STATQAT: STATISTICAL QUANTIZER OPTIMIZATION FOR DEEP NETWORKS

Anonymous authorsPaper under double-blind review

ABSTRACT

Quantization is essential for reducing the computational cost and memory usage of deep neural networks, enabling efficient inference on low-precision hardware. Despite the growing adoption of uniform and floating-point quantization schemes, selecting optimal quantization parameters remains a key challenge, particularly for diverse data distributions encountered during training and inference. This work presents a novel statistical error analysis framework for uniform and floating-point quantization, providing theoretical insight into error behavior across quantization configurations. Building on this analysis, we propose iterative quantizers designed for arbitrary data distributions and analytic quantizers tailored for Gaussian-like weight distributions. These methods enable efficient, low-error quantization suitable for both activations and weights. We incorporate our quantizers into quantization-aware training and evaluate them across integer and floating-point formats. Experiments demonstrate improved accuracy and stability, highlighting the effectiveness of our approach for training low-precision neural networks. ¹

1 Introduction

Quantization is a key optimization for deploying large deep learning models efficiently (1). By reducing parameter precision, quantization decreases model size and improves memory bandwidth utilization through reduced data movement (2; 3). When activations are also quantized, forward-pass computations can leverage specialized low-precision hardware units such as INT8, FP8, and FP4 matrix multipliers, which are widely used in linear and convolution layers (4; 5). These optimizations collectively enable low-latency, high-throughput inference (4).

Quantization schemes are typically categorized as non-uniform, uniform, or floating-point. Non-uniform quantization offers maximum flexibility by allowing arbitrary placement of quantization levels, but its hardware incompatibility limits its use primarily to weight-only quantization (6). Uniform quantization constrains levels to be equally spaced, enabling efficient implementation with integer formats such as INT8/4/2 or UINT8/4/2 (7; 8). Floating-point quantization, such as FP8 or FP4, follows the standard floating-point representation, where levels are dense near zero and expand exponentially, making it well-suited for parameters clustered around zero (9; 10).

In this work, we introduce a statistical framework for analyzing and minimizing quantization error in uniform and floating-point schemes. Our approach enables the selection of optimal quantization parameters for weights and activations. We propose iterative quantizers tailored to arbitrary data distributions observed in activations, and analytic quantizers optimized for Gaussian-distributed weights, a common assumption during training (11). We focus on quantization-aware training (QAT) to validate these quantizers across integer and floating-point formats. Importantly, our analysis extends to floating-point formats such as FP4, which are gaining native hardware support in modern accelerators like NVIDIA Blackwell. To our knowledge, this is the first closed-form analytic solution for FP quantizers in QAT, offering both theoretical novelty and practical hardware relevance.

¹Code is submitted with the paper. Upon acceptance, it will be publicly available.

2 BACKGROUND

The quantization operator is a many-to-one function, formulated in its general form as follows:

$$Q(x|\mathbf{l},\mathbf{t}) = l_k, \quad \text{if } t_k < x \le t_{k+1}, \tag{1}$$

where $x \in \mathbb{R}$ denotes a scalar input to be quantized, $t \in \mathbb{R}^N$ is a list containing N sorted thresholds defining the predetermined quantization intervals, and $l \in \mathbb{R}^{N-1}$ contains corresponding quantization levels for each interval. Given x, the operator compares its value to N-1 intervals defined by N thresholds t and assigns it to the corresponding level l_k .

2.1 Non-uniform quantization

The non-uniform quantization scheme does not constrain the choice of levels and thresholds, making it highly flexible. There are 2N-1 parameters, including thresholds and levels. In practice, thresholds are often chosen as the midpoints between consecutive levels, such as $t_k=(l_k+l_{k-1})/2$ for k=1:N-2, and $t_0=-\infty$, $t_{N-1}=\infty$, to cover the full input range. This choice minimizes the mean-squared quantization error given a fixed set of levels (12).

2.2 Uniform quantization

In this scheme, \boldsymbol{l} is constrained to have evenly spaced N-1 quantization levels, which can be determined using a scale parameter s and a shift parameter z. The levels and thresholds are then defined as: $l_k = sk + z$ for k = 0: N-2 and $t_k = s(k-1/2) + z$ for k = 1: N-2. This scheme underlies integer quantization methods such as INT8, INT4, and INT2, which are widely deployed in hardware-efficient inference engines (4; 3).

When z is set to -(N/2-1)s, the levels are symmetric around zero: $l_k = [-(N/2-1)s, \ldots, (N/2-1)s]$, and the thresholds become $t = [-\infty, -(N-3)s/2, \ldots, (N-3)s/2, \infty]$. This symmetric setting reduces parameter count at the expense of increased quantization error. Appendix A provides implementation details of uniform quantization.

2.3 FLOAT QUANTIZATION

A float number comprises a sign bit, E exponent bits, and M mantissa bits. Its value is computed as $(-1)^s \times 2^{(e-b)} \times (f_0 + \sum_{m=1}^M f_m 2^{-m})$, where b is a bias term for the exponent. E bits define the exponent value e, and $\{f_m\}_{m=1:M}$ define the fraction. This representation introduces non-uniform spacing, where smaller values are represented with higher resolution, making float formats such as FP8 and FP4 particularly suited for data clustered around zero (10; 11).

The floating-point standard includes subnormal and normal regions (13). When $f_0 = 0$, subnormal representation is active to handle very small magnitudes. When $f_0 = 1$, the normal grid is active. The positive half of the grid points are denoted as:

$$g_p = \begin{cases} p2^{1-M-b}, & \text{if } p \in [0, 2^M - 1] \\ 2^{\left\lfloor \frac{p}{2^M} \right\rfloor - b} (1 + (p \mod 2^M)2^{-M}), & \text{if } p \in [2^M, 2^{M+E} - c - 1] \end{cases}$$
 (2)

for $p=0:2^{M+E}-c-1$ where c accounts for special values like NaNs or INFs. This grid doubles its spacing every 2^M points. Including the negative half, the total number of grid points becomes $2^{M+E+1}-2c-1$. With scale s and shift z parameters, the quantization levels l are defined as

$$l_k = \begin{cases} z - sg_{-k+2^{M+E}-c-1}, & \text{if } k \in [0, 2^{M+E}-c-2] \\ z + sg_{k-2^{M+E}+c+1}, & \text{if } k \in [2^{M+E}-c-1, 2^{M+E+1}-2c-2] \end{cases}$$
(3)

where s and z are parameters to be determined. See Appendix B for the practical implementation.

2.4 QUANTIZED TRAINING

The training objective is to minimize the negative log-likelihood and a regularization term:

$$\mathcal{L}(\boldsymbol{\theta}) = -\log p(\boldsymbol{\theta}) - \sum_{n=1}^{N} \log p(\boldsymbol{y}_n | \boldsymbol{x}_n, \boldsymbol{\theta}), \tag{4}$$

where $\theta \in \mathbb{R}^M$ is the set of model parameters. The prior is often factorized across layers: $\log p(\theta) = \sum_{l=1}^{L} \log p(\theta_l)$. Per-layer quantization can then be enforced via an additional regularization term:

$$\mathcal{L}_r(\boldsymbol{\theta_l}) = \inf_{\boldsymbol{\theta_0} \in \boldsymbol{l}} ||\boldsymbol{\theta_l} - \boldsymbol{\theta_0}||_p = ||\boldsymbol{\theta_l} - Q(\boldsymbol{\theta_l}|\boldsymbol{l})||_p$$
 (5)

which penalizes the p-norm difference between the original parameters and their quantized counterparts. Since $Q(\cdot)$ is non-differentiable (14), the straight-through estimator (STE) (15) is commonly used to approximate gradients of the regularization term. Particularly, the gradients are computed at the quantized parameters with the loss function in Eq. 4. Subsequently, the parameter updates are performed in unconstrained space.

Additionally, activations can be quantized in quantization-aware training for low precision deployment (16; 8). A linear layer using quantized weights θ and activations x computes $\theta^T x \approx Q(\theta|l_{\theta})Q(x|l_{x})$.

3 ERROR-DRIVEN QUANTIZER PARAMETER OPTIMIZATION

The previous section introduced uniform and floating-point quantizers, each parameterized by a scale and shift to define quantization levels and thresholds. We now formulate the quantization error analytically for both schemes and propose iterative and analytic quantizers to optimize these parameters. For completeness, we include optimization details for non-uniform quantizers in Appendix E.

3.1 Uniform Quantizers: Error Model and Optimization

In the uniform quantization scheme, levels are defined as $l_k = sk + z$ for k = 0: N-2, and thresholds are $t_k = s(k-\frac{1}{2}) + z$ for k = 1: N-2, with $t_0 = -\infty$ and $t_{N-1} = \infty$. The mean-squared quantization error function is given by:

$$\mathbb{E}[e^2] = \mathbb{E}[(x - Q(x|\mathbf{l}, \mathbf{t}))^2] = \sum_{k=0}^{N-2} \int_{t_k}^{t_{k+1}} (x - l_k)^2 p(x) dx,$$
 (6)

where the errors corresponding to each quantization level are integrated over N-1 intervals given the data distribution p(x). We decompose this function as the sum of two terms, clipping E_c and stepping E_s error functions. Clipping error function is defined as the error due to the saturated regions:

$$E_c = \int_{-\infty}^{-s/2+z} (x-z)^2 p(x) dx + \int_{s(N-3/2)+z}^{\infty} (x-s(N-2)-z)^2 p(x) dx, \tag{7}$$

whereas the stepping error function E_s is defined as the power of stepping error $e_s = x - Q(x|\boldsymbol{l},\boldsymbol{t})$, which is a random variable uniformly distributed on the interval $[-\frac{s}{2},\frac{s}{2}]$ according to the stochastic uniform error model (17; 18). Thus $p(e_s) = \frac{1}{s}\mathbb{I}(-s/2 \le e_s \le s/2)$ with the mean value being zero due to symmetry. The mean squared value of e_s , which corresponds to the stepping error function E_s , is computed by

$$E_s = \mathbb{E}[e_s^2] = \int e_s^2 p(e_s) de_s = \frac{1}{s} \int_{-\frac{s}{2}}^{\frac{s}{2}} e_s^2 de_s = \frac{s^2}{12}.$$
 (8)

The total quantization error is $\mathbb{E}[e^2] = E_c + E_s$. Since reducing s shrinks E_s but increases E_c , this tradeoff defines an optimization problem over s and z.

3.1.1 ITERATIVE UNIFORM QUANTIZER

When there is no prior information about the distribution of the data, we minimize Eq. 6 via alternating optimization over s and z. Taking the derivatives of the error and setting them to zero yields:

$$s = \frac{\sum_{k=0}^{N-2} k \int_{t_k}^{t_{k+1}} (x-z) p(x) dx}{\sum_{k=0}^{N-1} k^2 \int_{t_k}^{t_{k+1}} p(x) dx},$$
(9)

$$z = \frac{\sum_{k=0}^{N-2} \int_{t_k}^{t_{k+1}} (x - sk) p(x) dx}{\sum_{k=0}^{N-1} \int_{t_k}^{t_{k+1}} p(x) dx}.$$
 (10)

Then, the thresholds are updated as $t_k = s(k-1/2) + z$ for k=1:N-2 where $t_0 = -\infty$ and $t_{N-1} = \infty$. First, the data points are quantized given the previous values of s and z, and then, s and z are updated given the quantized data. This process resembles a modified 1D k-means algorithm, constrained so that cluster centers (quantization levels) are evenly spaced.

3.1.2 NORMAL-OPTIMAL UNIFORM ANALYTIC QUANTIZER

Assuming the data is normally distributed (a reasonable prior for weights due to L_2 norm regularization (11; 19)), we derive an analytic uniform quantizer for parameter quantization. The input space is unbounded when the data distribution is normal. Hence, a clipping point is determined, which results in a clipping error. Denoting the clipping point by C, the clipping error function for zero-mean normally distributed data is computed analytically as follows:

$$E_c = 2(\sigma^2 + C^2)\mathcal{Q}\left(\frac{C}{\sigma}\right) - 2C\frac{\sigma}{\sqrt{2\pi}}\exp\left(-\frac{C^2}{2\sigma^2}\right),\tag{11}$$

where $\mathcal{Q}(z)=\frac{1}{\sqrt{2\pi}}\int_z^\infty \exp^{-\frac{u^2}{2}}du$ is the \mathcal{Q} function (see Appendix C for the derivation). Given a C value, the step size is computed as $s=\frac{2C}{N-1}$ due to uniformity. Given that the stepping error is uniformly distributed, the stepping error function is given by

$$E_s = \frac{C^2}{3(N-1)^2}. (12)$$

The goal is to maximize the signal-to-noise ratio:

$$SNR = \frac{\sigma^2}{E_c + E_s} = \left[2(1 + \frac{C^2}{\sigma^2}) \mathcal{Q}(\frac{C}{\sigma}) - \frac{C}{\sigma} \sqrt{\frac{2}{\pi}} e^{-\frac{C^2}{2\sigma^2}} + \frac{C^2}{3\sigma^2(N-1)^2} \right]^{-1}$$
(13)

Given N and σ , finding C that maximizes SNR does not have a closed form, so we perform a numerical search. Different from the iterative method, this search is done offline without requiring any data. In practice, we set $\sigma^2=1$ and find the optimal clipping C_{opt} that maximizes SNR with numerical search given the total number of levels N. Then, during the quantization process, the mean and variance of the data are computed, and the scale and shift parameters are determined accordingly as $s=2C_{opt}\sigma_{(x)}/(N-1)$ and $z=-s(N/2-1)+m_{(x)}$ where $\sigma_{(x)}$ and $m_{(x)}$ are empirical standard deviation and mean of the data points, respectively.

3.2 FLOATING POINT QUANTIZERS: ERROR MODEL AND OPTIMIZATION

Floating-point quantization introduces exponentially spaced levels that better match the distributions centered around zero (9; 10; 11). The quantization error function for floating-point schemes is analogous to the uniform case (Eq. 6), but uses the level definitions from Eq. 2, and sets the number of levels as $N=2^{M+E+1}-2c-1$.

As in uniform quantization, the total mean-squared quantization error is decomposed into a clipping error E_c and a stepping error E_s . The clipping error E_c is defined identically as in Eq. 7. The key challenge lies in computing E_s , as the floating-point grid consists of non-uniform intervals, with spacing that doubles every 2^M points. We define R such grid regions (subintervals) and denote each region by \mathcal{R}_r . The number of such regions is $R = 2 \left\lfloor \frac{k_{max}}{2^M} \right\rfloor + 2\mathbb{I}(k_{max} \mod 2^M > 0) - 3$ where $k_{max} = 2^{M+E} - c - 1$ is the index of the last representable grid point in the positive half-space.

Each region \mathcal{R}_r corresponds to an interval:

$$\mathcal{R}_{r} = \begin{cases}
[-s2^{\lfloor R/2 \rfloor - r + 3 - b} + z, -s2^{\lfloor R/2 \rfloor - r + 2 - b} + z], & r \leq \lfloor R/2 \rfloor \\
[-s2^{2 - b} - z, s2^{2 - b} + z], & r = \lfloor R/2 \rfloor + 1 \\
[s2^{r - \lfloor R/2 \rfloor - b} + z, s2^{r - \lfloor R/2 \rfloor + 1 - b} + z], & r > \lfloor R/2 \rfloor + 1
\end{cases}$$
(14)

The quantization step size within region r is denoted as $v_r = s2^{|r-1-\lfloor R/2\rfloor|+1-M-b}$.

Following the stochastic uniform error model, we model the stepping error in each region as conditionally uniform. Thus, the overall error distribution becomes a mixture of uniform distributions:

$$p(e_s) = \sum_{r=1}^{2R-1} \frac{1}{v_r} \mathbb{I}\left(-\frac{v_r}{2} \le e \le \frac{v_r}{2}\right) \cdot p(x \in \mathcal{R}_r), \tag{15}$$

where $p(x \in \mathcal{R}_r)$ denotes the probability mass of region r under p(x). (Appendix F gives an illustration of the stepping error). The expected stepping error is then:

$$E_s = \sum_{r=1}^{2R-1} \frac{v_r^2}{12} \cdot p(x \in \mathcal{R}_r).$$
 (16)

3.2.1 Iterative Floating-Point Quantizer

To optimize the parameters s and z, we employ an alternating optimization approach again. The scale update is $s = N_s(x)/D_s(x)$ where:

$$N_s(x) = \sum_{k=0}^{N/2-2} g_{-k+(N-3)/2} \int_{t_k}^{t_{k+1}} (x-z)p(x)dx + \sum_{k=N/2-1}^{N-2} g_{k-(N-3)/2} \int_{t_k}^{t_{k+1}} (x-z)p(x)dx,$$
(17)

$$D_s(x) = \sum_{k=0}^{N/2-2} g_{-k+(N-3)/2}^2 \int_{t_k}^{t_{k+1}} p(x)dx + \sum_{k=N/2-1}^{N-2} g_{k-(N-3)/2}^2 \int_{t_k}^{t_{k+1}} p(x)dx.$$
 (18)

The shift z is updated as $z = \frac{N_z(x)}{D_z(x)}$ with:

$$N_z(x) = \sum_{k=0}^{N/2-2} \int_{t_k}^{t_{k+1}} (x - sg_{-k+(N-3)/2}) p(x) dx + \sum_{k=N/2-1}^{N-2} \int_{t_k}^{t_{k+1}} (x - sg_{k-(N-3)/2}) p(x) dx,$$
(19)

$$D_z(x) = \sum_{k=0}^{N-2} \int_{t_k}^{t_{k+1}} p(x) dx.$$
 (20)

where the denominator $D_z(x)$ sums up to the total number of points I. Then, we update the thresholds as $t_k = s(l_k + l_{k+1})/2 + z$ for k = 1 : N - 2 where $t_0 = -\infty$ and $t_{N-1} = \infty$. In summary, in the first step, the data points are quantized given the previous values of s and z, and in the second step, s and z are updated given the quantized data. This procedure resembles a constrained 1D k-means clustering, but where cluster centers must lie on a scaled floating-point grid.

3.2.2 NORMAL-OPTIMAL ANALYTIC FLOAT QUANTIZER

Assuming $x \sim \mathcal{N}(0, \sigma^2)$, we can evaluate $p(x \in \mathcal{R}_r)$ analytically:

$$p(x \in \mathcal{R}_r) = F(\mathcal{R}_r^h) - F(\mathcal{R}_r^l), \tag{21}$$

where F(x) is the Gaussian CDF and \mathcal{R}_r^l , \mathcal{R}_r^h denote region bounds. From Eqs. 16 and 11, the total signal-to-noise ratio becomes:

$$SNR = \left[2\left(1 + \frac{C^2}{\sigma^2}\right)\mathcal{Q}\left(\frac{C}{\sigma}\right) - \frac{C}{\sigma}\sqrt{\frac{2}{\pi}}e^{-\frac{C^2}{2\sigma^2}} + \sum_{r=1}^{2R-1}\frac{v_r^2}{12\sigma^2}(F(\mathcal{R}_r^h) - F(\mathcal{R}_r^l))\right]^{-1}$$
(22)

As in the uniform case, we fix $\sigma^2=1$ and numerically find C_{opt} offline that maximizes SNR. We then compute scale as $s=C_{opt}\sigma_{(x)}/g_{2^{M+E}-c-1}$, and derive levels using Eq. 2 and Eq. 3.

3.3 QUANTIZED TRAINING RECIPE

In quantization-aware training (QAT), quantization operations are inserted into the forward pass to simulate inference-time behavior, while gradients are propagated using the straight-through estimator (STE) (15). However, fully optimizing quantization parameters at each training step is computationally prohibitive, particularly for iterative methods that require multiple passes to converge.

To address this, we adopt a single-step update scheme where the parameters at time step t, denoted s^t and z^t , are updated using the values from the previous step (s^{t-1}, z^{t-1}) via a single iteration. This incremental update approximates convergence while maintaining training efficiency. The procedure

Algorithm 1 Iterative Quantization Step $Q_k(x \mid s, z)$ assigns level indices; $Q(x \mid s, z)$ returns dequantized values. See Eqs. (26)–(30).

1: **Input:** $x \in \mathbb{R}^I$, previous scale s^{t-1} , shift z^{t-1}

2: Quantize: $Q_k(x \mid s^{t-1}, z^{t-1})$

3: Update scale

$$s^t = \frac{\sum_{j} (x_j - z^{t-1}) \cdot Q_k(x_j \mid s^{t-1}, z^{t-1})}{\sum_{j} Q_k(x_j \mid s^{t-1}, z^{t-1})^2}$$

4: Update shift:

$$z^{t} = \frac{1}{I} \sum_{i} x_{i} - s^{t-1} \cdot Q_{k}(x_{i} \mid s^{t-1}, z^{t-1})$$

5: Dequantize: $Q(\boldsymbol{x} \mid s^t, z^t)$

6: Output: $Q(\boldsymbol{x} \mid s^t, z^t), s^t, z^t$

Algorithm 2 Analytic Quantization Step Parameters computed via analytic method. Q_k and Q as defined in Eqs. (26)–(30).

1: **Input:** $x \in \mathbb{R}^I$, optimal clipping point C_{opt}

2: Compute mean: $\mu_x = \frac{1}{I} \sum_j x_j$

3: Compute variance: $\sigma_x^2 = \frac{1}{I} \sum_j (x_j - \mu_x)^2$

4: Set shift: $z = \mu_x$

5: Set scale: $s = 2C_{\text{opt}} \cdot \sigma_x / (N-1)$

6: Quantize: $Q_k(\boldsymbol{x} \mid s, z)$

7: Dequantize: $Q(\boldsymbol{x} \mid s, z)$

8: Output: $Q(\boldsymbol{x} \mid s, z)$

is outlined in Algorithm 1 for iterative quantizers. In contrast, analytic methods do not require convergence, as the updates are in closed form. Algorithm 2 illustrates a single quantization step employed at time step t. See Appendix G for the initialization of the iterative quantizer, and Appendix H for the relative update speeds of the quantizers.

4 EXPERIMENTS

4.1 SIGNAL-TO-NOISE RATIO (SNR) ANALYSIS

We first present an error analysis of different data formats based on the error models derived in Section 3 for zero-mean, unit-variance Gaussian distributed data. This enables direct computation of the signal-to-noise ratio (SNR) using the analytic formulas in Eq. 13 and Eq. 22. Figure 1 illustrates the resulting SNR curves as a function of the clipping point for 4-bit floating-point and 4/3/2-bit uniform quantizers. Results indicate that the 4-bit uniform quantizer theoretically can outperform the E2M1 FP4 floating-point quantizer when Gaussian distribution data is quantized. However, a floating-point quantizer offers more robustness in case clipping points are suboptimal. On the other hand, Figure 1 shows that E2M1 FP4 format achieves higher peak SNR than E3M0, but E3M0 shows resilience to suboptimal clipping due to its larger dynamic range. This analysis suggests that E2M1 FP4 and INT4/3/2 are proper choices when weight-only quantization-aware training is performed under L2 regularization, which intrinsically forces normally distributed data.

4.2 QUANTIZATION-AWARE TRAINING (QAT)

We experiment with end-to-end quantization-aware training, focusing on ResNet, MobileLLM, and Llama models. We use ResNet in our ablation study to show the effects of quantizer choice during training. Then, we use MobileLLM and Llama models to highlight that the proposed quantizers can achieve state-of-the-art accuracy. All experiments are conducted in 4-bit, representing the lowest precision currently supported by standardized floating-point formats.

4.2.1 ABLATION STUDY WITH RESNET

For the ablation study, we use the Resnet-18 model, and train it with CIFAR-10 in the QAT framework in which both activations and weights are quantized. Training recipe includes a stochastic gradient descent optimizer with a momentum of 0.9 and a weight decay coefficient of $5e^-4$.

The baseline training scheme is FP32. We compare the performance of min-max (FP4/INT4-minmax), analytic (FP4/INT4-analytic), and iterative quantization schemes (FP4/INT4-iterative). In the min-max training scheme, min-max quantizers are used for activations and weights in all quantizable

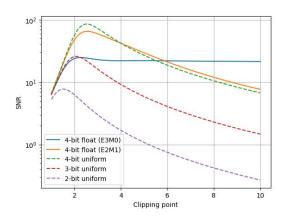


Figure 1: Signal-to-noise ratio versus clipping point for 4-bit float and 4/3/2-bit uniform quantizers, assuming zero-mean unit variance Gaussian input.

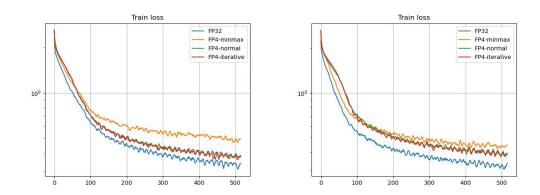


Figure 2: Loss functions with 4-bit uniform (left) and float (right) quantization across 40 epochs.

layers, including linear and 2D convolution layers in ResNet. The iterative quantization schemes use iterative quantizers for both activations and weights. Due to the arbitrary distributions of activations, the FP4/INT4-analytic scheme uses iterative quantizers for activations and analytic quantizers for the weights. The granularity is consistently tensor-wise in all experiments.

Figure 2.a shows training curves (log-scale) for 40 epochs for uniform quantizers. Likewise, Figure 2.b shows training curves for float quantizers. The loss curves consistently indicate the inferior performance of the min-max quantizer (See Appendix D for a discussion). We observe that analytic and iterative quantization schemes perform similarly. The final trained models result in the following loss values on the test set: [.0063, .0057, .0058] for minmax, analytic, and iterative with INT4, respectively. Likewise, [.0065, .0059, .0055] for the float quantizers. The full precision model gives .0050 test performance. The experiment indicates that analytic and iterative quantizers can perform similarly, especially in the INT4 setting. The performance of the analytic quantizer deteriorates slightly in FP4, which aligns with the SNR analysis in the previous section.

4.2.2 LLM PERFORMANCE

We compare StatQAT with state-of-the-art QAT frameworks, LLM-QAT (10) and ParetoQ (20), on MobileLLM (125M/600M) and Llama 3.2 (1B/3B) models. Evaluation was conducted on zero-shot common-sense reasoning tasks and WikiText-2.

Each model was fine-tuned using AdamW with a learning rate of 2e-5 with a cosine scheduler, L2 regularization of 0.01, and a batch size of 2 per GPU across 6 A100s. Epoch durations range from <1 minute (125M) to 9 minutes (3B) on WikiText-2. To match baseline settings, we apply symmetric weight-only quantization, using E2M1 for FP4 based on prior SNR analysis.

Table 1 compiles the results of the proposed StatQAT quantizers compared to baselines. The analytic quantizer has a comparable performance to the iterative quantizer in all cases. This justifies that estimating quantization parameters based on a Gaussian assumption holds in real-world deep learning models. We observe that baseline algorithms degrade significantly in tensor-wise granularity, whereas StatQAT quantizers can maintain performance better. The impact of optimal clipping is inherently greater in the per-tensor setting, where a single clipping value must handle a broader dynamic range, including outliers. This effect is well known and consistent across all algorithms, as also observed in our experiments. Our method shows large gains in the per-tensor case, demonstrating its ability to correctly and efficiently determine optimal clipping points.

In channel-wise granularity, StatQAT quantizers outperform LLM-QAT in the majority of cases and ParetoQ in all cases. We observe that the performance of LLM-QAT and ParetoQ is not competitive in large models compared to StatQAT when the granularity is tensor-wise. In the per-channel setting, where each channel has its own clipping value and thus a narrower distribution, the optimization problem is inherently easier, and performance differences are expected to be smaller. Nonetheless, our results show consistent improvements, albeit marginal, across all tested models. These gains, while modest, were reproducible across multiple datasets, indicating statistical significance.

5 RELATED WORK

Quantization reduces memory, compute, and bandwidth requirements in deep neural networks by lowering numerical precision. Early PTQ methods relied on heuristic techniques such as min–max scaling or k-means clustering (21; 3), which often failed at low bit-widths.

To overcome these limitations, quantization-aware training (QAT) was introduced, where quantization effects are simulated during training to improve robustness (4). Frameworks such as DoReFa-Net (7), PACT (8), and LSQ (16) proposed gradient approximations and parameter calibration tailored to specific quantizers, enabling accuracy retention in the 4-bit regime and below. More recent QAT approaches, including LLM-QAT (10) and ParetoQ (20), extend these ideas to large language models (LLMs), achieving state-of-the-art performance.

From a theoretical perspective, minimizing mean-squared quantization error has inspired more principled methods, such as Lloyd–Max optimization (22) and differentiable quantization objectives (23). Other works explore learned or data-driven quantizers (24; 25), but these approaches often incur significant overhead due to iterative optimization within each training step (26).

In parallel, the rise of LLMs has driven PTQ strategies specialized for transformer architectures. GPTQ (27) employs Hessian-based error-aware quantization to achieve high-fidelity INT4 quantization. AWQ (28) improves GPTQ by decoupling weight scales and applying per-channel clipping, while SmoothQuant (29) shifts scaling from activations to weights, enabling effective INT8 quantization for both weights and activations. These methods highlight the importance of error compensation, scale decoupling, and outlier-aware clipping for preserving LLM performance.

Meanwhile, FP8 and FP4 formats have emerged as promising alternatives for training and inference (30; 31). Kuzmin et al. (9) treat the clipping point as a trainable parameter, similar to LSQ, while Liu et al. (10) propose FP4 quantization in a PTQ setting, optimizing clipping points via online search. However, such iterative schemes are impractical for QAT, where clipping must be updated at every training step. To our knowledge, we present the first closed-form analytic solution for FP quantizers in QAT, offering both theoretical novelty and practical hardware relevance.

Our work introduces a unified statistical framework for uniform and floating-point quantization. The proposed analytic and iterative quantizers are lightweight and suitable for QAT. The latency of our analytic quantizer is slightly better than min–max scaling (Appendix H). This complements existing LLM quantization techniques and provides a principled foundation for large-scale QAT deployments.

6 CONCLUSION

We presented a comprehensive statistical framework for analyzing and minimizing quantization error in uniform and floating-point quantization schemes. Building on this foundation, we proposed novel iterative and analytic quantizers effective for activations and weights. Our methods support accurate,

| Model | Type | Granularity | Method | Arc-e | Arc-c | BoolQ | PIQA | SIQA | HellaSwag | OBQA | WinoGrande | Avg. | Wik |
|-------------------|-------|-------------|---------------------------------------|--------------|--------------|--------------|------|--------------|--------------|------|--------------|---------------------|-----|
| MobileLLM 125M | | Tensor | LLM-QAT | 37.5 | 20.5 | 45.8 | 56.7 | 36.8 | 29.9 | 13.6 | 50.8 | 36.5 | 13 |
| | | | ParetoQ | 37.2 | 19.8 | 45.0 | 56.8 | 36.6 | 30.1 | 14.6 | 50.0 | 36.3 | 13. |
| | | 1011501 | StatQAT-iterative | 44.0 | 20.5 | 56.3 | 61.1 | 36.6 | 31.3 | 17.4 | 52.9 | 40.0 | 10. |
| | FP4 | | StatQAT-analytic | 42.8 | 19.4 | 58.0 | 61.4 | 37.5 | 31.4 | 17.0 | 52.4 | 40.0 | 10. |
| | 114 | Channel | LLM-QAT | 44.9 | 20.0 | 56.1 | 62.6 | 38.1 | 32.0 | 17.6 | 51.8 | 40.4 | 9.8 |
| | | | ParetoQ | 44.2 | 19.2 | 60.5 | 62.0 | 37.9 | 31.9 | 18.4 | 51.7 | 40.7 | 10. |
| | | | StatQAT-iterative | 46.0 | 19.9 | 59.1 | 62.3 | 37.7 | 32.0 | 18.4 | 52.5 | 41.0 | 9.8 |
| | | | StatQAT-analytic | 44.9 | 21.1 | 60.9 | 62.0 | 36.9 | 31.6 | 17.6 | 54.1 | 41.1 | 10. |
| | INT4 | Tensor | LLM-QAT | 29.3 | 20.1 | 40.3 | 53.0 | 34.5 | 27.2 | 14.8 | 51.5 | 33.8 | 26. |
| | | | ParetoQ | 29.8 | 19.6 | 51.8 | 53.7 | 34.8 | 27.0 | 13.4 | 53.3 | 35.4 | 34. |
| | | | StatQAT-iterative | 42.3 | 20.8 | 44.6 | 61.5 | 36.4 | 31.4 | 16.0 | 51.9 | 38.1 | 10. |
| | | | StatQAT-analytic | 41.5 | 21.2 | 55.4 | 60.6 | 36.7 | 31.1 | 17.4 | 51.7 | 39.5 | 10. |
| | | Channel | LLM-QAT | 44.3 | 20.1 | 56.7 | 62.6 | 37.4 | 32.0 | 17.8 | 51.8 | 40.3 | 10. |
| | | | ParetoQ | 42.2 | 18.6 | 52.1 | 61.0 | 36.1 | 30.9 | 17.6 | 51.9 | 38.8 | 10. |
| | | | StatQAT-iterative | 45.6 | 20.9 | 53.8 | 61.8 | 37.2 | 32.1 | 17.4 | 52.2 | 40.1 | 10. |
| | | | StatQAT-analytic | 44.9 | 20.8 | 53.5 | 61.9 | 36.9 | 31.7 | 18.8 | 52.6 | 40.1 | 10. |
| | FP4 | Tensor | LLM-QAT | 44.8 | 23.5 | 47.3 | 62.2 | 36.7 | 35.1 | 15.6 | 53.4 | 39.8 | 9.1 |
| | | | ParetoQ | 45.0 | 23.6 | 46.8 | 61.4 | 36.9 | 35.0 | 15.8 | 53.3 | 39.7 | 9.8 |
| | | | StatQAT-iterative | 57.9 | 26.5 | 59.3 | 69.6 | 37.7 | 40.7 | 20.4 | 58.3 | 46.3 | 7.8 |
| | | | StatQAT-analytic | 57.6 | 27.5 | 59.4 | 69.3 | 38.8 | 40.3 | 21.2 | 56.9 | 46.4 | 7.3 |
| | | Channel | LLM-QAT | 59.1 | 26.3 | 63.1 | 69.7 | 40.3 | 42.3 | 24.2 | 58.3 | 47.9 | 7. |
| | | | ParetoQ | 57.6 | 25.6 | 61.8 | 68.7 | 39.5 | 41.8 | 22.0 | 58.3 | 46.9 | 7. |
| MobileLLM 600M | | | StatQAT-iterative | 59.8 | 26.7 | 62.5 | 69.4 | 40.6 | 42.1 | 25.2 | 58.6 | 48.1 | 7.3 |
| | | | StatQAT-analytic | 60.1 | 27.9 | 56.1 | 68.4 | 38.6 | 41.2 | 22.8 | 58.1 | 46.7 | 7. |
| | | Tensor | LLM-QAT | 28.7 | 20.1 | 40.3 | 54.7 | 34.0 | 27.5 | 13.8 | 50.7 | 33.7 | 25 |
| | | | ParetoQ | 27.6 | 20.2 | 38.9 | 51.4 | 32.8 | 26.5 | 14.4 | 48.5 | 32.5 | 43. |
| | | | StatQAT-iterative | 57.9 | 26.5 | 61.3 | 67.7 | 38.0 | 39.7 | 20.8 | 56.8 | 46.1 | 8.0 |
| | DATE: | | StatQAT-analytic | 56.1 | 26.3 | 58.3 | 67.6 | 38.2 | 39.2 | 21.6 | 56.0 | 45.4 | 8. |
| | INT4 | Channel | LLM-QAT | 58.5 | 25.2 | 62.0 | 69.8 | 39.9 | 41.7 | 21.8 | 58.7 | 47.2 | 7. |
| | | | ParetoQ | 56.8 | 26.5 | 62.6 | 69.0 | 41.0 | 40.9 | 21.4 | 55.9 | 46.8 | 7. |
| | | | StatQAT-iterative | 59.1 | 25.6 | 62.7 | 68.9 | 39.6 | 41.9 | 22.6 | 59.6 | 47.5 | 7.4 |
| | | | StatQAT-analytic | 57.8 | 26.5 | 57.2 | 67.9 | 37.4 | 40.3 | 22.2 | 58.4 | 46.0 | 7.3 |
| | | Tensor | LLM-QAT | 27.2 | 20.3 | 37.8 | 51.7 | 34.1 | 26.0 | 13.8 | 49.8 | 32.6 | 194 |
| | | | ParetoQ | 26.1 | 19.9 | 37.8 | 50.8 | 34.4 | 26.0 | 13.4 | 49.8 | 32.3 | 203 |
| | | | StatQAT-iterative | 60.7 | 29.5 | 45.4 | 70.9 | 40.9 | 42.5 | 21.6 | 55.6 | 45.9 | 9.0 |
| | | | StatQAT-analytic | 59.8 | 28.7 | 56.3 | 69.7 | 41.0 | 42.6 | 21.4 | 55.6 | 46.9 | 9. |
| | FP4 | Channel | LLM-QAT | 59.8 | 30.5 | 49.5 | 70.9 | 39.0 | 43.7 | 22.6 | 56.2 | 46.5 | 9.5 |
| | | | ParetoQ | 60.5 | 28.9 | 55.2 | 70.3 | 38.6 | 42.9 | 19.8 | 55.6 | 46.5 | 10. |
| | | | StatQAT-iterative | 61.9 | 30.3 | 51.0 | 72.2 | 40.2 | 44.6 | 23.2 | 57.2 | 47.6 | 8.9 |
| Llama 3.2 | | | StatQAT-analytic | 60.6 | 30.4 | 54.5 | 72.0 | 40.1 | 43.0 | 24.4 | 55.7 | 47.6 | 9. |
| 1B | | | LLM-QAT | 26.3 | 19.5 | 37.8 | 51.2 | 33.6 | 26.1 | 11.2 | 51.1 | 32.1 | 389 |
| 110 | | | ParetoQ | 26.2 | 20.5 | 37.8 | 50.3 | 34.0 | 26.1 | 12.4 | 49.3 | 32.1 | 493 |
| | | | StatQAT-iterative | 56.8 | 27.6 | 52.4 | 70.6 | 41.5 | 42.2 | 21.2 | 53.2 | 45.7 | 9.0 |
| | INT4 | | StatQAT-nerative StatQAT-analytic | 56.1 | 27.2 | 56.5 | 69.5 | 40.0 | 40.7 | 19.8 | 55.2 | 45.6 | 9. |
| | | Channel | LLM-QAT | 59.4 | 29.5 | 58.0 | 70.3 | 39.2 | 41.8 | 21.0 | 56.4 | 46.9 | 10. |
| | | | ParetoQ | 52.4 | 27.0 | 52.6 | 66.1 | 37.8 | 38.5 | 20.8 | 54.5 | 43.7 | 12. |
| | | | StatQAT-iterative | 61.0 | 29.9 | 56.1 | 71.3 | 40.6 | 43.6 | 24.8 | 55.6 | 47.9 | 9. |
| | | | StatQAT-nerative StatQAT-analytic | 59.4 | 29.7 | 59.2 | 70.1 | 41.0 | 42.0 | 23.4 | 56.5 | 47.7 | 9. |
| Llama 3.2 3B | | Tensor | LLM-QAT | 30.6 | 17.7 | 37.9 | 55.3 | 33.8 | 27.3 | 13.4 | 49.9 | 33.2 | 25. |
| | | | ParetoQ | 31.8 | | 38.0 | 54.5 | 33.6 | 27.4 | 13.4 | 48.1 | 33.1 | 24. |
| | | | | 58.0 | 18.6 30.9 | 38.0 46.4 | 69.0 | 39.2 | 42.5 | 18.2 | 59.7 | 45.5 | 8. |
| | | | StatQAT-iterative StatQAT-analytic | 58.0 64.3 | 33.4 | 46.4 49.6 | 71.2 | 39.2 41.2 | 42.5 43.6 | 21.6 | 59.7 59.0 | 45.5 48.0 | 8. |
| | FP4 | Channel | | | 41.2 | | | | | 29.2 | | | 7. |
| | | | LLM-QAT | 73.6 | | 65.5 | 75.8 | 42.2 | 52.1 | | 65.8 | 55.7 | |
| | | | ParetoQ | 72.3 | 40.7 | 61.5 | 75.0 | 42.4 | 51.2 | 28.0 | 65.2 | 54.5 | 7. |
| | | | StatQAT-iterative | 73.6 | 39.4 | 62.5 | 75.8 | 43.9 | 52.2 | 30.0 | 66.7 | 55.5 | 7. |
| | | | StatQAT-analytic | 69.7 | 38.5 | 61.7 | 74.6 | 42.6 | 49.1 | 27.0 | 63.1 | 53.3 | 7. |
| | | Tensor | LLM-QAT | 26.4 | 21.8 | 37.8 | 50.7 | 34.8 | 26.0 | 12.6 | 49.6 | 32.5 | 473 |
| | | | ParetoQ | 26.8 | 20.8 | 37.8 | 50.9 | 34.8 | 26.1 | 12.4 | 50.2 | 32.5 | 533 |
| | | | StatQAT-iterative | 53.9 | 29.4 | 60.9 | 66.2 | 38.8 | 40.6 | 19.0 | 55.2 | 45.5 | 8. |
| | INT4 | | StatQAT-analytic | 55.4 | 28.6 | 56.3 | 67.9 | 38.5 | 40.8 | 19.6 | 57.1 | 45.5 | 8. |
| | IN14 | Channel | LLM-QAT | 72.1 | 39.3 | 63.9 | 74.6 | 42.5 | 50.5 | 29.4 | 66.1 | 54.8 | 7. |
| | | | ParetoQ | 68.1 | 35.7 | 70.8 | 72.7 | 41.4 | 48.5 | 25.4 | 65.3 | 53.5 | 8. |
| | | | StatQAT-iterative | 68.9 | 36.9 | 62.9 | 75.3 | 43.6 | 51.4 | 30.6 | 66.6 | 54.5 | 7. |
| | | | StatQAT-analytic | 67.6 | 35.8 | 54.5 | 72.3 | 40.6 | 46.4 | 24.2 | 61.2 | 50.3 | 8. |

Table 1: Accuracy comparison of proposed StatQAT quantizers with baselines on benchmark datasets. Our test setup included the implementation of the baselines for a fair comparison. ± 0.1 variance is accounted for when bolding the best performance for statistical significance.

low-overhead quantization-aware training by enabling efficient estimation of optimal quantization parameters. Our error analysis and experimental results on ResNet and LLMs demonstrate that our quantizers achieve state-of-the-art performance. Notably, analytic quantizers achieve performance comparable to iterative ones at a fraction of the cost, making them highly practical for training. Our findings lay the groundwork for extending statistical quantizer design to more complex distributions and adaptive training scenarios. While our study focuses on 4-bit quantization due to current hardware support for formats like FP4, extending our framework to ultra-low precision remains an important direction for future research.

REFERENCES

- [1] Amir Gholami, Sehoon Kim, Zhen Dong, Zhewei Yao, Michael W Mahoney, and Kurt Keutzer. A survey of quantization methods for efficient neural network inference. In *Low-power computer vision*, pages 291–326. Chapman and Hall/CRC, 2022.
- [2] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *International Conference on Learning Representations (ICLR)*, 2016.
- [3] Scott Migacz. 8-bit inference with tensorrt. https://developer.nvidia.com/blog/int8-inference-autonomous-vehicles-tensorrt/, 2017. GPU Technology Conference.
- [4] 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. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [5] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in neural information processing systems*, 35:30318–30332, 2022.
- [6] Yoshitaka Gongyo et al. Learning non-uniform step sizes for neural network quantization. In *Asian Conference on Computer Vision (ACCV)*, 2024.
- [7] 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, 2016.
- [8] Yoonho Choi, Mostafa El-Khamy, and Jungwon Lee. Pact: Parameterized clipping activation for quantized neural networks, 2018.
- [9] Andrey Kuzmin, Mart Van Baalen, Yuwei Ren, Markus Nagel, Jorn Peters, and Tijmen Blankevoort. Fp8 quantization: The power of the exponent. *Advances in Neural Information Processing Systems*, 35:14651–14662, 2022.
- [10] Shih-yang Liu, Zechun Liu, Xijie Huang, Pingcheng Dong, and Kwang-Ting Cheng. Llm-fp4: 4-bit floating-point quantized transformers. *arXiv preprint arXiv:2310.16836*, 2023.
- [11] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in neural information processing systems*, 36:10088–10115, 2023.
- [12] Stuart P Lloyd. Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.
- [13] IEEE Computer Society. Ieee standard for floating-point arithmetic. *IEEE Std 754-2019* (*Revision of IEEE 754-2008*), 2019.
- [14] Christos Louizos, Charles Blundell, and Max Welling. Relaxed quantization for discretized neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [15] Yoshua Bengio et al. Estimating or propagating gradients through stochastic neurons for conditional computation. In *arXiv preprint arXiv:1308.3432*, 2013.
- [16] Steven K Esser, Jeffrey L McKinstry, Arun Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. In *International Conference on Learning Representations (ICLR)*, 2020.
- [17] Michael Bernhard, David Rörich, Thomas Handte, and Joachim Speidel. Analytical and numerical studies of quantization effects in coherent optical ofdm transmission with 100 gbit/s and beyond. *ITG-Fachtagung Photonische Netze*, pages 34–40, 2012.

[18] Henning Ehm, Sebastian Winter, and Robert Weigel. Analytic quantization modeling of ofdm signals using normal gaussian distribution. In 2006 Asia-Pacific Microwave Conference, pages 847–850. IEEE, 2006.

- [19] Charles Blundell et al. Weight uncertainty in neural networks. In *International Conference on Machine Learning (ICML)*, 2015.
- [20] Zechun Liu, Changsheng Zhao, Hanxian Huang, Sijia Chen, Jing Zhang, Jiawei Zhao, Scott Roy, Lisa Jin, Yunyang Xiong, Yangyang Shi, et al. Paretoq: Scaling laws in extremely low-bit llm quantization. *arXiv preprint arXiv:2502.02631*, 2025.
- [21] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)*, 2016.
- [22] Robert M Gray and David L Neuhoff. Quantization. IEEE Transactions on Information Theory, 44(6):2325–2383, 1998.
- [23] Stefan Uhlich et al. Differentiable quantization of deep neural networks. In *International Conference on Learning Representations (ICLR)*, 2020.
- [24] Christos Louizos, Charles Blundell, and Max Welling. Relaxed quantization for discretized neural networks. In *International Conference on Learning Representations (ICLR)*, 2019.
- [25] Sangil Jung, Byeongho Choi, Junha Kim, and Nojun Kwak. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *CVPR*, pages 4350–4359, 2019.
- [26] Charbel Sakr, Steve Dai, Rangha Venkatesan, Brian Zimmer, William Dally, and Brucek Khailany. Optimal clipping and magnitude-aware differentiation for improved quantization-aware training. In *International conference on machine learning*, pages 19123–19138. PMLR, 2022.
- [27] Elias Frantar, Pierre Stock, and Dan Alistarh. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323*, 2022.
- [28] Ji Lin, Zhenhua Tang, Yujun Liu, and Song Han. Awq: Activation-aware weight quantization for llm compression and acceleration. *arXiv preprint arXiv:2306.00978*, 2023.
- [29] Yuxuan Xiao, Yuhui Ren, Yifan Sun, Haotong Wang, Zheng Wang, et al. Smoothquant: Accurate and efficient post-training quantization for large language models. *arXiv* preprint arXiv:2306.03078, 2023.
- [30] Paulius Micikevicius, Dusan Stosic, Neil Burgess, Marius Cornea, Pradeep Dubey, Richard Grisenthwaite, Sangwon Ha, Alexander Heinecke, Patrick Judd, John Kamalu, et al. Fp8 formats for deep learning. *arXiv preprint arXiv:2209.05433*, 2022.
- [31] Xiao Sun, Naigang Wang, Chia-Yu Chen, Jiamin Ni, Ankur Agrawal, Xiaodong Cui, Swagath Venkataramani, Kaoutar El Maghraoui, Vijayalakshmi Viji Srinivasan, and Kailash Gopalakrishnan. Ultra-low precision 4-bit training of deep neural networks. *Advances in Neural Information Processing Systems*, 33:1796–1807, 2020.

A PRACTICAL IMPLEMENTATION OF UNIFORM QUANTIZATION

In real-world applications, instead of storing high-precision l_k values corresponding to each data point, only the indexes are stored in integer format for reduced data size. Updating Eq.1 to return indexes instead of levels as

$$Q_k(x|\mathbf{l}, \mathbf{t}) = k, \quad \text{if } t_k < x \le t_{k+1}, \tag{23}$$

we obtain the unsigned integer representation of point x. For signed integer representation, the function returns k - N/2 + 1.

We can alternatively compute the index for uniform unsigned integer quantization as follows:

$$Q_k(x|s,z) = clip(\lfloor \frac{x-z}{s} \rceil, n, m)$$
 (24)

where [] is the integer rounding, and clip(x, n, m) is the clipping operator where n, m defines the boundaries such that n = 0, m = N - 2. In this case, the level is computed as:

$$Q(x|s,z) = s\hat{k} + z \tag{25}$$

where $\hat{k} = Q_k(x|s,z)$ is the quantization index of x.

For signed integer quantization, the index is computed as follows

$$Q_k(x|s,z) = clip(\lfloor \frac{x-z}{s} \rfloor, n, m) - \frac{N}{2} + 1$$
 (26)

and the levels as $Q(x|s,z) = s(\hat{k} + N/2 - 1) + z$.

B PRACTICAL IMPLEMENTATION OF FLOAT QUANTIZATION

Given a higher precision floating point number x, we can quantize it to a lower precision floating point number by first computing the step size v given the scaled and shifted input value:

$$v = \begin{cases} 2^{\lfloor \log_2 |\frac{x-z}{s}| \rfloor - M}, & \text{if } \lfloor \log_2 |\frac{x-z}{s}| + b \rfloor \geq 1 \\ 2^{1-M-b}, & \text{otherwise} \end{cases}$$

Then, we quantize x with this step size

$$Q_k(x|s,z) = v * \operatorname{clip}(\lfloor \frac{x-z}{sv} \rceil, n, m)$$
 (27)

When we scale-shift back, we obtain dequantized data:

$$Q(x|s,z) = sQ_k(x|s,z) + z$$
(28)

where $m=-n=2^{2^E-b-2}(2-2^{-M})$. Note the input dependency of the step size due to the non-uniform normal grid. $Q_k(x|s,z)$ is the grid point in lower precision, which we store for reduced data size or low precision computation.

C DERIVATION OF CLIPPING ERROR FUNCTION

The clipping error is defined as

$$N_c = 2 \int_C^\infty (x - C)^2 p(x) dx \tag{29}$$

where

$$p(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{x^2}{2\sigma^2}} \tag{30}$$

is the centered normal data distribution with σ^2 variance. We can decompose the error function as follows:

$$N_{c} = 2 \int_{C}^{\infty} x^{2} p(x) dx + 2C^{2} \int_{C}^{\infty} p(x) dx - 4C \int_{C}^{\infty} x p(x) dx$$
 (31)

The integral in the first term is computed by integration by parts:

$$\int_{C}^{\infty} x^{2} p(x) dx = \frac{1}{\sigma \sqrt{2\pi}} \int_{C} x^{2} e^{-\frac{x^{2}}{2\sigma^{2}}} dx$$

$$= \frac{\sigma^{2}}{2} erf\left(\sqrt{\frac{x}{2\sigma^{2}}}\right) - \frac{\sigma}{\sqrt{2\pi}} x e^{-\frac{x^{2}}{2\sigma^{2}}} \Big|_{x=C}^{x=\infty}$$

$$= \sigma^{2} \left(\frac{1}{2} - \frac{1}{2} erf\left(\frac{C}{\sigma\sqrt{2}}\right)\right) + \frac{C\sigma}{\sqrt{2\pi}} e^{-\frac{C^{2}}{2\sigma^{2}}}$$

$$= \sigma^{2} Q\left(\frac{C}{\sigma}\right) + \frac{C\sigma}{\sqrt{2\pi}} e^{-\frac{C^{2}}{2\sigma^{2}}}$$
(32)

The integral in the second term is computed by the definition of the erf function:

$$\int_{C}^{\infty} p(x)dx = \frac{1}{\sigma\sqrt{2\pi}} \int_{C}^{\infty} e^{-\frac{x^{2}}{2\sigma^{2}}} dx$$

$$= \frac{1}{2} erf(\sqrt{\frac{1}{2\sigma^{2}}}x)\Big|_{x=C}^{x=\infty}$$

$$= \frac{1}{2} (1 - erf(\sqrt{\frac{1}{2\sigma^{2}}}x))$$

$$= Q(\frac{C}{\sigma})$$
(33)

The integral in the third term is computed by integration by parts:

$$\int_{C}^{\infty} xp(x)dx = \frac{1}{\sigma\sqrt{2\pi}} \int_{C}^{\infty} xe^{-\frac{x^{2}}{2\sigma^{2}}} dx$$

$$= -\frac{C\sigma}{\sqrt{2\pi}} e^{-\frac{x^{2}}{2\sigma^{2}}} \Big|_{x=C}^{x=\infty}$$

$$= \frac{C\sigma}{\sqrt{2\pi}} e^{-\frac{C^{2}}{2\sigma^{2}}}$$
(34)

Plugging Eq. 32, Eq. 33, and Eq. 34 into Eq. 31 results in

$$N_c = 2(\sigma^2 + C^2)Q(\frac{C}{\sigma}) - 2C\frac{\sigma}{\sqrt{2\pi}}e^{-\frac{C^2}{2\sigma^2}}$$
 (35)

D RELATION TO MIN-MAX QUANTIZER AND CASTING

D.1 UNIFORM QUANTIZATION

Assume the data is uniformly distributed, $p(x) = Unif(x|-\alpha,\alpha)$. The clipping error is then computed analytically as follows:

$$E_c = 2 \int_C^\infty (x - C)^2 p(x) dx = \frac{1}{\alpha} \int_C^\infty (x - C)^2 dx$$
$$= -\frac{C^3}{3\alpha} + C^2 - C\alpha + \frac{\alpha^2}{3}$$
(36)

while E_s is the same as in Eq. 8. It can be observed that E_c decreases cubically with C until α , while E_s increases quadratically. Therefore, in the case of uniformly distributed data, the minimum error is achieved when C is set to α , resulting in zero clipping error. This indicates that the min-max quantizer is the optimal uniform quantizer in terms of mean squared quantizer error only when the data is uniformly distributed, which is not the case in deep learning model parameters or activations. Therefore, while this quantizer is commonly used in quantization tools, it has also been the primary source of quantization error due to high stepping error, especially in low-bit settings.

D.2 FLOAT QUANTIZATION

Directly casting the input data corresponds to using quantization levels without scaling and shifting. This method introduces large quantization errors if the data values are much smaller than the largest grid point (stepping error) or if the data values are larger than the largest grid (clipping error). These errors become more significant when we use lower bit quantization, such as in FP4 data types.

To avoid clipping errors, we can shift and scale the levels so that they match the data range using min-max quantization with $s=\frac{\alpha-\beta}{2l_{max}}$ and $z=\beta+sl_{max}$ where l_{max} is the largest grid point floating data type can represent, $\alpha=\max(\boldsymbol{X})$ and $\beta=\min(\boldsymbol{X})$. Then we can transform the data as (x-z)/s and do the typical float casting. This method zeros out the clipping error but introduces a large stepping error when there are outliers in the dataset. Unlike uniform quantization, the min-max float quantizer is not optimal for uniform distributed data. Indeed, there is no such tractable data distribution that this quantizer is optimal for due to base2 grid spacing.

E OPTIMIZATION OF NON-UNIFORM QUANTIZERS

Mean-squared quantization error is formulated as follows:

$$\mathbb{E}[e^2] = \mathbb{E}[(x - Q(x|\boldsymbol{l}, \boldsymbol{t}))^2] = \sum_{k=0}^{N-2} \int_{t_k}^{t_{t+1}} (x - l_k)^2 p(x) dx$$
(37)

where the errors corresponding to each quantization level are integrated over N-1 intervals given the data distribution p(x). The objective is to determine the quantization levels \boldsymbol{l} and thresholds \boldsymbol{t} to minimize the error function.

E.1 ITERATIVE METHOD

In non-uniform quantization, there are no constraints on the parameters. This allows for iterative optimization of the objective. We present an alternating optimization algorithm for this purpose. By taking the derivative of the quantization error in Eq. 37 with respect to l_k and setting it to zero, we can derive the following update for the levels:

$$l_k = \mathbb{E}[x|t_k < x \le t_{k+1}] \tag{38}$$

for k = 0: N - 2. This expectation is computed using the empirical mean of the data points in the interval. Taking the derivative with respect to t_k yields the following update for the thresholds:

$$t_k = \frac{1}{2}(l_k + l_{k-1}) \tag{39}$$

for k=1:N-2. Then, $t_0=-\infty$, and $t_N=\infty$ are set to cover the unbounded data space. In summary, in the first step of the algorithm, input data is compared to the thresholds determined at the previous iteration and quantized accordingly, and in the second step, the levels and thresholds are updated given the recently quantized data. This algorithm converges to optimal mean-squared quantization error regardless of the data distribution when initialized properly. One can recognize that the iterative solution is a 1-dimensional k-means clustering solution, which has been used in (2) for deep learning model quantization.

E.2 NORMAL-OPTIMAL ANALYTIC QUANTIZER

The deep learning models have approximately normally distributed weights (11). Based on this assumption, we propose an analytic quantizer, which is fast and as accurate as the iterative method as long as the distribution assumption holds. The update equation 38 in the iterative method relies on computing the expectation of the data given an interval. Given the data is zero-mean normally distributed with unit variance, this expectation can be computed analytically without any samples via the first moment of the two-sided truncated normal distribution:

$$\mathbb{E}[x|t_k < x \le t_{k+1}] = -\frac{p(t_k + 1) - p(t_k)}{\mathcal{F}(t_{k+1}) - \mathcal{F}(t_k)}$$
(40)

Iterating through Eq. 40 and Eq. 39 converges to the optimal non-uniform quantizer for data with a centered unit variance normal distribution, in terms of mean squared quantization error. Unlike the iterative method, this optimization is performed offline without any data, and the computed optimal levels \boldsymbol{l}_{opt} and thresholds \boldsymbol{t}_{opt} are used afterward for all layers during the quantization process as follows: the mean and variance of each quantizable weight tensor are computed, and the quantizer levels are updated accordingly as $\boldsymbol{l} = \sigma_{(x)} \boldsymbol{l}_{opt} + m_{(x)}$, where $\sigma_{(x)}$ and $m_{(x)}$ are the empirical standard deviation and mean, respectively. This makes the online quantization process much faster, as only basic statistics are computed at that stage.

E.3 RELATION TO QUANTILE QUANTIZATION

Another non-iterative method for non-uniform quantization is quantile quantization, which computes the quantiles of N equally spaced probabilities and uses them as quantization levels (11). To find the quantile function, we first compute empirical CDF, then use its inverse to determine threshold levels as follows:

$$t_k = \mathcal{F}_x^{-1}(k/(N-1)) \tag{41}$$

for k=1:N-2, $t_0=-\infty$ and $t_{N-1}=\infty$. Then we estimate levels as $l_k=\mathbb{E}[x|t_k< x< t_{k+1}]$ for k=0:N-2. In practice, we sort the input data points and assign the data points with indexes k(I-1)/(N-1) as thresholds, where I is the total number of data points. One can notice that if we use only levels to quantize by finding the closest levels using squared distance, we don't exactly equalize the histogram. Ideally, one should use thresholds to determine the index and then assign the level accordingly by Eq. 1 for improved equalization. Empirically, using levels with squared distance works better in terms of mean squared quantization error. Equalizing the bins does not necessarily result in lower quantization error. Indeed, this is true only for uniform distributed data, which is not the case for deep learning model weights.

E.4 COMPARISON OF QUANTIZATION LEVELS

Non-uniform quantizers have no constraints on choosing quantization levels. Here we show how non-uniform quantizers compare to each other given a simulated input tensor of shape [256x256] from a zero-mean unit variance normal distribution. As seen in Figure 3, iterative and analytic quantizers converge to the same points for the normally distributed data. The quantizer errors are MSE = [0.21, 0.21, 0.25], and MAE = [0.41, 0.41, 0.41] for iterative, analytic, and quantile quantization, respectively. Quantile quantizer has a higher MSE than others, although all quantizers give a comparable MAE.

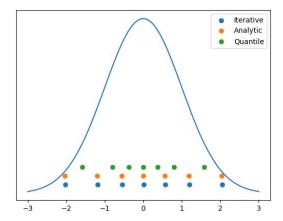


Figure 3: Inferred quantization levels of 3-bit non-uniform quantizers for a normally distributed input data.

F STEPPING ERROR DISTRIBUTION

We define total quantization error as the sum of clipping and stepping errors. Clipping error is obtained with quadratic error integrated over the data distribution. However, the stepping error depends on the quantization scheme. Figure 4 illustrates the differences in the stepping error distribution of uniform and float quantization schemes given a set of normally distributed input data with zero mean and unit variance. Regardless of the data distribution, quantizer type, and total bits, we obtain uniform distributed stepping error in a uniform quantization scheme. The error is bounded in [-s/2, s/2] with probability 1/s as shown in Figure 4.a-c. From this observation, we compute the stepping error power as in Eq. 8. On the other hand, the error distribution is a mixture of uniforms in floating point quantization schemes due to the exponential spacing of every 2^M point. Thus, the error space widens at each consecutive region as shown in Figure 4.b-d. Hence, the stepping error power is computed by Eq. 16 since now the data distribution affects the mixture coefficient, which changes the stepping error power.

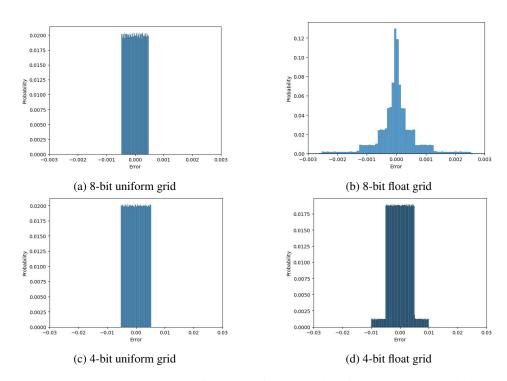


Figure 4: Quantization error distribution for float and uniform quantization grids

G INITIALIZATION OF STATQAT-ITERATIVE

 Quantizer initialization is critical to avoid early-stage gradient instabilities. For weights, if the initialization follows a normal distribution, as is typical under common initialization schemes and regularization, the analytic quantizer derived under Gaussian assumptions is a natural choice. For weights initialized from uniform distributions, the min-max method provides a better match to the initial data distribution.

Activation quantizers are initialized during the first forward pass. Since activation distributions can be highly non-Gaussian and vary across tasks and layers, the ideal approach is to fit an iterative quantizer on the first minibatch. However, this can be computationally expensive. Empirically, we find that analytic quantizers still provide competitive initialization performance and allow stable training from the outset.

H ITERATION SPEEDS OF QUANTIZERS

We compare the iteration rate of the min-max, analytic, and iterative quantizers in Table 2 for FP4 and INT4 datatypes and channel and tensor-wise granularity on A100 GPU. We use 2048×2048 tensor as input. In all configurations, the iterative quantizer is the slowest by an order of magnitude, while the analytic quantizer is slightly faster than min-max. These results confirm that the analytic quantizer is light enough to be used as a drop-in replacement for the min-max quantizer for QAT.

| Туре | Quantizer | Granularity | 1k iter/sec |
|------|-----------|-------------|-------------|
| INT4 | Min-max | Channel | 21.6 |
| INT4 | Iterative | Channel | 4.4 |
| INT4 | Analytic | Channel | 25.1 |
| FP4 | Min-max | Channel | 20.1 |
| FP4 | Iterative | Channel | 1.9 |
| FP4 | Analytic | Channel | 24.6 |
| INT4 | Min-max | Tensor | 21.2 |
| INT4 | Iterative | Tensor | 4.6 |
| INT4 | Analytic | Tensor | 24.2 |
| FP4 | Min-max | Tensor | 18.8 |
| FP4 | Iterative | Tensor | 1.9 |
| FP4 | Analytic | Tensor | 23.0 |

Table 2: Iteration rate of quantizers.