

---

# Low-Precision Training in Logarithmic Number System using Multiplicative Weight Update

---

**Jiawei Zhao\***  
Caltech  
jiawei@caltech.edu

**Steve Dai**  
NVIDIA  
sdai@nvidia.com

**Rangharajan Venkatesan**  
NVIDIA  
rangharajanv@nvidia.com

**Ming-Yu Liu**  
NVIDIA  
mingyul@nvidia.com

**Brucek Khailany**  
NVIDIA  
bkhailany@nvidia.com

**Bill Dally**  
NVIDIA  
bdally@nvidia.com

**Anima Anandkumar**  
Caltech  
NVIDIA  
anima@caltech.edu

## Abstract

Training large-scale deep neural networks (DNNs) currently requires a significant amount of energy, leading to serious environmental impacts. One promising approach to reduce the energy costs is representing DNNs with low-precision numbers. While it is common to train DNNs with forward and backward propagation in low-precision, training directly over low-precision weights, without keeping a copy of weights in high-precision, still remains to be an unsolved problem. This is due to complex interactions between learning algorithms and low-precision number systems. To address this, we jointly design a low-precision training framework involving a logarithmic number system (LNS) and a multiplicative weight update training method, termed LNS-Madam. LNS has a high dynamic range even in a low-bitwidth setting, leading to high energy efficiency and making it relevant for on-board training in energy-constrained edge devices. We design LNS to have the flexibility of choosing different bases for weights and gradients, as they usually require different quantization gaps and dynamic ranges during training. By drawing the connection between LNS and multiplicative update, LNS-Madam ensures low quantization error during weight update, leading to a stable convergence even if the bitwidth is limited. Compared to using a fixed-point or floating-point number system and training with popular learning algorithms such as SGD and Adam, our joint design with LNS and LNS-Madam optimizer achieves better accuracy while requiring smaller bitwidth. Notably, with only 5-bit for gradients, the proposed training framework achieves accuracy comparable to full-precision state-of-the-art models such as ResNet-50 and BERT. To verify the efficiency of our framework, we also conduct energy estimations by analyzing the math datapath units during training. We calculate that our design achieves over 60x energy reduction compared to FP32 on BERT models. For full training of ResNet-50 on ImageNet, our design reduces the carbon emissions by 98% around.

---

\*Work done during an internship at NVIDIA Research.

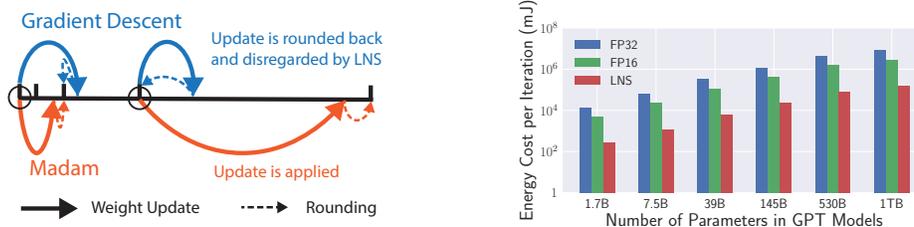


Figure 1: **Left:** an illustration for updating weights using Gradient Descent (GD) and Madam under logarithmic representation. Each coordinate represents a number stored in LNS. Assume the weights at two circles receive the same gradient. The updates generated from GD are disregarded as the weights move larger, whereas the updates generated by Madam are adjusted with the weights. **Right:** energy costs over a range of GPT models from 1 billion to 1 trillion parameters. The models are scaled by a throughput efficient method proposed by Narayanan et al. [11]. The energy calculation method is introduced in Section 5.3.

## 1 Introduction

Deep neural networks (DNNs) have been widely used in many applications, including image classification and language processing. However, training and deploying DNNs usually require significant computational costs and energy usage, emitting a tremendous amount of carbon dioxide ( $CO_2$ ) and severely impacting our environment [1]. To save energy, previous studies have tried to represent DNNs with low-precision numbers. Traditionally, numbers in neural networks are represented using floating-point (32-bit) numbers, which leads to large arithmetic and memory footprint and hence, large energy consumption. However, recent studies suggest that the high-precision number format is redundant, and the models can be quantized in low-precision with little loss in accuracy [2, 3]. Low-precision numbers only require low-bitwidth computational units, leading to better computational efficiency and less memory requirement.

**LNS-based Low-Precision Training.** With the advent of large-scale networks, training and fine-tuning increasingly become energy expensive, which brings the need for training in low-precision. Low-precision training is built upon the techniques for low-precision inference. One of the most popular inference quantization methods is quantization-aware training (QAT), which directly trains the quantized model using straight-through estimators (STE) [4]. However, QAT does not accelerate the training phase since the gradients are still in full-precision. To speed up the training phase, several studies focus on low-precision training where gradients are also quantized during QAT [3, 5]. In this case, both forward and backward propagation can be implemented efficiently on low-bitwidth computational units like NVIDIA Tensor Cores [6].

While the above low-precision training methods generally reduce computational costs, the energy efficiency can be further improved by choosing a logarithmic number system (LNS) as the backend of low-precision training. LNS achieves higher computational efficiency by transforming expensive multiplication operations in linear space to inexpensive additions in logarithmic space. Previous works also demonstrate that logarithmic representations are suitable for representing the dynamics of training for DNNs [7]. This is because a logarithmic representation attains a wide dynamic range and has a non-uniform distribution. This is also in line with the weight distribution of neural networks and anatomical findings of biological synapses [8, 9].

Although previous works demonstrate that it is feasible to train networks in low-precision using LNS, larger datasets and state-of-the-art models using LNS have not yet shown good results [7, 10]. One of the reasons is the inflexibility of log-base settings in LNS. Standard LNS used in previous works fixes log-base to be exactly two. However, a smaller and more flexible log-base setting is needed as the weights and gradients usually require different quantization gaps during training. Although this flexibility better approximates the dynamics of training and thus leads to better accuracy, it also introduces additional hardware overhead because the conversion operation between logarithmic and linear formats requires a larger size look-up table (LUT) implemented as Read Only Memory (ROM). This motivates us to design a LNS that leverages log-base flexibility while retaining log-to-linear conversion efficiency.

**Training with Low-Precision Weight Update.** While low-precision training significantly reduces the computational costs, its overall efficiency is hampered by the high-precision requirement for weight update process. Conventional low-precision training methods require high-precision weight gradients and weights to maintain optimization stability. In fact, most recent studies even use a full-precision (FP32) copy of weights [7, 10]. However, training with hybrid numerical formats induces

additional costs, and expensive FP32 arithmetic is not available especially in cheap energy-constrained edge devices.

This high-precision requirement for weight update is due to complex interactions between learning algorithms and number systems, which is usually ignored in previous studies. However, this interaction plays an important role especially in low-precision training based on LNS. For example, as illustrated in Figure 1, updates generated by stochastic gradient descent (SGD) are disregarded more frequently by LNS when the weights become larger. This suggests that conventional learning algorithms may not be suitable for LNS. On the other hand, due to lack of this understanding, previous LNS studies simply apply SGD by default and thus require high-precision weights to avoid numerical instabilities [7, 12].

To directly update the weights in low-precision, we employ a learning algorithm tailored to LNS. Recently, Bernstein et al. [13] proposes a learning algorithm Madam that can train the network over a discrete logarithmic weight space. Specifically, Madam updates the weights in a multiplicative way, which is similar to updating the weights additively in logarithmic space. We also illustrate the update mechanism of Madam in Figure 1. Although previous works on Madam consider its connections to logarithmic representation, they still employ a full-precision training without considering low-precision LNS. This brings us to design a learning algorithm that is based on Madam and is also tailored to LNS for low-precision training.

**Our contributions are summarized as follows:**

1. We propose a low-precision training framework that considers forward propagation, backward propagation, and weight update process all in LNS. We also use a multiplicative weight update algorithm LNS-Madam for optimizing weights directly in their logarithmic representation.
2. We design a multi-base LNS where the log-base can be powers of two. Multi-base LNS provides additional flexibility of choosing different bases for weights and gradients, as they require different quantization gaps and dynamic ranges during training. In addition, we propose an approximation for addition arithmetic in LNS to further improve its energy efficiency.
3. With only 5-bit for backward propagation, multi-base LNS achieves accuracy comparable to full-precision state-of-the-art models such as ResNet-50 and BERT. This is the first time that such large-scale language models are trained in low-precision under LNS.
4. We propose a new optimizer LNS-Madam to optimize weights directly in their logarithmic representation without any conversion. Given both theoretical analyses and empirical evaluations, we verify that LNS-Madam achieves significantly lower quantization error compared to SGD or Adam.
5. Empirical results also validate that LNS-Madam always maintains high accuracy even when the precision of weight update is severely limited. For BERT model on SQuAD and GLUE benchmarks, LNS-Madam is 20% better than Adam with respect to F-1 score, when the weight update is in 10-bit.
6. We present an energy efficiency analysis for multi-base LNS on various neural networks. By calculating the energy consumed by math datapath units in forward and backward propagation, the analysis shows that LNS achieves 60x energy reduction compared to 32-bit floating-point (FP32) on BERT models. Notably, for a full training of ResNet-50 on ImageNet, our design reduces the total carbon emissions by 98% around.

For the following, we first introduce how to apply multi-base LNS to a standard quantized training where the weight update process is in full-precision. Then we include quantized weight update into our consideration and propose LNS-Madam optimizer.

## 2 Multi-Base Logarithmic Number System

In this section, we introduce our multi-base logarithmic number system (LNS), including the corresponding number representation and arithmetic operations.

We start with our mathematical formulation through this paper. We assume the DNN  $F(\cdot, W)$  is composed of  $L$  layers with learnable weights  $W$  and activations  $X$  across the layers.  $\mathcal{L}(W)$  denotes the training loss. The forward propagation is defined as:  $X_l = f_l(X_{l-1}, W_l)$ , where  $\forall l \in [1, L]$  is any layer.  $\nabla_{X_l} = \frac{\partial \mathcal{L}(W)}{\partial X_l}$  and  $\nabla_{W_l} = \frac{\partial \mathcal{L}(W)}{\partial W_l}$  denote gradients with respect to activations and weights, respectively. For number system, we define  $\mathcal{B}$  as the bitwidth,  $x$  as any number, and  $x^q$  as its quantized format.

**Multi-base Logarithmic Representation.** Unlike traditional works that use exact two as the base of logarithmic representation, we propose a multi-base logarithmic representation that allows the base to be two with a fractional exponent, such that:

$$x = \text{sign} \times 2^{\tilde{x}/\gamma} \quad \tilde{x} = 0, 1, 2, \dots, 2^{\mathcal{B}-1} - 1,$$

where  $\tilde{x}$  is an integer, and we term  $\gamma$  as a base factor. We restrict  $\gamma$  be powers of two to achieve better efficiency in hardware.  $\gamma$  controls the quantization gap, which is the distance between successive representable values within the number system. Previous works have already demonstrated that logarithmic quantized neural networks achieve better performance when relaxing  $\gamma$  from 1 to 2 [7]. However, a smaller and more flexible log-base setting is needed as the weights and gradients usually require different quantization gaps during training. Therefore we generalize this setting by providing the flexibility of controlling the quantization gap in order to approximate the training dynamics more accurately.

**Arithmetic Operations.** One of the benefits for using LNS stems from the low computational cost of its arithmetic operations. We use dot product operations as an example since they are prevalent during training. Consider two vectors  $\mathbf{a} \in \mathbb{R}^n$  and  $\mathbf{b} \in \mathbb{R}^n$  are represented by their integer exponents  $\tilde{\mathbf{a}}$  and  $\tilde{\mathbf{b}}$  in LNS. A dot product operation between them can be represented as follows:

$$\mathbf{a}^T \mathbf{b} = \sum_{i=1}^n \text{sign}_i \times 2^{\tilde{a}_i/\gamma} \times 2^{\tilde{b}_i/\gamma} = \sum_{i=1}^n \text{sign}_i \times 2^{(\tilde{a}_i + \tilde{b}_i)/\gamma} = \sum_{i=1}^n \text{sign}_i \times 2^{\tilde{p}_i/\gamma}, \quad (1)$$

where  $\text{sign}_i = \text{sign}(\mathbf{a}_i) \oplus \text{sign}(\mathbf{b}_i)$ . In this dot product operation, each element-wise multiplication is computed as an addition between integer exponents, which significantly reduces the computational cost by requiring adders instead of multipliers. However, the accumulation is difficult to be computed efficiently as it requires first converting from logarithmic to linear format and then performing the addition operation. The conversion between these formats is expensive as it requires computing  $2^{\tilde{p}_i/\gamma}$  using polynomial expansion.

**Conversion Approximation.** To address the cost of the conversion, we propose an efficient hybrid approximation method. Our method is based on Mitchell approximation [14]:  $2^{\tilde{x}/2^b} \approx (1 + \tilde{x}/2^b)$ , where the logarithmic format can be efficiently approximated to the linear format. In order to reduce the approximation error, rather than directly using Mitchell approximation, we design a hybrid way that leverages both Mitchell approximation and exact conversion. In addition, since the approximation serves as an additional non-linear operation in neural networks, we find the approximated training does not damage accuracy in practice. We present a detailed description of our approximation in the appendix.

### 3 Quantized Training by Multi-Base LNS

In this section, we introduce how to apply multi-base LNS to quantized training, as illustrated in Figure 2.

**Logarithmic Quantization.** To realize reduced precision for values and arithmetic during training, we define a logarithmic quantization function  $Q_{\log} : \mathbb{R} \rightarrow \mathbb{R}$ , which quantizes a real number into a sign and an integer exponent using a limited number of bits.  $Q_{\log}$  is defined as follows:

$$Q_{\log}(x) = \text{sign}(x) \times s \times 2^{(\tilde{x}/\gamma)}, \quad (2)$$

where  $\tilde{x} = \text{clamp}(\text{round}(\log_2(|x|/s) \times \gamma), 0, 2^{\mathcal{B}-1} - 1)$ , and  $s \in \mathbb{R}$  denotes a scale factor.  $Q_{\log}$  first brings scaled numbers  $|x|/s$  into their logarithmic space, magnifies them by the base factor  $\gamma$  and then performs rounding and clamping functions to convert them into desired integer exponents  $\tilde{x}$ .

**Quantization for Forward Propagation.** We apply quantization-aware training (QAT) for quantizing weights and activations during forward propagation. Each quantizer is associated with a STE to allow the gradients to directly flow through the non-differentiable quantizer during backward pass [4]. Because QAT views each quantization function as an additional non-linear operation in the networks, the deterministic quantization error introduced by any quantizer in the forward pass is implicitly reduced through training. We define weight quantization function as  $Q_W$  and activation quantization function as  $Q_A$  for each layer during forward propagation, where  $W_l^q = Q_W(W_l)$  and  $X_l^q = Q_A(f_l(X_{l-1}^q, W_l^q))$ .

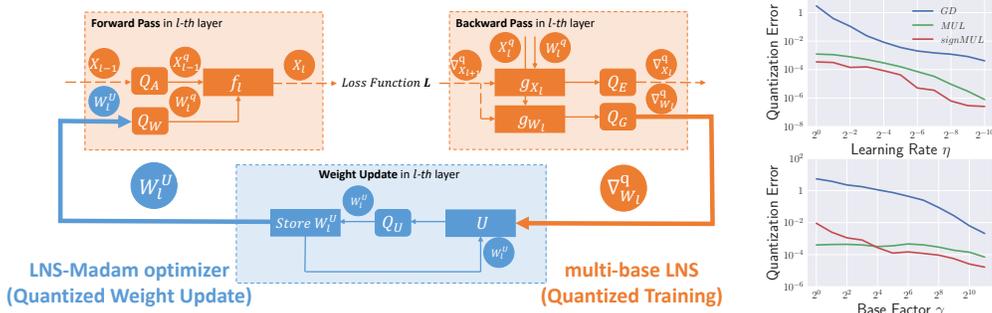


Figure 2: **Left:** illustration of our end-to-end low-precision training framework. Quantized training includes quantizing weights  $W$  and activations  $X$  in forward propagation, and weight gradients  $\nabla_W$  and activation gradients  $\nabla_X$  in backward propagation.  $g_X$  and  $g_W$  denote the functions to compute gradients. Quantized weight update applies a quantization function  $Q_U$  over weights after any learning algorithm  $U$  updates them. The quantized weights  $W^U$  are the actual numbers stored in the system. **Right:** quantization error from different learning algorithms on ImageNet. The errors are averaged over all iterations in the first epoch. The results suggest that multiplicative algorithms introduce significantly lower errors compared to the gradient descent, which are also in line with our theoretical results in Section 4.

**Quantization for Backward Propagation.** In order to accelerate training in addition to inference, gradients also need to be quantized into low-precision numbers. As shown by recent studies, the distribution of gradients resembles a Gaussian or Log-Normal distribution [15, 16]. This suggests that logarithmic representation may be more suitable than fixed-point or floating-point representations when quantizing gradients. We quantize the activation gradients using quantization function  $Q_E$ :  $\nabla_{X_l}^q = Q_E(\nabla_{X_l})$ . We also quantize the weight gradients using quantization function  $Q_G$ :  $\nabla_{W_l}^q = Q_G(\nabla_{W_l})$ . Because the activation gradients are involved in the most of the backward computations, the bitwidth of  $Q_E$  significantly affects the computational costs in backward propagation. Therefore, in this work we focus on reducing the bitwidth requirement of activation gradients.

## 4 Logarithmic Quantized Weight Update and LNS-Madam

Although logarithmic quantized training significantly improves training efficiency, its overall efficiency is hampered by the high-precision requirement of weights during weight update. Previous works generally assume the weights are updated over a full-precision weight space, in other words, a full-precision copy of weights is maintained [7, 12]. However, this offsets the efficiency improvement achieved by quantized training, and expensive FP32 arithmetic is not available especially in cheap energy-constrained edge devices.

Therefore, in this work, we consider quantized weight update, where the weights are updated over a discrete logarithmic space instead of a full-precision one. We note that quantized weight update is orthogonal to quantized training due to different objectives. Quantized training tries to maintain the fidelity of weight gradients, which aims to provide accurate gradients for the weight update. On the other hand, quantized weight update aims to reduce gaps between updated weights and their quantized counterparts.

**Quantized Weight Update.** To better understand this problem, we first define a generalized form of a weight update as:  $W_{t+1} = U(W_t, \nabla_{W_t})$ , where  $U$  represents any learning algorithm. For example, gradient descent (GD) algorithm takes  $U_{GD} = W - \eta \nabla_W$ , where  $\eta$  is learning rate.

Because the weights need to be represented in a quantized format, it is necessary to consider the effect of logarithmic quantization during weight update. We define *logarithmic quantized weight update* as follows:

$$W_{t+1}^U = Q_{\log}(U(W_t, \nabla_{W_t})). \quad (3)$$

In this case,  $W_{t+1}^U$  can be directly stored in a logarithmic format without using FP32. For simplicity, we assume weight gradients  $\nabla_W$  are exact as quantized training is orthogonal to this problem. Switching to the approximated gradient estimates will not affect our theoretical results.

**Quantization Error.** Because of the logarithmic quantization, quantized weight update inevitably introduces a mismatch between the quantized weights and their full-precision counterparts. To preserve the reliability of the optimization, we aim to reduce the quantization error (i.e., the mismatch) to a lower level. For the following, we take a theoretical point of view to discuss how different

learning algorithms affect the quantization error under LNS. Detailed assumptions and proofs can be found in the appendix.

Due to the logarithmic structure, we focus on minimizing a quantization error  $r_t = \|\log_2 |W_{t+1}^U| - \log_2 |W_{t+1}|\|^2$ , which measures the L2 norm of the difference between the weights and their quantized counterparts in logarithmic space. Because  $r_t$  quantify a relative difference between  $|W_{t+1}^U|$  and  $|W_{t+1}|$ , minimizing  $r_t$  is equivalent to minimizing relative quantization error  $\|(W_{t+1} - W_{t+1}^U)/W_{t+1}\|^2$ .

We assume a simplified logarithmic quantization where the scale factor and the clamping function are ignored. This ensures our focus is on the effect of the quantization gap determined by  $\gamma$  instead of the dynamic range. We also replace the deterministic rounding with a stochastic counterpart  $SR$  where  $\mathbb{E} SR(x) = x$  for any real number. Although  $SR$  helps us establish the theoretical results, in practice  $SR$  requires random generators that induce additional costs, and thus are not suitable for energy-efficient training.

Given everything we need, we use gradient descent as an example to discuss why traditional learning algorithms are not suited for LNS-based quantized weight update. The theorem is stated as follows:

**Theorem 1.** *The quantization error  $r_{t,GD}$  introduced by logarithmic quantized gradient descent at iteration  $t$  can be bounded in expectation, as:*

$$\mathbb{E} r_{t,GD} \leq \frac{\sqrt{d}}{\gamma} \|\log_2 (|W_t| - \eta_1 \nabla_{W_t})\|, \quad (4)$$

where  $d$  is the dimension of  $W$  and  $\eta_1$  is the learning rate of  $U_{GD}$ .

Theorem 1 suggests that  $r_{t,GD}$  is magnified when the magnitudes of weights become larger. This is because the updates  $\eta_1 \nabla_{W_t}$  generated by GD are not proportional to the magnitudes of weights.  $\eta_1 \nabla_{W_t}$  can be orders of magnitude smaller than the quantization gaps as weights become larger, and thus these updates often are disregarded by quantization function  $Q_{\log}$ . We intuitively illustrate this problem in Figure 1.

To ensure the updates are proportional to the weights, a direct way is to update the weights multiplicatively. Because the weights are represented in LNS, we further consider a special multiplicative learning algorithm tailored to LNS, which updates the weights directly over their logarithmic space:

$$U_{MUL} = \text{sign}(W) \odot 2^{\tilde{W} - \eta \nabla_{W} \odot \text{sign}(W)}$$

where  $\tilde{W} = \log_2 |W|$  are the exponents of the magnitude of weights, and  $\odot$  denotes element-wise multiplication.  $U_{MUL}$  makes sure the magnitude of each element  $k$  of the weights increases when the sign  $\text{sign}(W_k)$  and  $\nabla_{W_k}$  agree and decreases otherwise. The quantization error wrt.  $U_{MUL}$  is stated as follows:

**Theorem 2.** *The quantization error  $r_{t,MUL}$  introduced by logarithmic quantized multiplicative weight update at iteration  $t$  can be bounded in expectation, as:*

$$\mathbb{E} r_{t,MUL} \leq \frac{\sqrt{d} \eta_2}{\gamma} \|\nabla_{W_t}\|, \quad (5)$$

where  $d$  is the dimension of  $W$  and  $\eta_2$  is the learning rate of  $U_{MUL}$ .

Theorem 2 indicates that  $r_{t,MUL}$  does not depend on the magnitudes of weights, and thus the quantization error is not magnified when the weights become larger. This is also illustrated at Figure 1.

Interestingly, we find that the quantization error  $r_{t,MUL}$  can be further reduced by regularizing the information of gradients for the learning algorithm  $U_{MUL}$ :

**Lemma 1.** *Assume the multiplicative learning algorithm  $U_{MUL}$  only receives the sign information of gradients where  $U_{MUL} = \tilde{W} - \eta_2 \text{sign}(\nabla_W) \odot \text{sign}(W)$ . The upper bound on quantization error  $r_{t,MUL}$  becomes:*

$$\mathbb{E} r_{t,MUL} \leq \frac{d \eta_2}{\gamma}. \quad (6)$$

The result in Lemma 1 suggests that  $r_{t,MUL}$  can be independent of both weights and gradients when only taking sign information of gradients during weight update. We denote this special learning algorithm as  $U_{\text{sign}MUL}$ .  $U_{\text{sign}MUL}$  is a multiplicative version of signSGD, which has been studied widely [16, 17].

To verify our theoretical results, we empirically measure the quantization errors for the three aforementioned learning algorithms over a range of  $\eta$  and  $\gamma$ . As shown in Figure 2, the empirical findings are in line with our theoretical results. Although all learning algorithms introduce less errors when  $\eta$  and  $\gamma$  become smaller, the multiplicative algorithms introduce significantly lower errors compared to the gradient descent.

Because  $U_{signMUL}$  further reduces the error by regularizing the information of gradients, this motivates us to design a learning algorithm based on the rules of  $U_{signMUL}$ . However,  $U_{signMUL}$  still fails to minimize the loss without showing it updates the weights in a descent direction. This brings us the need to discuss  $U_{signMUL}$  under an optimization point of view. Interestingly, because signSGD is a special case of Adam optimizer, we notice that  $U_{signMUL}$  is also a special case of a recently proposed optimizer Madam, which is a multiplicative version of Adam [13].

Madam updates the weights multiplicatively using normalized gradients:

$$U_{Madam} = W \odot e^{-\eta sign(W) \odot g^*} \quad g^* := g^t / \sqrt{g_2}, \quad (7)$$

where  $g$  represents the gradient vector  $\nabla_W$ , and  $g^*$  denote a normalized gradient, which is the fraction between  $g$  and the square root of its second moment estimate  $\sqrt{g_2}$ . Bernstein et al. [13] demonstrates that Madam achieves state-of-the-art accuracy over multiple tasks with a relatively fixed learning rate  $\eta$ . They also theoretically prove the descent property of Madam that ensures its convergence. Although Bernstein et al. [13] further shows the possibility of applying Madam over a discrete logarithmic weight space, they still employ a full-precision training without considering low-precision LNS.

By considering both  $U_{signMUL}$  and  $U_{Madam}$ , we propose a new optimizer LNS-Madam that ensures fast convergence while introducing small quantization error. LNS-Madam directly optimizes the weights over their base-2 logarithmic space using the gradient normalization technique described in Equation 7. Details of LNS-Madam are shown in Algorithm 1.

Because LNS-Madam directly updates base-2 exponents of weights in LNS, there is no need for linear-to-log conversion when the weights are in LNS, which further reduces the energy cost during weight update.

## 5 Experiments

In this section, we empirically evaluate multi-base LNS and LNS-Madam on large-scale datasets and state-of-the-art models. In the end, we present an energy analysis to show the efficiency of multi-base LNS.

We design an end-to-end quantization training system to simulate both quantized training and quantized weight update. We benchmark our framework on various tasks including ResNet models on CIFAR-10 and ImageNet, and BERT-base and BERT-large language models on SQuAD and GLUE. Detailed descriptions of the system design, datasets, and models can be found in the appendix.

For the following, we first evaluate multi-base LNS under a standard quantized training setting, and then we include quantized weight update into our evaluation to demonstrate the effectiveness of the LNS-Madam optimizer.

### 5.1 Benchmarking Quantized Training by Multi-Base LNS

**Bitwidth and Base Factor Selection.** As the energy costs of training are usually dominated by backward propagation, in this work we put our focus on minimizing the bitwidth requirement of gradients. We fix weights and activations in 8-bit as a standard practice [3]. Although we note that LNS could train networks with a bitwidth setting lower than 8-bit, we want to rule out their effects in order to focus on the backward propagation. We also maintain weight gradients in 8-bit, which is lower than 16-bit or 32-bit settings used in other studies [10, 7].

To find an appropriate combination of bitwidth  $\mathcal{B}$  and base factor  $\gamma$ , we design an efficient searching strategy that first finds a desired dynamic range and then tests different  $\gamma$  to find a minimal required

---

#### Algorithm 1 LNS-Madam

---

**Require:** Base-2 weight exponents  $\tilde{W}$ , where  $\tilde{W} = \log_2(W)$ , learning rate  $\eta$ , momentum  $\beta$   
Initialize  $g_2 \leftarrow 0$   
**repeat**  
 $g \leftarrow \text{StochasticGradient}()$   
 $g_2 \leftarrow (1 - \beta)g^2 + \beta g_2$   
 $g^* \leftarrow g / \sqrt{g_2}$   
 $\tilde{W} \leftarrow \tilde{W} - \eta g^* \odot \text{sign}(W)$   
**until** converged

---

Table 1: Benchmarking multi-base LNS on various datasets and models. We compare multi-base LNS with INT and FP32 baselines. A unified bitwidth setting is used across tasks: 5-bit for activation gradients  $Q_E$ , and 8-bit for weights  $Q_W$ , activations  $Q_A$  and weight gradients  $Q_G$ . Each task uses a uniform number of training epochs. Energy reduction compares both the absolute and relative training costs between LNS and FP32 (the higher the better).

Dataset	Task	Multi-Base LNS	INT	Full-Precision	Energy Reduction
CIFAR-10	ResNet-18	93.37% $\pm$ 0.17	93.28% $\pm$ 0.15	93.51% $\pm$ 0.23	2.7e-3 KWh (98.4%)
ImageNet	ResNet-50	75.89% $\pm$ 0.05	73.82% $\pm$ 0.06	76.38% $\pm$ 0.05	5.7e-1 KWh (98.4%)
SQuAD	BERT-base	87.90% $\pm$ 0.21	23.7% $\pm$ 0.49	88.36% $\pm$ 0.19	6.8e-3 KWh (98.3%)
SQuAD	BERT-large	90.34% $\pm$ 0.17	16.1% $\pm$ 0.62	90.80% $\pm$ 0.22	2.4e-2 KWh (98.3%)
GLUE	BERT-base	88.07% $\pm$ 0.45	81.03% $\pm$ 0.32	88.92% $\pm$ 0.39	4.2e-4 KWh (98.3%)
GLUE	BERT-large	89.15% $\pm$ 0.41	87.87% $\pm$ 0.36	89.35% $\pm$ 0.44	1.5e-3 KWh (98.3%)

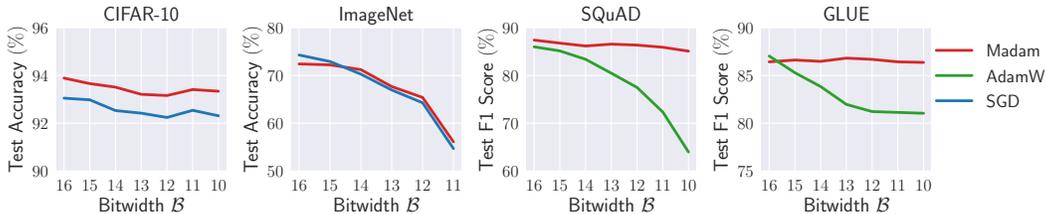


Figure 3: Comparing LNS-Madam with SGD and Adam optimizers under multi-base LNS with logarithmic quantized weight update. The bitwidth of the weights during the update  $Q_U$  is varied from 16-bit to 10-bit.

quantization gap. After applying the searching strategy across a wide range of tasks, we identify the best setting where 5-bit with  $\gamma = 1$  is applied for activation gradients and 8-bit with  $\gamma = 8$  is applied for the rest of operands. A detailed selection strategy and results are discussed in the appendix.

**Scaling Techniques.** Scale techniques play an essential role in accommodating the dynamic range of the values, and therefore significantly impact quantization error and network accuracy. Although studying the scaling techniques is orthogonal to the focus of our work, sophisticated scaling techniques are necessary for training with ultra-low bitwidth settings. In this work, we adopt an effective per-vector scaling technique recently proposed by Dai et al., which allocates multiple scale factors within a single dimension of a tensor. As we apply the per-vector scaling for every low-precision experiment, this technique should not bias our comparison. Although we find LNS with a simple per-layer scaling achieves almost no loss accuracy on ImageNet even in 4-bit, we observe a large degradation when applying the same setting on BERT fine-tuning tasks. Additional experiments studying the effects of different scaling techniques can be found in the appendix.

**Comparisons.** Given the bitwidth setting and per-vector scaling mentioned earlier, we first benchmark multi-base LNS under a quantized training setting with full-precision weight update. As shown in Table 1, multi-base LNS consistently outperforms fixed-point number system (INT) and achieves accuracy comparable to the FP32 counterpart. In addition, compared to FP32, multi-base LNS achieves large energy reductions across tasks. As INT usually requires bitwidth larger than 5-bit to represent the gradients, we observe large performance gaps between LNS and INT on SQuAD benchmarks.

We also compare multi-base LNS with related works in the field. Recently, Chen et al. [5] proposes block householder quantizer (BHQ) that establishes new state-of-the-art results for training ResNet-50 on ImageNet with low-precision gradients. BHQ is built upon fixed-point representation and aims to reduce the variance of gradients. We compare multi-base LNS with BHQ over a range of bitwidth settings of activation gradients. As shown in table 3, LNS matches the results of BHQ and even yields better accuracy at 5-bit and 7-bit. Furthermore, since BHQ requires a fixed-point number system, multi-base LNS naturally has better energy efficiency compared to it.

## 5.2 Benchmarking Quantized Weight Update by LNS-Madam

After showing strong performance for low-precision multi-base LNS under quantized training, we further include the quantized weight update into consideration to benchmark LNS-Madam optimizer. We compare LNS-Madam with popular optimizers SGD and AdamW under a multi-base LNS, which is configured with the same bitwidth setting as before.

We empirically search the best learning rate  $\eta$  for LNS-Madam from  $2^{-4}$  to  $2^{-10}$ , and we find  $\eta = 2^{-7}$  works uniformly well across tasks. This is in line with the robustness of the learning rate for Madam as discussed in Bernstein et al. [13]. We also use well-tuned learning rate settings for SGD and AdamW, as provided by default.

As shown in Figure 3, we vary the bitwidth of  $Q_U$  from 16-bit to 10-bit to test their performance over a wide range. The results suggest compared to other optimizers, LNS-Madam always maintains higher accuracy when precision is severely limited. Notably, for BERT model on SQuAD and GLUE benchmarks, LNS-Madam is 20% better than Adam with respect to F-1 score, when the weight update is in 10-bit. We observe large degradation for both LNS-Madam and SGD on ImageNet training, and we believe this is because the weights in some layers inevitably require higher precision settings during training. We leave it as future work to explore LNS-Madam under a customized precision setting.

### 5.3 Energy Efficiency Analysis

Finally, we analyze the energy efficiency of multi-base LNS. We implemented datapath units in sub-16nm state-of-the-art process technology using the methodology adopted by the MAGNet DNN accelerator [19]. To align with the bitwidth setting of LNS we used, the hardware implementation has two separate datapaths specialized for executing different steps of neural network training. A datapath with  $\mathcal{B} = 8, \gamma = 8$  is used to perform computations in the forward pass, while a datapath with  $\mathcal{B} = 5, \gamma = 1$  is used for the backward pass. Also, we consider an 8-bit fixed point (INT8) datapath, a FP16 datapath, and a FP32 datapath for the baseline. As shown in Table 2, LNS achieves 1.47x to 1.53x energy reduction over INT8, more than 20x improvements over FP16, and more than 60x improvements over FP32. In Figure 1, we also show the energy costs over a range of GPT models. Given the energy results, we also calculate expected carbon emissions using a method provided by Patterson et al. [1]. We describe our calculation methods and present additional results in the appendix.

## 6 Related Works

**Low-Precision Training.** To achieve good accuracy at reduced precision, quantization-aware training (QAT) is commonly applied to directly train the quantized model using straight-through estimators [4, 20–23]. To accelerate the training phase, several studies suggest quantizing the gradients during backward propagation [24, 10, 3]. To maintain the fidelity of the gradient accumulation, some low-precision training methods assume a full-precision copy for weights during the weight update [24, 5]. Other studies reduce the precision for the weight update by using high-precision gradient accumulator [25], stochastic rounding [26, 3] or additionally quantizing the residual part of weights [27, 28]. However, they mostly apply SGD or Adam during the weight update without considering the relationship between the precision of the weights and the underlying learning algorithms.

**Logarithmic Number System.** Previous works demonstrate the effectiveness of using logarithmic representation for DNNs [7, 29–31]. Furthermore, some studies suggest using multiple levels of log-base to reduce the quantization error [7, 12]. However, few of them address the additional computational cost induced by this multi-base design nor scale the system to state-of-the-art neural networks for both training and inference. From the perspective of hardware design, a few studies focus on improving the efficiency of LNS by utilizing the significant cost reduction of multiplications [32–34, 30, 31].

**Multiplicative Weight Update.** Multiplicative algorithms, such as exponentiated gradient algorithm and Hedge algorithm in AdaBoost framework [35, 36], have been well studied in the field of machine learning. In general, multiplicative updates are applied to problems where the optimization domain’s geometry is described by relative entropy, such as probability distribution [35]. Recently,

Table 2: Energy efficiency for different models and numeric formats. The energy costs per-iteration (mJ) are shown below.

Model	LNS	INT8	FP16	FP32
<b>ResNet-18</b>	0.15	0.23	3.34	9.43
<b>ResNet-50</b>	0.28	0.42	6.15	17.35
<b>BERT-Base</b>	2.33	3.43	49.82	140.55
<b>BERT-Large</b>	8.12	11.98	173.73	490.14

Table 3: Compare LNS and BHQ over a range of activation gradients bitwidth  $\mathcal{B}_E$  on ImageNet. The results of test accuracy (%) are shown below.

	LNS	INT	BHQ
<b>8-bit</b>	76.04	75.72	<b>76.23</b>
<b>7-bit</b>	<b>76.16</b>	75.12	76.14
<b>6-bit</b>	76.06	74.41	<b>76.21</b>
<b>5-bit</b>	<b>75.89</b>	73.82	75.70
<b>4-bit</b>	73.73	44.03	<b>74.04</b>

[13] proposes an optimizer Madam that focuses on optimization domains described by any relative distance measure instead of only relative entropy. Madam shows great performance in training large-scale neural networks. However, Madam requires full-precision training without considering its connection to LNS-based low-precision training.

## 7 Discussions and Broader Impacts

In this work, we propose an end-to-end low-precision training framework using multi-base LNS and learning algorithm LNS-Madam. One of the most important applications of our low-precision framework is learning neural networks over energy-constrained edge devices. This is fundamental for intelligent edge devices to easily adapt to the changing and non-stationary environments by learning on the device. Also, our work has positive societal outcomes for energy conservation and environmental cleaning. Based on our energy analysis, our multi-base LNS reduces the carbon emissions by 98% around for full training of ResNet-50 on ImageNet. This shows a promising opportunity for using LNS-based hardware to conduct environmental-friendly deep learning research in the future.

## Acknowledgements

We would like to sincerely thank Mohammad Shoeybi and Yuanyuan Shi for meaningful discussions and suggestions. A. Anandkumar gratefully acknowledges the support from Bren endowed chair at Caltech.

## References

- [1] David Patterson, Joseph Gonzalez, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David So, Maud Texier, and Jeff Dean. Carbon emissions and large neural network training, 2021.
- [2] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, 2015.
- [3] Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. In *Neural Information Processing Systems*, 2018.
- [4] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- [5] Jianfei Chen, Yu Gai, Zhewei Yao, Michael W Mahoney, and Joseph E Gonzalez. A statistical framework for low-bitwidth training of deep neural networks. *arXiv preprint arXiv:2010.14298*, 2020.
- [6] NVIDIA Corporation. Nvidia a100 tensor core gpu architecture. *NVIDIA*, 2020.
- [7] Daisuke Miyashita, Edward H. Lee, and Boris Murmann. Convolutional neural networks using logarithmic data representation, 2016.
- [8] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018.
- [9] Thomas M. Bartol, Jr., Cailey Bromer, Justin Kinney, Michael A. Chirillo, Jennifer N. Bourne, Kristen M. Harris, and Terrence J. Sejnowski. Nanoconnectomic upper bound on the variability of synaptic plasticity. *eLife*, 2015.
- [10] 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. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1796–1807. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/13b919438259814cd5be8cb45877d577-Paper.pdf>.

- [11] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Anand Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. Efficient large-scale language model training on gpu clusters, 2021.
- [12] Sebastian Vogel, Mengyu Liang, Andre Guntoro, Walter Stechele, and Gerd Ascheid. Efficient hardware acceleration of cnns using logarithmic data representation with arbitrary log-base. In *Proceedings of the International Conference on Computer-Aided Design, ICCAD '18*, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450359504. doi: 10.1145/3240765.3240803. URL <https://doi.org/10.1145/3240765.3240803>.
- [13] Jeremy Bernstein, Jiawei Zhao, Markus Meister, Ming-Yu Liu, Anima Anandkumar, and Yisong Yue. Learning compositional functions via multiplicative weight updates. *arXiv preprint arXiv:2006.14560*, 2020.
- [14] J. N. Mitchell. Computer multiplication and division using binary logarithms. *IRE Transactions on Electronic Computers*, EC-11(4):512–517, 1962. doi: 10.1109/TEC.1962.5219391.
- [15] Brian Chmiel, Liad Ben-Uri, Moran Shkolnik, Elad Hoffer, Ron Banner, and Daniel Soudry. Neural gradients are near-lognormal: improved quantized and sparse training. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=EoFNy62JGd>.
- [16] Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signSGD with majority vote is communication efficient and fault tolerant. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=BJxhijAcY7>.
- [17] Jeremy Bernstein, Jiawei Zhao, Kamyar Azizzadenesheli, and Anima Anandkumar. signsgd with majority vote is communication efficient and fault tolerant, 2019.
- [18] Steve Dai, Rangha Venkatesan, Haoxing Ren, Brian Zimmer, William Dally, and Brucek Khailany. VS-Quant: Per-vector Scaled Quantization for Accurate Low-Precision Neural Network Inference. *arXiv preprint*, 2021.
- [19] R. Venkatesan, Y. S. Shao, M. Wang, J. Clemons, S. Dai, M. Fojtik, B. Keller, A. Klinefelter, N. Pinckney, P. Raina, Y. Zhang, B. Zimmer, W. J. Dally, J. Emer, S. W. Keckler, and B. Khailany. Magnet: A modular accelerator generator for neural networks. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 1–8, Nov 2019. doi: 10.1109/ICCAD45719.2019.8942127.
- [20] 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.
- [21] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pages 525–542. Springer, 2016.
- [22] Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental network quantization: Towards lossless cnns with low-precision weights. *arXiv preprint arXiv:1702.03044*, 2017.
- [23] Benoît 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*, pages 2704–2713, 2018.
- [24] Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable methods for 8-bit training of neural networks. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, pages 5145–5153. Curran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/file/e82c4b19b8151ddc25d4d93baf7b908f-Paper.pdf>.
- [25] Charbel Sakr and Naresh Shanbhag. Per-tensor fixed-point quantization of the back-propagation algorithm. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rkxaNjA9Ym>.

- [26] Shuang Wu, Guoqi Li, Feng Chen, and Luping Shi. Training and inference with integers in deep neural networks. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=HJGXzmspb>.
- [27] Xiao Sun, Jungwook Choi, Chia-Yu Chen, Naigang Wang, Swagath Venkataramani, Vijayalakshmi Viji Srinivasan, Xiaodong Cui, Wei Zhang, and Kailash Gopalakrishnan. Hybrid 8-bit floating point (hfp8) training and inference for deep neural networks. *Advances in neural information processing systems*, 32:4900–4909, 2019.
- [28] Christopher De Sa, Megan Leszczynski, Jian Zhang, Alana Marzoev, Christopher R. Aberger, Kunle Olukotun, and Christopher Ré. High-accuracy low-precision training, 2018.
- [29] E. H. Lee, D. Miyashita, E. Chai, B. Murmann, and S. S. Wong. Lognet: Energy-efficient neural networks using logarithmic computation. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5900–5904, 2017. doi: 10.1109/ICASSP.2017.7953288.
- [30] Jeff Johnson. Rethinking floating point for deep learning. *arXiv preprint arXiv:1811.01721*, 2018.
- [31] J. Johnson. Efficient, arbitrarily high precision hardware logarithmic arithmetic for linear algebra. In *2020 IEEE 27th Symposium on Computer Arithmetic (ARITH)*, pages 25–32, 2020. doi: 10.1109/ARITH48897.2020.00013.
- [32] Hassaan Saadat, Haseeb Bokhari, and Sri Parameswaran. Minimally biased multipliers for approximate integer and floating-point multiplication. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 37(11):2623–2635, 2018.
- [33] Hassaan Saadat, Haris Javaid, and Sri Parameswaran. Approximate integer and floating-point dividers with near-zero error bias. In *2019 56th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2019.
- [34] Hassaan Saadat, Haris Javaid, Aleksandar Ignjatovic, and Sri Parameswaran. Realm: reduced-error approximate log-based integer multiplier. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1366–1371. IEEE, 2020.
- [35] Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Information and Computation*, 1997.
- [36] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 1997.
- [37] M.C. McFarland, A.C. Parker, and R. Camposano. The high-level synthesis of digital systems. *Proceedings of the IEEE*, 78(2):301–318, Feb. 1990.
- [38] Hao Wu, Patrick Judd, Xiaojie Zhang, Mikhail Isaev, and Paulius Micikevicius. Integer quantization for deep learning inference: Principles and empirical evaluation. *arXiv preprint arXiv:2004.09602*, 2020.
- [39] Nvidia v100 tensor core gpu. URL <https://www.nvidia.com/en-us/data-center/v100/>.
- [40] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition*, 2016.
- [41] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [42] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [43] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009.

- [44] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1264. URL <https://www.aclweb.org/anthology/D16-1264>.
- [45] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium, November 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-5446. URL <https://www.aclweb.org/anthology/W18-5446>.

## A Quantization Error Analysis

Here we present proofs of theorems and lemmas presented in the main paper, as well as some additional details of the empirical evaluations.

### A.1 Proofs

Before presenting the proofs, we want to clarify the error definition and assumptions introduced in the main paper. Previously, we claim that minimizing  $r_t = \|\log_2 |W_{t+1}^U| - \log_2 |W_{t+1}|\|^2$  is equivalent to minimizing relative quantization error  $\|(W_{t+1} - W_{t+1}^U)/W_{t+1}\|^2$ . To better understand it, we transform the form of relative quantization error as follows:

$$\begin{aligned} \|(W_{t+1} - W_{t+1}^U)/W_{t+1}\|^2 &= \|(I - W_{t+1}^U)/W_{t+1}\|^2 \\ &= \|(I - |W_{t+1}^U|)/|W_{t+1}|\|^2 \quad (\text{sign}(W_{t+1}^U) = \text{sign}(W_{t+1})) \\ &= \|(I - 2^{\log_2 |W_{t+1}^U| - \log_2 |W_{t+1}|})\|^2 \quad (\text{transfer to base-2 logarithmic space}) \end{aligned}$$

This relaxation suggests that minimizing  $r_t$  is equivalent to minimizing the relative quantization error.

We start to introduce the simplified logarithmic quantization we used for the analysis. The stochastic rounding (SR) is defined as follows:

$$\text{SR}(x) = \begin{cases} \lfloor x \rfloor + 1 & \text{for } p \leq x - \lfloor x \rfloor, \\ \lfloor x \rfloor & \text{otherwise,} \end{cases} \quad (8)$$

where  $p \in [0, 1]$  is generated by a uniform random number generator. SR makes sure the rounded number is an unbiased estimate of its full-precision counterpart:  $\mathbb{E} \text{SR}(x) = x$ , which is an important property for the analysis.

Equipped with SR, we define the simplified logarithmic quantization function:

$$\text{Q}_{\log}(x) = \text{sign}(x) \times 2^{\tilde{x}/\gamma}, \quad (9)$$

where  $\tilde{x} = \text{SR}(\log_2 |x| \times \gamma)$ . We ignore the scale factor and the clamping function to ensure our focus is on the effect of the quantization gap instead of the dynamic range.

Before proving our main results, we want to introduce an important proposition that describes the error introduced by stochastic rounding.

**Proposition 1.** *For any vector  $x$ , the quantization error introduced by stochastic rounding  $r = \text{SR}(x) - x$  can be bounded in expectation, as:*

$$\mathbb{E} \|r\|^2 \leq \sqrt{d} \|x\|, \quad (10)$$

where  $d$  is the dimension of  $x$ .

*Proof.* Let  $r_i$  denotes the  $i$ th element of  $r$  and let  $q_i = x_i - \lfloor x_i \rfloor$ .  $r_i$  can be represented as follows:

$$\begin{aligned} r_i &= \begin{cases} \lfloor x_i \rfloor + 1 - x_i & \text{for } p \leq x_i - \lfloor x_i \rfloor, \\ \lfloor x_i \rfloor - x_i & \text{otherwise,} \end{cases} \\ &= \begin{cases} -q_i + 1 & \text{for } p \leq q_i, \\ -q_i & \text{otherwise.} \end{cases} \end{aligned}$$

$r_i$  can be bounded by expectation, as:

$$\begin{aligned} \mathbb{E} r_i^2 &\leq (-q_i + 1)^2 q_i + (-q_i)^2 (1 - q_i) \\ &= q_i (1 - q_i) \\ &\leq \min\{q_i, 1 - q_i\} \\ &= \min\{x_i - \lfloor x_i \rfloor, 1 - x_i + \lfloor x_i \rfloor\} \\ &\leq |x_i|. \end{aligned}$$

Therefore, by summing over index  $i$ , we can get:

$$\begin{aligned} \mathbb{E} \|r\|^2 &\leq \|x\|_1 \\ &\leq \sqrt{d} \|x\|. \end{aligned}$$

□

Now we start to prove Theorem 1 given  $U_{GD} = W - \eta \nabla_W$ .

**Theorem 1.** *The quantization error  $r_{t,GD}$  introduced by logarithmic quantized gradient descent at iteration  $t$  can be bounded in expectation, as:*

$$\mathbb{E} r_{t,GD} \leq \frac{\sqrt{d}}{\gamma} \|\log_2(|W_t| - \eta_1 \nabla_{W_t})\|, \quad (4)$$

where  $d$  is the dimension of  $W$  and  $\eta_1$  is the learning rate of  $U_{GD}$ .

*Proof.*

$$\mathbb{E} r_{t,GD} = \|\log_2 |Q_{\log}(W_t - \eta_1 \nabla_{W_t})| - \log_2 |W_t - \eta_1 \nabla_{W_t}|\|^2. \quad (11)$$

By replacing  $Q_{\log}$  with Equation 9, we can get:

$$\log_2 |Q_{\log}(W_t - \eta_1 \nabla_{W_t})| = \frac{1}{\gamma} \text{SR}(\gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|).$$

Plug it back to Equation 11:

$$\mathbb{E} r_{t,GD} = \frac{1}{\gamma^2} \|\text{SR}(\gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|) - \gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|\|^2.$$

Given Proposition 1, we can upper bound the quantization error introduced by stochastic rounding:

$$\|\text{SR}(\gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|) - \gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|\|^2 \leq \sqrt{d} \|\gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|\|.$$

Therefore, we can get:

$$\begin{aligned} \mathbb{E} r_{t,GD} &\leq \frac{\sqrt{d}}{\gamma^2} \|\gamma \log_2 |W_t - \eta_1 \nabla_{W_t}|\| \\ &\leq \frac{\sqrt{d}}{\gamma} \|\log_2 |W_t - \eta_1 \nabla_{W_t}|\|. \end{aligned}$$

□

Given  $U_{MUL} = \text{sign}(W) \odot 2^{\tilde{W} - \eta \nabla_{W_t} \odot \text{sign}(W)}$ , Theorem 2 follows a similar proof as Theorem 1.

**Theorem 2.** *The quantization error  $r_{t,MUL}$  introduced by logarithmic quantized multiplicative weight update at iteration  $t$  can be bounded in expectation, as:*

$$\mathbb{E} r_{t,MUL} \leq \frac{\sqrt{d} \eta_2}{\gamma} \|\nabla_{W_t}\|, \quad (5)$$

where  $d$  is the dimension of  $W$  and  $\eta_2$  is the learning rate of  $U_{MUL}$ .

*Proof.*

$$\mathbb{E} r_{t,MUL} = \|\log_2 |Q_{\log}(2^{\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)})| - \log_2 |2^{\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)}|\|^2. \quad (12)$$

By replacing  $Q_{\log}$  with Equation 9, we can get:

$$\log_2 |Q_{\log}(2^{\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)})| = \frac{1}{\gamma} \text{SR}(\gamma (\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t))).$$

Plug it back to Equation 12:

$$\mathbb{E} r_{t,MUL} = \frac{1}{\gamma^2} \|\text{SR}(\gamma (\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t))) - \gamma (\tilde{W}_t - \eta_2 \nabla_{W_t} \odot \text{sign}(W_t))\|^2.$$

Because  $\tilde{W}_t$  is already an integer,  $\text{SR}(\gamma \tilde{W}_t) - \gamma \tilde{W}_t = 0$ , and thus we can eliminate  $\tilde{W}_t$  in the equation:

$$\mathbb{E} r_{t,MUL} = \frac{1}{\gamma^2} \|\text{SR}(-\gamma \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)) + \gamma \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)\|^2.$$

Similar to the proof of Theorem 1, we can upper bound it using Proposition 1, and get:

$$\begin{aligned} \mathbb{E} r_{t,MUL} &\leq \frac{\sqrt{d}}{\gamma^2} \|\gamma \eta_2 \nabla_{W_t} \odot \text{sign}(W_t)\| \\ &\leq \frac{\sqrt{d} \eta_2}{\gamma} \|\nabla_{W_t}\|. \end{aligned}$$

□

**Lemma 1.** *Assume the multiplicative learning algorithm  $U_{MUL}$  only receives the sign information of gradients where  $U_{MUL} = \tilde{W} - \eta_2 \text{sign}(\nabla_{W_t}) \odot \text{sign}(W)$ . The upper bound on quantization*

error  $r_{t,MUL}$  becomes:

$$\mathbb{E} r_{t,MUL} \leq \frac{d \eta_2}{\gamma}. \quad (6)$$

*Proof.* We can simply replace  $\nabla_{W_t}$  with  $sign(\nabla_{W_t})$  in the result of Theorem 2, and show:

$$\frac{\sqrt{d} \eta_2}{\gamma} \|sign(\nabla_{W_t})\| \leq \frac{d \eta_2}{\gamma}.$$

□

## A.2 Evaluations

As shown in Figure 2, we evaluate empirical quantization errors from different learning algorithms when training ResNet-50 on ImageNet. The quantization error is computed at each iteration by  $\|\log_2 |W_{t+1}^U| - \log_2 |W_{t+1}|\|^2$ . We run each experiment with a full epoch and average the quantization error over iterations. When varying learning rate  $\eta$ , we fix the base factor  $\gamma$  as  $2^{10}$ . We also fix  $\eta$  as  $2^{-6}$  when varying  $\gamma$ .

## B Multi-Base LNS

### B.1 Conversion Approximation

We first recap the dot product operation we defined before.

$$\mathbf{a}^T \mathbf{b} = \sum_{i=1}^n sign_i \times 2^{\tilde{a}_i/\gamma} \times 2^{\tilde{b}_i/\gamma} = \sum_{i=1}^n sign_i \times 2^{(\tilde{a}_i + \tilde{b}_i)/\gamma} = \sum_{i=1}^n sign_i \times 2^{\tilde{p}_i/\gamma}, \quad (13)$$

where  $sign_i = sign(\mathbf{a}_i) \oplus sign(\mathbf{b}_i)$ .

To understand how we approximate the conversion, we first introduce how ordinary conversion is computed in LNS. Let  $\tilde{p}_{iq}$  and  $\tilde{p}_{ir}$  be positive integers representing quotient and remainder of the intermediate result  $\tilde{p}_i/\gamma$  in Equation 13, and let  $v_r = 2^{\tilde{p}_{ir}/\gamma}$ . Therefore,

$$\begin{aligned} 2^{\tilde{p}_i/\gamma} &= 2^{\tilde{p}_{iq}/\gamma} \times 2^{\tilde{p}_{ir}/\gamma} = 2^{\tilde{p}_{iq}} \times 2^{\tilde{p}_{ir}/\gamma} \\ &= (v_r \ll \tilde{p}_{iq}), \end{aligned} \quad (14)$$

where  $\ll$  is left bit-shifting. This transformation enables fast conversion by applying efficient bit-shifting over  $v_r$  whose value is bounded by the remainder. The different constant values of  $v_r = 2^{\tilde{p}_{ir}/\gamma}$  can be pre-computed and stored in a hardware look-up table (LUT), where the remainder  $\tilde{p}_{ir}$  is used to select the constant for  $v_r$ . The quotient  $\tilde{p}_{iq}$  then determines how far to shift the constant. Furthermore, because  $\gamma = 2^b$ , the least significant bits (LSB) of the exponent are the remainder and the most significant bits (MSB) are the quotient. As the size of the LUT grows, the computational overhead from conversion increases significantly. Typically, the LUT is required to contain  $2^b$  entries for storing all possible values of  $v_r$ , which can be a large overhead for large values of  $b$ .

A straightforward solution for reducing the size of LUT is utilizing Mitchell approximation [14]:  $v_r = 2^{\tilde{p}_{ir}/2^b} = (1 + \tilde{p}_{ir}/2^b)$ . However, if  $v_r$  is far away from zero or one, the approximation error induced by Mitchell approximation will be significant. To alleviate this error, we propose a hybrid approximation that trades off efficiency and approximation error. Specifically, we split  $\tilde{p}_{ir}$  into  $\tilde{p}_{irM}$  and  $\tilde{p}_{irL}$  to represent the MSB and LSB of the remainder, respectively. LSB values  $2^{\tilde{p}_{irL}}$  are approximated using Mitchell approximation, and MSB values  $2^{\tilde{p}_{irM}}$  are pre-computed and stored using LUT, such that:

$$\begin{aligned} v_r &= 2^{\tilde{p}_{ir}/2^b} = 2^{\tilde{p}_{irM}/2^b} \times 2^{\tilde{p}_{irL}/2^b} \\ &= (1 + \tilde{p}_{irL}/2^b) \times 2^{\tilde{p}_{irM}/2^b}, \end{aligned} \quad (15)$$

where  $\tilde{p}_{irM}$  and  $\tilde{p}_{irL}$  represent  $b_m$  MSB and  $b_l$  LSB bits of  $\tilde{p}_{ir}$ . This reduces the size of LUT to  $2^{b_m}$  entries. For efficient hardware implementation, we use  $2^{b_m}$  registers to accumulate different partial sum values and then multiply with constants from the LUT.

Table 4: Benchmarking conversion approximation for multi-base LNS. Best prediction performance for each task is highlighted.

	LUT=1	LUT=2	LUT=4	LUT=8
<b>CIFAR-10</b> (Accuracy)	92.58	92.54	92.68	<b>93.43</b>
<b>ImageNet</b> (Accuracy)	75.80	75.85	75.94	<b>76.05</b>
<b>SQuAD</b> (F-1)	87.57	87.11	<b>88.00</b>	87.82
<b>GLUE</b> (F-1)	84.89	85.48	86.93	<b>88.15</b>
<b>Energy Cost</b> (fJ / op)	12.29	14.71	17.24	19.02

## B.2 Approximation-Aware Training

We demonstrate that the approximation error induced by the conversion approximation does not affect the overall accuracy significantly. The deterministic approximators can be viewed as additional non-linear layers in networks, and thus they can be learned during training. Because this process is similar to quantization-aware training, we denote it as approximation-aware training. To verify our hypothesis, we simulate the proposed conversion approximation in LNS. The approximators are only applied to the forward propagation to allow approximation-aware training. After training, an approximated model can also be deployed for fast inference. With 8-bit and base factor  $\gamma = 8$ , we evaluate the approximation setting from  $LUT = 1$  to  $LUT = 8$ . We use the BERT-base model for SQuAD and GLUE in this experiment, and we also report the energy cost per operation for each approximation setting. As shown in Table 4, the approximated networks achieve almost no loss of accuracy while reducing energy cost by 35% maximally.

## B.3 Energy Analysis

Logarithmic datapath was coded in C++ and synthesized using a commercial HLS tool to map untimed C++ code to cycle-accurate RTL [37]. Once the RTL is generated, a standard logic synthesis flow is used to obtain gate-level netlist and estimate area. To evaluate energy consumption, standard power analysis tools use gate-level simulation results from post-synthesis netlist. We perform our analysis in a sub-16nm state-of-the-art process technology. Table 5 shows the design tools used in the evaluation.

Table 5: Experimental Setup

Design Tools	
HLS Compiler	Mentor Graphics Catapult HLS
Verilog simulator	Synopsys VCS
Logic synthesis	Synopsys Design Compiler Graphical
Power Analysis	Synopsys PT-PX

# C Experiments

## C.1 End-To-End Hardware Simulation System

We simulate multi-base LNS using a Pytorch-based neural network quantization library that implements a set of common neural network layers (e.g., convolution, fully-connected) for training and inference in both full and quantized modes [38]. The library supports integer quantization in fixed-point number system originally, and we further extend it to support logarithmic number system as well. The library also provides utilities for scaling values to the representable integer range of the specific number format.

With this library, a typical quantized layer consists of a conventional layer implemented in floating-point preceded by a weight quantizer and an input quantizer that convert the weights and inputs of

the layer to the desired quantized format. For the backward pass, after the gradients pass through the STE in each quantizer, they will also be quantized by  $Q_{\log}$ .

We run the experiments using our internal computing cluster equipped with 8 V100 Tensor Core GPUs for each instance [39].

## C.2 Datasets and Models

**ResNet Models** We use residual networks for benchmarks on image datasets [40]. We quantize all fully-connected and convolutional layers in ResNet, including both forward and backward propagation. Besides, we leave batch-norm layers at full-precision for simplicity. SGD optimizer is applied by default with its standard learning rate schedule.

**BERT Models** We perform quantization on pre-trained BERT models for language fine-tuning tasks. BERT models are the state-of-the-art language representation models which include 110M parameters in the BERT-base model and 320M parameters in the BERT-large model [41]. We quantize all GEMM operations for both models, which consist of 99% of all parameters. AdamW optimizer is applied by default.

**CIFAR-10** We use ResNet-18 to evaluate different quantization settings on CIFAR-10 dataset [42]. CIFAR-10 consists of 60,000 images in 10 different classes. The network is trained for 300 epochs, and we use a fixed learning rate decay schedule that decayed every 100 epochs. Besides, we use a tuned SGD optimizer by default, where the initial learning rate is 0.1, weight decay is 0.0001, and momentum is 0.9.

**ImageNet** The ILSVRC2012 ImageNet dataset consists of 1.2 million images belonging to 1,000 classes [43]. We use Resnet-50 as the base model, and the network is trained for 90 epochs for all settings. Similarly, we use a tuned SGD optimizer with a learning rate warmup by default, where the configuration is the same as the one in the CIFAR-10 experiment. For LNS-Madam optimizer, as mentioned in previous studies, multiplicative learning algorithms may enjoy a weight initialization different from the standard practice. Therefore, on the ImageNet benchmark, we use SGD as a warm-up for the first 10 epochs to mitigate this initialization effect.

**SQuAD** The Stanford Question Answering Dataset (SQuAD v1.1) is a collection of 100k crowd-sourced question/answer pairs [44]. We evaluate our framework on SQuAD and use the generic pre-trained BERT and BERT-large models for the fine-tuning ([41]). The maximum sequence length and document stride are also set to be 384 and 128, respectively. We fine-tune the network for 2 epochs and use a tuned AdamW optimizer by default, where its learning rate starts from 0.00003.

**GLUE** The General Language Understanding Evaluation (GLUE) benchmark is a collection of diverse natural language understanding tasks [45]. We use a pre-processing setting similar to the setting for SQuAD and fine-tunes using both BERT-base and BERT-large models for 2 epochs. AdamW optimizer is applied by default, where its initial learning rate is 0.00002.

## C.3 Base Factor Selection

As mentioned earlier, to find an appropriate combination of bitwidth  $\mathcal{B}$  and base factor  $\gamma$ , we design an efficient searching strategy that first finds a desired dynamic range and then tests different  $\gamma$  to find a minimal required quantization gap. Specifically, our strategy first decides an appropriate log dynamic range:  $(0, (2^{\mathcal{B}-1} - 1)/\gamma)$  and then varies the combinations of base factor and bitwidth within the range.

As shown in Table 6, we fix the bitwidth as 8-bit and vary the base factor  $\gamma$  to find the appropriate dynamic ranges for forward and backward quantization. According to the results, we find a dynamic range around  $(0, 32)$  that uniformly works across  $Q_W, Q_A, Q_E$ , and  $Q_G$ .

This dynamic range determines  $O(M)$  combinations of base factor and bitwidth where  $M$  is the number of bitwidth settings considered. In this case, we only need to verify  $O(M)$  combinations to find the best one, which is more computationally efficient. For example, the first two columns in Table 7 show some combinations we considered varying from 8-bit to 5-bit. After evaluating our framework over a wide range of learning tasks, we identify the best setting that uniformly works well, where 5-bit with  $\gamma = 1$  is applied for activation gradients  $Q_E$ . Given the fixed 8-bit setting for

$Q_W, Q_A$ , and  $Q_G$  as we mentioned in Section 5.1, we apply  $\gamma = 8$  for them to match their dynamic range requirements.

We do note that LNS could train networks with a bitwidth setting lower than 8-bit. As shown in Table 7, LNS with a unified 5-bit setting could achieve almost no loss of accuracy. However, we do find the accuracy degradation for BERT fine-tuning tasks under the unified 5-bit setting. This suggests a more fine-grained scaling technique or a higher bitwidth setting is needed for the forward quantization under LNS. Because we focus on quantizing gradients in this work, we rule out the effects of the ultra-low forward quantization by fixing their bitwidth as 8-bit uniformly. From the perspective of energy efficiency, the conversion approximation can be applied for the forward quantization with  $\gamma = 8$  seamlessly to improve the efficiency for both training and inference.

Table 6: Base Factor Selection. We fix the bitwidth as 8-bit and vary the base factor  $\gamma$  from 1 to 32. Quant Forward or Quant Backward denotes the settings where either forward propagation or backward propagation is quantized while leaving the rest of computation in full-precision. The results of test accuracy (%) are shown below.

Base Factor $\gamma$	Log Dynamic Range	Quant Forward	Quant Backward
1	(0,127)	NaN	NaN
2	(0,63.5)	75.81	75.79
4	(0,31.8)	75.96	76.07
8	(0,15.9)	75.88	76.23
16	(0,7.9)	76.32	63.67
32	(0,4.0)	68.15	20.71

Table 7: Unified bitwidth setting for  $B_W, B_A$  and  $B_E$  on ImageNet. The combinations of  $B$  and  $\gamma$  follow a fixed log dynamic range selected in Table 6. We compare LNS and INT with different scaling techniques. The results of test accuracy (%) are shown below.

Bitwidth $B$	Base Factor $\gamma$	LNS (per-vector)	LNS (per-layer)	INT (per-vector)	INT (per-layer)
<b>8-bit</b>	8	76.05	75.96	75.69	75.30
<b>7-bit</b>	4	76.10	76.24	75.01	72.65
<b>6-bit</b>	2	75.83	75.83	74.25	65.20
<b>5-bit</b>	1	75.55	75.13	73.59	43.97

#### C.4 Scaling Techniques

To study the effects of different scaling techniques, we compare per-vector scaling with a simple per-layer scaling over a range of bitwidth settings. The results are shown in Table 7. Because LNS naturally enjoys a higher dynamic range, LNS performs much better than INT when the bitwidth is strongly limited. When applying per-vector scaling with a default vector size of 16, every setting gets improved compared to its per-layer counterpart.