ACIQ: ANALYTICAL CLIPPING FOR INTEGER QUANTIZATION OF NEURAL NETWORKS

Anonymous authors
Paper under double-blind review

ABSTRACT

We analyze the trade-off between quantization noise and clipping distortion in low precision networks. We identify the statistics of various tensors, and derive exact expressions for the mean-square-error degradation due to clipping. By optimizing these expressions, we show marked improvements over standard quantization schemes that normally avoid clipping. For example, just by choosing the accurate clipping values, more than 40% accuracy improvement is obtained for the quantization of VGG16-BN to 4-bits of precision. Our results have many applications for the quantization of neural networks at both training and inference time.

1 INTRODUCTION

A significant drawback of deep learning models is their computational costs. Low precision is one of the key techniques being actively studied recently to conquer the problem. With hardware support, low precision training and inference can compute more operations per second, reduce memory bandwidth and power consumption, and allow fit bigger networks into a device. In general, a low-precision scheme involves a floating-point to integer conversion, which introduces quantization noise into the network. This quantization noise is strongly linked to the dynamic range, defined as the range between the largest and smallest values that we are need to quantized. For a given N-bit integer representation, a smaller dynamic range leads to a smaller spacing between the 2^N quantization levels, enabling improved resolution and smaller quantization noise. To reduce this quantization noise, the dynamic range can be limited by clipping the values in the tensor. This clipping process introduces an additional noise because of the loss of information that otherwise would be carried by the clipped portion of the tensor. Hence, a trade-off between clipping and quantization effects exist. To find the best clippig value we need to minimize the information loss.

In this paper, we study the effect of clipping with the aim of improving overall quantization noise. To this end, we first study the distribution of values within these tensors. In all our measurements, the statistical distributions of weights and activations are observed to follow a bell-curve. This indicates that large values occur very rarely compared to small values, and suggests that the loss of information due to the clipping process might be compensated by improving the resolution of the more common smaller values.

To optimize this process further, it is essential to understand the underlying distributions of tensors before applying the clipping. By running a few statistical tests, we were able to see on a variety of convolution models that activation tensors follow either a Gaussian or Laplacian distributions with a high degree of certainty (p-value < 0.01). This modeling of activation tensors enable a clear formulation of the mean-square quantization error and constitute the first step for its minimization. In Section 4, we provide a rigorous formulation to optimize the quantization effect of activation tensors using clipping by analyzing both the Gaussian and the Laplace priors. This formulation is henceforth refered to as Analytical Clipping for Integer Quantization (ACIQ).

These analytical results have many applications for the quantization of neural networks at both training and inference time. For example, a straightforward quantization of the weights and activations to 8-bit fixed point representation has been shown to have a negligible effect on the model accuracy. Yet, in the majority of the applications, further reduction of precision quickly degrades performance, calling for an optimal clipping scheme to minimize information-loss during quantization. On a more general level, exploiting the statistics of the activation tensors to minimize their quantization toll is orthogonal to other techniques for network quantization. It can work in synergy with other schemes

to achieve more than they could have by working individually. Finally, it is easy to implement and requires only the adjustment of clipping value according to an analytical formula.

We further demonstrate the applicability and usability of our analytic terms on the following challenging problem. Given a pre-trained network, we would like to quantize the weights to 8-bit of precision and most activations to 4-bits of precision without any further processing (e.g., re-training). This specific setting is of a particular interest due to quantization of activations to 4-bits, which alleviates a major bottleneck in terms of memory bandwidth. Prior attempts using standard techniques (Krishnamoorthi, 2018) show severe degradation on accuracy. While several recent works were able to overcome this issue by additional re-training (McKinstry et al., 2018), this is not feasible in many practical settings, e.g., we often do not have the dataset on which the network is working on. By using smart clipping, we were able to significantly improve validation accuracy for configurations with 8-bit weight and 4-bit activation as summarized in Table 1.

Model	**Dataset**	**Accuracy Improvement**
VGG16	ImageNet	16.4%
VGG16-BN	ImageNet	42.8%
ResNet-18	ImageNet	16.0%
ResNet-50	ImageNet	22.8%
ResNet-101	ImageNet	26.2%
Inception v3	ImageNet	34.5%

Table 1: The effect of ACIQ on validation error (8-bit weights, 4-bit activations)

.

The methods introduced in this work may be additionally useful to current and future applications, such as the attempts to fully train in a low precision setting (Banner et al., 2018). In this scenario, intermediate activations and weights are known to have their statistics change throughout the training process, further highlighting the need for a precise and careful scale determination.

2 PREVIOUS WORK

In many cases, taking a model trained for full precision and directly quantizing it to 8-bit precision, without any re-training, can result in a relatively low loss of accuracy (Jacob et al., 2017a; Gong et al., 2018; Krishnamoorthi, 2018). More recently, (Banner et al., 2018; Wu et al., 2018) has shown that 8-bit precision is sufficeint not just for running trained models, but also for training them. Yet, naively quantizing a full precision model below 8-bits usually incurs significant accuracy degradation.

Many attempts have been made to diminish this effect, but they usually employ training of the model with quantization constraints or modifying network structure (Lin et al., 2017). Recent work has been able to quantize weights and activations to 4-bits, but requires a long training time of 110 epochs (McKinstry et al., 2018). Other concepts to enable quantization below 8-bits have been proposed (Rastegari et al., 2016; Zhou et al., 2016; Choi et al., 2018). Per-channel scale has shown to be important both for inference (Rastegari et al., 2016) and for training (Wu et al., 2018). The concept of non-uniform quantization has recently been suggested by (Park et al., 2018) and (Baskin et al., 2018).

To the best of our knowledge, there have been only a few attempts to clip activations before. Nvidia proposes an iterative method to search for a good clipping threshold based on the Kullback-Leibler divergence measure (Migacz, 2017). More recently, (Choi et al., 2018; Jung et al., 2018) have proposed an activation clipping parameter that is optimized during training. Finally, Wu et al. (2018) has proposed a heuristic approach to adjust the clipping values according to the desired precision. As contrast to previous works, we propose a simple and efficient method to find the optimal clipping threshold analytically from the distribution of the tensor, by minimizing MSE error. Our method takes into account not only the target precision but also the statistics of the tensor. We prove analytically that our method finds an optimal solution for threshold that minimizes MSE error and demonstrate its usefulness on several Imagenet topologies.

3 PROBABILITY DISTRIBUTION FITTING FOR ACTIVATIONS

In this section, we construct an estimate of the underlying probability density function of the activation tensors in deep learning models. Traditionally, no clipping was made in standard GEMMLWOP. Hence, quantization levels are uniformity spaced between the largest and the smallest values in the tensor. Yet, this approach is non-optimal, due to the fact that the activation tensors have bell-shaped distributions. Therefore, to reduce quantization noise (or increase resolution), it might be desirable to clip the values above a certain threshold. In order to find the best clipping value, that on the one hand, maintains low clipping rate, but on the other hand improves resolution, we must understand the underlying data distribution.

To that end, we collect data of various tensors at different layers. We observe that data is in general symmetrically distributed around a mean. We next estimate the goodness of fit to several bell-shaped distributions. This is done by measuring the static distance (largest vertical line) between the cumulative distribution function (CDF) of the empirical observed distribution and the CDF of the reference distribution (also know by Kolmogorov-Smirnov test (Lopes, 2011). By considering the activations of all layers of ResNet50 on ImageNet, we obtain the average static distance to each of the following distributions with p-value< 0.01:

Distribution	**Static Distance**
Laplace	0.070
Normal	0.053
Logistic	0.150
Cauchy	0.142
Uniform	0.490
Loglaplace	0.505
Lognorm	0.540

As one can see, the best fit is established by the Laplace and Normal distributions. In figure 3 (see Appendix), we plot the normalized distributions of the activation tensors at different layers for Resnet50, and compare them against both priors.

4 AN ANALYTICAL MODEL OF TENSOR QUANTIZATION

In this section, we provide a detailed analysis for establishing the optimal clipping values under either Gaussian or Laplace distributions. We first derive a generic expression for any given distribution for the expected MSE as a function of clipping value. We then use this expression to develop a specific expression for each distribution. Finally, we establish the optimal clipping values by solving the equations for which the derivative with respect to the clipping value are set to zero.

Let X be a high precision tensor with a probability density function $f(x)$. Without loss of generality, we assume a prepossessing step has been made so that the average value in the tensor zero i.e., $\overline{X} = \mu = 0$ (we do not lose generality since we can always subtract and add this mean). Assuming bit-width M, we would like to quantize the values in the tensors uniformally to 2^M discrete values.

Commonly (e.g., in GEMMLOWP), integer tensors are uniformly quantized in the range $[\alpha, -\alpha]$, where α is determined by the tensor maximal absolute value. In the following we show that the this choice of α is suboptimal, and suggest a model where the tensor values are clipped to reduce quantization noise. For any $x \in \mathbb{R}$, we define the clipping function $\text{clip}(x, \alpha)$ as follows

$$\text{clip}(x, \alpha) = \begin{cases} x & \text{if } |x| \leq \alpha \\ \text{sign}(x) \cdot \alpha & \text{if } |x| > \alpha \end{cases} \tag{1}$$

Denoting by α the clipping value, the range $[\alpha, -\alpha]$ is partitioned to 2^M equal quantization regions. Hence, the quantization step Δ between two adjacent quantized values is established as follows:

$$\Delta = \frac{2\alpha}{2^M} \tag{2}$$

Our model assumes values are rounded to the midpoint of the region (bin) i.e., for every index $i \in [0, 2^M - 1]$ all values that fall in $[-\alpha + i \cdot \Delta, -\alpha + (i+1) \cdot \Delta]$ are rounded to $-\alpha + (2i+1)\frac{\Delta}{2}$.

Then, the expected mean-square-error between X and its quantized version $Q(X)$ can be written as follows:

$$\begin{aligned}
&E[(X - Q(X))^2] = \\
&= \int_{-\infty}^{-\alpha} f(x) \cdot (x - \alpha)^2 dx + \\
&+ \sum_{i=0}^{2^M - 1} \int_{-\alpha + i \cdot \Delta}^{-\alpha + (i+1) \cdot \Delta} f(x) \cdot \left[x - \left(-\alpha + \frac{(2i+1)\Delta}{2} \right) \right]^2 dx + \\
&+ \int_{\alpha}^{\infty} f(x) \cdot (x - \alpha)^2 dx
\end{aligned} \tag{3}$$

Equation 3 is composed of three parts. The first and last terms quantify the contribution of $\mathrm{clip}(x, \alpha)$ to the expected mean-square-error. Note that for symmetrical distributions around zero (e.g., Gaussian $N(0, \sigma)$ or Laplace $\mathrm{Laplace}(0, b)$) these two terms are equal and their sum can therefore be evaluated by multiplying any of the terms by 2.

The second term corresponds to the expected mean-square-error when the range $[-\alpha, \alpha]$ is quantized uniformly to 2^M discrete levels. We take the common additive orthogonal noise model for the analysis of quantization error. The validity of this model is well established in cases of sufficient resolution. For more details see Marco & Neuhoff (2005). Accordingly, given a quantizaiton level $i \leq 2^M$, we approximate the second term by replacing the probability density function $f(x)$ by the uniform density function $\frac{1}{\Delta}$, as follows:

$$\begin{aligned}
&\int_{-\alpha + i \cdot \Delta}^{-\alpha + (i+1) \cdot \Delta} f(x) \cdot \left(x + \alpha - \frac{(2i+1)\Delta}{2} \right)^2 dx \approx \\
&\approx \frac{1}{2^M} \int_{-\alpha + i \cdot \Delta}^{-\alpha + (i+1) \cdot \Delta} \frac{1}{\Delta} \cdot \left(x + \alpha - \frac{(2i+1)\Delta}{2} \right)^2 dx = \\
&= \frac{1}{3 \cdot 2^M \cdot \Delta} \cdot \left(x + \alpha - \frac{(2i+1)\Delta}{2} \right)^3 \Bigg|_{-\alpha + i \cdot \Delta}^{-\alpha + (i+1) \cdot \Delta} = \\
&= \frac{\Delta^2}{12 \cdot 2^M}
\end{aligned}$$

Substituting $\Delta = \frac{2\alpha}{2^M}$, provides an evaluation for the second term in Equation 3 as follows:

$$\begin{aligned}
&\sum_{i=0}^{2^M - 1} \int_{-\alpha + i \cdot \Delta}^{-\alpha + (i+1) \cdot \Delta} f(x) \cdot \left(x + \alpha - \frac{(2i+1)\Delta}{2} \right)^2 dx \approx \\
&\approx \sum_{i=0}^{2^M - 1} \frac{\Delta^2}{12 \cdot 2^M} = 2^M \cdot \frac{\Delta^2}{12 \cdot 2^M} \\
&= \frac{\alpha^2}{3 \cdot 2^{2M}}
\end{aligned} \tag{4}$$

By substituting Equation 4 into Equation 3, and using the symmetrical argument mentioned above, Equation 3 can be simplified for symmetrical distributions as follows:

$$E[(X - Q(X))^2] = \frac{\alpha^2}{3 \cdot 2^{2M}} + 2 \cdot \int_{\alpha}^{\infty} f(x) \cdot (x - \alpha)^2 dx \tag{5}$$

In the following we provide a closed form solution for the case where the density probability distribution function $f(x)$ is either Gaussian $N(0, \sigma)$ or Laplace $\mathrm{Laplace}(0, b)$.

4.1 Laplace Clipping

In the following we develop an expression based on Equation 5 for the Laplace case. Since we assume that $\mu = 0$, we have the following Laplace density function $f(x) = \frac{1}{2b} \mathrm{e}^{-\frac{|x|}{b}}$. In order to

derive a closed form solution for Equation 5, we need to evaluate

$$2 \cdot \int_{\alpha}^{\infty} f(x) \cdot (x-\alpha)^2 dx \tag{6}$$

Let $\Psi(x)$ represent the expression below:

$$\Psi(x) = -\frac{1}{2}\left[x^2 + 2\left(b-\alpha\right)x + 2b^2 - 2\alpha b + \alpha^2\right] \mathrm{e}^{-\frac{x}{b}} + C \tag{7}$$

By taking the derivative of $\Psi(x)$ with respect to x, it is easy to see that $\Psi(x)$ is the correct antiderivative of the integrand in equation 6. Hence,

$$\int_{\alpha}^{\infty} f(x) \cdot (x-\alpha)^2 dx = \Psi(\inf) - \Psi(\alpha) = b^2 \cdot e^{-\frac{\alpha}{b}}$$

We can finally state Equation 5 for the laplace case as follows.

$$E[(X - Q(X))^2] \approx 2 \cdot b^2 \cdot e^{\frac{-\alpha}{b}} + \frac{\alpha^2}{3 \cdot 2^{2M}} \tag{8}$$

Finally, in order to find the α that give the minimum MSE, the corresponding derivative with respect to α is set equal to zero as follows:

$$\frac{\partial E[(X - Q(X))^2]}{\partial \alpha} = \frac{2\alpha}{3 \cdot 2^{2M}} - 2b\mathrm{e}^{-\frac{\alpha}{b}} = 0 \tag{9}$$

4.2 Gaussian Clipping

We now turn to evaluate Equation 5 for the Gaussian case. Given a Gaussian random variable $X \sim N(0, \sigma)$, we define $\Psi(x)$ to represent the expression below:

$$\Psi(x) = \frac{\left(\alpha^2 + \sigma^2\right) \operatorname{erf}\left(\frac{x}{\sqrt{2}\sigma}\right)}{2} - \frac{\left(x\sigma - 2\alpha\sigma\right) \mathrm{e}^{-\frac{x^2}{2\sigma^2}}}{\sqrt{2\pi}} \tag{10}$$

As in subsection 4.1, one can observe that by taking the derivative of $\Psi(x)$ with respect to x, it is easy to show that $\Psi(x)$ is the correct antiderivative of Equation 6 for the case where f represents the Gaussian density function i.e., $f(x) = \frac{1}{\sqrt{2\pi}\sigma}\mathrm{e}^{-\frac{x^2}{2\sigma^2}}$. Next, we use $\Psi(x)$ on the range $[\alpha, \infty]$ to evaluate Equation 6 for the Gaussian case as follows:

$$\begin{aligned}
&\int_{\alpha}^{\infty} f(x) \cdot (x-\alpha)^2 dx = \Psi(\infty) - \Psi(\alpha) \\
&= \frac{\alpha^2 + \sigma^2}{2} \cdot \left[1 - \operatorname{erf}\left(\frac{\alpha}{\sqrt{2}\sigma}\right)\right] - \frac{\alpha \cdot \sigma \cdot \mathrm{e}^{-\frac{\alpha^2}{2 \cdot \sigma^2}}}{\sqrt{2\pi}}
\end{aligned}$$

Equation 5 can thus be written for the case of Gaussian distribution as follows:

$$E[(X - Q(X))^2] \approx \frac{\alpha^2 + \sigma^2}{2} \cdot \left[1 - \operatorname{erf}\left(\frac{\alpha}{\sqrt{2}\sigma}\right)\right] - \frac{\alpha \cdot \sigma \cdot \mathrm{e}^{-\frac{\alpha^2}{2 \cdot \sigma^2}}}{\sqrt{2\pi}} + \frac{\alpha^2}{3 \cdot 2^{2M}} \tag{11}$$

In order to find the optimal clipping values for which mean-square-error is minimized, we need to differentiate $E[(X - Q(X))^2]$ with respect to α and set the derivative equal zero as follows.

$$\frac{\partial E[(X - Q(X))^2]}{\partial \alpha} = \alpha\left[1 - \operatorname{erf}\left(\frac{\alpha}{\sqrt{2}\sigma}\right)\right] - \frac{\sigma^2 \mathrm{e}^{-\frac{\alpha^2}{2\sigma^2}}}{\sqrt{2\pi}\sigma} - \frac{\sigma \mathrm{e}^{-\frac{\alpha^2}{2\sigma^2}}}{\sqrt{2\pi}} + \frac{2\alpha}{3 \cdot 2^{2M}} \tag{12}$$

5 Experiments

We evaluate the proposed ideas of analytical clipping on multiple models, and demonstrate its usefulness by suggesting a fast method to quantize neural network aggressively and without a significant

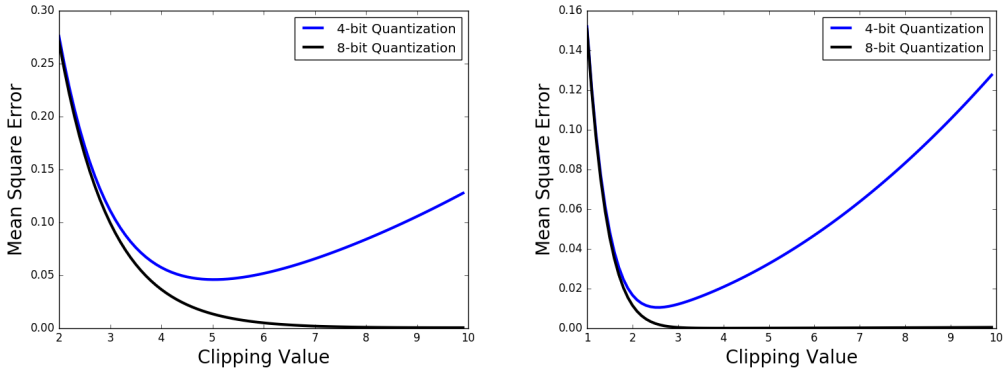

(a) Laplace ($\mu = 0$ and $b = 1$)

(b) Gaussian ($\mu = 0$ and $\sigma = 1$)

Figure 1: Expected mean-square-error as a function of clipping value for 4-bit and 8-bit quantization as predicted by equations 8 and 11

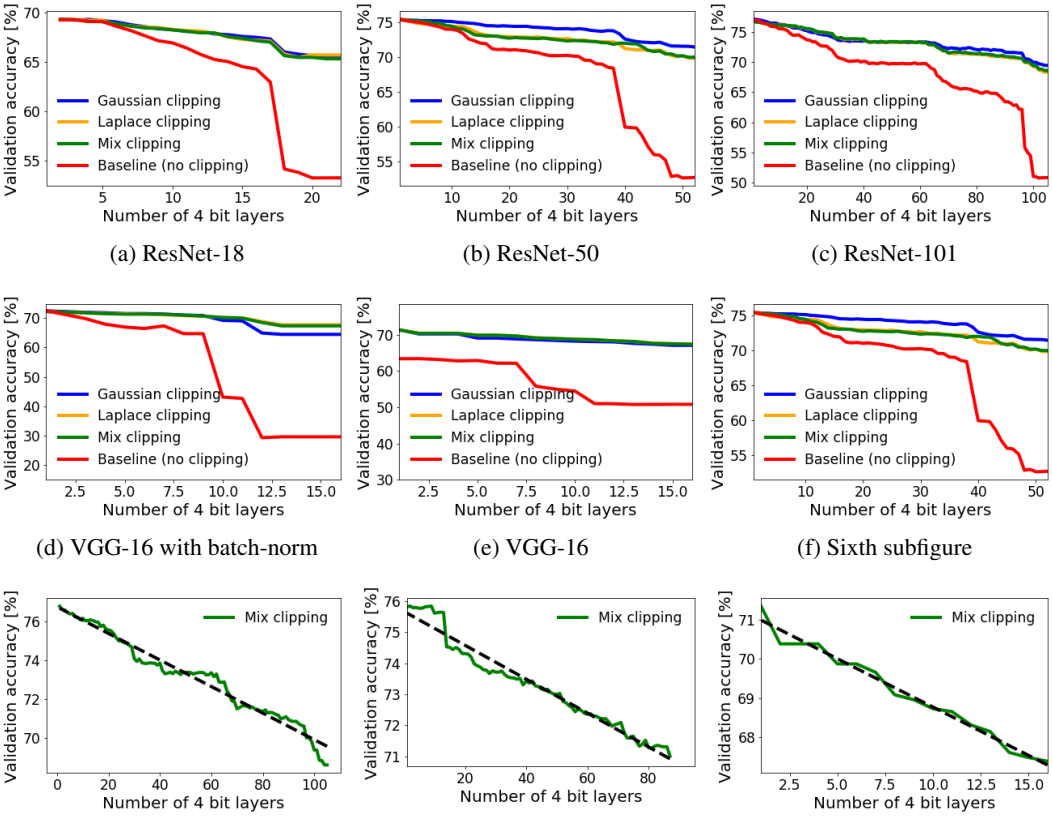

(a) ResNet-18

(b) ResNet-50

(c) ResNet-101

(d) VGG-16 with batch-norm

(e) VGG-16

(f) Sixth subfigure

(g) ResNet-101

(h) InceptioinV3

(i) VGG-16

Figure 2: Accuracy degradation as a function of quantization level: much smaller and predictable accuracy degradation with optimal clipping. Red curve represents the standard GEMMLOWP quantization method; yellow curve represents the accuracy of our method, assuming Laplace distribution; blue curve represents the accuracy of our method, assuming Gaussian distribution; green curve represents the accuracy of our method when distribution is selected at according to MSE estimation.

accuracy drop. Our method is simple and efficient, and does not require re-training. During quantization we only adjust the clipping values according to our analytical formula, and compare it against the standard GEMMLOWP quantization (Jacob et al., 2017b) that avoids clipping. We consider a

mixed precision scheme of 4 and 8 bits quantization. All weights are kept at 8-bit precision. The goal is to quantize as many activation layers as possible to 4-bits of precision without a significant accuracy degradation. This setting is important as activation tensors constitute a major memory bottleneck for inference in large models. The code to replicate all our experiments is available online.[1]

For each particular activation tensor, we use Equations 8 and 11 to evaluate the expected mean-square-error (MSE). Based on the expected MSE and a quantization requirement (i.e., the number of activation layers that need to be quantized to 4-bits), we decide whether to keep the tensor at 8-bits of precision or quantize the tensor into a 4-bit representation. Our analysis in Section 4 suggests that no clipping is required when a tensor is quantized to 8 bits of precision (see figure 1). On the other hand, clipping is important when the tensor is quantized to 4-bits of precision. In that case, we use the optimal clipping values stored in a lookup table. Given a tensor with estimated statistical parameters (e.g., std or mean), we search for the optimal clipping value that was pre-calculated offline (by solving either equation 9 or equation 12) and stored for fast retrieval.

We ran experiments to see how many activation layers can be quantized from 8 to 4 bits of precision without a significant accuracy drop. These experiments have been made with ResNet-18, ResNet-50, ResNet-101, VGG-16, VGG-16 with batch normalization and Inception-v3 on the ImageNet dataset. For each method, we select N layers (tensors) to be quantized to 4 bits using our optimized clipping method and compare it against the standard GEMMLOWP approach. In figure 2 we present this accuracy-quantization tradeoff. Our results clearly show the significant potential of our optimized clipping method. For example, just by choosing the correct clipping values, more than half of the layers can be quatized to 4 bits with an accuracy drop of less than 1% for Res18, Res50 and both versions of Vgg16. Also, unlike the standard approach, degradation is much more predictable and follow a roughly linear trend. Therefore, by adjusting the number of layers, we can conveniently and linearly control the accuracy-quantization trade-off, as can be seen by the linear relation shown in Figures 2g, 2h and 2i.

6 Conclusion

We introduce ACIQ - an optimized clipping framework for improved quantization of neural networks. Optimized clipping is shown to have a drastic impact on quantization in a variety of models. The underlying reason lies in the statistical dispersion of activations, where large values occur very rarely. We show the bell-curve statistics of activations are best fit as either Laplace or Gaussian distributions, and formulate the clipping process as an optimization problem. The solution to this optimization problem constitutes a polynomial-exponential equation that can be calculated numerically for a variety of statistical parameters, and stored in a lookup table for fast retrieval. This scheme is very simple and easy to implement either in software or in hardware.

While results are very encouraging, this work is only the first step on the ladder for successful deployment of clipping in neural networks. First, our main focus in this work is quantization of activations, while similar evaluation still needs to be done for weights. On a more general level, our framework is not restricted to the inference settings and can be extended to training. For example, our preliminary results show that quantization of gradients might benefit from clipping of the small values (i.e., sparsification). Establishing the correct threshold for gradients is yet another important direction for future work. While much work still needs to be done with regards to optimized clipping, we believe our work clearly demonstrates the major importance of this concept for the quantization of neural networks.

References

Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable methods for 8-bit training of neural networks. *arXiv preprint arXiv:1805.11046*, 2018.

Chaim Baskin, Eli Schwartz, Evgenii Zheltonozhskii, Natan Liss, Raja Giryes, Alex M Bronstein, and Avi Mendelson. Uniq: Uniform noise injection for the quantization of neural networks. *arXiv preprint arXiv:1804.10969*, 2018.

[1] https://github.com/submission2019/AnalyticalScaleForIntegerQuantization

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

Jiong Gong, Haihao Shen, Guoming Zhang, Xiaoli Liu, Shane Li, Ge Jin, Niharika Maheshwari, Evarist Fomenko, and Eden Segal. Highly efficient 8-bit low precision inference of convolutional neural networks with intelcaffe. *arXiv preprint arXiv:1805.08691*, 2018.

Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *arXiv preprint arXiv:1712.05877*, 2017a.

Benoit Jacob et al. gemmlowp: a small self-contained low-precision gemm library.(2017), 2017b.

Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Youngjun Kwak, Jae-Joon Han, and Changkyu Choi. Joint training of low-precision neural network with quantization interval parameters. *arXiv preprint arXiv:1808.05779*, 2018.

Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.

Xiaofan Lin, Cong Zhao, and Wei Pan. Towards accurate binary convolutional neural network. In *Advances in Neural Information Processing Systems*, pp. 345–353, 2017.

Raul HC Lopes. Kolmogorov-smirnov test. In *International encyclopedia of statistical science*, pp. 718–720. Springer, 2011.

Daniel Marco and David L Neuhoff. The validity of the additive noise model for uniform scalar quantizers. *IEEE Transactions on Information Theory*, 51(5):1739–1755, 2005.

Jeffrey L McKinstry, Steven K Esser, Rathinakumar Appuswamy, Deepika Bablani, John V Arthur, Izzet B Yildiz, and Dharmendra S Modha. Discovering low-precision networks close to full-precision networks for efficient embedded inference. *arXiv preprint arXiv:1809.04191*, 2018.

S Migacz. 8-bit inference with tensorrt. In *GPU Technology Conference*, 2017.

Eunhyeok Park, Sungjoo Yoo, and Peter Vajda. Value-aware quantization for training and inference of neural networks. *arXiv preprint arXiv:1804.07802*, 2018.

Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European Conference on Computer Vision*, pp. 525–542. Springer, 2016.

Shuang Wu, Guoqi Li, Feng Chen, and Luping Shi. Training and inference with integers in deep neural networks. *arXiv preprint arXiv:1802.04680*, 2018.

Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.

Appendices

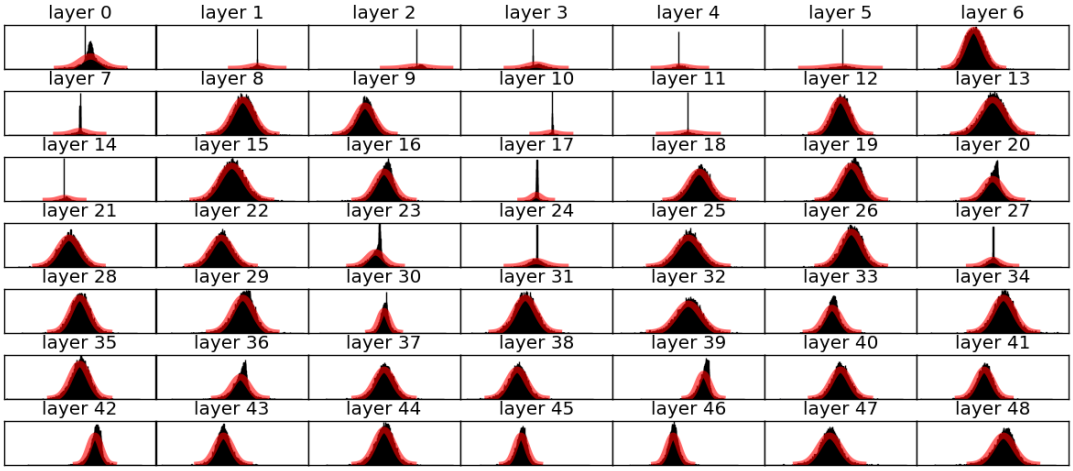

(a) Fitting the activations with a Gaussian Distribution

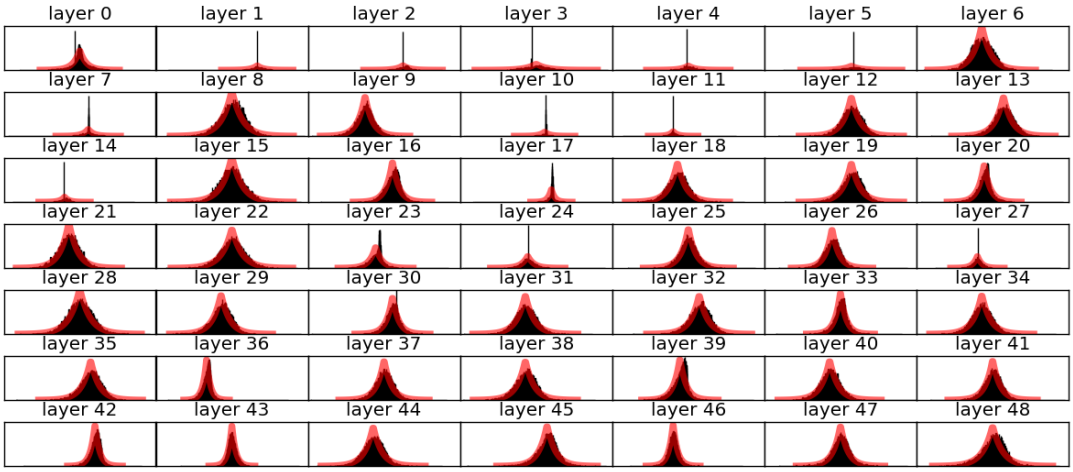

(b) Fitting the activations with a Laplace Distribution

Figure 3: Fitting activation tensors to Gauss and Laplace distribution at different ResNet50 layers. The statistical fit to both models is acceptable and similar.