2234–2242.
- [25] N. Keskar, D. Mudigere, J. Nocedal, M. Smelyanskiy, and P. Tang, "On large-batch training for deep learning: Generalization gap and sharp minima," *arXiv preprint arXiv:1609.04836*, 2016.
- [26] J. Ba, J. Kiros, and G. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [27] T. Salimans and D. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 901–909.
- [28] Y. LeCun, L. Bottou, G. Orr, and K.-R. Müller, "Efficient backprop," in *Neural Networks: Tricks of the Trade, this book is an outgrowth of a 1996 NIPS workshop*. Springer-Verlag, 1998, pp. 9–50.
- [29] D. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (ELUs)," *CoRR*, vol. abs/1511.07289, 2015. [Online]. Available: <http://arxiv.org/abs/1511.07289>
- [30] S. Santurkar, D. Tsipras, A. Ilyas, and A. Madry, "How Does Batch Normalization Help Optimization? (No, It Is Not About Internal Covariate Shift)," *ArXiv e-prints*, May 2018.
- [31] S. Zagoruyko and N. Komodakis, "Wide residual networks," 2016, arxiv.1605.07146.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Identity mappings in deep residual networks," Microsoft Research, Tech. Rep., 2016, arxiv.1603.05027.
- [33] T. Devries and G. W. Taylor, "Improved regularization of convolutional neural networks with cutout," *CoRR*, vol. abs/1708.04552, 2017. [Online]. Available: <http://arxiv.org/abs/1708.04552>
- [34] G. Hinton, O. Vinyals, and J. Dean, "Distilling the Knowledge in a Neural Network," *ArXiv e-prints*, Mar. 2015.
- [35] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," *CoRR*, vol. abs/1706.04599, 2017. [Online]. Available: <http://arxiv.org/abs/1706.04599>
- [36] B. Moons, D. Bankman, L. Yang, B. Murmann, and M. Verhelst, "Bina-ryeye: An always-on energy-accuracy-scalable binary CNN processor with all memory on chip in 28nm CMOS," in *IEEE Custom Integrated Circuits Conference, CICC, San Diego, April 8-11, 2018*, pp. 1–4.
- [37] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks." in *ISSCC*. IEEE, 2016, pp. 262–263.
- [38] K. Ando, K. Ueyoshi, K. Orimo, H. Yonekawa, S. Sato, H. Nakahara, S. Takamaeda-Yamazaki, M. Ikebe, T. Asai, T. Kuroda, and M. Motomura, "Brein memory: A single-chip binary/ternary reconfigurable in-memory deep neural network accelerator achieving 1.4 TOPS at 0.6 W," *IEEE Journal of Solid-State Circuits*, 12 2017.