
Fixed Point Quantization of Deep Convolutional Networks

Darryl D. Lin

Qualcomm Research, San Diego, CA 92121, USA

DARRYL.DLIN@GMAIL.COM

Sachin S. Talathi

Qualcomm Research, San Diego, CA 92121, USA

TALATHI@GMAIL.COM

V. Sreekanth Annapureddy

NetraDyne Inc., San Diego, CA 92121, USA

SREEKANTHAV@GMAIL.COM

Abstract

In recent years increasingly complex architectures for deep convolution networks (DCNs) have been proposed to boost the performance on image recognition tasks. However, the gains in performance have come at a cost of substantial increase in computation and model storage resources. Fixed point implementation of DCNs has the potential to alleviate some of these complexities and facilitate potential deployment on embedded hardware. In this paper, we propose a quantizer design for fixed point implementation of DCNs. We formulate and solve an optimization problem to identify optimal fixed point bit-width allocation across DCN layers. Our experiments show that in comparison to equal bit-width settings, the fixed point DCNs with optimized bit width allocation offer $> 20\%$ reduction in the model size without any loss in accuracy on CIFAR-10 benchmark. We also demonstrate that fine-tuning can further enhance the accuracy of fixed point DCNs beyond that of the original floating point model. In doing so, we report a new state-of-the-art fixed point performance of 6.78% error-rate on CIFAR-10 benchmark.

1. Introduction

Recent advances in the development of deep convolution networks (DCNs) have led to significant progress in solving non-trivial machine learning problems involving image recognition (Krizhevsky et al., 2012) and speech recognition (Deng et al., 2013). Over the last two years several

advances in the design of DCNs (Zeiler & Fergus, 2014; Simonyan & Zisserman, 2014; Szegedy et al., 2014; Chatfield et al., 2014; He et al., 2014; Ioffe & Szegedy, 2015) have not only led to a further boost in achieved accuracy on image recognition tasks but also have played a crucial role as a feature generator for other machine learning tasks such as object detection (Krizhevsky et al., 2012) and localization (Sermanet et al., 2013), semantic segmentation (Girshick et al., 2014) and image retrieval (Krizhevsky et al., 2012; Razavian et al., 2014). These advances have come with an added cost of computational complexity, resulting from DCN designs involving any combinations of: increasing the number of layers in the DCN (Szegedy et al., 2014; Simonyan & Zisserman, 2014; Chatfield et al., 2014), increasing the number of filters per convolution layer (Zeiler & Fergus, 2014), decreasing stride per convolution layer (Sermanet et al., 2013; Simonyan & Zisserman, 2014) and hybrid architectures that combine various DCN layers (Szegedy et al., 2014; He et al., 2014; Ioffe & Szegedy, 2015).

While increasing computational complexity has afforded improvements in the state-of-the-art performance, the added burden of training and testing makes these networks impractical for real world applications that involve real time processing and for deployment on mobile devices or embedded hardware with limited power budget. One approach to alleviate this burden is to increase the computational power of the hardware used to deploy these networks. An alternative approach that may be cost efficient for large scale deployment is to implement DCNs in fixed point, which may offer advantages in reducing memory bandwidth, lowering power consumption and computation time as well as the storage requirements for the DCNs.

In general, there are two approaches to designing a fixed point DCN: (1) convert a pre-trained floating point DCN model into a fixed point model without training, and (2) train a DCN model with fixed point constraint. While

the second approach may produce networks with superior accuracy numbers (Rastegari et al., 2016; Lin & Talathi, 2016), it requires tight integration between the network design, training and implementation, which is not always feasible. In this paper, we will mainly focus on the former approach. In many real-world applications a pre-trained DCN is used as a feature extractor, followed by a domain specific classifier or a regressor. In these applications, the user does not have access to the original training data and the training framework. For these types of use cases, our proposed algorithm will offer an optimized method to convert any off-the-shelf pre-trained DCN model for efficient run time in fixed point.

The paper is organized as follows: In Section 2, we present a literature survey of the related works. In Section 3, we develop quantizer design for fixed point DCNs. In Section 4 we formulate an optimization problem to identify optimal fixed point bit-width allocation per layer of DCNs to maximize the achieved reduction in complexity relative to the loss in the classification accuracy of the DCN model. Results from our experiments are reported in Section 5 followed by conclusions in the last section.

2. Related work

Fixed point implementation of DCNs has been explored in earlier works (Courbariaux et al., 2014; Gupta et al., 2015). These works primarily focused on training DCNs using low precision fixed-point arithmetic. More recently, Lin et al. (2015) showed that deep neural networks can be effectively trained using only binary weights, which in some cases can even improve classification accuracy relative to the floating point baseline.

The works above all focused on the approach of designing the fixed point network during training. The works of Kyu Yeon & Sung (2014); Sajid et al. (2015) more closely resemble our work. In Kyu Yeon & Sung (2014), the authors propose a floating point to fixed point conversion algorithm for fully-connected networks. The authors used an exhaustive search strategy to identify optimal fixed point bit-width for the entire network. In a follow-up paper (Sajid et al., 2015), the authors applied their proposed algorithm to DCN models where they analyzed the quantization sensitivity of the network for each layer and then manually decide the quantization bit-widths. Other works that are somewhat closely related are Vanhoucke et al. (2011); Gong et al. (2014). Vanhoucke et al. (2011) quantized the weights and activations of pre-trained deep networks using 8-bit fixed-point representation to improve inference speed. Gong et al. (2014) on the other hand applied code-book based on scalar and vector quantization methods in order to reduce the model size.

In the spirit of Sajid et al. (2015), we also focus on optimizing DCN models that are pre-trained with floating point precision. However, as opposed to exhaustive search method adopted by Sajid et al. (2015), our objective is to convert the pre-trained DCN model into a fixed-point model using an optimization strategy based on signal-to-quantization-noise-ratio (SQNR). In doing so, we aim to improve upon the inference speed of the network and reduce storage requirements. The benefit of our approach as opposed to the brute force method is that it is grounded in a theoretical framework and offers an analytical solution for bit-width choice per layer to optimize the SQNR for the network. This offers an easier path to generalize to networks with significantly large number of layers such as the one recently proposed by He et al. (2015).

Other approaches to handle complexity of deep networks include: (a) leveraging high complexity networks to boost performance of low complexity networks, as proposed in Hinton et al. (2014), (b) compressing neural networks using hashing (Chen et al., 2015), and (c) combining pruning and quantization during training to reduce the model size without affecting the accuracy (Han et al., 2015). These methods are complementary to our proposed approach and the resulting networks with reduced complexity can be easily converted to fixed point using our proposed method. In fact, the DCN model that we performed experiments with and report results on, (see Section 5.2), was trained under the dark knowledge framework by using the inception network (Ioffe & Szegedy, 2015) trained on ImageNet as the master network.

3. Floating point to fixed point conversion

In this section, we will propose an algorithm to convert a floating point DCN to fixed point. For a given layer of DCN the goal of conversion is to represent the input activations, the output activations, and the parameters of that layer in fixed point. This can be seen as a process of quantization.

3.1. Optimal uniform quantizer

There are three inter-dependent parameters to determine for the fixed point representation of a floating point DCN: bit-width, step-size (resolution), and dynamic range. These are related as:

$$\text{Range} \approx \text{Stepsize} \cdot 2^{\text{Bitwidth}} \quad (1)$$

Given a fixed bit-width, the trade-off is between having large enough range to reduce the chance of overflow and small enough resolution to reduce the quantization error. The problem of striking the best trade-off between overflow error and quantization error has been extensively studied in the literature. Table 1 below shows the step sizes of the

optimal symmetric uniform quantizer for uniform, Gaussian, Laplacian and Gamma distributions. The quantizers are optimal in the sense of minimizing the SQNR.

Table 1. Step-sizes of optimal symmetric uniform quantizer for various input distributions (Shi & Sun, 2008)

Bit-width β	Uniform	Gaussian	Laplacian	Gamma
1	1.0	1.596	1.414	1.154
2	0.5	0.996	1.087	1.060
3	0.25	0.586	0.731	0.796
4	0.125	0.335	0.456	0.540

For example, suppose the input is Gaussian distributed with zero mean and unit variance. If we need a uniform quantizer with bit-width of 1 (i.e. 2 levels), the best approach is to place the quantized values at -0.798 and 0.798. In other words, the step size is 1.596. If we need a quantizer with bit-width of 2 (i.e. 4 levels), the best approach is to place the quantized values at -1.494, -0.498, 0.498, and 1.494. In other words, the step size is 0.996.

In practice, however, even though a symmetric quantizer is optimal for a symmetric input distribution, sometimes it is desirable to have 0 as one of the quantized values because of the potential savings in model storage and computational complexity. This means that for a quantizer with 4 levels, the quantized values could be -0.996, 0.0, 0.996, and 1.992.

Assuming an optimal uniform quantizer with ideal input, the resulting SQNR as a function of the bit-width is shown in Figure 1. It can be observed that the quantization efficiency decreases as the Kurtosis of the input distribution increases.

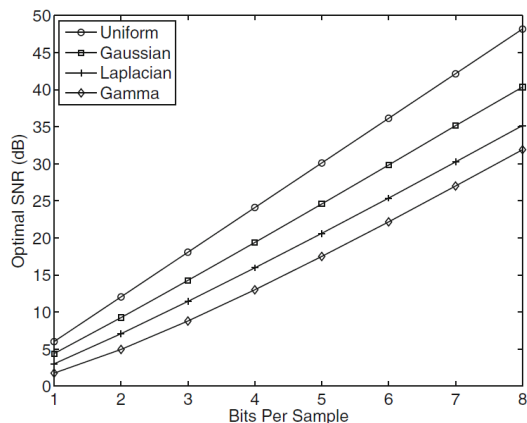


Figure 1. Optimal SQNR achieved by uniform quantizer for uniform, Gaussian, Laplacian and Gamma distributions (You, 2010)

Another take-away from this figure is that there is an ap-

proximately linear relationship between the bit-width and resulting SQNR:

$$\gamma_{\text{dB}} \approx \kappa \cdot \beta \quad (2)$$

where $\gamma_{\text{dB}} = 10 \log_{10}(\gamma)$ is the SQNR in dB, κ is the quantization efficiency, and β is the bit-width. Note that the slopes of the lines in Figure 1 depict the optimal quantization efficiency for ideal distributions. The quantization efficiency for uniform distribution is the well-known value of 6dB/bit (Shi & Sun, 2008), while the quantization efficiency for Gaussian distribution is about 5dB/bit (You, 2010). Actual quantization efficiency for non-ideal inputs can be significantly lower. Our experiments show that the SQNR resulting from uniform quantization of the actual weights and activations in the DCN is between 2 to 4dB/bit.

3.2. Empirical distributions in a pre-trained DCN

Figure 2(a) and Figure 2(b) depict the empirical distributions of weights and activations, respectively for the convolutional layers of the DCN we designed for CIFAR-10 benchmark (see Section 5). Note that the activations plotted here are before applying the activation functions.

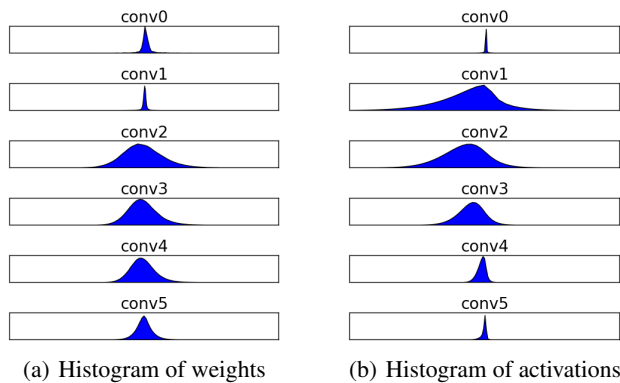


Figure 2. Distribution of weights & activations in a DCN design for CIFAR-10 benchmark.

Given the similarity of these distributions to the Gaussian distribution, in all our analysis we have assumed Gaussian distribution for both weights and activations. However, we also note that the distribution of weights and activations for some layers are less Gaussian-like. It will therefore be of interest to experiment with step-sizes for other distributions (see Table 1), which is beyond the scope of present work.

3.3. Model conversion

Any floating point DCN model can be converted to fixed point by following these steps:

- Run a forward pass in floating point using a large set of typical inputs and record the activations.

- Collect the statistics of weights, biases and activations for each layer.
- Determine the fixed point formats of the weights, biases and activations for each layer.

Note that determining the fixed point format is equivalent to determining the resolution, which in turn means identifying the number of fractional bits it requires to represent the number. The following equations can be used to compute the number of fractional bits:

- Determine the effective standard deviation of the quantity being quantized: ξ .
- Calculate step size via Table 1: $s = \xi \cdot \text{Stepsize}(\beta)$.
- Compute number of fractional bits: $n = -\lceil \log_2 s \rceil$.

In these equations,

- ξ is the effective standard deviation of the quantity being quantized, an indication of the width of the distribution we want to quantize. For example, if the quantized quantities follow an ideal zero mean Gaussian distribution, then $\xi = \sigma$, where σ is the true standard deviation of quantized values. If the actual distribution has longer tails than Gaussian, which is often the case as shown in Figure 2(a) and 2(b), then $\xi > \sigma$. In our experiments in Section 5, we set $\xi = 3\sigma$.
- $\text{Stepsize}(\beta)$ is the optimal step size corresponding to quantization bit-width of β , as listed in Table 1.
- n is the number of fractional bits in the fixed point representation. Equivalently, 2^{-n} is the resolution of the fixed point representation and a quantized version of s . Note that $\lceil \cdot \rceil$ is one choice of a rounding function and is not unique.

4. Bit-width optimization across a deep network

In the absence of model fine-tuning, converting a floating point deep network into a fixed point deep network is essentially a process of introducing quantization noise into the neural network. It is well understood in fields like audio processing or digital communications that as the quantization noise increases, the system performance degrades. The effect of quantization can be accurately captured in a single quantity, the SQNR.

In deep learning, there is not a well-formulated relationship between SQNR and classification accuracy. However, it is reasonable to assume that in general higher quantization noise level leads to worse classification performance. Given that SQNR can be approximated theoretically and analyzed layer by layer, we focus on developing a theoretical framework to optimize for the SQNR. We then conduct empirical investigations into how the proposed optimization for SQNR affect classification accuracy of the DCN.

4.1. Impact of quantization on SQNR

In this section, we will derive the relationship between the quantization of the weight, bias and activation values respectively, and the resulting output SQNR.

4.1.1. QUANTIZATION OF INDIVIDUAL VALUES

Quantization of individual values in a DCN, whether it is an activation or weight value, readily follows the quantizer discussion in Section 3.1. For instance, for weight value w , the quantized version \tilde{w} can be written as:

$$\tilde{w} = w + n_w, \quad (3)$$

where n_w is the quantization noise. As illustrated in Figure 1, if w approximately follows a uniform, Gaussian, Laplacian or Gamma distribution, the SQNR, γ_w , as a result of the quantization process can be written as:

$$10 \log(\gamma_w) = 10 \log \frac{\mathbb{E}[w^2]}{\mathbb{E}[n_w^2]} \approx \kappa \cdot \beta, \quad (4)$$

where κ is the quantization efficiency and β is the quantizer bit-width.

4.1.2. QUANTIZATION OF BOTH ACTIVATIONS AND WEIGHTS

Consider the case where weight w is multiplied by activation a , where both w and a are quantized with quantization noise n_w and n_a , respectively. The product can be approximated, for small n_w and n_a , as follows:

$$\begin{aligned} \tilde{w} \cdot \tilde{a} &= (w + n_w) \cdot (a + n_a) \\ &= w \cdot a + w \cdot n_a + n_w \cdot a + n_w \cdot n_a \\ &\cong w \cdot a + w \cdot n_a + n_w \cdot a. \end{aligned} \quad (5)$$

The last equality holds if $|n_a| \ll |a|$ and $|n_w| \ll |w|$. A very important observation is that the SQNR of the product, $w \cdot a$, as a result of quantization, satisfies

$$\frac{1}{\gamma_{w \cdot a}} = \frac{1}{\gamma_w} + \frac{1}{\gamma_a}. \quad (6)$$

This is characteristic of a linear system. The defining benefit of this realization is that introducing quantization noise to weights and activations independently is equivalent to adding the total noise after the product operation in a normalized system. This property will be used in later analysis.

4.1.3. FORWARD PASS THROUGH ONE LAYER

In a DCN with multiple layers, computation of the i th activation in layer $l+1$ of the DCN can be expressed as follows:

$$a_i^{(l+1)} = \sum_{j=1}^N w_{i,j}^{(l+1)} a_j^{(l)} + b_i^{(l+1)}, \quad (7)$$

where (l) represents the l th layer, N represents number of additions, $w_{i,j}$ represents the weight and b_i represents the bias.

Ignoring the bias term for the time being, since $a_i^{(i+1)}$ is simply a sum of terms like $w_{i,j}^{(i+1)} a_j^{(i)}$, which when quantized all have the same SQNR $\gamma_{w^{(i+1)} \cdot a^{(i)}}$. Assuming the product terms $w_{i,j}^{(i+1)} a_j^{(i)}$ are independent, it follows that the value of $a_i^{(i+1)}$, before further quantization, has inverse SQNR that equals

$$\frac{1}{\gamma_{w_{i,j}^{(i+1)} a_j^{(i)}}} = \frac{1}{\gamma_{w_{i,j}^{(i+1)}}} + \frac{1}{\gamma_{a_j^{(i)}}} = \frac{1}{\gamma_{w^{(i+1)}}} + \frac{1}{\gamma_{a^{(i)}}} \quad (8)$$

After $a_i^{(i+1)}$ is quantized to the assigned bit-width, the resulting inverse SQNR then becomes $\frac{1}{\gamma_{a^{(i+1)}}} + \frac{1}{\gamma_{w^{(i+1)}}} + \frac{1}{\gamma_{a^{(i)}}}$. We are not considering the biases in this analysis because, assuming that the biases are quantized at the same bit-width as the weights, the SQNR is dominated by the product term $w_{i,j}^{(i+1)} a_j^{(i)}$. Note that Equation 8 matches rather well with experiments as shown in Section 5.3, even though the independence assumption of $w_{i,j}^{(i+1)} a_j^{(i)}$ does not always hold.

4.1.4. FORWARD PASS THROUGH ENTIRE NETWORK

Equation 8 can be generalized to all the layers in a DCN (although we have found empirically that the approximation applies better for convolutional layers than fully-connected layers). Consider layer L in a deep network, where all the activations and weights are quantized. Extending Equation 8, we obtain the SQNR (γ_{output}) at the output of layer L as:

$$\frac{1}{\gamma_{\text{output}}} = \frac{1}{\gamma_{a^{(0)}}} + \frac{1}{\gamma_{w^{(1)}}} + \frac{1}{\gamma_{a^{(1)}}} + \dots + \frac{1}{\gamma_{w^{(L)}}} + \frac{1}{\gamma_{a^{(L)}}} \quad (9)$$

In other word, the SQNR at the output of a layer in DCN is the *Harmonic Mean* of the SQNRs of all preceding quantization steps. This simple relationship reveals some very interesting insights:

- All the quantization steps contribute equally to the overall SQNR of the output, regardless if it's the quantization of weights, activations, or input, and irrespective of where it happens (at the top or bottom of the network).
- Since the output SQNR is the harmonic mean, the network performance will be dominated by the worst quantization step. For example, if the activations of a particular layer has a much smaller bit-width than other layers, it will be the bottleneck of network performance, because based on Equation 9, $\gamma_{\text{output}} \leq \gamma_{a^{(l)}}$ for all l .

- Depth makes quantization more challenging, but not exceedingly so. The rest being the same, doubling the depth of a DCN will half the output SQNR (3dB loss). But this loss can be readily recovered by adding 1 bit to the bit-width of all weights and activations, assuming the quantization efficiency is more than 3dB/bit. However, this theoretical prediction will need to be empirically verified in future works.

4.1.5. EFFECTS OF OTHER NETWORK COMPONENTS

- **Batch normalization:** Batch normalization (Ioffe & Szegedy, 2015) improves the speed of training a deep network by normalizing layer inputs. After the network is trained, the batch normalization layer is a fixed linear transformation and can be absorbed into the neighboring convolutional layer or fully-connected layer. Therefore, the quantization effect due to batch normalization does not need to be explicitly modeled.
- **ReLU:** In Equation 7, for simplicity we omitted the activation function applied to $a_j^{(l)}$. When the activation function is ReLU and the quantization noise is small, all the positive values at the input to the activation function will have the same SQNR at the output, and the negative values become zero (effectively reducing the number of additions, N). In other words,

$$\begin{aligned} a_i^{(l+1)} &= \sum_{j=1}^N w_{i,j}^{(l+1)} g(a_j^{(l)}) + b_i^{(l+1)} \\ &= \sum_{j=1}^M w_{i,j}^{(l+1)} a_j^{(l)} + b_i^{(l+1)}, \end{aligned} \quad (10)$$

where $g(\cdot)$ is the ReLU function and $M \leq N$ is the number of $a_j^{(l)}$'s that are positive.

In this case, the ReLU function has little impact on the SQNR of $a_i^{(l+1)}$. ReLU only starts to affect SQNR calculation when the perturbation caused by quantization is sufficiently large to alter the sign of $a_j^{(l)}$. Therefore, our analysis may become increasingly inaccurate as the bit-width becomes too small (quantization noise too large).

- **Non-ReLU activations:** Other nonlinear activation functions such as tanh, sigmoid, PReLU functions are much harder to model and analyze. However, in Section 5.1 we will see that applying the analysis in this section to a network with PReLU activation functions still yields useful enhancements.

4.2. Cross-layer bit-width optimization

From Equation 9, it is seen that trade-offs can be made between quantizers of different layers to produce the same

γ_{output} . That is to say, we can choose to use smaller bit-widths for some layers by increasing bit-widths for other layers. For example, this may be desirable because of the following reasons:

- Some layers may require a large number of computations (multiply-accumulate operations). Reducing the bit-widths for these layers would reduce the overall network computation load.
- Some layers may contain a large number of network parameters (weights). Reducing the weight bit-widths for these layers would reduce the overall model size.

Interestingly, such objectives can be formulated as an optimization problem and solved exactly. Suppose our goal is to reduce model size while maintaining a minimum SQNR at the DCN output. We can use ρ_i as the scaling factor at quantization step i , which in this case represents the number of parameters being quantized in the quantization step. The problem can be written as:

$$\min_{\gamma_i} \sum_i \rho_i \left(\frac{10 \log \gamma_i}{\kappa} \right), \quad \text{s.t.} \quad \sum_i \frac{1}{\gamma_i} \leq \frac{1}{\gamma_{\min}} \quad (11)$$

where $10 \log \gamma_i$ is the SQNR in dB domain, and $(10 \log \gamma_i)/\kappa$ is the bit-width in the i th quantization step according to Equation 2. γ_{\min} is the minimum output SQNR required to achieve a certain level of accuracy. The optimization constraint follows from Equation 9 that the output SQNR is the harmonic mean of the SQNR of intermediate quantization steps.

Substituting by $\lambda_i = \frac{1}{\gamma_i}$ and removing the constant scalars from the objective function, the problem can be reformulated as:

$$\min_{\lambda_i} - \sum_i \rho_i \log \lambda_i, \quad \text{s.t.} \quad \sum_i \lambda_i \leq C \quad (12)$$

where the constant $C = \frac{1}{\gamma_{\min}}$. This is a classic convex optimization problem with the *water-filling* solution (Boyd & Vandenberghe, 2004): $\frac{\rho_i}{\lambda_i} = \text{constant}$.

Or equivalently,

$$\rho_i \gamma_i = \text{constant} \quad (13)$$

Recognizing that $10 \log \gamma_i = \kappa \beta_i$ based on Equation 2, the solution can be rewritten as:

$$\frac{10 \log \rho_i}{\kappa} + \beta_i = \text{constant} \quad (14)$$

In other words, the difference between the optimal bit-widths of two quantization steps is inversely proportional to the difference of ρ 's in dB, scaled by the quantization efficiency.

$$\beta_i - \beta_j = \frac{10 \log(\rho_j/\rho_i)}{\kappa} \quad (15)$$

This is a surprisingly simple and insightful relationship. For example, assuming $\kappa = 3\text{dB/bit}$, the bit-widths β_i and β_j would differ by 1 bit if ρ_j is 3dB (or 2x) larger than ρ_i . More specifically, for model size reduction, layers with more parameters should use relatively lower bit-width, as it leads to better model compression under the overall SQNR constraint.

5. Experiments

In this section we study the effect of reduced bit-width for both weights and activations versus traditional 32-bit single-precision floating point approach (16-bit half-precision floating point is expected to produce comparable results as single-precision in most cases). In particular, we will implement the fixed point quantization algorithm described in Section 3 and investigate the effectiveness of the bit-width optimization algorithm in Section 4. In addition, using the quantized fixed point network as the starting point, we will further fine-tune the fixed point network within the restricted alphabets of weights and activations.

5.1. Bit-width optimization for CIFAR-10 classification

We evaluate our proposed cross-layer bit-width optimization algorithm on the CIFAR-10 benchmark using the algorithm prescribed in Section 4.2.

Table 2. Parameters per layer in our CIFAR-10 network

Layer	Input channels	Output img size	Filter dim	Params ($\times 10^6$)
conv0	3	22 \times 22	3 \times 3	0.007
conv1	256	12 \times 12	3 \times 3	0.295
conv2	128	10 \times 10	3 \times 3	0.295
conv3	256	8 \times 8	3 \times 3	0.590
conv4	256	5 \times 5	3 \times 3	0.590
conv5	256	2 \times 2	7 \times 7	1.606
fc0	128	-	-	0.005

In Table 2, we compute the number of parameters in each layer of our CIFAR-10 network. Consider the objective of minimizing the overall model size. Provided that the quantization efficiency $\kappa = 3\text{dB/bit}$, our derivation in Section 4.2 shows that the optimal bit-width of layer conv0 and conv1 would differ by $10 \log(0.295/0.007)/\kappa = 5\text{bits}$. Similarly, assuming the bit-width for layer conv0 is β_0 , the subsequent convolutional layers should have bit-width values as indicated in Table 3.

In our experiment in this section, we will ignore the fully-connected layer and assume a fixed bit-width of 16 for both weights and activations. This is because fully-connected layers have different SQNR characteristics and need to be optimized separately. Although the fully-connected lay-

Table 3. Optimal bit-width allocation in our CIFAR-10 network, assuming the bit-width of layer conv0 is β_0

Layer	conv1	conv2	conv3	conv4	conv5
Bit-width	$\beta_0 - 5$	$\beta_0 - 5$	$\beta_0 - 6$	$\beta_0 - 6$	$\beta_0 - 8$

ers can often be quantized more aggressively than convolutional layers, since the number of parameters of fc0 is very small in this experiment, we will set the bit-width to a large value to eliminate the impact of quantizing fc0 from the analysis, knowing that the large bit-width of fc0 has very little impact on the overall model size. We will also set the activation bit-widths of all the layers to a large value of 16 because they do not affect the model size.

Figure 3 displays the model size vs. error rate in a scatter plot, we can clearly see the advantage of cross-layer bit-width optimization. When the model size is large (bit-width is high), the error rate saturates at around 6.9%. When the model size reduces below approximately 25Mbits, the error rate starts to increase quickly as the model size decreases. In this region, cross-layer bit-width optimization offers $> 20\%$ reduction in model size for the same performance.

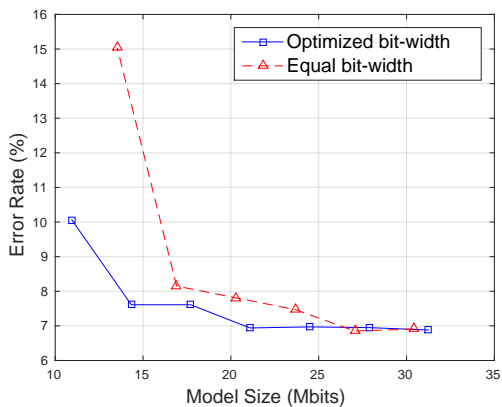


Figure 3. Model size vs. error rate with and without cross-layer bit-width optimization (CIFAR-10)

5.2. Bit-width optimization for ImageNet classification

In Section 5.1, we performed cross-layer bit-width optimization for our CIFAR-10 network with the objective of minimizing the overall model size while maintaining accuracy. Here we carry out a similar exercise for an AlexNet-like DCN that is trained on ImageNet-1000. The DCN architecture is described in Table 4.

For setting the bit-width of convolutional layers of this DCN, we follow the steps in Section 5.1 with the assumption that the bit-width for layer conv1 is β_1 . The resulting

Table 4. Parameters per layer in our AlexNet implementation

Layer	Input channels	Output img size	Filter dim	Params ($\times 10^6$)
conv1	3	112×112	7×7	0.014
conv2	96	28×28	5×5	0.384
conv3	160	14×14	3×3	0.277
conv4	192	14×14	3×3	0.332
conv5	192	14×14	3×3	0.277
fc1	160	-	-	16.056
fc2	2048	-	-	4.194

bit-width allocation for all convolutional layers is summarized in Table 5.

Table 5. Optimal bit-width allocation in our AlexNet-like network, assuming the bit-width of layer conv1 is β_1

Layer	conv2	conv3	conv4	conv5
Bit-width	$\beta_1 - 5$	$\beta_1 - 4$	$\beta_1 - 5$	$\beta_1 - 4$

For fully-connected layers we first keep the network as floating point and quantize the weights of fully-connected layers only. We then reduce bit-width of fully-connected layers until the classification accuracy starts to degrade. We found that the minimum bit-width for the fully-connected layers before performance degradation occurs is 6.

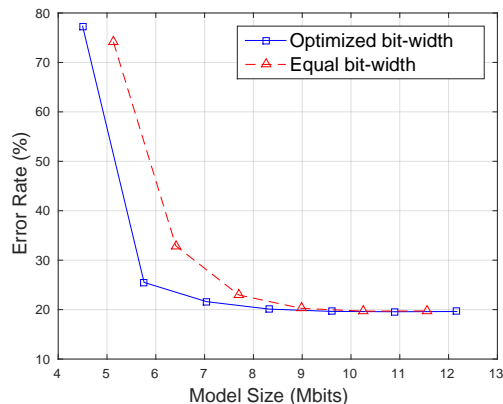


Figure 4. Model size vs. top-5 error rate with and without cross-layer bit-width optimization (ImageNet)

Figure 4 depicts the convolutional layer model size vs. top-5 error rate tradeoff for both optimized bit-width and equal bit-width scenarios. Similar to our CIFAR-10 network, there is a clear benefit of cross-layer bit-width optimization in terms of model size reduction. In some regions the saving can be upto 1Mbits.

However, unlike our CIFAR-10 network where the convo-

lutional layers make up most of the model size, in AlexNet-like DCN the fully-connected layers dominate in terms of number of parameters. With bit-width of 6, the overall size of fully-connected layers is more than 100Mbits. This means that the saving of 1Mbits brings in less than 1% of overall model size reduction! This is in clear contrast to the 20% model size reduction we reported for our CIFAR-10 network. Therefore, it is worth nothing that the proposed cross-layer layer bit-width optimization algorithm is most effective when the network size is dominated by convolutional layers, and is less effective otherwise.

The experiments on our CIFAR-10 network and AlexNet-like network demonstrate that the proposed cross-layer bit-width optimization offers clear advantage over the equal bit-width scheme. As summarized in Table 3 and Table 5, a simple offline computation of the inter-layer bit-width relationship is all that is required to perform the optimization.

However, in the absence of customized design, the implementation of the optimized bit-widths can be limited by the software or hardware platform on which the DCN operates. Often, the optimized bit-width needs to be rounded up to the next supported bit-width, which may in turn impact the network classification accuracy and model size.

5.3. Validation for SQNR prediction

To verify that our SQNR calculation presented in Section 4.1 is valid, we will conduct a small experiment. More specifically, we will focus on the optimized networks in Figure 3 and compare the measured SQNR per layer to the SQNR predictions according to Equation 4 and 8.

Table 6. Predicated SQNR vs. measured SQNR (in dB) in our CIFAR-10 network

	Example 1		Example 2	
	Predicted	Measured	Predicted	Measured
conv1	23.83	24.93	20.85	20.12
conv2	20.9	23.75	17.91	16.82
conv3	17.93	22.62	14.94	17.83
conv4	16.16	21.81	13.2	14.74
conv5	12.54	18.01	9.55	9.26

Table 6 contains the comparison between the theoretical SQNR and the measured SQNR (in dB) for layers conv1 to conv5 for two of the optimized networks. We observe that while the two SQNR values do not match numerically, they follow similar decreasing trend as the activations propagate deeper into the network. It should be noted that our theoretical SQNR predictions are based purely on the weight and activation bit-widths of each layer as well as the quantization efficiency κ . The theory does not rely on any in-

formation related to the network parameters or the data it is tested on.

5.4. Model fine-tuning

Although our focus is fixed point implementation without training, our quantizer design can also be used as a starting point for further model fine-tuning when the training model and training parameters are available.

Table 7. CIFAR-10 classification error rate with different bit-width combinations

Activation Bit-width	Weight Bit-width			
	4	8	16	Float
4	8.30	7.50	7.40	7.44
8	7.58	6.95	6.95	6.78
16	7.58	6.82	6.92	6.83
Float	7.62	6.94	6.96	6.98

Table 7 contains the classification error rate (in %) for the CIFAR-10 network after fine-tuning the model for 30 epochs. We experiment with different weight and activation bit-width combinations, ranging from floating point to 4-bit, 8-bit, and 16-bit fixed point. It is shown that even the (4b, 4b) bit-width combination works well (8.30% error rate) when the network is fine-tuned after quantization. In addition, the (float, 8b) setting generates an error rate of 6.78%, which is the new state-of-the-art result even though the activations are only 8-bit fixed point values. This may be attributed to the regularization effect of the added quantization noise (Lin et al., 2015; Luo & Yang, 2014).

6. Conclusions

Fixed point implementation of deep networks is important for real world embedded applications, which involve real time processing with limited power budget. In this paper, we develop a principled approach to converting a pre-trained floating point DCN model to its fixed point equivalent. We show that the naive method of quantizing all the layers in the DCN with uniform bit-width value results in DCN networks with subpar performance in terms of error rates relative to our proposed approach of SQNR based optimization of bit-widths. Specifically, we present results for a floating point DCN trained CIFAR-10 benchmark, which on conversion to its fixed point counter-part results in >20 % reduction in model size without any loss in accuracy. We note that our proposed algorithm facilitates easy conversion of any off-the-shelf DCN model for efficient real world on-device application. Finally, through fine-tuning experiments we demonstrate that our quantizer design methodology is a useful starting point for further model fine-tuning after the floating-point-to-fixed-point conversion.

Acknowledgements

We would like to acknowledge fruitful discussions and valuable feedback from our colleagues: David Julian, Anthony Sarah, Daniel Fontijne, Somdeb Majumdar, Aniket Vartak, Blythe Towal, and Mark Staskauskas.

References

- Boyd, Stephen and Vandenberghe, Lieven. *Convex Optimization*. Cambridge University Press, 2004.
- Chatfield, K., Simonyan, S., Vedaldi, V., and Zisserman, A. Return of the devil in the details: Delving deep into convolutional nets. *arXiv:1405:3531*, 2014.
- Chen, W., Wilson, J. T., Tyree, S., Weinberger, K. Q., and Chen, Y. Compressing neural networks with the hashing trick. In *ICML*, 2015.
- Courbariaux, M., Bengio, Y., and David, J. Low precision arithmetic for deep learning. *arXiv:1412.7024*, 2014.
- Deng, L., G.E., Hinton, and Kingsbury, B. New types of deep neural network learning for speech recognition and related applications: an overview. In *IEEE International Conference on Acoustic, Speech and Signal Processing*, pp. 8599–8603, 2013.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014.
- Gong, Y., Liu, L., Yang, M., and Bourdev, L. Compressing deep convolutional networks using vector quantization. *arXiv:1412.6115*, 2014.
- Gupta, S., Agrawal, A., Gopalakrishnan, K., and Narayanan, P. Deep learning with limited numerical precision. *arXiv:1502.02551*, 2015.
- Han, S., Mao, H., and Dally, W. J. A deep neural network compression pipeline: Pruning, quantization, Huffman encoding. *arXiv:1510.00149*, 2015.
- He, K., Zhang, X., Ren, S., and Sun, J. Spatial pyramid pooling in deep convolutional networks for visual recognition. *arXiv:1406:4729*, 2014.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. *arXiv:1512.03385*, 2015.
- Hinton, G., Vinyals, O., and Dean, J. Distilling knowledge in a neural network. In *Proc. Deep Learning and Representation Learning NIPS Workshop*, 2014.
- Ioffe, S. and Szegedy, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv:1502:03167*, 2015.
- Krizhevsky, A., Sutskever, I., and Hinton, G.E. ImageNet classification with deep convolutional neural networks. In *NIPS*, 2012.
- Kyuyeon, H. and Sung, W. Fixed-point feedforward deep neural network design using weights +1, 0 and -1. In *IEEE Workshop on Signal Processing Systems (SiPS)*, 2014.
- Lin, D. D. and Talathi, S. S. Overcoming challenges in fixed point training of deep convolutional networks. In *ICML Workshop*, 2016.
- Lin, Z., Courbariaux, M., Memisevic, R., and Bengio, Y. Neural networks with few multiplications. *arXiv:1510.03009*, 2015.
- Luo, Y. and Yang, F. Deep learning with noise. <http://www.andrew.cmu.edu/user/fanyang1/DLWN.pdf>, 2014.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. XNOR-Net: ImageNet classification using binary convolutional neural networks. *arXiv:1603.05279*, 2016.
- Razavian, A. S., Azizpour, H., Sullivan, J., and S., Carlsson. CNN features off-the-shelf: An astounding baseline for recognition. In *CVPR*, 2014.
- Sajid, A., Kyuyeon, H., and Sung, W. Fixed point optimization of deep convolutional neural networks for object recognition. In *IEEE International Conference on Acoustic, Speech and Signal Processing*, 2015.
- Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., and LeCun, Y. OverFeat: Integrated recognition, localization and detection using convolutional networks. *arXiv:1312:6229*, 2013.
- Shi, Yun Q. and Sun, Huifang. *Image and Video Compression for Multimedia Engineering: Fundamentals, Algorithms, and Standards*. CRC Press, 2008.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv:1409:1556*, 2014.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. *arXiv:1409.4842*, 2014.
- Vanhoucke, V., Senior, A., and Mao, M. Improving the speed of neural networks on CPUs. In *Proc. Deep Learning and Unsupervised Feature Learning NIPS Workshop*, 2011.
- You, Yuli. *Audio Coding: Theory and Applications*. Springer, 2010.

Zeiler, M.D. and Fergus, R. Visualizing and understanding convolutional neural networks. In *ECCV*, 2014.