

REVISITING BFloat16 TRAINING

Anonymous authors

Paper under double-blind review

ABSTRACT

State-of-the-art generic low-precision training algorithms use a mix of 16-bit and 32-bit precision, creating the folklore that 16-bit precision alone is not enough to maximize model accuracy. As a result, deep learning accelerators are forced to support both 16-bit and 32-bit compute units which is more costly than only using 16-bit units for hardware design. We ask *can we do pure 16-bit training which requires only 16-bit compute units, while still matching the model accuracy attained by 32-bit training*. Towards this end, we study pure 16-bit training algorithms on the widely adopted BFloat16 compute unit. While these units conventionally use *nearest rounding* to cast output to 16-bit precision, we show that nearest rounding for model weight updates can often cancel small updates, which degrades the convergence and model accuracy. Motivated by this, we identify two simple existing techniques, stochastic rounding and Kahan summation, to remedy the model accuracy degradation in pure 16-bit training. We empirically show that these two techniques can enable up to 7% absolute validation accuracy gain in pure 16-bit training. This leads to 0.1% lower to 0.2% higher matching validation accuracy compared to 32-bit precision training across seven deep learning applications.

1 INTRODUCTION

Recently there has been an explosion in the compute resources required for training deep learning models (Shoeybi et al., 2019; Rajbhandari et al., 2019; Real et al., 2019). As a result, there has been broad interest in leveraging low-precision (< 32 -bit) training algorithms to reduce the required compute resources (De Sa et al., 2017; Hubara et al., 2017; Gupta et al., 2015). Among these algorithms, mixed-precision training—in which model activations and gradients are stored using a 16-bit floating point format while model weights and optimizer states use 32-bit precision—is commonly used when training generic deep learning models (Micikevicius et al., 2017; Kalamkar et al., 2019). While there is a wide body of literature showing that low-precision training can minimally impact accuracy on specific models (Wang et al., 2018b; De Sa et al., 2015; Zhang et al., 2017), conventional wisdom suggests that at least some 32-bit computation is required as a fail-safe in generic deep learning training. As such, new accelerator architectures for deep learning are forced to support both 32-bit and 16-bit compute units. This is much more costly in terms of area, power, and speed when compared to hardware with only 16-bit compute units (Horowitz, 2014; Galal et al., 2013).

In this paper we question if 32-bit compute units are truly needed for new deep learning hardware accelerators. Namely, can we match the model accuracy of 32-bit-precision algorithms while leveraging *only* 16-bit compute units? To answer this question, we study *pure 16-bit training* algorithms, ones which use only 16-bit compute units and which store activations, gradients, model weights, and optimizer states all in a 16-bit precision. Specifically, we focus on training with the BFloat16 compute unit which is widely adopted in modern deep learning accelerators (Jouppi et al., 2017; Burgess et al., 2019). Such units take 16-bit inputs, perform computation, and then round the results to a 16-bit output. BFloat16 compute units can provide $3\times$ higher power efficiency, $1.5\times$ lower latency, and $1.5\times$ less chip area than 32-bit units (Horowitz, 2014; Galal et al., 2013). In addition, pure 16-bit training algorithms can reduce the memory footprint and bandwidth consumption of model weights and optimizers by $2\times$ compared to mixed precision or 32-bit precision training, especially for large models with billions of weights (Shoeybi et al., 2019; Rajbhandari et al., 2019). Developing reliable pure 16-bit training algorithms will enable hardware designers to realize these advantages.

The simplest approach to pure 16-bit training is to take a 32-bit baseline and “make it low-precision” by replacing all the 32-bit numbers with 16-bit numbers and replacing each 32-bit floating-point op-

eration with its 16-bit analog, using nearest rounding¹ to quantize as necessary: we call this approach the *standard* algorithm. *Unfortunately, we show empirically that standard pure 16-bit training does not match 32-bit training on model accuracy across deep learning models.* For example, the standard pure 16-bit training algorithm one would run on conventional hardware attains 16% and 7% lower training and validation accuracies than a 32-bit baseline. Motivated by this observation, we start by analyzing what factors limit the model accuracy of this standard pure 16-bit algorithm.

The goal of our analysis is to inspire a clean, minimal set of simple techniques that allow pure 16-bit training to attain strong model accuracy for state-of-the-art deep learning models across application domains. Towards this end, we derive insights from a simple least-squares regression model in Section 3. Using this least-squares regression model, we reveal that nearest rounding of compute unit outputs causes significant convergence degradation and consequent model accuracy loss. More concretely, we show a key theoretical insight hidden in existing work: when running stochastic gradient descent on a least-squares regression model, nearest rounding while updating model weights ignores small updates. This phenomenon significantly degrades the convergence of stochastic gradient descent when model updates become small relative to model weights, which is also what we observe when training deep learning models. In comparison, nearest rounding in the forward and backward pass of backpropagation has a negligible impact on convergence. These insights lead us to consider two simple existing techniques to achieve high-accuracy pure 16-bit training. First, we can use *stochastic rounding* instead of nearest rounding for the model weight updates. Here, the rounded weights become an unbiased estimate of the precise weights without rounding: thus, regardless of the magnitude of updates, the expectation of rounded weights converges at the same speed as the precise weights. Second, we can use the well-known Kahan summation algorithm (Kahan, 1965) to accumulate model updates while still keeping nearest rounding for all operations. This method tracks and compensates weight rounding errors across iterations with auxiliary 16-bit values, which avoids catastrophic cancellation of many small model weight updates.

Empirically, in Section 4 we first validate that, as suggested by our theory, nearest rounding for model weight updates is the sole bottleneck for convergence and model accuracy on several deep learning models. We then demonstrate that pure 16-bit training using stochastic rounding or Kahan summation on model weight updates can match 32-bit training in model accuracy across a wide range of applications (He et al., 2016; Amodei et al., 2016; Devlin et al., 2018; Naumov et al., 2019). To validate that nearest rounding for model weight updates is the cause of the accuracy degradation, we show that if we store model weights in 32-bit precision without rounding during weight updates, and we keep using 16-bits and nearest rounding for all other operations, then the attained model accuracy matches full 32-bit precision training. Next, we demonstrate that 16-bit training with stochastic rounding for weight updates attains model accuracy matching 32-bit training for five out of seven applications in our study. Note that while it works most of the time, this is not a silver bullet, as using stochastic rounding alone could not fully match 32-bit training on all models. To address this, we show that Kahan summation for model weight updates closes remaining gaps on all the models we consider; this Kahan summation comes with a trade off, as it requires $2\times$ weight memory, but achieves up to 0.2% higher validation accuracy than stochastic rounding. *Our results suggest that deep learning accelerators using only 16-bit compute units are feasible if stochastic rounding and Kahan summation are supported respectively by the hardware and the software stack.*

2 PRELIMINARY

In this section we establish the background and notation for our study and present the preliminary observations that motivate our work. We focus on the case of stochastic gradient descent (SGD), which is the primary workhorse used to train deep learning models. SGD computes gradients from a subset of training samples, and uses them to update the model weights so as to decrease the loss in expectation. In the classic supervised learning setting, let (\mathbf{X}, \mathbf{y}) be a dataset where $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n] \in \mathbb{R}^{n \times d}$ and $\mathbf{y} = (y_1, y_2, \dots, y_n) \in \mathbb{R}^n$. On this dataset, we use stochastic gradient descent to optimize a loss function $f(\mathbf{w}) = 1/n \sum_{i=1}^n f_i(\mathbf{w}, \mathbf{x}_i, y_i)$ defined by the model. At the t -th iteration, we sample an index subset $\sigma(t) \subset \{1, 2, \dots, n\}$ and compute a sample gradient $\nabla f_{\sigma(t)}(\mathbf{w}_t)$ as an unbiased estimate of the full gradient $\nabla f(\mathbf{w})$. In deep learning, model training can be described as a compute graph where the compute graph operators such as addition and ma-

¹This nearest rounding is the standard rounding mode for compute unit output commonly supported across hardware platforms (Intel, 2018; Nvidia, 2020).

trix multiplication are the nodes. For example, the model weight update operator is defined as the subtraction in the operation $\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \nabla f_{\sigma(t)}(\mathbf{w}_t)$ which updates the model weight \mathbf{w} .

16-bit Compute Units On modern hardware, numerical operators in the compute graph are supported by fused multiply-and-accumulation (FMAC) compute units. These units work by computing $a \leftarrow a + (x \times y)$, where x and y are input floating point numbers, and a is an accumulator that is part of the FMAC unit. Importantly, for a 16-bit FMAC unit, the accumulator a has higher-than-16-bit precision. This higher precision accumulator is inexpensive compared to the multiply in terms of chip area and energy consumption in FMAC units, but is critical to the numerical accuracy of operations such as matrix multiplication and convolution. Thus 16-bit FMAC units with higher precision accumulator is standard for modern hardware including TPUs and GPUs (Chao & Saeta, 2019; Markidis et al., 2018; Stephens, 2019); this will likely continue to be standard in emerging accelerators. Because of the higher precision accumulator, the result in the accumulator then needs to be rounded to 16-bits before it is output from the FMAC unit (e.g. before writing to memory). FMAC units use the same hardware implementation to support all operators from simple additions to computationally intensive convolutions, so this output-rounding step happens for all the operators in a compute graph.

Nearest Rounding FMAC output rounding is widely implemented with *nearest rounding* as the standard mode, due to its efficient support across hardware platforms. “Nearest rounding” means rounding a higher precision floating point number to the closest low-precision representable value. *Given that the add step already uses accurate higher precision accumulators, this nearest rounding is the primary source of numerical errors for training using 16-bit FMAC units.* In this context of 16 bit FMAC units and nearest rounding, we discuss the following training-precision approaches.

- In *32-bit precision training*, all the compute graph operators read and write memory using a 32-bit precision. These operators require 32-bit compute units, which constrains the compute and memory efficiency.
- In *mixed precision training*, model weights are stored in 32-bit precision while activations and gradients use 16-bit precision. Thus, new accelerators customized to maximize efficiency for mixed precision training still require 32-bit compute units for operators involving 32-bit weights as the input; this has lower efficiency in power, speed and chip area than only using 16-bit units.
- In *pure 16-bit training*, activation, gradients and model weights are all stored in 16-bit precision. All operators use pure 16-bit input and write out 16-bit output after rounding. Thus, aside from just saving memory, pure 16-bit training can eliminate the requirement for 32-bit compute units. This opens the possibility for highly-efficient accelerators using only 16-bit compute units. In spite of this favorable efficiency, we now show that it can be surprisingly challenging for standard pure 16-bit training to match 32-bit training on model accuracy.

Motivating Observations Although recent works have shown that certain models are robust to numerical error during training (Wang et al., 2018b; De Sa et al., 2015; Zhang et al., 2017), surprisingly, we observe that it is challenging for pure 16-bit training to attain the same accuracy as 32-bit precision training on several state-of-the-art deep learning models. To demonstrate this, we compare 32-bit precision training and standard pure 16-bit training (using nearest rounding for all its operator outputs). For example, Figure 1 illustrates that for a BERT-Base model for natural language inference, the standard pure 16-bit training algorithm demonstrates 16% and 7% lower training and validation accuracies than 32-bit precision training². This gap suggests that nearest rounding is the primary source of numerical error in pure 16-bit algorithms, significantly degrading the convergence and model accuracy. To alleviate this problem, in Section 3, we study how nearest rounding impacts convergence, and we expose insights which lead to simple techniques to improve the model accuracy in pure 16-bit training.




Figure 1: Standard pure 16-bit training shows lower training accuracy compared to 32-bit training on a BERT model.

²We note that 32-bit and standard 16-bit training in Figure 1 start from the same initialization with the same model accuracy. The gap at the beginning of the two curves is due to curve smoothing for visualization clarity. We refer to Appendix D.1 for the unsmoothed curves which start from the same model accuracy.

3 PRECISE MODEL WEIGHT UPDATES ARE ALL YOU NEED

To understand how to improve the model accuracy attained by pure 16-bit training, in this section we analyze the impact of nearest rounding, the primary source of numerical error in the standard 16-bit training algorithm. In Section 3.1, we show that when model updates are small relative to the model weights, nearest rounding for model weight updates ignores small updates and consequently impedes convergence of stochastic gradient descent. In contrast, we show that nearest rounding in the forward and backward compute can have a much weaker impact on the convergence throughout training. These insights emphasize the importance of more precise model weight updates for pure 16-bit training. To achieve such precise updates using pure 16-bit training, in Section 3.2 we consider using two simple existing numerical techniques, stochastic rounding and Kahan summation, for model weight updates. Although we use a least-squares regression model for its simplicity in exposing insights, we will show in Section 4 that our insights transfer empirically to representative deep learning models.

3.1 UNDERSTANDING THE IMPACT OF NEAREST ROUNDING

In this section, we use a simple least-squares regression model $\frac{1}{2n} \sum_{i=1}^n \|\mathbf{x}_i^T \mathbf{w} - y_i\|^2$ with batch size 1 as a proxy to expose the impact of numerical errors due to nearest rounding. First, we discuss the impact of nearest rounding for model weight updates. Then, we show that nearest rounding for forward and backward compute has comparatively much weaker influence on accuracy. [More concretely, we focus on underdetermined least-squares regression problems with high dimensional model weights and input features; this is intended to reflect the overparameterized setting in modern deep learning models \(Li et al., 2018; Jacot et al., 2018\).](#) In this setting, the model has the capacity to perfectly fit the dataset with $y_i = \mathbf{x}_i^T \mathbf{w}^*$ where \mathbf{w}^* is the minimizer of $f(\mathbf{w})$ (A1). To ensure convergence with a bounded gradient variance, we assume bounded input data $\|\mathbf{x}_i\|^2 \leq L$ (A2). We let ϵ denote the machine epsilon of our floating point format, such that if u and v are adjacent representable numbers in our floating point format, $\epsilon|u| \leq |u - v| \leq 2\epsilon|u|$. Under this standard assumption for numerical analysis on floating point numbers (Stoer & Bulirsch, 2013), nearest rounding $\mathbf{Q}(\cdot)$ will have a bounded error of $|\mathbf{Q}(u) - u| \leq \epsilon|u|$ for any u in range. To simplify the presentation, we ignore overflow and underflow in our analysis here, and disregard factors of ϵ^2 (as is standard in analyses of floating point error).

Nearest Rounding for Model Weight Updates When stochastic gradient descent updates model weights with nearest rounding, the model weights evolve as $\mathbf{w}_{t+1} = \mathbf{Q}(\mathbf{w}_t - \alpha \nabla f_{\sigma(t)}(\mathbf{w}_t))$. For a weight dimension i , if the model update $[\alpha \nabla f_{\sigma(t)}(\mathbf{w}_t)]_i$ is smaller than half of the difference between $[\mathbf{w}_t]_i$ and its neighboring representable value in a certain precision format, nearest rounding cancels this model update. This often emerges in the mid-to-late training stage when the magnitude of gradient becomes small or learning rate decays small. Formally, we show that nearest rounding can cancel updates across all dimensions for a least-squares regression model when approaching the optimal weights in Theorem 1; we defer the proof to Appendix A.

Theorem 1. *Consider running one step of SGD on a least-squares regression model under assumptions A1 and A2. The model weight update will be entirely canceled by nearest rounding if*

$$\|\mathbf{w} - \mathbf{w}^*\| \leq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|, \quad (1)$$

where w_j^* denotes the j -th dimension of the optimal solution $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w})$. Additionally, if we run multiple steps of SGD using nearest-rounded weight updates and fixed learning rate $\alpha \leq 1/L$, then the distance of the weights \mathbf{w}_t at any timestep t from the optimum is bounded by

$$\|\mathbf{w}_t - \mathbf{w}^*\| \geq \min \left(\frac{\epsilon(1 - \alpha L)}{\alpha L + \epsilon} \cdot \min_j |w_j^*|, \|\mathbf{w}_0 - \mathbf{w}^*\| \right).$$

Theorem 1 reveals that for the least-squares regression model, nearest rounding cancels the entire model updates when the distance towards the optimal solution \mathbf{w}^* is small relative to the magnitude of \mathbf{w}^* . Thus in the late stage of training, the model weights can halt in a region with radius $\frac{\epsilon}{\alpha L + \epsilon} \min_j |w_j^*|$ around \mathbf{w}^* . Our lower bound shows that this region limits the convergence of SGD with nearest rounding on model weight updates for least-squares regression models. In this bound, the key property is the dependency on the step size: as the step size becomes small, this error lower bound becomes *worse*, which is the opposite of the usual effect of diminishing the step size in SGD.

Given that \mathbf{w}^* can be arbitrarily far from the zero vector, our lower bound also reveals a substantial convergence degradation for stochastic gradient descent using floating point numbers. This specific degradation, which depends on the magnitude of \mathbf{w}^* , does not emerge in training with fixed point numbers which is a setting widely used for analyzing the impact of rounding (Hou et al., 2018; Li et al., 2017). Because the lower bound on $\|\mathbf{w}_0 - \mathbf{w}^*\|$ is also in the order of $\mathcal{O}(\epsilon)$, this bound is worse for lower precision formats with a large ϵ value. In Section 4, we will empirically show that these insights from the least-squares regression model can also generalize to deep learning models, which explains the convergence and model accuracy degradation due to the small updates cancellation in model weight updates.

Nearest Rounding for Forward and Backward Compute In contrast to the significant convergence degradation imposed by nearest rounding on model weight updates, we show that the nearest rounding in the gradient computation (in the forward and backward passes of backpropagation) can impact convergence minimally. To show this, we consider stochastic gradient descent with nearest rounding only for compute operations which generate activations and gradients. Here to compute the gradient for least-squares regression models, the linear layer passes the rounded activation $a_i = \mathbf{Q}(\mathbf{x}_i^T \mathbf{w} - y_i)$ to the loss layer. (We see no quantization error within the dot product $\mathbf{x}_i^T \mathbf{w}$ itself, as all accumulation here is done with the higher-precision accumulator of the FMAC.) In the backward stage, the loss layer feeds the rounded activation gradients $g_{a,i} = \mathbf{Q}(a_i)$ back to the linear layer. The weight gradient is then computed as $\nabla_{\mathbf{Q}} f_i(\mathbf{w}) := \mathbf{Q}(g_{a,i} \mathbf{x}_i) = \mathbf{Q}(\mathbf{Q}(\mathbf{Q}(\mathbf{x}_i^T \mathbf{w} - y_i)) \mathbf{x}_i)$. To isolate the impact of nearest rounding for activations and gradients, we do not round model weights in this setting. Formally, we show that SGD with activation and gradient rounding allows for an upper bound on $\|\mathbf{w}_t - \mathbf{w}^*\|$ which can be substantially smaller than the lower bound in Theorem 1.

Theorem 2. *Consider running multiple steps of SGD on a least-squares regression model under assumptions A1 and A2, using nearest rounding for only forward and backward compute, but exact arithmetic for model weight updates. Then if the step size is small enough that $\alpha \leq 1/L$, the distance of the weights \mathbf{w}_t at any timestep t from the optimum will be bounded by*

$$\mathbf{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2] \leq \exp\left(-\alpha \mu t \left(1 - \frac{4\epsilon L}{\mu}\right)\right) \cdot \|\mathbf{w}_0 - \mathbf{w}^*\|^2,$$

where μ is the smallest eigenvalue of the data covariance matrix $\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T$.

As t can be made arbitrarily large, this bound guarantees us substantially more accurate solutions than the lower bound attained by using nearest rounding only for model weight updates in Theorem 1. This shows that rounding for the weight updates is the primary source of error, as even with all other operations quantized, the algorithm is guaranteed to converge closer to the optimum than is even possible with just the weight updates rounded with nearest rounding. Note that the bound in Theorem 2 is able to guarantee arbitrarily accurate solutions because we ignore underflow here. In practice, precision would eventually be limited by underflow even in the setting of Theorem 2; however, the underflow threshold for BFloat16 is small enough that this represents a level of error that deep learning applications are generally able to tolerate. We refer to Appendix A for the proof.

Theory Validation To validate our insights, we compare the impact of nearest rounding for model weight updates against that of nearest rounding in forward and backward compute on a synthetic 10-dimensional least-squares regression problem. Specifically, the input data are sampled from a zero-mean unit-variance normal distribution while the model weight is generated from a uniform distribution in the range of $[0, 100)$. We perturb the label with a zero-mean normal distribution with standard deviation 0.5. As shown in Figure 2, when using a learning rate 0.01 and 16-bit nearest rounding for model weight updates, we observe that the training loss saturates at a magnitudes higher level than stochastic gradient descent without rounding because of updates cancellation. Meanwhile, when using nearest rounding only for forward and backward compute, the loss saturates at a level close to that attained by training without rounding. These observations align with our insights on the relative impact of nearest rounding for model weight updates and for forward and backward compute.




Figure 2: **Theory validation.** On a least square regression model, (smoothed) training losses with 16-bit nearest rounding for weight updates saturate at a higher level than 32-bit training. With only using nearest rounding for forward and backward compute, the losses saturate much closer to 32-bit training.

3.2 HIGH-ACCURACY PURE 16-BIT TRAINING

In Section 3.1, we showed that nearest rounding for model weight updates is the bottleneck for convergence in the standard pure 16-bit training algorithm; this is because it cancels small model updates which degrades the model weight precision. This motivates us to consider two existing techniques, *stochastic rounding* and *Kahan summation* (Kahan, 1965) for improving weight updates. These techniques have been reliably applied in different numerical domains (Hopkins et al., 2020; Antonana et al., 2017) and can hypothetically enable high-accuracy pure 16-bit training. [We present details on how to integrate these techniques into SGD and AdamW optimizers with pure 16-bit model weights and optimizer states such as momentum in Appendix B.](#)

Stochastic Rounding Stochastic rounding for floating point numbers has been used in training certain model components (Zhang et al., 2018) and can potentially improve the model accuracy of pure 16-bit training for general models. Specifically, let \mathbb{S} be the set of all values representable by a limited precision format: the upper and lower neighboring values for $a \in \mathbb{R}$ are $a_u = \min_{x \geq a, x \in \mathbb{S}} x$ and $a_l = \max_{x \leq a, x \in \mathbb{S}} x$. Stochastic rounding randomly rounds a up to a_u with probability $(a - a_l)/(a_u - a_l)$ and otherwise rounds down to a_l . We consider pure 16-bit training using stochastic rounding only for the subtraction output in the model update $\mathbf{w}_t - \alpha \nabla f_{\sigma(i)}(\mathbf{w}_t)$. We keep nearest rounding for all the other compute operations. Here, the rounded model weight is an unbiased estimate of the precise value, so it will still make progress in expectation; this prevents the halting effect from nearest rounding on model weight updates. [We note that in modern hardware, stochastic rounding can be implemented without any expensive multiplication or division arithmetic \(De Sa et al., 2017\).](#) Thus using stochastic rounding for model weight updates adds minimal overhead when training modern deep learning models; we discuss explicitly how to achieve this in Appendix B.1.

Kahan Summation The Kahan summation algorithm uses an auxiliary variable to track numerical errors and to compensate the accumulation results. In the context of pure 16-bit training, we use a 16-bit auxiliary variable $\mathbf{c}_t \in \mathbb{R}^d$ to track the error in model weights. To ensure pure 16-bit data paths, we keep nearest rounding for all operators in compute graphs, including those during Kahan accumulation in Algorithm 1. At iteration t , we first compensate the current model

Algorithm 1 SGD updates with Kahan summation

- 1: Auxiliary value $\mathbf{c}_0 \leftarrow 0$ at initialization
 - 2: **Input:** Model updates $-\alpha \nabla f_{\sigma(i)}(\mathbf{w}_t)$ at iter. t
 - 3: $\mathbf{u}_{t+1} \leftarrow -\alpha \nabla f_{\sigma(i)}(\mathbf{w}_t)$
 - 4: $\mathbf{y}_{t+1} \leftarrow \mathbf{u}_{t+1} - \mathbf{c}_t$ \triangleright Compensate updates
 - 5: $\mathbf{s}_{t+1} \leftarrow \mathbf{w}_t + \mathbf{y}_{t+1}$ \triangleright Accumulate updates
 - 6: $\mathbf{c}_{t+1} \leftarrow (\mathbf{s}_{t+1} - \mathbf{w}_t) - \mathbf{y}_{t+1}$ \triangleright Measure errors
 - 7: $\mathbf{w}_{t+1} \leftarrow \mathbf{s}_{t+1}$
 - 8: **Return:** \mathbf{w}_{t+1}
-

update \mathbf{u}_{t+1} by subtracting the previous error \mathbf{c}_t . We then compute the new model weights by adding the compensated updates \mathbf{y}_{t+1} to the current weights \mathbf{w}_t . We reversely subtract previous model weights \mathbf{w}_t and the compensated updates \mathbf{y}_{t+1} to acquire the new numerical error \mathbf{c}_{t+1} in the updated weights \mathbf{w}_{t+1} . For small updates \mathbf{u}_t which would cause no change in the weights after nearest rounding, this reverse subtraction records the canceled updates in the error \mathbf{c}_{t+1} . Across iterations, small updates can be accumulated in \mathbf{c}_{t+1} until \mathbf{c}_{t+1} grow large enough to affect the model weights; this allows convergence to continue when it would otherwise halt due to nearest-rounding effects. [In spite of the additional auxiliary value, pure 16-bit training with Kahan summation for model weight updates can still have advantages in terms of throughput and memory consumption compared to 32-bit and mixed precision training; we refer to Appendix B.2 for details.](#)

4 EXPERIMENTS IN DEEP LEARNING

Our theory in Section 3 reveals that nearest rounding on model weight updates is the primary source of numerical error during training. This motivates us to suggest using stochastic rounding and Kahan summation in pure 16-bit training for improved model accuracy. To first validate our theory, in this section we start by demonstrating that by ablating nearest rounding on model weight updates from the standard 16-bit training algorithm, the model accuracy gap compared to 32-bit precision training can be closed on deep learning models. Next, we show empirically that with either stochastic rounding or Kahan summation on model weight updates, pure 16-bit training can match the accuracy of 32-bit precision training across representative deep learning applications.

Experiment Setup To validate the accuracy bottleneck, we consider three representative models: ResNet-18 (He et al., 2016) on the CIFAR10 image classification (Krizhevsky et al., 2009), BERT-

Table 1: **Model accuracy bottleneck for the standard pure 16-bit training algorithm.** This algorithm shows validation accuracy gap compared to 32-bit training. In an ablation of this algorithm, we use 32-bit model weights and turn off nearest rounding only on model weight updates. This eliminates the gap, suggesting that nearest rounding on model weight updates is the accuracy bottleneck.

Model	Dateset (Metric)	32-bit	Standard 16-bit	Standard 16-bit & 32-bit weights
ResNet-18	CIFAR10 (Acc%)	95.45 ± 0.07	94.23 ± 0.12	95.40 ± 0.05
DLRM	Kaggle (AUC%)	80.27 ± 0.01	78.49 ± 0.08	80.26 ± 0.01
BERT-Base	MNLI (Acc%)	84.26 ± 0.08	77.53 ± 0.07	84.34 ± 0.04




Figure 3: **Training accuracy gap imposed by the standard pure 16-bit training.** The standard algorithm fails to match the training accuracy of 32-bit training, especially in the middle-to-late stage. We close this accuracy gap by ablating nearest rounding for weight updates from the standard algorithm. This indicates that nearest rounding for model weight update is the accuracy bottleneck.

Base (Devlin et al., 2018) on the MNLI natural language inference (Wang et al., 2018a), and DLRM model (Naumov et al., 2019) on the Kaggle Advertising Challenge (CriteoLabs, 2014). To extensively evaluate pure 16-bit training with stochastic rounding and Kahan summation, we additionally consider larger datasets and more applications: ResNet-50 on the ImageNet (Deng et al., 2009), BERT-Base on the Wiki103 language model³ (Merity et al., 2016), DLRM model on the Criteo Terabyte dataset (CriteoLabs, 2018), and Deepspeech2 (Amodei et al., 2016) on the LibriSpeech datasets (Panayotov et al., 2015). As there is no publicly available accelerator with the software and hardware support necessary for our study, we simulate pure 16-bit training using the QPyTorch simulator (Zhang et al., 2019). QPyTorch models PyTorch kernels such as matrix multiplication as compute graph operators, and effectively simulates FMAC units with 32-bit accumulators⁴. For all training algorithms, we use the same hyperparameters as their original papers or code repositories. We report statistically meaningful results with averaged metrics and standard deviations across runs with 3 random seeds. We refer to Appendices C and D for experiment details and extended results.

The Model Accuracy Bottleneck To validate our insights from Section 3, we first show empirically that nearest rounding on the model weights is the primary model accuracy bottleneck on several deep learning models. To do this, we keep the model weights in 32-bit precision and turn off nearest rounding on the model weight updates while keeping nearest rounding for all other operators in the compute graph. Figure 3 shows that the standard pure 16-bit training algorithm with nearest rounding on all operators has up to 16% training accuracy gap compared to 32-bit training. Although this gap can be small in the early training phase, it grows larger in later stages. In contrast, by ablating nearest rounding on model weight updates, the standard algorithm can fully match the training accuracy attained by 32-bit training. We notice in Table 1 that this ablation can also close the 1.2% to 6.7% validation accuracy gap when comparing the standard pure 16-bit training to 32-bit training. These observations validate our insights from Section 3.1 and motivate the use of stochastic rounding and Kahan summation on model weight updates.

³We subsample 25% of the Wiki103 and 100 hours of Librispeech training set because of the training time.

⁴Following the convention in mixed precision training (Micikevicius et al., 2017), our simulator uses fused operators for computationally inexpensive activation and normalization layers.

Table 2: **Pure 16-bit training can match 32-bit training on model accuracy.** With stochastic rounding or Kahan summation for model weight updates, pure 16-bit training can attain 0.1% lower to 0.2% higher absolute value for validation accuracy metrics across applications.

Model	Dateset (Metric)	32-bit	16-bit Stochastic	16-bit Kahan	Standard 16-bit
ResNet-18	CIFAR10 (Acc%)	95.45 \pm 0.07	95.33 \pm 0.08	95.36 \pm 0.07	94.23 \pm 0.12
ResNet-50	ImageNet (Acc%)	75.70 \pm 0.05	75.45 \pm 0.03	75.61 \pm 0.14	67.10 \pm 0.24
DLRM	Kaggle (AUC%)	80.27 \pm 0.01	80.18 \pm 0.02	80.26 \pm 0.01	78.49 \pm 0.08
	Terabyte (AUC%)	80.32 \pm 0.00	80.25 \pm 0.00	80.32 \pm 0.00	78.79 \pm 0.02
BERT	MNLI (Acc%)	84.26 \pm 0.08	84.35 \pm 0.12	84.45 \pm 0.03	77.53 \pm 0.07
	Wiki103 (PPL)	5.50 \pm 0.50	5.84 \pm 0.53	5.45 \pm 0.51	56.88 \pm 1.77
DeepSpeech2	Librispeech (WER)	62.71 \pm 0.07	62.85 \pm 0.07	62.87 \pm 0.18	69.42 \pm 0.22




Figure 5: **Training accuracy for pure 16-bit training.** With stochastic rounding or Kahan summation enabled for model weight updates, pure 16-bit training matches 32-bit precision training in terms of training accuracy with negligible differences across the applications in our experiments.

High-accuracy Pure 16-bit Training Next, we validate empirically that enabling stochastic rounding or Kahan summation for model weight updates allows pure 16-bit training to attain matching model accuracy as 32-bit training. In Table 2, we first show that by using stochastic rounding for model weight updates, pure 16-bit training matches the validation accuracy of 32-bit training with at most 0.1% difference on the CIFAR10, Kaggle, Terabyte, MNLI and Librispeech datasets, a majority of the applications in our experiments. For applications where stochastic rounding still shows a non-negligible accuracy gap with more than 0.1% discrepancy, we show that Kahan summation for model weight updates can enable pure 16-bit training to match the model accuracy of 32-bit training algorithms. We show that Kahan summation for model weight updates can boost pure 16-bit training to higher validation accuracy than using stochastic rounding. More concretely, the Kahan summation for model weight updates shows 0.2% higher top-1 accuracy and 0.1% higher AUC respectively for ResNet-50 on ImageNet and for recommendation on Terabyte than using stochastic rounding. Consequently as shown in Table 2, by using Kahan summation for model weight updates, pure 16-bit training match the model accuracy attained by 32-bit precision training across all the applications in our experiments. This validates that stochastic rounding and Kahan summation can enable pure 16-bit training algorithms to match the model accuracy of 32-bit training.

Memory efficiency and model accuracy trade-off Additionally, we show that stochastic rounding and Kahan summation can be combined for pure 16-bit training, which exposes a memory efficiency and model accuracy trade-off for practitioners to exploit. As an example, in Figure 4 we demonstrate this trade-off by incrementally replacing stochastic rounding with Kahan summation on various model weights in the DLRM model on the Kaggle dataset. As we apply Kahan summation to more model weights, the weight memory cost increases by up to $2\times$. As this cost increases we also observe up to 0.04% improvement in AUC. This exploits a memory efficiency and model accuracy trade-off that users should consider when deciding which technique to leverage.




Figure 4: **Efficiency and accuracy trade-off.** With stochastic rounding and Kahan summation on different parts of the DLRM-Kaggle, it attains higher model accuracy at the cost of more weight memory.

5 RELATED WORK

There is a plethora of research work on low-precision training for deep learning models. On certain specific models such as convolutional or recurrent neural networks, training with fixed point or floating point precisions lower than 16-bit has been shown to be feasible when using customized techniques (Wang et al., 2018b; Zhou et al., 2016; Hubara et al., 2017; Courbariaux et al., 2014; Ott et al., 2016; Sun et al., 2019). *Instead of proposing new techniques for specific model types using lower than 16-bit precision, we focus on finding the minimal set of simple techniques required by generic mode training on future general-purpose deep learning accelerators requiring only modern 16-bit compute units. Such emerging accelerators have the potential to unlock substantially improved hardware efficiency compare to those still requiring 32-bit compute units.*

Recent work shows that in low precision training, stochastic gradient descent with stochastic rounding for model weight updates only degrades worst-case upper bounds on convergence minimally; this shows stochastic rounding is an effective technique to attain strong model accuracy in low precision training (Li et al., 2017; Hou et al., 2018). Our analysis is orthogonal and complementary to these upper bounds. Specifically, we prove a lower bound on convergence when using standard nearest rounding for model weight updates. This lower bound shows that nearest rounding for weight updates can substantially degrade the convergence even in the most optimistic case no matter how learning rates are tuned. We use this lower bound together with the previous upper bounds to inform future accelerator designers that only supporting nearest rounding is not enough to maximize model accuracy; to alleviate this problem, stochastic rounding for model weight updates is one of the minimal supports required in future training accelerators.

6 CONCLUSION

In this paper we study pure 16-bit training algorithms that require only 16-bit compute units. We show that nearest rounding on model weight updates is the primary cause of convergence and model accuracy degradation in standard pure 16-bit training. To alleviate this issue, we apply two existing techniques: stochastic rounding and Kahan summation. With these techniques, we demonstrate that pure 16-bit training can match the model accuracy of 32-bit precision training across many deep learning models. Our study suggests that it is feasible to design high-accuracy deep learning accelerators using only 16-bit compute units if stochastic rounding and Kahan algorithm are supported.

REFERENCES

- Dario Amodei, Sundaram Ananthanarayanan, Rishita Anubhai, Jingliang Bai, Eric Battenberg, Carl Case, Jared Casper, Bryan Catanzaro, Qiang Cheng, Guoliang Chen, et al. Deep speech 2: End-to-end speech recognition in english and mandarin. In *International conference on machine learning*, pp. 173–182, 2016.
- Mikel Antonana, Joseba Makazaga, and Ander Murua. Reducing and monitoring round-off error propagation for symplectic implicit runge-kutta schemes. *Numerical Algorithms*, 76(4):861–880, 2017.
- Neil Burgess, Jelena Milanovic, Nigel Stephens, Konstantinos Monachopoulos, and David Mansell. Bfloat16 processing for neural networks. In *2019 IEEE 26th Symposium on Computer Arithmetic (ARITH)*, pp. 88–91. IEEE, 2019.
- C Chao and B Saeta. Cloud tpu: Codesigning architecture and infrastructure. In *Hot Chips*, volume 31, 2019.
- Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Training deep neural networks with low precision multiplications. *arXiv preprint arXiv:1412.7024*, 2014.
- CriteoLabs. Display advertising challenge, 2014. URL <https://labs.criteo.com/2014/02/kaggle-display-advertising-challenge-dataset/>.
- CriteoLabs. Criteo 1tb click logs dataset, 2018. URL <https://ailab.criteo.com/criteo-1tb-click-logs-dataset/>.

- Christopher De Sa, Matthew Feldman, Christopher Ré, and Kunle Olukotun. Understanding and optimizing asynchronous low-precision stochastic gradient descent. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 561–574, 2017.
- Christopher M De Sa, Ce Zhang, Kunle Olukotun, and Christopher Ré. Taming the wild: A unified analysis of hogwild-style algorithms. In *Advances in neural information processing systems*, pp. 2674–2682, 2015.
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.
- 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.
- Sameh Galal, Ofer Shacham, John S Brunhaver II, Jing Pu, Artem Vassiliev, and Mark Horowitz. Fpu generator for design space exploration. In *2013 IEEE 21st Symposium on Computer Arithmetic*, pp. 25–34. IEEE, 2013.
- Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. Deep learning with limited numerical precision. In *International Conference on Machine Learning*, pp. 1737–1746, 2015.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Michael Hopkins, Mantas Mikaitis, Dave R Lester, and Steve Furber. Stochastic rounding and reduced-precision fixed-point arithmetic for solving neural ordinary differential equations. *Philosophical Transactions of the Royal Society A*, 378(2166):20190052, 2020.
- Mark Horowitz. 1.1 computing’s energy problem (and what we can do about it). In *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp. 10–14. IEEE, 2014.
- Lu Hou, Ruiliang Zhang, and James T Kwok. Analysis of quantized models. In *International Conference on Learning Representations*, 2018.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- Intel. Bfloat16-hardware numerics definition, 2018. URL <https://software.intel.com/content/www/us/en/develop/download/bfloat16-hardware-numerics-definition.html>.
- Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pp. 8571–8580, 2018.
- Norman P Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pp. 1–12, 2017.
- W Kahan. Further remarks on reducing truncation errors, *commun. Assoc. Comput. Mach.*, 8:40, 1965.
- Dhiraj Kalamkar, Dheevatsa Mudigere, Naveen Mellempudi, Dipankar Das, Kunal Banerjee, Sasikanth Avancha, Dharna Teja Vooturi, Nataraj Jammalamadaka, Jianyu Huang, Hector Yuen, et al. A study of bfloat16 for deep learning training. *arXiv preprint arXiv:1905.12322*, 2019.
- Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. The cifar-10 dataset, 2009. URL <https://www.cs.toronto.edu/~kriz/cifar.html>.

- Hao Li, Soham De, Zheng Xu, Christoph Studer, Hanan Samet, and Tom Goldstein. Training quantized nets: A deeper understanding. In *Advances in Neural Information Processing Systems*, pp. 5811–5821, 2017.
- Yuanzhi Li, Tengyu Ma, and Hongyang Zhang. Algorithmic regularization in over-parameterized matrix sensing and neural networks with quadratic activations. In *Conference On Learning Theory*, pp. 2–47. PMLR, 2018.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Stefano Markidis, Steven Wei Der Chien, Erwin Laure, Ivy Bo Peng, and Jeffrey S Vetter. Nvidia tensor core programmability, performance & precision. In *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pp. 522–531. IEEE, 2018.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843*, 2016.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. *arXiv preprint arXiv:1710.03740*, 2017.
- Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Malleevich, Ilia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. Deep learning recommendation model for personalization and recommendation systems. *CoRR*, abs/1906.00091, 2019. URL <https://arxiv.org/abs/1906.00091>.
- Nvidia. Floating point and ieee 754 compliance for nvidia gpus, 2020. URL <https://docs.nvidia.com/cuda/floating-point/index.html>.
- Joachim Ott, Zhouhan Lin, Ying Zhang, Shih-Chii Liu, and Yoshua Bengio. Recurrent neural networks with limited numerical precision. *arXiv preprint arXiv:1608.06902*, 2016.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimization towards training a trillion parameter models. *arXiv preprint arXiv:1910.02054*, 2019.
- Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pp. 4780–4789, 2019.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using gpu model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- N Stephens. Bfloat16 processing for neural networks on armv8-a. See https://community.arm.com/developer/ip-products/processors/b/ml-ip-blog/posts/bfloat16-processing-for-neural-networks-on-armv8_2d00_a (accessed 14 October 2019), 2019.
- Josef Stoer and Roland Bulirsch. *Introduction to numerical analysis*, volume 12. Springer Science & Business Media, 2013.
- 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. In *Advances in Neural Information Processing Systems*, pp. 4900–4909, 2019.

- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018a.
- Naigang Wang, Jungwook Choi, Daniel Brand, Chia-Yu Chen, and Kailash Gopalakrishnan. Training deep neural networks with 8-bit floating point numbers. In *Advances in neural information processing systems*, pp. 7675–7684, 2018b.
- Hantian Zhang, Jerry Li, Kaan Kara, Dan Alistarh, Ji Liu, and Ce Zhang. Zipml: Training linear models with end-to-end low precision, and a little bit of deep learning. In *International Conference on Machine Learning*, pp. 4035–4043, 2017.
- Jian Zhang, Jiyan Yang, and Hector Yuen. Training with low-precision embedding tables. In *Systems for Machine Learning Workshop at NeurIPS*, volume 2018, 2018.
- Tianyi Zhang, Zhiqiu Lin, Guandao Yang, and Christopher De Sa. QPyTorch: A low-precision arithmetic simulation framework. *arXiv preprint arXiv:1910.04540*, 2019.
- 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.

A THEORY PROOF

In this section, we present the proof for Theorem 1 and Theorem 2.

A.1 PROOF OF THEOREM 1

Proof. We first prove that when

$$\|\mathbf{w} - \mathbf{w}^*\| \leq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|,$$

the model updates across all the dimensions will be canceled in a *single stochastic gradient descent step* on the least-squares regression model. We then prove the lower bound for $\|\mathbf{w}_t - \mathbf{w}^*\|$ after using *stochastic gradient descent across many iterations*.

We assume that floating-point representation has the property that for any representable number $u \in \mathbb{R}$, all other representable numbers $v \in \mathbb{R}$ have $|Q(u) - v| \leq \epsilon|u|$. Also in our bounded input data assumption (A2), we know that the data magnitude is bounded such that $\|\mathbf{x}_i\|^2 \leq L$. Additionally we also assume the overparameterization setting where $\mathbf{x}_i \mathbf{w}^* = y_i$ (A1) with $\mathbf{w}^* = \arg \min_{\mathbf{w} \in \mathbb{R}^d} f(\mathbf{w})$.

On the condition to cancel updates for all the dimensions For the least-squares regression model, the sample gradient at model weight \mathbf{w} is

$$\nabla f_i(\mathbf{w}) = \mathbf{x}_i(\mathbf{x}_i^T \mathbf{w} - y_i) = \mathbf{x}_i(\mathbf{x}_i^T \mathbf{w} - \mathbf{x}_i^T \mathbf{w}^*) = \mathbf{x}_i \mathbf{x}_i^T (\mathbf{w} - \mathbf{w}^*).$$

for index $i \in \{1, 2, \dots, n\}$. When the model weight update gets canceled by nearest rounding Q , we have

$$Q(\mathbf{w} - \alpha \nabla f_i(\mathbf{w})) = \mathbf{w}. \quad (2)$$

To make this hold, it suffices to show that, for all indices $j \in \{1, \dots, d\}$,

$$|[\mathbf{w} - \alpha \nabla f_i(\mathbf{w})]_j - w_j| \leq \epsilon |w_j|.$$

with w_j being the j -th dimension of \mathbf{w} . This is equivalent to show that

$$\alpha \left| [\mathbf{x}_i \mathbf{x}_i^T (\mathbf{w} - \mathbf{w}^*)]_j \right| \leq \epsilon |w_j|. \quad (3)$$

In order to prove Equation (3), we first observe that by Cauchy-Schwarz inequality,

$$\begin{aligned} \alpha \left| [\mathbf{x}_i \mathbf{x}_i^T (\mathbf{w} - \mathbf{w}^*)]_j \right| &= \alpha \left| [\mathbf{x}_i]_j \right| \cdot \left| \mathbf{x}_i^T (\mathbf{w} - \mathbf{w}^*) \right| \\ &\leq \alpha \|\mathbf{x}_i\| \cdot \|\mathbf{x}_i\| \cdot \|\mathbf{w} - \mathbf{w}^*\| \\ &\leq \alpha L \cdot \|\mathbf{w} - \mathbf{w}^*\|. \end{aligned} \quad (4)$$

We also observe that

$$|w_j| = |[\mathbf{w} - \mathbf{w}^*]_j + w_j^*| \geq |w_j^*| - |[\mathbf{w} - \mathbf{w}^*]_j| \geq \min_j |w_j^*| - \|\mathbf{w} - \mathbf{w}^*\|. \quad (5)$$

with w_j^* being the j -th dimension of \mathbf{w}^* .

Equation (4) and Equation (5) shows that, to cancel weight updates, namely satisfying Equation (2), it suffices to prove that

$$\alpha L \cdot \|\mathbf{w} - \mathbf{w}^*\| \leq \epsilon \min_j |w_j^*| - \epsilon \|\mathbf{w} - \mathbf{w}^*\|. \quad (6)$$

When we have

$$\|\mathbf{w} - \mathbf{w}^*\| \leq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|,$$

Equation (6) is satisfied. This proves that the weight updates will be canceled with

$$Q(\mathbf{w} - \alpha \nabla f_i(\mathbf{w})) = \mathbf{w}.$$

Lower bound on convergence *Case I:* the model is initialized far away from the optimum \mathbf{w}^* with $\|\mathbf{w}_0 - \mathbf{w}^*\| \geq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|$.

To prove the lower bound

$$\|\mathbf{w}_t - \mathbf{w}^*\| \geq \frac{\epsilon(1 - \alpha L)}{\alpha L + \epsilon} \cdot \min_j |w_j^*|. \quad (7)$$

for any $t \in \{1, 2, \dots\}$, we discuss two possible situations.

In the first situation, if for any t , \mathbf{w}_t never steps into the quantization noise ball with center \mathbf{w}^* and radius $\frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|$. Then the bound in Equation (7) directly holds.

In the second situation, we assume model weights step into the quantization noise ball at a certain iteration $t + 1$ for the first time. Thus at iteration t we have that

$$\begin{aligned} \|\mathbf{w}_{t+1} - \mathbf{w}^*\| &= \|\mathbf{w}_t - \mathbf{w}^* - \alpha \nabla f_{\sigma(t)}(\mathbf{w}_t)\| \\ &\geq \|\mathbf{w}_t - \mathbf{w}^*\| - \alpha \|\nabla f_{\sigma(t)}(\mathbf{w}_t)\| \\ &= \|\mathbf{w}_t - \mathbf{w}^*\| - \alpha \|\mathbf{x}_i \mathbf{x}_i^T (\mathbf{w}_t - \mathbf{w}^*)\| \\ &\geq \|\mathbf{w}_t - \mathbf{w}^*\| - \alpha \|\mathbf{x}_i\|^2 \|\mathbf{w}_t - \mathbf{w}^*\| \\ &\geq (1 - \alpha L) \|\mathbf{w}_t - \mathbf{w}^*\|. \end{aligned}$$

Because \mathbf{w}_t is still outside of the quantization noise ball at iteration t , we have $\|\mathbf{w}_t - \mathbf{w}^*\| \geq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|$. Thus we have

$$\|\mathbf{w}_{t+1} - \mathbf{w}^*\| \geq \frac{\epsilon(1 - \alpha L)}{\alpha L + \epsilon} \cdot \min_j |w_j^*|. \quad (8)$$

Case II: the model is initialized in the quantization noise ball with $\|\mathbf{w}_0 - \mathbf{w}^*\| \leq \frac{\epsilon}{\alpha L + \epsilon} \cdot \min_j |w_j^*|$. Because inside the quantization noise ball, the model weights halt, we have that

$$\|\mathbf{w}_t - \mathbf{w}^*\| \geq \min \left(\frac{\epsilon(1 - \alpha L)}{\alpha L + \epsilon} \cdot \min_j |w_j^*|, \|\mathbf{w}_0 - \mathbf{w}^*\| \right). \quad (9)$$

□

A.2 PROOF OF THEOREM 2

Proof. Under the overparameterization assumption $y_i = \mathbf{x}_i^T \mathbf{w}^*$ (A1) and bounded data assumption $\|\mathbf{x}_i\|^2 \leq L$ (A2), we first observe that the error from the activation quantization is bounded by

$$\begin{aligned} |Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) - (\mathbf{x}_i^T \mathbf{w}_t - y_i)| &\leq \epsilon |\mathbf{x}_i^T \mathbf{w}_t - y_i| \\ &= \epsilon |\mathbf{x}_i^T (\mathbf{w}_t - \mathbf{w}^*)| \\ &\leq \epsilon \sqrt{L} \|\mathbf{w}_t - \mathbf{w}^*\|. \end{aligned}$$

Note that here, we are assuming no numerical errors happen inside the dot product because we suppose it is computed by a single FMAC with a higher-precision accumulator. It follows that

$$\begin{aligned} \|Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i - (\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i\| &\leq \|\mathbf{x}_i\| \cdot |Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) - (\mathbf{x}_i^T \mathbf{w}_t - y_i)| \\ &\leq \epsilon L \|\mathbf{w}_t - \mathbf{w}^*\|. \end{aligned} \quad (10)$$

The second quantization is redundant in the least-squares regression case, because the already quantized activation will be leveraged as the activation gradient. And the additional activation gradient quantization would not take any effect to introduce new numerical error. For the third quantization, we observe that its error in the j -th coordinate will be bounded by

$$|Q(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j - (Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j| \leq \epsilon |(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j|,$$

which implies that

$$\begin{aligned}
& \|Q(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i) - (Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)\| \\
&= \sqrt{\sum_{j=1}^n |Q(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j - (Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j|^2} \\
&\leq \sqrt{\sum_{j=1}^n \epsilon^2 |(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)_j|^2} \\
&= \epsilon \|Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i\| \\
&\leq \epsilon \|\mathbf{x}_i \mathbf{x}_i^T (\mathbf{w}_t - \mathbf{w}^*)\| + \mathcal{O}(\epsilon^2) \\
&\leq \epsilon L \|\mathbf{w}_t - \mathbf{w}^*\| + \mathcal{O}(\epsilon^2).
\end{aligned}$$

where the second last inequality is a consequence of Equation (10). If we disregard the ϵ^2 term, we have that

$$\|Q(Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i) - (Q(\mathbf{x}_i^T \mathbf{w}_t - y_i) \mathbf{x}_i)\| \leq 2\epsilon L \|\mathbf{w}_t - \mathbf{w}^*\|.$$

This means that we can write our gradient update step as

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \alpha \mathbf{x}_{\sigma(t)} \mathbf{x}_{\sigma(t)}^T (\mathbf{w}_t - \mathbf{w}^*) + \boldsymbol{\eta}_t,$$

where $\sigma(t)$ is an index from the dataset sampled uniformly at random, and $\boldsymbol{\eta}_t$ is an error term bounded by $\|\boldsymbol{\eta}_t\| \leq 2\alpha\epsilon L \|\mathbf{w}_t - \mathbf{w}^*\|$. This gives

$$\mathbf{w}_{t+1} - \mathbf{w}^* = \left(\mathbf{I} - \alpha \mathbf{x}_{\sigma(t)} \mathbf{x}_{\sigma(t)}^T \right) (\mathbf{w}_t - \mathbf{w}^*) + \boldsymbol{\eta}_t,$$

Taking the norm and the expected value gives

$$\begin{aligned}
\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &= \mathbf{E} \left[(\mathbf{w}_t - \mathbf{w}^*)^T \left(\mathbf{I} - \alpha \mathbf{x}_{\sigma(t)} \mathbf{x}_{\sigma(t)}^T \right)^2 (\mathbf{w}_t - \mathbf{w}^*) \right. \\
&\quad \left. + 2\boldsymbol{\eta}_t^T \left(\mathbf{I} - \alpha \mathbf{x}_{\sigma(t)} \mathbf{x}_{\sigma(t)}^T \right) (\mathbf{w}_t - \mathbf{w}^*) + \|\boldsymbol{\eta}_t\|^2 \right].
\end{aligned}$$

If we let Σ denote the expected value

$$\Sigma = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^T,$$

then

$$\mathbf{E} \left[(\mathbf{x}_{i_t} \mathbf{x}_{i_t}^T) (\mathbf{x}_{i_t} \mathbf{x}_{i_t}^T)^T \right] = \mathbf{E} \left[\|\mathbf{x}_{\sigma(t)}\|^2 \mathbf{x}_{\sigma(t)} \mathbf{x}_{\sigma(t)}^T \right] \preceq L \Sigma.$$

Because of this, we have that

$$\begin{aligned}
\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 &\leq \mathbf{E} \left[(\mathbf{w}_t - \mathbf{w}^*)^T (\mathbf{I} - 2\alpha \Sigma + \alpha^2 L \Sigma) (\mathbf{w}_t - \mathbf{w}^*) \right. \\
&\quad \left. + 2\boldsymbol{\eta}_t^T (\mathbf{I} - \alpha \Sigma) (\mathbf{w}_t - \mathbf{w}^*) + \|\boldsymbol{\eta}_t\|^2 \right].
\end{aligned}$$

Since we assumed $\alpha L \leq 1$, with an additional application of Cauchy-Schwarz, we can simplify this to

$$\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \leq \mathbf{E} \left[(\mathbf{w}_t - \mathbf{w}^*)^T (\mathbf{I} - \alpha \Sigma) (\mathbf{w}_t - \mathbf{w}^*) + 2\|\boldsymbol{\eta}_t\| \|\mathbf{w}_t - \mathbf{w}^*\| + \|\boldsymbol{\eta}_t\|^2 \right].$$

If we let μ denote the smallest eigenvalue of Σ , then this can be bounded by

$$\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \leq \mathbf{E} \left[(1 - \alpha\mu) \|\mathbf{w}_t - \mathbf{w}^*\|^2 + 2\|\boldsymbol{\eta}_t\| \|\mathbf{w}_t - \mathbf{w}^*\| + \|\boldsymbol{\eta}_t\|^2 \right].$$

Substituting our bound on $\boldsymbol{\eta}$ gives

$$\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \leq \mathbf{E} \left[(1 - \alpha\mu) \|\mathbf{w}_t - \mathbf{w}^*\|^2 + 4\alpha\epsilon L \|\mathbf{w}_t - \mathbf{w}^*\|^2 + 4\alpha^2 \epsilon^2 L^2 \|\mathbf{w}_t - \mathbf{w}^*\|^2 \right].$$

Disregarding the ϵ^2 term, we have

$$\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \leq \mathbf{E} \left[(1 - \alpha\mu) \|\mathbf{w}_t - \mathbf{w}^*\|^2 + 4\alpha\epsilon L \|\mathbf{w}_t - \mathbf{w}^*\|^2 \right],$$

which simplifies to

$$\mathbf{E} \|\mathbf{w}_{t+1} - \mathbf{w}^*\|^2 \leq (1 - \alpha\mu + 4\alpha\epsilon L) \mathbf{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2].$$

In other words, as long as $\epsilon L \ll \mu$ (i.e. ϵ is small relative to the inverse of the condition number $\kappa = L/\mu$ of the problem, we will still converge at a linear rate. More explicitly,

$$\mathbf{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2] \leq (1 - \alpha\mu + 4\alpha\epsilon L)^t \cdot \|\mathbf{w}_0 - \mathbf{w}^*\|^2.$$

Or,

$$\mathbf{E} [\|\mathbf{w}_t - \mathbf{w}^*\|^2] \leq \exp(-\alpha\mu t (1 - 4\epsilon\kappa)) \cdot \|\mathbf{w}_0 - \mathbf{w}^*\|^2.$$

□

B HIGH-ACCURACY PURE 16-BIT TRAINING

In this section, we present the algorithm for SGD and AdamW (Loshchilov & Hutter, 2017) when combined with stochastic rounding or Kahan summation on the model weight updates in Algorithms 2 to 5; these four algorithms describe the optimizers in our experiments. In these algorithms, all the scalars and tensors such as model weights and optimizer states are in BFloat16 precision. On top of these values, all floating point arithmetic operators takes BFloat16 inputs and uses nearest rounding for the output unless noticed otherwise. For SGD and AdamW with stochastic rounding on the model weight updates, we define operator \ominus as a subtraction with BFloat16 inputs and stochastic rounding for the output. Across our experiments using these optimizers, to calculate mini-batch gradient $\nabla f_{\sigma(t)}(\mathbf{w}_t)$, forward and backward compute operators consumes BFloat16 input and perform nearest rounding for the output.

To demonstrate the hardware efficiency of Pure 16-bit training algorithms in Algorithms 2 to 5, we discuss the hardware overhead of using stochastic rounding and Kahan summation for model weight updates in Appendices B.1 and B.2 respectively.

B.1 SYSTEM EFFICIENCY OF STOCHASTIC ROUNDING

In modern hardware, training with stochastic rounding for model weight updates can be implemented efficiently. It has been demonstrated that stochastic rounding can be implemented without any expensive multiplication or division arithmetic De Sa et al. (2017). Instead, it only requires 1) generating random bit sequences with a shift register, 2) adding random sequence to the lower mantissa bits and 3) truncating; these operations are inexpensive relative to other optimizer operations regardless of the model size. Secondly, we note that the minimal technique in our study only requires stochastic rounding for model weight updates while even cheaper nearest rounding remains for forward, backward and optimizer operations other than weight updates. Because of the above two reasons, using stochastic rounding for model weight updates only adds minimal overhead compared to standard 16-bit training fully with nearest rounding.

B.2 SYSTEM EFFICIENCY OF KAHAN SUMMATION

Despite the fact that 16-bit Kahan accumulation requires an additional 16-bit auxiliary value, it can still demonstrates advantages over 32-bit and mixed precision training in terms of throughput and memory consumption. In more details, optimizers in both 32-bit and mixed precision training operate on 32-bit weights (for mixed precision, the weights here refer to the master copy) and 32-bit optimizer states such as momentum. On the other hand, in 16-bit training with Kahan summation, the optimizers leverage fully 16-bit weights, optimizer states and auxiliary variables.

Regarding the throughput, optimizers in 32-bit and mixed precision training require compute units with 32-bit multiply, while optimizers in 16-bit training with Kahan summation can leverage units with 16-bit multiply. Given that 2) modern MAC units with 16-bit multiply can be implemented with 1.5X throughput compared to those with 32-bit multiply [10] and 3) Kahan summation only introduce 3 additional add/subtract operations which are inexpensive relative to multiply (e.g. Adam has 9 major multiply operations), optimizers in 16-bit training with Kahan summation can demonstrate meaningful speedup.

Regarding the memory efficiency, using Adam optimizer as an example, 16-bit training with Kahan summation saves 33% and 43% memory for weights plus optimizer states respectively over 32-bit training and mixed precision training (mixed precision training has both 16-bit and 32-bit weights in

memory). We note that these benefits here are especially pronounced for training large SOTA models with billions of model weights using optimizers with sophisticated optimizer states (Shoeybi et al., 2019; Rajbhandari et al., 2019).

Algorithm 2 SGD with Stochastic Rounding

```

1: INPUT: learning rate  $\eta_t$ , momentum  $\mu$ , weight decay  $d$ , gradient  $\nabla f_{\sigma(t)}(\mathbf{w}_t)$  at iteration  $t$ 
2: while  $t < T$  do
3:    $\mathbf{g}_{t+1} \leftarrow \nabla f_{\sigma(t)}(\mathbf{w}_t) + d * \mathbf{w}_t$ 
4:    $\mathbf{m}_{t+1} \leftarrow \mu * \mathbf{m}_t + \mathbf{g}_{t+1}$ 
5:    $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t \ominus (\eta_t * \mathbf{m}_{t+1})$ 
6:    $t \leftarrow t + 1$ 
7: end while
8: RETURN:  $\mathbf{w}_T$ 

```

Algorithm 3 SGD with Kahan Summation

```

1: INPUT: learning rate  $\eta_t$ , momentum  $\mu$ , weight decay  $d$ , gradient  $\nabla f_{\sigma(t)}(\mathbf{w}_t)$  at iteration  $t$ 
2:  $\mathbf{c}_0 \leftarrow 0$ 
3: while  $t < T$  do
4:    $\mathbf{g}_{t+1} \leftarrow \nabla f_{\sigma(t)}(\mathbf{w}_t) + d * \mathbf{w}_t$ 
5:    $\mathbf{m}_{t+1} \leftarrow \mu * \mathbf{m}_t + \mathbf{g}_{t+1}$ 
6:    $\mathbf{u}_{t+1} \leftarrow -(\eta_t * \mathbf{m}_{t+1})$ 
7:    $\mathbf{y}_{t+1} \leftarrow \mathbf{u}_{t+1} - \mathbf{c}_t$ 
8:    $\mathbf{s}_{t+1} \leftarrow \mathbf{w}_t + \mathbf{y}_{t+1}$ 
9:    $\mathbf{c}_{t+1} \leftarrow (\mathbf{s}_{t+1} - \mathbf{w}_t) - \mathbf{y}_{t+1}$ 
10:   $\mathbf{w}_{t+1} \leftarrow \mathbf{s}_{t+1}$ 
11:   $t \leftarrow t + 1$ 
12: end while
13: RETURN:  $\mathbf{w}_T$ 

```

Algorithm 4 AdamW with Stochastic Rounding

```

1: INPUT: learning rate  $\eta_t$ , beta  $(\beta_1, \beta_2)$ , weight decay  $d$ , gradient  $\nabla f_{\sigma(t)}(\mathbf{w}_t)$  at iteration  $t$ 
2:  $c_{1,0} \leftarrow 1, c_{2,0} \leftarrow 1$ 
3: while  $t < T$  do
4:    $\mathbf{g}_{t+1} \leftarrow \nabla f(\mathbf{w}_t, \mathbf{x}_t)$ 
5:    $\mathbf{m}_{t+1} \leftarrow \beta_1 * \mathbf{m}_t + (1 - \beta_1) * \mathbf{g}_t$ 
6:    $\mathbf{v}_{t+1} \leftarrow \beta_2 * \mathbf{v}_t + (1 - \beta_2) * \mathbf{g}_t^2$ 
7:    $c_{1,t+1} \leftarrow c_{1,t} * \beta_1$ 
8:    $c_{2,t+1} \leftarrow c_{2,t} * \beta_2$ 
9:    $\hat{\mathbf{m}}_{t+1} \leftarrow \mathbf{m}_{t+1} / (1 - c_{1,t+1})$ 
10:   $\hat{\mathbf{v}}_{t+1} \leftarrow \sqrt{\mathbf{v}_{t+1} / (1 - c_{2,t+1})}$ 
11:   $\mathbf{w}_{t+1} \leftarrow \mathbf{w}_t \ominus (\eta_t * \hat{\mathbf{m}}_{t+1} / (\hat{\mathbf{v}}_{t+1} + \epsilon) + \eta_t * d * \mathbf{w}_t)$ 
12:   $t \leftarrow t + 1$ 
13: end while
14: RETURN:  $\mathbf{w}_T$ 

```

C EXPERIMENT DETAILS

In this section, we first discuss the details in the experiment setup. We first discuss the hyperparameter values we use in Appendix C.1. We then present the infrastructure details used in our experiments in Appendix C.2.

C.1 HYPERPARAMETERS

For each model in our experiment, we use the hyperparameter values acquired from the original paper or code repository. In the rest of this section, we discuss the detailed hyperparameter values for all the seven deep learning applications in our experiments.

Algorithm 5 AdamW with Kahan Summation

```

1: INPUT: learning rate  $\eta_t$ , betas  $(\beta_1, \beta_2)$ , weight decay  $d$ , gradient  $\nabla f_{\sigma(t)}(\mathbf{w}_t)$  at iteration  $t$ 
2:  $c_{1,0} \leftarrow 1, c_{2,0} \leftarrow 1$ 
3:  $\mathbf{c}_0 \leftarrow 0$ 
4: while  $t < T$  do
5:    $\mathbf{g}_t \leftarrow \nabla f(\mathbf{w}_t, \mathbf{x}_t)$ 
6:    $\mathbf{m}_{t+1} \leftarrow \beta_1 * \mathbf{m}_t + (1 - \beta_1) * \mathbf{g}_t$ 
7:    $\mathbf{v}_{t+1} \leftarrow \beta_2 * \mathbf{v}_t + (1 - \beta_2) * \mathbf{g}_t^2$ 
8:    $c_{1,t+1} \leftarrow c_{1,t} * \beta_1$ 
9:    $c_{2,t+1} \leftarrow c_{2,t} * \beta_2$ 
10:   $\hat{\mathbf{m}}_{t+1} \leftarrow \mathbf{m}_{t+1} / (1 - c_{1,t+1})$ 
11:   $\hat{\mathbf{v}}_{t+1} \leftarrow \sqrt{\mathbf{v}_{t+1} / (1 - c_{2,t+1})}$ 
12:   $\mathbf{u}_{t+1} \leftarrow -(\eta_t * \hat{\mathbf{m}}_{t+1} / (\hat{\mathbf{v}}_{t+1} + \epsilon) + \eta_t * d * \mathbf{w}_t)$ 
13:   $\mathbf{y}_{t+1} \leftarrow \mathbf{u}_{t+1} - \mathbf{c}_t$ 
14:   $\mathbf{s}_{t+1} \leftarrow \mathbf{w}_t + \mathbf{y}_{t+1}$ 
15:   $\mathbf{c}_{t+1} \leftarrow (\mathbf{s}_{t+1} - \mathbf{w}_t) - \mathbf{y}_{t+1}$ 
16:   $\mathbf{w}_{t+1} \leftarrow \mathbf{s}_{t+1}$ 
17:   $t \leftarrow t + 1$ 
18: end while
19: RETURN:  $\mathbf{w}_T$ 

```

ResNet-18-CIFAR10⁵ We train the model for 350 epochs using the SGD optimizer with mini-batch size, weight decay, and momentum values of 128, 5×10^{-4} , and 0.9 respectively. The learning rate starts from 0.1 and We manually divide it by 10 at epochs 150 and 250. We summarize these hyperparameters in Table 3.

Table 3: Training hyperparameter for ResNet-18-CIFAR10

Hyperparameter	Value
Optimizer	SGD
Batchsize	128
Training epochs	350
Learning rate	0-150 epochs: 0.1 150-250 epochs: 0.01 250-350 epochs: 0.001
weight decay	5×10^{-4}
momentum	0.9

ResNet-50-ImageNet⁶ We train the model for 90 epochs using the SGD optimizer with mini-batch size, weight decay, and momentum values of 256, 1×10^{-4} , and 0.9 respectively. The learning rate starts from 0.1 and we manually divide it by 10 after each 30 epochs. We summarize these hyperparameters in Table 4.

BERT-MNLI⁷ We fine-tune the model for 3 epochs using the AdamW optimizer with mini-batch size and first order momentum values of 256, 0.9 respectively. The learning rate starts from 2×10^{-5} and linearly decays to 0 during the training. Note in pure 16-bit training algorithms, we use $\beta_2 = 0.997$ for pure 16-bit training algorithms. This is because 0.999 is considered as 1 in BFloat16 format; we use the 0.997 which is a BFloat16 representable value that is the closest to 0.999 and is smaller than 1. We summarize these hyperparameters in Table 5.

⁵<https://github.com/kuangliu/pytorch-cifar>

⁶<https://github.com/pytorch/examples/tree/master/imagenet>

⁷<https://github.com/huggingface/transformers>

⁸<https://github.com/pytorch/fairseq/blob/master/examples/roberta/README.pretraining.md>

Table 4: Training hyperparameter for ResNet-50-ImageNet

Hyperparameter	Value
Optimizer	SGD
Batchsize	256
Training epochs	90
Learning rate	0-30 epochs: 0.1 30-60 epochs: 0.01 60-90 epochs: 0.001
Weight decay	1×10^{-4}
Momentum	0.9

Table 5: Training hyperparameter for BERT-MNLI

Hyperparameter	Value
Model type	Base
Optimizer	AdamW
Batchsize	64
Training epochs	3
Learning rate	2×10^{-5}
Decay rate for 1st moment β_1 (32-bit and 16-bit)	0.9
Decay rate for 2nd moment β_2 (32-bit)	0.999
Decay rate for 2nd moment β_2 (16-bit)	0.997

BERT-Wiki103⁸ We pre-train the model for 31250 iterations using the AdamW optimizer with mini-batch size, weight decay, first order momentum, and second order momentum values of 512, 0.01, 0.9, and 0.98 respectively; compared to the original training steps, we scale down the training steps proportionally because we subsampled the datasets. The first 8% of training is used for learning rate warm-up and the peak learning rate value is 1×10^{-4} . Throughout the rest of the training (remaining 92%), learning rate decays to 0 linearly. We summarize these hyperparameters in Table 6.

Table 6: Training hyperparameter for BERT-wiki103

Hyperparameter	Value
Model type	Base
Sub-sample rate	25%
Optimizer	AdamW
Batchsize	512
Peak learning rate	1×10^{-4}
Total iterations	31250
Warmup iterations	2500
Weight decay	0.01
Decay rate for 1st moment β_1	0.9
Decay rate for 2nd moment β_2	0.98

DLRM-Kaggle⁹ We fine-tune the model for 1 epoch using the SGD optimizer with mini-batch size of 128. The learning rate has constant value of 0.1 for the whole training. We summarize these hyperparameters in Table 7.

⁹<https://github.com/facebookresearch/dlrm>

Table 7: Training hyperparameter for DLRM-Criteo Kaggle

Hyperparameter	Value
Sparse feature size	16
Bottom MLP	13-512-256-64-16
Top MLP	512-256-1
Optimizer	SGD
Batchsize	128
Training epochs	1
Learning rate	0.1
Loss function	BCE

DLRM-Terabyte For this model, we obtain hyperparameters from NVIDIA¹⁰ DLRM repository while we use the Facebook⁹ implementation. The reason is that we want to keep the implementation the same as DLRM-Kaggle experiment; however, the NVIDIA hyperparameters presents the state-of-the-art time-to-accuracy metric. We fine-tune the model for 1 epoch using the SGD optimizer with mini-batch size of 32768. The first 5% of training is used for learning rate warm-up and the peak learning rate value is 28. The learning rate decay starts at the middle (50% of iterations) of the training, and throughout the rest of the training (remaining 50%), learning rate decays to 0 linearly. We summarize these hyperparameters in Table 8.

Table 8: Training hyperparameter for DLRM-Criteo Terabyte

Hyperparameter	Value
Sparse feature size	16
Bottom MLP	512-256-128
Top MLP	1024-1024-512-256-2
Optimizer	SGD
Batchsize	32768
Training epochs	1
Learning rate	28
Warmup iterations	6400
Learning rate decay start Step	64000
Learning rate decay steps	80000
Loss function	BCE

DLRM-Kaggle¹¹ We train the model for 60 epochs using the SGD optimizer with mini-batch size, weight decay, and first order momentum values of 64, 1×10^{-5} , and 0.9 respectively. The learning rate starts from 3×10^{-4} and is reduced by 1% after each epoch. We summarize these hyperparameters in Table 9.

C.2 INFRASTRUCTURE DETAILS

We use AWS p3.16xlarge instances for the large applications including BERT-Wiki103, DeepSpeech2-LibriSpeech and ResNet-50-ImageNet. For the rest of the applications, we train on p3.2xlarge. The AWS p3.2xlarge instance has one NVIDIA’s Tesla V100 GPU, and p3.16xlarge has 8 of them. For all the experiments, we use Python version 3.8.1, PyTorch version 1.5, and CUDA version 10.2.

¹⁰<https://github.com/NVIDIA/DeepLearningExamples/blob/master/PyTorch/Recommendation/DLRM/dlrn/scripts/main.py>

¹¹<https://github.com/SeanNaren/deepspeech.pytorch/releases/tag/v2.0>

Table 9: Training hyperparameter for Deepspeech2-Librispeech

Hyperparameter	Value
RNN type	LSTM
Sub-sampled rate	10%
Optimizer	SGD
Batchsize	64
Training epochs	60
Learning rate	3×10^{-4}
Momentum	0.9
Weight decay	1×10^{-5}
Max-norm	400

D EXTENDED EMPIRICAL RESULTS

In this section, we discuss additional results complementary to those in Section 4. In Appendix D.1, we discuss the validation accuracy curves in our ablation study in Section 4 where we validate the model accuracy bottlenecks. In Appendix D.2, we present the validation accuracy curves for pure 16-bit training algorithms with stochastic rounding or Kahan summation for model weight updates. Finally in Appendix D.3, we present additional results on the following four aspects: 1) The percentage of model weight updates which are canceled during training for a representative deep learning model; 2) The model accuracy attained with lower than 16-bit precision format when stochastic rounding or Kahan summation is enabled for model weight updates; 3) The model accuracy attained by pure 16-bit training when we combine stochastic rounding and Kahan summation on all the model weights; 4) the model accuracy of pure 16-bit training using Float16 instead of BFloat16 with stochastic rounding or Kahan summation enabled.

D.1 THE MODEL ACCURACY BOTTLENECK

In Figure 7, we present the validation accuracy curves in our ablation study to validate the accuracy bottleneck. We can observe that the standard pure 16-bit training algorithm demonstrates lower validation accuracy metric values compared to 32-bit training across the applications. In the ablation, we save the model weights in 32-bit, turn off nearest rounding on model weight updates and keep nearest rounding for all the other compute operators. This ablation isolates the influence of nearest rounding on model weight updates. We can see that this ablated algorithm can match the validation accuracy attained by 32-bit precision training. These results together with the ones in Section 4 shows that nearest rounding on model weight updates is the primary cause for model accuracy gap in training deep learning models.

In our experiments, 32-bit and standard 16-bit training start from the same initialization with the same model accuracy. For the natural language inference task in Figure 1, the gap at the beginning of the curves is due to aggressive smoothing for the clarify in visualization. In Figure 6, we can observe that the unsmoothed curves for 32-bit training and standard 16-bit training indeed start from the same training accuracy level at the beginning.

D.2 HIGH-ACCURACY PURE 16-BIT TRAINING

In Figure 8, we present the validation accuracy curves attained by pure 16-bit training with stochastic rounding or Kahan summation for model weight updates. First we observe that with nearest rounding for all compute operations, the standard pure 16-bit training algorithm demonstrate validation accuracy gap compared to 32-bit training across training stages. In the contrast, we can observe that by using stochastic rounding or Kahan summation for model weight updates, pure 16-bit training achieves validation accuracy curves matching that of 32-bit training.

D.3 OTHER EXPERIMENT RESULTS

Nearest rounding cancels weight updates In this experiment, we show that a large portion of the model weight updates can be canceled due to nearest rounding on weight updates for a representa-




Figure 6: Unsmoothed training accuracy at different number of training iterations for BERT-MNLI.




Figure 7: **Validation accuracy gap imposed by the standard pure 16-bit training.** The standard algorithm fail to match the validation accuracy from 32-bit precision training. By ablating nearest rounding for model weight update from the standard method, we can recover the accuracy gap. This indicates the dominating impact of nearest rounding for model weight updates.

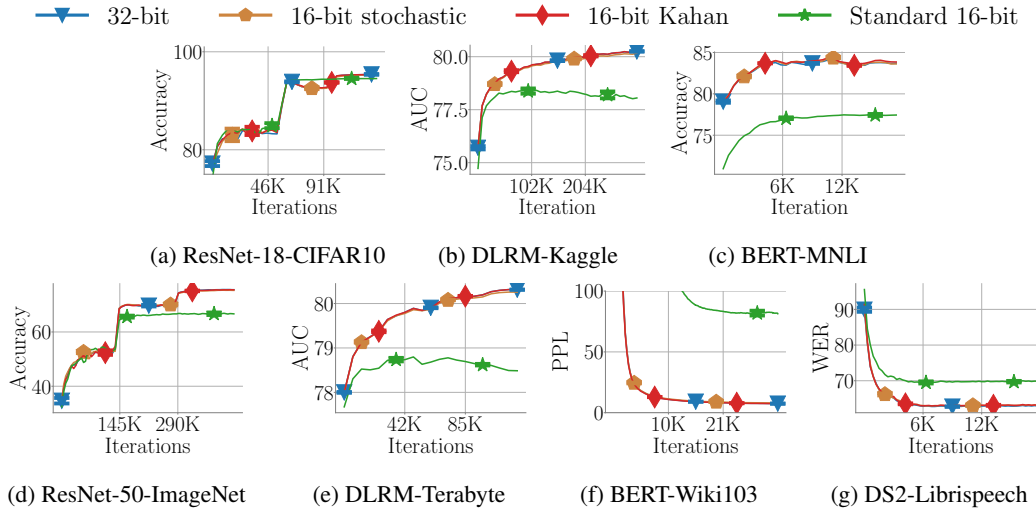


Figure 8: **Validation accuracy for pure 16-bit training.** With stochastic rounding or Kahan summation enabled for model weight updates, pure 16-bit training matches 32-bit precision training in terms of validation accuracy with negligible differences across the applications in our experiments.

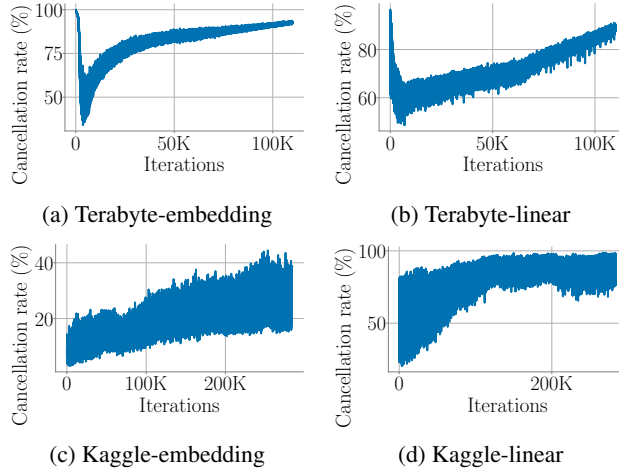


Figure 9: **The percentage of weight updates that are non-zero but get canceled in the standard pure 16-bit training.** For both embedding and linear layers, the percentage increases in the mid-to-late training stage for the DLRM model on the Kaggle and Terabyte datasets. The high percentage of cancellation aligns with our insights on the impact of nearest rounding on model weight updates.

tive deep learning model; this aligns with our theoretical insights on the impact of weight updates cancellation from the least-squares regression model. Specifically, we train DLRM model using the standard 16-bit training. In this experiment, we record the percentage of weight updates that are non-zero and get canceled by nearest rounding on model weight updates at each training iteration. Because the DLRM model consists of two different layer types which are embeddings and linear layers in MLPs, we present the results in Figure 9 on one embedding layer and one linear layer for each of the dataset we use for the DLRM model; the behavior on these layers is representative for many different embedding layers and linear layers in the DLRM model. As it is shown in Figure 9, on both the Kaggle and the Terabyte datasets, the percentage of weight updates that are non-zero but get canceled increases as the training progresses. For both the embedding and linear layers, the percentage of canceled weight updates can be larger than 80% in the mid-to-late training stage. For the Kaggle dataset, the increasing percentage of cancellation is due to the shrinking gradient magnitude because we use a constant learning rate on this dataset. On the other hand, we use a decaying learning rate on the Terabyte dataset. Thus the increasing trend on the percentage of weight update cancellation is a consequence of the compound effect from decaying gradient magnitude and decaying learning rates. These observations align with our insights on the least-squares regression model where we observe severe model weight updates cancellation in the mid-to-late training stage.

Going below 16-bit We train DLRM model on Criteo Kaggle dataset with precisions lower than 16-bit. For the lower precision representations, we keep 8 exponent bits as BFloat16 and only lower the number of mantissa bits to ensure enough number dynamic range for stable training. In Figure 10, we demonstrate the model accuracy attained with 10, 12 and 14-bit precision. Besides using 14-bit precision with Kahan summation for model weight updates, the other configuration all demonstrate lower training AUC value compared to 32-bit training and pure 16-bit training with stochastic rounding or Kahan summation. This implies that to go lower than 16-bit precision for training generic deep learning model, additional techniques might be required.




Figure 11: **Pure 16-bit training using stochastic rounding and Kahan summation simultaneously.** By using the Kahan summation and stochastic rounding simultaneously for model weight updates, pure 16-bit training can also robustly attain matching model accuracy compared to 32-bit precision training.




Figure 10: **Training accuracy with precisions lower than 16-bit.** Besides 14-bit training with Kahan summation for model weight updates, training with lower precision than 16-bit shows lower training accuracy compared to either 16-bit or 32-bit training on DLRM-Kaggle.

Combining two techniques To test the robustness of the two numerical techniques for pure 16-bit training, we discuss the model accuracy attained by applying stochastic rounding and Kahan summation simultaneously on all the model weights. In Figure 11, we can observe that stochastic rounding and Kahan summation can work together robustly and attain validation accuracy matching that of 32-bit training.

Using Float16 instead of BFloat16 We also consider doing pure 16-bit training with Float16 precision. We apply stochastic rounding and Kahan summation for model weights in pure 16-bit training with Float16 precision. We consider the minimal algorithms without any loss scaling techniques as in mixed precision training. As shown in Figure 12, we can observe that pure 16-bit training with Float16 precision (even with the stochastic rounding or Kahan summation technique) demonstrate significant model accuracy degradation compared to 32-bit precision training. Because Float16 has more mantissa bits but less exponent bits than BFloat16, the model accuracy gap is mostly due to the small dynamic range of Float16.

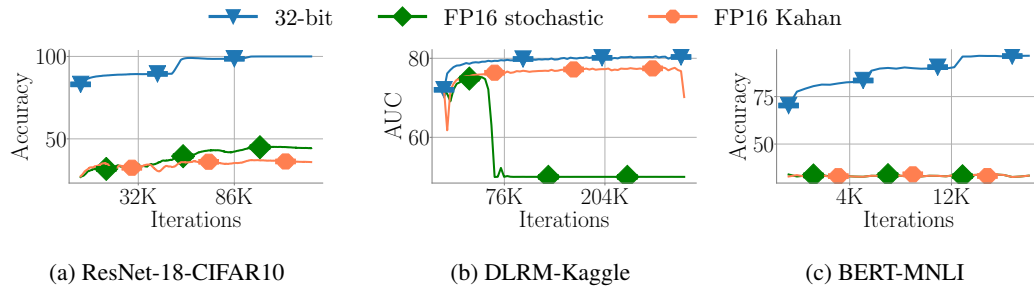


Figure 12: **Training accuracy gap imposed by pure 16-bit training with Float16** . Despite the fact that we enabled stochastic rounding or Kahan summation for the model weight updates, pure 16-bit training with Float16 still demonstrate significant model accuracy gap compared to 32-bit training. This is mostly due to the small dynamic range of the Float16 precision.