

# LOW-RANK TRAINING OF DEEP NEURAL NETWORKS FOR EMERGING MEMORY TECHNOLOGY

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

The recent success of neural networks for solving difficult decision tasks has incentivized incorporating smart decision making “at the edge.” However, this work has traditionally focused on neural network *inference*, rather than *training*, due to memory and compute limitations, especially in emerging non-volatile memory systems, where writes are energetically costly and reduce lifespan. Yet, the ability to train at the edge is becoming increasingly important as it enables real-time adaptability to device drift and environmental variation, user customization, and federated learning across devices. In this work, we address two key challenges for training on edge devices with non-volatile memory: low write density and low auxiliary memory. We present a low-rank training scheme that addresses these challenges while maintaining computational efficiency. We then demonstrate the technique on a representative convolutional neural network across several adaptation problems, where it out-performs standard SGD both in accuracy and in number of weight writes.

## 1 INTRODUCTION

Deep neural networks have shown remarkable performance on a variety of challenging inference tasks. As the energy efficiency of deep-learning inference accelerators improves, some models are now being deployed directly to edge devices to take advantage of increased privacy, reduced network bandwidth, and lower inference latency. Despite edge deployment, training happens predominately in the cloud. This limits the privacy advantages of running models on-device and results in static models that do not adapt to evolving data distributions in the field.

Efforts aimed at on-device training address some of these challenges. Federated learning aims to keep data on-device by training models in a distributed fashion (Konecný et al., 2016). On-device model customization has been achieved by techniques such as weight-imprinting (Qi et al., 2018), or by retraining limited sets of layers. On-chip training has also been demonstrated for handling hardware imperfections (Zhang et al., 2017; Gonugondla et al., 2018). Despite this progress with small models, on-chip training of larger models is bottlenecked by the limited memory size and compute horsepower of edge processors.

Emerging non-volatile (NVM) memories such as resistive random access memory (RRAM) have shown great promise for energy and area-efficient inference (Yu, 2018). However, on-chip training requires a large number of writes to the memory, and RRAM writes cost significantly more energy than reads (e.g., 10.9 pJ/bit versus 1.76 pJ/bit (Wu et al., 2019)). Additionally, RRAM endurance is on the order of  $10^6$  writes (Grossi et al., 2019), shortening the lifetime of a device due to memory writes for on-chip training.

In this paper, we present an online training scheme amenable to NVM memories to enable next generation edge devices. Our contributions are (1) an algorithm called Streaming Kronecker Sum Approximation (SKS), and its analysis, which addresses the two key challenges of low write density and low auxiliary memory; (2) two techniques “gradient max-norm” and “streaming batch norm” to help training specifically in the online setting; (3) a suite of adaptation experiments to demonstrate the advantages of our approach.

## 2 RELATED WORK

**Efficient training for resistive arrays.** Several works have aimed at improving the efficiency of training algorithms on resistive arrays. Of the three weight-computations required in training (forward, backprop, and weight update), weight updates are the hardest to parallelize using the array structure. Stochastic weight updates (Gokmen & Vlasov, 2016) allow programming of all cells in a crossbar at once, as opposed to row/column-wise updating. Online Manhattan rule updating (Zamanidoost et al., 2015) can also be used to update all the weights at once. Several works have proposed new memory structures to improve the efficiency of training (Soudry et al., 2015; Ambrogio et al., 2018). The number of writes has also been quantified in the context of chip-in-the-loop training (Yu et al., 2016).

**Distributed gradient descent.** Distributed training in the data center is another problem that suffers from expensive weight updates. Here, the model is replicated onto many compute nodes and in each training iteration, the mini-batch is split across the nodes to compute gradients. The distributed gradients are then accumulated on a central node that computes the updated weights and broadcasts them. These systems can be limited by communication bandwidth, and compressed gradient techniques (Aji & Heafield, 2017) have therefore been developed. In Lin et al. (2017), the gradients are accumulated over multiple training iterations on each compute node and only gradients that exceed a threshold are communicated back to the central node. In the context of on-chip training with NVM, this method helps reduce the number of weight updates. However, the gradient accumulator requires as much memory as the weights themselves, which negates the density benefits of NVM.

**Low-Rank Training.** Our work draws heavily from previous low-rank training schemes that have largely been developed for use in recurrent neural networks to uncouple the training memory requirements from the number of time steps inherent to the standard truncated backpropagation through time (TBPTT) training algorithm. Algorithms developed since then to address the memory problem include Real-Time Recurrent Learning (RTRL) (Williams & Zipser, 1989), Unbiased Online Recurrent Optimization (UORO) (Tallec & Ollivier, 2017), Kronecker Factored RTRL (KF-RTRL) (Mujika et al., 2018), and Optimal Kronecker Sums (OK) (Benzing et al., 2019). These latter few techniques rely on the weight gradients in a weight-vector product looking like a sum of outer products (i.e., Kronecker sums) of input vectors with backpropagated errors. Instead of storing a growing number of these sums, they can be approximated with a low-rank representation involving fewer sums.

## 3 TRAINING NON-VOLATILE MEMORY

The meat of most deep learning systems are many weight matrix - activation vector products  $\mathbf{W} \cdot \mathbf{a}$ . Fully-connected (dense) layers use them explicitly:  $\mathbf{a}^{[\ell]} = \sigma(\mathbf{W}^{[\ell]} \mathbf{a}^{[\ell-1]} + \mathbf{b}^{[\ell]})$  for layer  $\ell$ , where  $\sigma$  is a non-linear activation function (more details are discussed in detail in Appendix C.1). Recurrent neural networks use one or many matrix-vector products per recurrent cell. Convolutional layers can also be interpreted in terms of matrix-vector products by unrolling the input feature map into strided convolution-kernel-size slices. Then, each matrix-vector product takes one such input slice and maps it to all channels of the corresponding output pixel (more details are discussed in Appendix C.2).

The ubiquity of matrix-vector products allows us to adapt the techniques discussed in “Low-Rank Training” of Section 2 to other network architectures. Instead of reducing the memory across time steps, we can reduce the memory across training samples in the case of a traditional feedforward neural network. However, in traditional training (e.g., on a GPU), this technique does not confer advantages. Traditional training platforms often have ample memory to store a batch of activations and backpropagated gradients, and the weight updates  $\Delta \mathbf{W}$  can be applied directly to the weights  $\mathbf{W}$  once they are computed, allowing temporary activation memory to be deleted. The benefits of low-rank training only become apparent when looking at the challenges of proposed NVM devices:

**Low write density (LWD).** In NVM, writing to weights at every sample is costly in energy, time, and endurance. These concerns are exacerbated in multilevel cells, which require several steps of an iterative write-verify cycle to program the desired level. We therefore want to minimize the number of writes to NVM.

**Low auxiliary memory (LAM).** NVM is the densest form of memory. In 40nm technology, RRAM 1T-1R bitcells @ 0.085  $\mu\text{m}^2$  (Chou et al., 2018) are 2.8x smaller than 6T SRAM cells @ 0.242  $\mu\text{m}^2$  (TSMC, 2019). Therefore, NVM should be used to store the memory-intensive weights. By the same

token, no other on-chip memory should come close to the size of the on-chip NVM. In particular, if our  $b$ -bit NVM stores a weight matrix of size  $n_o \times n_i$ , we should use at most  $r(n_i + n_o)b$  auxiliary non-NVM memory, where  $r$  is a small constant. Despite these space limitations, the reason we might opt to use auxiliary (large, high endurance, low energy) memory is because there are places where writes are frequent, violating LWD if we were to use NVM.

In the traditional minibatch SGD setting with batch size  $B$ , an upper limit on the write density per cell per sample is easily seen:  $1/B$ . However, to store such a batch of updates without intermediate writes to NVM would require auxiliary memory proportional to  $B$ . Therefore, a trade-off becomes apparent. If  $B$  is reduced, LAM is satisfied at the cost of LWD. If  $B$  is raised, LWD is satisfied at the cost of LAM. Using low-rank training techniques, the auxiliary memory requirements are decoupled from the batch size, allowing us to increase  $B$  while satisfying both LWD and LAM<sup>1</sup>. Additionally, because the low-rank representation uses so little memory, a larger bitwidth can be used, potentially allowing for gradient accumulation in a way that is not possible with low bitwidth NVM weights. In the next section, we elaborate on the low-rank training method.

## 4 LOW-RANK TRAINING METHOD

Let  $\mathbf{z}^{(i)} = \mathbf{W}\mathbf{a}^{(i)} + \mathbf{b}$  be the standard affine transformation building block of some larger network, e.g.,  $\mathbf{y}_p^{(i)} = f_{post}(\mathbf{z}^{(i)})$  and  $\mathbf{a}^{(i)} = f_{pre}(\mathbf{x}^{(i)})$  with prediction loss  $\mathcal{L}(\mathbf{y}_p^{(i)}, \mathbf{y}_t^{(i)})$ , where  $(\mathbf{x}^{(i)}, \mathbf{y}_t^{(i)})$  is the  $i^{\text{th}}$  training sample pair. Then weight gradient  $\nabla_{\mathbf{W}}\mathcal{L}^{(i)} = \mathbf{d}\mathbf{z}^{(i)} (\mathbf{a}^{(i)})^\top = \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)}$  where  $\mathbf{d}\mathbf{z}^{(i)} = \nabla_{\mathbf{z}^{(i)}}\mathcal{L}^{(i)}$ . A minibatch SGD weight update accumulates this gradient over  $B$  samples:  $\Delta\mathbf{W} = -\eta \sum_{i=1}^B \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)}$  for learning rate  $\eta$ .

For a rank- $r$  training scheme, approximate the sum  $\sum_{i=1}^B \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)}$  by iteratively updating two rank- $r$  matrices  $\tilde{\mathbf{L}} \in \mathbb{R}^{n_o \times r}$ ,  $\tilde{\mathbf{R}} \in \mathbb{R}^{n_i \times r}$  with each new outer product:  $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top \leftarrow \text{rankReduce}(\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top + \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)})$ . Therefore, at each sample, we convert the rank- $q = r + 1$  system  $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top + \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)}$  into the rank- $r$   $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top$ . In the next sections, we discuss how to compute *rankReduce*.

### 4.1 OPTIMAL KRONECKER SUM APPROXIMATION (OK)

One option for *rankReduce*( $\mathbf{X}$ ) to convert from rank  $q = r + 1$   $\mathbf{X}$  to rank  $r$  is a minimum error estimator, which is implemented by selecting the top  $r$  components of a singular value decomposition (SVD) of  $\mathbf{X}$ . However, a naive implementation is computationally infeasible and biased:  $\mathbb{E}[\text{rankReduce}(\mathbf{X})] \neq \mathbf{X}$ . Benzing et al. (2019) solves these problems by proposing a minimum variance unbiased estimator for *rankReduce*, which they call the OK algorithm<sup>2</sup>.

The OK algorithm can be understood in two key steps: first, an efficient method of computing the SVD of a Kronecker sum; second, a method of splitting the singular value matrix  $\Sigma$  into two rank- $r$  matrices whose outer product is a minimum-variance, unbiased estimate of  $\Sigma$ . Details can be found in their paper, however we include a high-level explanation in Sections 4.1.1 and 4.1.2 to aid our discussions. Note that our variable notation differs from Benzing et al. (2019).

#### 4.1.1 EFFICIENT SVD OF KRONECKER SUMS

Let  $\mathbf{L} = [\tilde{\mathbf{L}}, \mathbf{d}\mathbf{z}^{(i)}]$  and  $\mathbf{R} = [\tilde{\mathbf{R}}, \mathbf{a}^{(i)}]$  so that  $\mathbf{L}\mathbf{R}^\top = \tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top + \mathbf{d}\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)}$ . Recall that *rankReduce* should turn rank- $q$   $\mathbf{L}\mathbf{R}^\top$  into an updated rank- $r$   $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top$ .

QR-factorize  $\mathbf{L} = \mathbf{Q}_L\mathbf{R}_L$  and  $\mathbf{R} = \mathbf{Q}_R\mathbf{R}_R$  where  $\mathbf{Q}_L \in \mathbb{R}^{n_o \times q}$ ,  $\mathbf{Q}_R \in \mathbb{R}^{n_i \times q}$  are orthogonal so that  $\mathbf{L}\mathbf{R}^\top = \mathbf{Q}_L(\mathbf{R}_L\mathbf{R}_R^\top)\mathbf{Q}_R^\top$ . Let  $\mathbf{C} = \mathbf{R}_L\mathbf{R}_R^\top \in \mathbb{R}^{q \times q}$ . Then we can find the SVD of  $\mathbf{C} = \mathbf{U}_C\mathbf{\Sigma}_C\mathbf{V}_C^\top$  in  $\mathcal{O}(q^3)$  time (Cline & Dhillon, 2006), making it computationally feasible on small devices. Now we have:

<sup>1</sup>This can alternately be achieved by sub-sampling the training data by  $r/B$  where  $r$  is the OK rank. The purpose of using a low-rank estimate is that for the same memory cost, it is significantly more informational than the sub-sampled data, allowing for faster training convergence.

<sup>2</sup>Their target application differs slightly in that they handle matrix - vector Kronecker sums rather than vector - vector Kronecker sums.

$$\mathbf{L}\mathbf{R}^\top = \mathbf{Q}_L(\mathbf{U}_C\mathbf{\Sigma}\mathbf{V}_C^\top)\mathbf{Q}_R^\top = (\mathbf{Q}_L\mathbf{U}_C)\mathbf{\Sigma}(\mathbf{Q}_R\mathbf{V}_C)^\top \quad (1)$$

which gives the SVD of  $\mathbf{L}\mathbf{R}^\top$  since  $\mathbf{Q}_L\mathbf{U}_C$  and  $\mathbf{Q}_R\mathbf{V}_C$  are orthogonal and  $\mathbf{\Sigma}$  is diagonal. This SVD computation has a time complexity of  $\mathcal{O}((n_i+n_o+q)q^2)$  and a space complexity of  $\mathcal{O}((n_i+n_o+q)q)$ .

#### 4.1.2 MINIMUM VARIANCE, UNBIASED ESTIMATE OF $\mathbf{\Sigma}$

In Benzing et al. (2019), it is shown that the problem of finding a rank- $r$  minimum variance unbiased estimator of  $\mathbf{L}\mathbf{R}^\top$  can be reduced to the problem of finding a rank- $r$  minimum variance unbiased estimator of  $\mathbf{\Sigma}$  and plugging it in to (1).

Further, it is shown that such an optimal approximator for  $\mathbf{\Sigma} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_q)$ , where  $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_q$  will involve keeping the  $m-1$  largest singular values and mixing the smaller singular values  $\sigma_m, \dots, \sigma_q$  within their  $(k+1) \times (k+1)$  submatrix with  $m, k$  defined below. Let:

$$m = \min_i \text{ s.t. } (q-i)\sigma_i \leq \sum_{j=i}^q \sigma_j \quad k = q - m$$

$$\mathbf{x}_0 = \left( \sqrt{1 - \frac{\sigma_m k}{s_1}}, \dots, \sqrt{1 - \frac{\sigma_q k}{s_1}} \right)^\top \quad s_1 = \sum_{i=m}^q \sigma_i$$

Note that  $\|\mathbf{x}_0\|_2 = 1$ . Let  $\mathbf{X} \in \mathbb{R}^{(k+1) \times (k)}$  be orthogonal such that its left nullspace is the span of  $\mathbf{x}_0$ . Then  $\mathbf{X}\mathbf{X}^\top = \mathbf{I} - \mathbf{x}_0\mathbf{x}_0^\top$ . Now, let  $\mathbf{s} \in \{-1, 1\}^{(k+1) \times 1}$  be uniform random signs and define:

$$\mathbf{X}_s = (\mathbf{s} \odot \mathbf{X}_{:,1}, \dots, \mathbf{s} \odot \mathbf{X}_{:,k}) \quad \mathbf{Z} = \sqrt{\frac{s_1}{k}} \cdot \mathbf{X}_s$$

$$\tilde{\mathbf{\Sigma}}_L = \tilde{\mathbf{\Sigma}}_R = \text{diag}(\sqrt{\sigma_1}, \dots, \sqrt{\sigma_{m-1}}, \mathbf{Z}) \quad (2)$$

where  $\odot$  is an element-wise product. Then  $\tilde{\mathbf{\Sigma}}_L\tilde{\mathbf{\Sigma}}_R^\top = \tilde{\mathbf{\Sigma}}$  is a minimum variance, unbiased<sup>3</sup> rank- $r$  approximation of  $\mathbf{\Sigma}$ . Plugging  $\tilde{\mathbf{\Sigma}}$  into (1),

$$\mathbf{L}\mathbf{R}^\top = (\mathbf{Q}_L\mathbf{U}_C)\mathbf{\Sigma}(\mathbf{Q}_R\mathbf{V}_C)^\top \approx (\mathbf{Q}_L\mathbf{U}_C)\tilde{\mathbf{\Sigma}}(\mathbf{Q}_R\mathbf{V}_C)^\top = (\mathbf{Q}_L\mathbf{U}_C\tilde{\mathbf{\Sigma}}_L)(\mathbf{Q}_R\mathbf{V}_C\tilde{\mathbf{\Sigma}}_R)^\top \quad (3)$$

Thus,  $\tilde{\mathbf{L}} = \mathbf{Q}_L\mathbf{U}_C\tilde{\mathbf{\Sigma}}_L \in \mathbb{R}^{n_o \times r}$  and  $\tilde{\mathbf{R}} = \mathbf{Q}_R\mathbf{V}_C\tilde{\mathbf{\Sigma}}_R \in \mathbb{R}^{n_i \times r}$  gives us a minimum variance, unbiased, rank- $r$  approximation  $\tilde{\mathbf{L}}\tilde{\mathbf{R}}^\top$ .

## 4.2 STREAMING KRONECKER SUM APPROXIMATION (SKS)

Although the standalone OK algorithm presented by Benzing et al. (2019) has good asymptotic computational complexity, our vector-vector outer product sum use case permits further optimizations. In this section we present these optimizations, and we refer readers to the explicit implementation called Streaming Kronecker Sum Approximation (SKS) in Algorithm 1 of Appendix A.

### 4.2.1 MAINTAIN ORTHOGONAL $\mathbf{Q}_L, \mathbf{Q}_R$

The main optimization is a method of avoiding recomputing the QR factorization of  $\mathbf{L}$  and  $\mathbf{R}$  at every step. Instead, we keep track of orthogonal matrices  $\mathbf{Q}_L, \mathbf{Q}_R$ , and weightings  $\mathbf{c}_x$  such that  $\tilde{\mathbf{L}} = \mathbf{Q}_L \cdot \text{diag}(\sqrt{\mathbf{c}_x})_{[:r]}$  and  $\tilde{\mathbf{R}} = \mathbf{Q}_R \cdot \text{diag}(\sqrt{\mathbf{c}_x})_{[:r]}$ . Upon receiving a new sample, a single inner loop of the numerically-stable modified Gram-Schmidt (MGS) algorithm (Björck, 1967) can be used to update  $\mathbf{Q}_L$  and  $\mathbf{Q}_R$ . The orthogonal basis coefficients  $\mathbf{c}_L = \mathbf{Q}_L^\top \mathbf{d}\mathbf{z}^{(i)}$  and  $\mathbf{c}_R = \mathbf{Q}_R^\top \mathbf{a}^{(i)}$  computed during MGS can be used to find the new value of  $\mathbf{C} = \mathbf{c}_L\mathbf{c}_R^\top + \text{diag}(\mathbf{c}_x)$ .

<sup>3</sup>The fact that it is unbiased:  $\mathbb{E}[\tilde{\mathbf{\Sigma}}] = \mathbf{\Sigma}$  can be easily verified.

After computing  $\tilde{\Sigma}_L = \tilde{\Sigma}_R$  in (2), we can orthogonalize these matrices into  $\tilde{\Sigma}_L = \tilde{\Sigma}_R = \mathbf{Q}_x \mathbf{R}_x$ . Then from (3), we have  $\tilde{\mathbf{L}} \tilde{\mathbf{R}}^\top = (\mathbf{Q}_L \mathbf{U}_C \mathbf{Q}_x) (\mathbf{R}_x \mathbf{R}_x^\top) (\mathbf{Q}_R \mathbf{V}_C \mathbf{Q}_x)^\top$ . With this formulation, we can maintain orthogonality in  $\mathbf{Q}_L, \mathbf{Q}_R$  by setting:

$$\mathbf{Q}_L \leftarrow \mathbf{Q}_L \mathbf{U}_C \mathbf{Q}_x \quad \mathbf{Q}_R \leftarrow \mathbf{Q}_R \mathbf{V}_C \mathbf{Q}_x \quad \mathbf{c}_x \leftarrow \text{diag}(\mathbf{R}_x \mathbf{R}_x^\top)$$

These matrix multiplies require  $\mathcal{O}((n_i + n_o)q^2)$  multiplications, so this optimization does not improve asymptotic complexity bounds. This optimization may nonetheless be practically significant since matrix multiplies are easy to parallelize and would typically not be the bottleneck of the computation compared to Gram-Schmidt. The next section discusses how to orthogonalize  $\tilde{\Sigma}_L$  efficiently and why  $(\mathbf{R}_x \mathbf{R}_x^\top)$  is diagonal.

#### 4.2.2 ORTHOGONALIZATION OF $\tilde{\Sigma}_L$

Orthogonalization of  $\tilde{\Sigma}_L$  is relatively straightforward. From (2), the columns of  $\tilde{\Sigma}_L$  are orthogonal since  $\mathbf{Z}$  is orthogonal. However, they do not have unit norm. We can therefore pull out the norm into a separate diagonal matrix  $\mathbf{R}_x$  with diagonal elements  $\sqrt{\mathbf{c}_x}$ :

$$\mathbf{Q}_x = \begin{bmatrix} \mathbf{I}_{m-1} & 0 \\ 0 & \mathbf{X}_s \end{bmatrix} \quad \sqrt{\mathbf{c}_x} = (\sqrt{\sigma_1}, \dots, \sqrt{\sigma_{m-1}}, \underbrace{\sqrt{s_1/k}}_{q-m+1 \text{ times}})$$

#### 4.2.3 FINDING ORTHONORMAL BASIS $\mathbf{X}$

We generated  $\mathbf{X}$  by finding an orthonormal basis that was orthogonal to a vector  $\mathbf{x}_0$  so that we could have  $\mathbf{X} \mathbf{X}^\top = \mathbf{I} - \mathbf{x}_0 \mathbf{x}_0^\top$ . An efficient method of producing this basis is through Householder matrices  $(\mathbf{x}_0, \mathbf{X}) = \mathbf{I} - 2 \mathbf{v} \mathbf{v}^\top / \|\mathbf{v}\|^2$  where  $\mathbf{v} = \mathbf{x}_0 - \mathbf{e}^{(1)}$  and  $(\mathbf{x}_0, \mathbf{X})$  is a  $k+1 \times k+1$  matrix with first column  $\mathbf{x}_0$  and remaining columns  $\mathbf{X}$  (Householder, 1958; user1551, 2013).

#### 4.2.4 EFFICIENCY COMPARISONS TO STANDARD APPROACH

The OK/SKS methods require  $\mathcal{O}((n_i + n_o + q)q^2)$  operations per sample and  $\mathcal{O}(n_i n_o q)$  operations after collecting  $B$  samples, giving an amortized cost of  $\mathcal{O}((n_i + n_o + q)q^2 + n_i n_o q/B)$  operations per sample. Meanwhile, a standard approach expands the Kronecker sum at each sample, costing  $\mathcal{O}(n_i n_o)$  operations per sample. If  $q \ll B, n_i, n_o$  then the low rank method is superior to minibatch SGD in both memory and computational cost.

## 5 CONVEX CONVERGENCE

SKS introduces variance into the gradient estimates, so here we analyze the implications for online convex convergence. We analyze the case of strongly convex loss landscapes  $f^t(\mathbf{w}^t)$  for flattened weight vector  $\mathbf{w}^t$  and online sample  $t$ . In Appendix B, we show that with inverse squareroot learning rate, when the loss landscape Hessians satisfy  $0 < \mathbf{c} \mathbf{I} \preceq \nabla^2 f^t(\mathbf{w}^t)$  and under constraint (4) for the size of gradient errors  $\boldsymbol{\varepsilon}^t$ , where  $\mathbf{w}^*$  is the optimal offline weight vector, the online regret (5) is sublinear in the number of online steps  $T$ . We can approximate  $\|\boldsymbol{\varepsilon}\|$  and show that convex convergence is likely when (6) is satisfied in the biased, zero-variance case (equivalent to raw SVD, i.e., not applying Section 4.1.2), or when (7) is satisfied in the unbiased, minimum-variance case.

$$\|\boldsymbol{\varepsilon}^t\| \leq \frac{c}{2} \|\mathbf{w}^t - \mathbf{w}^*\| \quad (4) \quad R(T) = \sum_{t=1}^T f^t(\mathbf{w}^t) - \sum_{t=1}^T f^t(\mathbf{w}^*) \quad (5)$$

$$\sum_{i=1}^B \left( \sigma_q^{(t,i)} \right)^2 \leq \frac{c^2}{4} \|\mathbf{w}^t - \mathbf{w}^*\|^2 \quad (6) \quad \sum_{i=1}^B \sigma_r^{(t,i)} \sigma_q^{(t,i)} \leq \frac{c^2}{8} \|\mathbf{w}^t - \mathbf{w}^*\|^2 \quad (7)$$

Equations (6, 7) suggest conditions under which fast convergence may be more or less likely and also point to methods for improving convergence. We discuss these in more detail in Appendix B.3.

### 5.1 CONVERGENCE EXPERIMENTS

We validate (4) with several linear regression experiments on a static input batch  $\mathbf{X} \in \mathbb{R}^{1024 \times 100}$  and target  $\mathbf{Y}_t \in \mathbb{R}^{256 \times 100}$ . In Figure 1(a), Gaussian noise at different strengths (represented by different colors) is added to the true batch gradients at each update step. Notice that convergence slows significantly to the right of the dashed lines, which is the region where (4) no longer holds<sup>4</sup>.

In Figure 1(b), we validate Equations (4, 6, 7) by testing the SVD and SKS cases with rank  $r = 10$ . In these particular experiments, SKS adds too much variance, causing it to operate to the right of the dashed lines. However, both SVD and SKS can be seen to reduce their variance as training progresses. In the case of SVD, it is able to continue training as it tracks the right dashed line.

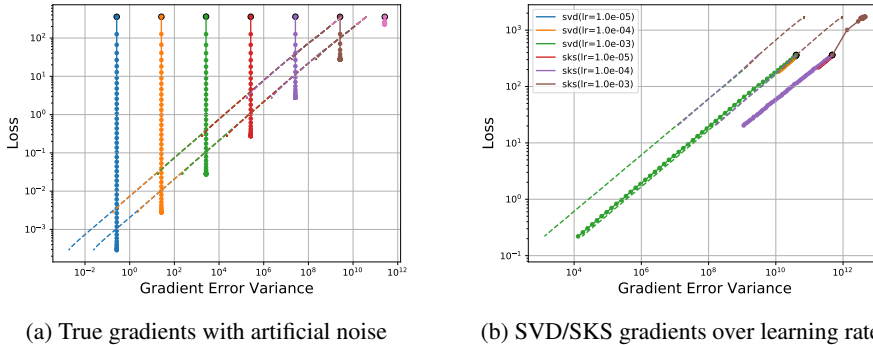


Figure 1: In both plots, the solid line with markers plots the loss vs. gradient error variance (LHS of (4)) across 50 steps of SGD for several different setups. The left dashed line represents the RHS of (4) and the right dashed line is the RHS with  $C$  instead of  $c$ .

## 6 IMPLEMENTATION DETAILS

**Quantization.** The NN is quantized in both the forward and backward directions with uniform power-of-2 quantization, where the clipping ranges are fixed at the start of training<sup>5</sup>. Weights are quantized to 8 bits between -1 and 1, biases to 16 bits between -8 and 8, activations to 8 bits between 0 and 2, and gradients to 8 bits between -1 and 1. Both the weights  $\mathbf{W}$  and weight updates  $\Delta\mathbf{W}$  are quantized to the same LSB so that weights cannot be used for accumulation beyond the fixed quantization dynamic range. This is in contrast to using high bitwidth (Zhou et al., 2016; Banner et al., 2018) or floating point accumulators. See Appendix D for more details on quantization.

**Gradient Max-Norming.** State-of-the-art methods in training, such as Adam (Kingma & Ba, 2014), use auxiliary memory per parameter to normalize the gradients. Unfortunately, we lack the memory budget to support these additional variables, especially if they must be updated every sample<sup>6</sup>. Instead, we propose dividing each gradient tensor by the maximum absolute value of its elements. This stabilizes the range of gradients across samples. See Appendix E for more details on gradient max-norming. In the experiments, we refer to this method as “max-norm” (opposite “no-norm”).

**Streaming Batch Normalization.** Batch normalization (Ioffe & Szegedy, 2015) is a powerful technique for improving training performance which has been suggested to work by smoothing the loss landscape (Santurkar et al., 2018). We hypothesize that this may be especially helpful when parameters are quantized as in our case. However, in the online setting, we receive samples one-at-a-time rather than in batches. We therefore propose a streaming batch norm that uses moving average statistics rather than batch statistics as described in detail in Appendix F.

<sup>4</sup>As discussed in Appendix B.1,  $B < n_i$ , so we substitute  $c$  in  $c/2\|\mathbf{w}^t - \mathbf{w}^*\|$  with the minimum non-zero Eigenvalue of the Hessian  $\tilde{c}$  when plotting the RHS of (4).

<sup>5</sup>Future work might look into how to change these clipping ranges, but this is beyond the scope of this paper.

<sup>6</sup>SKS could potentially approximate Adam. SKS on  $\mathbf{a}^2, \mathbf{dz}^2, \mathbf{a}, \mathbf{dz}$  allows for a low-rank approximation of the variance of the gradients, however, this is unlikely to work well because of numerical stability (e.g., estimated variances might be negative).

## 7 EXPERIMENTS

### 7.1 ADAPTATION EXPERIMENTS

To test the effectiveness of SKS, experiments are performed on a representative CNN with four  $3 \times 3$  convolution layers and two fully-connected layers. We generate “offline” and “online” datasets based on MNIST (see Appendix G), including one in which the statistical distribution shifts every 10k images. We then optimize an online SGD and rank-4 SKS model for fair comparison (see Appendix H). To see the importance of different training techniques, we run several ablations in Appendix I. Finally, we compare these different training schemes in different environments, meant to model real life. In these hypothetical scenarios, a model is first trained on the offline training set, and is then deployed to a number of devices at the edge that make supervised predictions (they make a prediction, then are told what the correct prediction would have been).

We present results on four hypothetical scenarios. First, a control case where both external/environment and internal/NVM drift statistics are exactly the same as during offline training. Second, a case where the input image statistical distribution shifts every 10k samples, selecting from augmentations such as spatial transforms and background gradients (see Section G). Third and fourth are cases where the NVM drifts from the programmed values, roughly modeling NVM memory degradation. In the third case, Gaussian noise is applied to the weights as if each weight was a single multi-level memory cell whose analog value drifted in a Brownian way. In the fourth case, random bit flips are applied as if each weight was represented by  $b$  memory cells (see Appendix G for details). For each hypothetical scenario, we plot five different training schemes: pure quantized inference (no training), bias-only training, standard SGD training, SKS training, and SKS training with max-normed gradients. In SGD training and for training biases, parameters are updated at every step in an online fashion. These are seen as different colored curves in Figure 2.

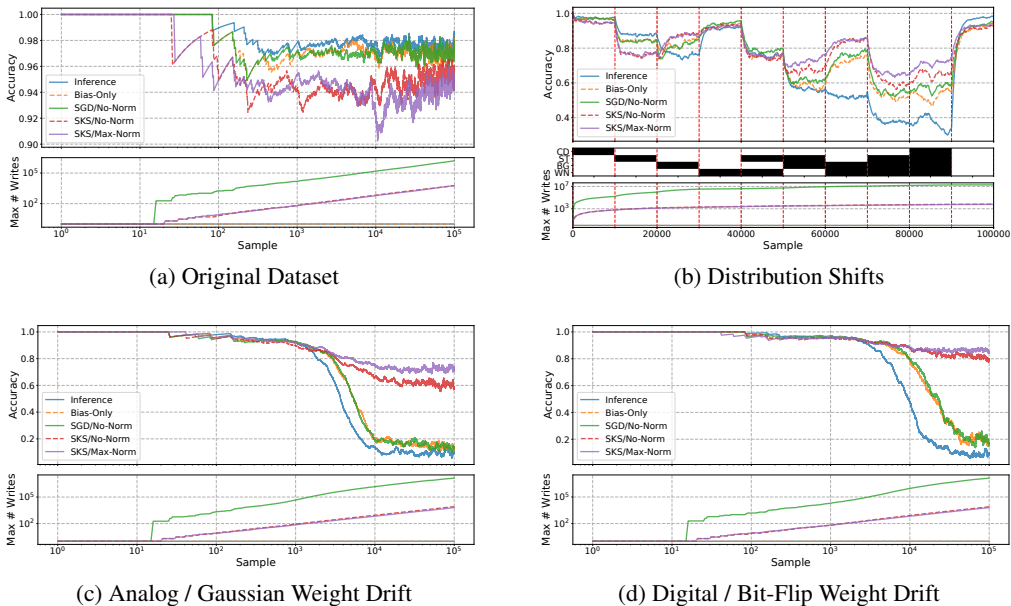


Figure 2: Adaptation of various training schemes over four different training environments (a) to (d). In each training environment, the top plot shows the exponential moving averages (0.999) of the per-sample online accuracy of the five training schemes, while the bottom plot shows the maximum number of updates applied to any given convolution or fully-connected kernel memory cell. For the distribution shifts in (b), the enabled augmentations at each contiguous 10k samples is shown (CD = class distribution, ST = spatial transforms, BG = background gradients, WN = white noise).

Inference does best in the control case, but does poorly in adaptation experiments. SGD doesn’t improve significantly on bias-only training, likely because SGD cannot accumulate gradients less than a weight LSB. SKS, on the other hand, shows significant improvement, especially after several

thousand samples in the weight drift cases. Additionally, SKS shows about three orders of magnitude improvement compared to SGD in the worst case number of weight updates. Much of this reduction is due to the convolutions, where updates are applied at each pixel. However, reduction in fully-connected writes is still important because of potential energy savings. SKS/max-norm performs best in terms of accuracy across all environments and has similar weight update cost to SKS/no-norm.

## 7.2 TRANSFER LEARNING AND ALGORITHM COMPARISONS

To test the broader applicability of low rank training techniques, we run several experiments on ImageNet with ResNet-34 (Deng et al., 2009; He et al., 2016), a potentially realistic target for dense NVM inference on-chip. For ImageNet-size images, updating the low-rank approximation at each pixel quickly becomes infeasible, both because of the single-threaded nature of the algorithm, and because of the increased variance of the estimate at larger batch sizes. Instead, we focus on training the final layer weights ( $1000 \times 512$ ). ResNet-34 weights are initialized to those from Paszke et al. (2017) and the convolution layers are used to generate feature vectors for 10k ImageNet training images<sup>7</sup>, which are quantized and fed to a one-layer quantized<sup>8</sup> neural network. To speed up experiments, the layer weights are initialized to the pretrain weights, modulated by random noise that causes inference top-1 accuracy to fall to  $52.7\% \pm 0.9\%$ . In Table 1, we see that the unbiased SKS has the strongest recovery accuracies, although biased SVD also does quite well. The high-variance UORO and true SGD have weak or non-existent recoveries.

Table 1: Accuracy recovery beyond inference (%), mean with standard deviation from 5 random seeds) between different algorithms (all with max-norm; effective batch size  $B = 100$  if applicable), tested at different ranks ( $r$ ), and learning rates ( $\eta$ ). Optimal learning rates are bolded.

Algorithm	$\eta$	0.003	0.010	0.030	0.100	0.300
	$r$					
SGD	-	$+0.3 \pm 0.2$	$+0.3 \pm 0.2$	$+0.3 \pm 0.2$	<b><math>+0.9 \pm 0.2</math></b>	$-3.9 \pm 0.8$
UORO	1	<b><math>+0.4 \pm 0.2</math></b>	$+0.3 \pm 0.4$	$-1.8 \pm 0.9$	$-7.6 \pm 1.6$	$-31.7 \pm 1.6$
SVD	1	$+1.9 \pm 0.2$	<b><math>+5.8 \pm 1.0</math></b>	$-3.4 \pm 1.0$	$-19.4 \pm 0.9$	$-40.7 \pm 1.1$
	2	$+1.4 \pm 0.4$	<b><math>+6.5 \pm 0.7</math></b>	$+6.3 \pm 0.6$	$-5.2 \pm 0.9$	$-36.3 \pm 0.9$
	4	$+1.3 \pm 0.4$	<b><math>+6.5 \pm 0.7</math></b>	$+5.2 \pm 0.8$	$-3.3 \pm 1.0$	$-33.8 \pm 0.8$
SKS	8	$+1.4 \pm 0.3$	<b><math>+5.6 \pm 0.8</math></b>	$+4.3 \pm 0.9$	$-2.4 \pm 1.0$	$-32.8 \pm 0.9$
	1	<b><math>+0.3 \pm 0.2</math></b>	$+0.3 \pm 0.2$	$-0.7 \pm 0.4$	$-2.7 \pm 1.7$	$-26.5 \pm 2.6$
	2	$+0.3 \pm 0.2$	$+0.4 \pm 0.3$	$-0.1 \pm 0.4$	<b><math>+1.3 \pm 0.9</math></b>	$-12.9 \pm 1.1$
	4	$+0.4 \pm 0.2$	$+0.6 \pm 0.2$	$+1.9 \pm 0.3$	<b><math>+8.0 \pm 1.1</math></b>	$-5.1 \pm 1.1$
	8	$+0.4 \pm 0.2$	$+1.1 \pm 0.2$	$+3.3 \pm 0.7$	<b><math>+4.8 \pm 1.5</math></b>	$-15.8 \pm 1.7$

## 8 CONCLUSION

We demonstrated the potential for SKS to solve the major challenges facing online training on NVM-based edge devices: low write density and low auxiliary memory. SKS is a computationally-efficient, memory-light algorithm capable of decoupling batch size from auxiliary memory, allowing larger effective batch sizes, and consequently lower write densities. Additionally, we noted that SKS may allow for training under severe weight quantization constraints as rudimentary gradient accumulations are handled by the  $\mathbf{L}$ ,  $\mathbf{R}$  matrices, which can have high bitwidths (as opposed to SGD, which may squash small gradients to 0). We found expressions for when SKS might have better convergence properties. Across a variety of online adaptation problems and a large-scale transfer learning demonstration, SKS was shown to match or exceed the performance of SGD while using a small fraction of the number of updates. Finally, we suspect that these techniques could be applied to a broader range of problems. Auxiliary memory minimization may be analogous to communication minimization in training strategies such as federated learning, where gradient compression is important.

<sup>7</sup>The decision to use training data is deliberate, however experiments on out-of-sample images, such as Recht et al. (2019) show similar behavior.

<sup>8</sup>Quantization ranges are chosen to optimize accuracy and are different from those in Section 7.1.



## REFERENCES

- Alham Fikri Aji and Kenneth Heafield. Sparse communication for distributed gradient descent. *arXiv preprint arXiv:1704.05021*, 2017.
- Stefano Ambrogio, Pritish Narayanan, Hsin-yu Tsai, Robert M. Shelby, Irem Boybat, Carmelo di Nolfo, Severin Sidler, Massimo Giordano, Martina Bodini, Nathan C. P. Farinha, Benjamin Killeen, Christina Cheng, Yassine Jaoudi, and Geoffrey W. Burr. Equivalent-accuracy accelerated neural-network training using analogue memory. *Nature*, 558(7708):60–67, June 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0180-5.
- Ron Banner, Itay Hubara, Elad Hoffer, and Daniel Soudry. Scalable methods for 8-bit training of neural networks. In *Advances in Neural Information Processing Systems*, pp. 5145–5153, 2018.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Frederik Benzing, Marcelo Matheus Gauy, Asier Mujika, Anders Martinsson, and Angelika Steger. Optimal Kronecker-sum approximation of real time recurrent learning. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 604–613, Long Beach, California, USA, 09–15 Jun 2019. PMLR. URL <http://proceedings.mlr.press/v97/benzing19a.html>.
- Åke Björck. Solving linear least squares problems by gram-schmidt orthogonalization. *BIT Numerical Mathematics*, 7(1):1–21, 1967.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- C. Chou, Z. Lin, P. Tseng, C. Li, C. Chang, W. Chen, Y. Chih, and T. J. Chang. An N40 256K×44 embedded RRAM macro with SL-precharge SA and low-voltage current limiter to improve read and write performance. In *2018 IEEE International Solid - State Circuits Conference - (ISSCC)*, pp. 478–480, February 2018. doi: 10.1109/ISSCC.2018.8310392.
- Alan Kaylor Cline and Inderjit S Dhillon. Computation of the singular value decomposition, 2006.
- J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- M. Ernestus. Elastic transformation of an image in python. <https://gist.github.com/erniejunior/601cdf56d2b424757de5>, 2016.
- Tayfun Gokmen and Yurii Vlasov. Acceleration of Deep Neural Network Training with Resistive Cross-Point Devices. *Frontiers in Neuroscience*, 10, July 2016. ISSN 1662-453X. doi: 10.3389/fnins.2016.00333.
- S. K. Gonugondla, M. Kang, and N. R. Shanbhag. A Variation-Tolerant In-Memory Machine Learning Classifier via On-Chip Training. *IEEE Journal of Solid-State Circuits*, 53(11):3163–3173, November 2018. doi: 10.1109/JSSC.2018.2867275.
- Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017.
- A. Grossi, E. Vianello, M. M. Sabry, M. Barlas, L. Grenouillet, J. Coignus, E. Beigne, T. Wu, B. Q. Le, M. K. Wootters, C. Zambelli, E. Nowak, and S. Mitra. Resistive ram endurance: Array-level characterization and correction techniques targeting deep learning applications. *IEEE Transactions on Electron Devices*, 66(3):1281–1288, March 2019. doi: 10.1109/TED.2019.2894387.
- S Haykin. *Adaptive filter theory, ser. always learning*, 2014.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pp. 1026–1034, 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.
- Alston S Householder. Unitary triangularization of a nonsymmetric matrix. *Journal of the ACM (JACM)*, 5(4):339–342, 1958.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Jakub Konečný, H. Brendan McMahan, Daniel Ramage, and Peter Richtárik. Federated optimization: Distributed machine learning for on-device intelligence. *CoRR*, abs/1610.02527, 2016. URL <http://arxiv.org/abs/1610.02527>.
- Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep gradient compression: Reducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887*, 2017.
- Fuyun Ling, Dimitris Manolakis, and John Proakis. A recursive modified gram-schmidt algorithm for least-squares estimation. *IEEE transactions on acoustics, speech, and signal processing*, 34(4): 829–836, 1986.
- Asier Mujika, Florian Meier, and Angelika Steger. Approximating real-time recurrent learning with random kronecker factors. In *Advances in Neural Information Processing Systems*, pp. 6594–6603, 2018.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- Hang Qi, Matthew Brown, and David G. Lowe. Low-Shot Learning with Imprinted Weights. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5822–5830, Salt Lake City, UT, June 2018. IEEE. ISBN 978-1-5386-6420-9. doi: 10.1109/CVPR.2018.00610.
- Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*, 2019.
- Jimmy SJ Ren and Li Xu. On vectorization of deep convolutional neural networks for vision tasks. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Shibani Santurkar, Dimitris Tsipras, Andrew Ilyas, and Aleksander Madry. How does batch normalization help optimization? In *Advances in Neural Information Processing Systems*, pp. 2483–2493, 2018.
- Patrice Y. Simard, Dave Steinkraus, and John Platt. Best practices for convolutional neural networks applied to visual document analysis. Institute of Electrical and Electronics Engineers, Inc., August 2003. URL <https://www.microsoft.com/en-us/research/publication/best-practices-for-convolutional-neural-networks-applied-to-visual-document-analysis/>.
- D. Soudry, D. Di Castro, A. Gal, A. Kolodny, and S. Kvatinsky. Memristor-Based Multilayer Neural Networks With Online Gradient Descent Training. *IEEE Transactions on Neural Networks and Learning Systems*, 26(10):2408–2421, October 2015. doi: 10.1109/TNNLS.2014.2383395.
- Corentin Tallec and Yann Ollivier. Unbiased online recurrent optimization. *arXiv preprint arXiv:1702.05043*, 2017.

- TSMC. 40nm Technology - Taiwan Semiconductor Manufacturing Company Limited. <https://www.tsmc.com/english/dedicatedFoundry/technology/40nm.htm>, 2019.
- user1551. Rotation matrix in arbitrary dimension to align vector. Mathematics Stack Exchange, 2013. URL <https://math.stackexchange.com/q/525587>. URL:<https://math.stackexchange.com/q/525587> (version: 2013-10-14).
- Ronald J Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280, 1989.
- T. F. Wu, B. Q. Le, R. Radway, A. Bartolo, W. Hwang, S. Jeong, H. Li, P. Tandon, E. Vianello, P. Vivet, E. Nowak, M. K. Wootters, H. . P. Wong, M. M. S. Aly, E. Beigne, and S. Mitra. 14.3 a 43pj/cycle non-volatile microcontroller with 4.7s shutdown/wake-up integrating 2.3-bit/cell resistive ram and resilience techniques. In *2019 IEEE International Solid- State Circuits Conference - (ISSCC)*, pp. 226–228, Feb 2019. doi: 10.1109/ISSCC.2019.8662402.
- S. Yu, Z. Li, P. Chen, H. Wu, B. Gao, D. Wang, W. Wu, and H. Qian. Binary neural network with 16 Mb RRAM macro chip for classification and online training. In *2016 IEEE International Electron Devices Meeting (IEDM)*, pp. 16.2.1–16.2.4, December 2016. doi: 10.1109/IEDM.2016.7838429.
- Shimeng Yu. Neuro-inspired computing with emerging nonvolatile memorys. *Proceedings of the IEEE*, 106(2):260–285, 2018.
- E. Zamanidoost, F. M. Bayat, D. Strukov, and I. Kataeva. Manhattan rule training for memristive crossbar circuit pattern classifiers. In *2015 IEEE 9th International Symposium on Intelligent Signal Processing (WISP) Proceedings*, pp. 1–6, May 2015. doi: 10.1109/WISP.2015.7139171.
- J. Zhang, Z. Wang, and N. Verma. In-Memory Computation of a Machine-Learning Classifier in a Standard 6T SRAM Array. *IEEE Journal of Solid-State Circuits*, 52(4):915–924, April 2017. ISSN 0018-9200. doi: 10.1109/JSSC.2016.2642198.
- 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.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 928–936, 2003.

## A SKS ALGORITHM

**Algorithm 1** Streaming Kronecker Sum Approximation

---

**State:**  $\mathbf{Q}_L \in \mathbb{R}^{n_o \times q}$ ;  $\mathbf{Q}_R \in \mathbb{R}^{n_i \times q}$ ;  $\mathbf{c}_x \in \mathbb{R}^{q \times 1}$   
**Input:**  $\mathbf{dz}^{(i)} \in \mathbb{R}^{n_o \times 1}$ ;  $\mathbf{a}^{(i)} \in \mathbb{R}^{n_i \times 1}$  for  $i \in [1, B]$   
**for**  $i = 1 \dots B$  **do**  
  {Modified Gram-Schmidt.}  
   $\mathbf{c}_L, \mathbf{c}_R \leftarrow 0^{q \times 1}$   
  **for**  $j = 1 \dots r$  **do**  
     $\mathbf{c}_{L,j} \leftarrow \mathbf{Q}_{L,j} \cdot \mathbf{dz}^{(i)}$ ;     $\mathbf{dz}^{(i)} \leftarrow \mathbf{dz}^{(i)} - \mathbf{c}_{L,j} \cdot \mathbf{Q}_{L,j}$   
     $\mathbf{c}_{R,j} \leftarrow \mathbf{Q}_{R,j} \cdot \mathbf{a}^{(i)}$ ;     $\mathbf{a}^{(i)} \leftarrow \mathbf{a}^{(i)} - \mathbf{c}_{L,j} \cdot \mathbf{Q}_{L,j}$   
  **end for**  
   $c_{L,q} \leftarrow \|\mathbf{dz}^{(i)}\|$ ;     $\mathbf{Q}_{L,q} \leftarrow \mathbf{dz}^{(i)} / c_{L,q}$   
   $c_{R,q} \leftarrow \|\mathbf{a}^{(i)}\|$ ;     $\mathbf{Q}_{R,q} \leftarrow \mathbf{a}^{(i)} / c_{R,q}$   
  {Generate  $\mathbf{C}$  and find its SVD.}  
   $\mathbf{C} \leftarrow \mathbf{c}_L \mathbf{c}_R^\top + \text{diag}(\mathbf{c}_x)$   
   $\mathbf{U}_C \cdot \text{diag}(\boldsymbol{\sigma}) \cdot \mathbf{V}_C^\top \leftarrow \text{SVD}(\mathbf{C})$   
  {Minimum-variance unbiased estimator for  $\boldsymbol{\Sigma}$ .}  
   $m \leftarrow \min j$  s.t.  $(q-j)\sigma_j \leq \sum_{\ell=j}^q \sigma_\ell$   
   $s_1 \leftarrow \sum_{i=m}^q \sigma_i$ ,     $k \leftarrow q - m$   
   $\mathbf{v} \leftarrow \sqrt{1 - k/s_1} \cdot \boldsymbol{\sigma}_{[m:]} - \mathbf{e}^{(1)}$   
   $\mathbf{s} \leftarrow \{-1, 1\}^{(k+1) \times 1}$  {Ind. uniform random signs.}  
   $\mathbf{X}_s \leftarrow (\mathbf{I} + (\mathbf{s} \odot \mathbf{v})(\mathbf{v}/v_1)^\top)_{[2:]}$  {Householder.}  
  {QR-factorization of  $\tilde{\boldsymbol{\Sigma}}_L$ .}  
   $\mathbf{Q}_x \leftarrow \begin{bmatrix} \mathbf{I} & 0 \\ 0 & \mathbf{X}_s \end{bmatrix} \in \mathbb{R}^{q \times r}$   
   $\mathbf{c}_x \leftarrow (\sigma_1, \dots, \sigma_{m-1}, \underbrace{s_1/k, \dots, s_1/k}_{q-m+1 \text{ times}})$   
  {Update the first  $r$  columns of  $\mathbf{Q}_L, \mathbf{Q}_R$ .}  
   $\mathbf{Q}_{L[:r]} \leftarrow \mathbf{Q}_L \cdot \mathbf{U}_C \cdot \mathbf{Q}_x$   
   $\mathbf{Q}_{R[:r]} \leftarrow \mathbf{Q}_R \cdot \mathbf{V}_C \cdot \mathbf{Q}_x$   
**end for**  
  {Compute final  $\tilde{\mathbf{L}}, \tilde{\mathbf{R}}$  where  $\nabla_{\mathbf{w}} \mathcal{L} \approx \tilde{\mathbf{L}} \tilde{\mathbf{R}}^\top$ .}  
   $\tilde{\mathbf{L}} \leftarrow (\mathbf{Q}_L \cdot \text{diag}(\sqrt{\mathbf{c}_x}))_{[:r]}$   
   $\tilde{\mathbf{R}} \leftarrow (\mathbf{Q}_R \cdot \text{diag}(\sqrt{\mathbf{c}_x}))_{[:r]}$

---

## B CONVEX CONVERGENCE

In this section we will attempt to bound the regret (defined below) of an SGD algorithm using noisy SKS estimates  $\tilde{\mathbf{g}} = \mathbf{g} + \boldsymbol{\varepsilon}$  in the convex setting, where  $\mathbf{g}$  are the true gradients and  $\boldsymbol{\varepsilon}$  are the errors introduced by the low rank SKS approximation. Here,  $\mathbf{g}$  is a vector of size  $N$  and can be thought of as a flattened/concatenated version of the gradient tensors (e.g.,  $N = n_i \cdot n_o$ ).

Our proof follows the proof in Zinkevich (2003). We define  $\mathbb{F}$  as the convex feasible set (valid settings for our weight tensors) and assume that  $\mathbb{F}$  is bounded with  $D = \max_{\mathbf{w}, \mathbf{v} \in \mathbb{F}} \|\mathbf{w} - \mathbf{v}\|$  being the maximum distance between two elements of  $\mathbb{F}$ . Further, assume a batch  $t$  of  $B$  samples out of  $T$  total batches corresponds to a loss landscape  $f^t(\mathbf{w}^t)$  that is strongly convex in weight parameters  $\mathbf{w}^t$ , so there are positive constants  $C \geq c > 0$  such that  $c\mathbf{I} \preceq \nabla^2 f^t(\mathbf{w}^t) \preceq C\mathbf{I}$  for all  $t$  (Boyd & Vandenberghe, 2004, Section 9.3). We define regret as  $R(T) = \sum_{t=1}^T f^t(\mathbf{w}^t) - \sum_{t=1}^T f^t(\mathbf{w}^*)$  where  $\mathbf{w}^* = \text{argmin}_{\mathbf{w}} \sum_{t=1}^T f^t(\mathbf{w})$  (i.e., it is an optimal offline minimizer of  $f^1, \dots, f^T$ ).

The gradients seen during SGD are  $\mathbf{g}^t = \nabla f^t(\mathbf{w}^t)$  and we assume they are bounded by  $G = \max_{\mathbf{w} \in \mathbb{F}, t \in [1, T]} \|\nabla f^t(\mathbf{w})\|$ . We also assume errors are bounded by  $\mathcal{E} = \max_{t \in [1, T]} \|\varepsilon^t\|$ . Therefore,  $\max_{t \in [1, T]} \|\tilde{\mathbf{g}}^t\| \leq \max_{t \in [1, T]} \|\mathbf{g}^t\| + \|\varepsilon^t\| \leq G + \mathcal{E}$  by the triangle inequality.

**Theorem 1.** *Assume SKS-based SGD is applied with learning rate  $\eta_t = 1/\sqrt{t}$ . Then, under the additional constraint  $\mathbf{g}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) - \frac{\varepsilon}{2} \|\mathbf{w}^t - \mathbf{w}^*\|_2^2 \leq \tilde{\mathbf{g}}^t \cdot (\mathbf{w}^t - \mathbf{w}^*)$ , we have sublinear regret:*

$$R(T) \leq \frac{D^2}{2} \sqrt{T} + (G + \mathcal{E})^2 (\sqrt{T} - 1/2)$$

*Proof.* From strong convexity  $c\mathbf{I} \preceq \nabla^2 f^t(\mathbf{w}^t) \preceq C\mathbf{I}$  for all  $t$ ,

$$f^t(\mathbf{w}) + \mathbf{g}^t \cdot (\mathbf{v} - \mathbf{w}) + \frac{c}{2} \|\mathbf{v} - \mathbf{w}\|_2^2 \leq f^t(\mathbf{v}) \quad \text{for all } \mathbf{v} \quad (8)$$

In particular, if we consider  $\mathbf{v} = \mathbf{w}^*$  and rearrange,

$$\begin{aligned} f^t(\mathbf{w}) - f^t(\mathbf{w}^*) &\leq \mathbf{g}^t \cdot (\mathbf{w} - \mathbf{w}^*) - \frac{c}{2} \|\mathbf{w} - \mathbf{w}^*\|_2^2 \\ f^t(\mathbf{w}) - f^t(\mathbf{w}^*) &\leq \tilde{\mathbf{g}}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) \end{aligned} \quad (9)$$

Consider a gradient update  $\mathbf{w}^{t+1} = P_{\mathbb{F}}(\mathbf{w}^t - \eta_t \tilde{\mathbf{g}}^t)$ , where  $P_{\mathbb{F}}$  projects the update back to  $\mathbb{F}$ . Then,

$$\begin{aligned} \|\mathbf{w}^{t+1} - \mathbf{w}^*\|_2^2 &= \|P(\mathbf{w}^t - \eta_t \tilde{\mathbf{g}}^t) - \mathbf{w}^*\|_2^2 \leq \|\mathbf{w}^t - \eta_t \tilde{\mathbf{g}}^t - \mathbf{w}^*\|_2^2 \\ &= \|\mathbf{w}^t - \mathbf{w}^*\|_2^2 - 2\eta_t (\mathbf{w}^t - \mathbf{w}^*) \cdot \tilde{\mathbf{g}}^t + \eta_t^2 \|\tilde{\mathbf{g}}^t\|_2^2 \\ &\leq \|\mathbf{w}^t - \mathbf{w}^*\|_2^2 - 2\eta_t (\mathbf{w}^t - \mathbf{w}^*) \cdot \tilde{\mathbf{g}}^t + \eta_t^2 (G + \mathcal{E})^2 \\ \tilde{\mathbf{g}}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) &\leq \frac{1}{2\eta_t} (\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{t+1} - \mathbf{w}^*\|_2^2) + \frac{\eta_t}{2} (G + \mathcal{E})^2 \end{aligned} \quad (10)$$

From (9, 10),

$$f^t(\mathbf{w}^t) - f^t(\mathbf{w}^*) \leq \frac{1}{2\eta_t} (\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{t+1} - \mathbf{w}^*\|_2^2) + \frac{\eta_t}{2} (G + \mathcal{E})^2 \quad (11)$$

We now bound the regret:

$$\begin{aligned} R(T) &= \sum_{t=1}^T [f^t(\mathbf{w}^t) - f^t(\mathbf{w}^*)] \\ &\leq \sum_{t=1}^T \left[ \frac{1}{2\eta_t} (\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 - \|\mathbf{w}^{t+1} - \mathbf{w}^*\|_2^2) + \frac{\eta_t}{2} (G + \mathcal{E})^2 \right] \\ &= \frac{\|\mathbf{w}^1 - \mathbf{w}^*\|_2^2}{2\eta_1} - \frac{\|\mathbf{w}^{T+1} - \mathbf{w}^*\|_2^2}{2\eta_T} + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{w}^t - \mathbf{w}^*\|_2^2 + \frac{(G + \mathcal{E})^2}{2} \sum_{t=1}^T \eta_t \\ &\leq \frac{\|\mathbf{w}^1 - \mathbf{w}^*\|_2^2}{2\eta_1} + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) \|\mathbf{w}^t - \mathbf{w}^*\|_2^2 + \frac{(G + \mathcal{E})^2}{2} \sum_{t=1}^T \eta_t \\ &\leq \frac{D^2}{2\eta_1} + \frac{1}{2} \sum_{t=2}^T \left( \frac{1}{\eta_t} - \frac{1}{\eta_{t-1}} \right) D^2 + \frac{(G + \mathcal{E})^2}{2} \sum_{t=1}^T \eta_t \\ &= \frac{D^2}{2\eta_T} + \frac{(G + \mathcal{E})^2}{2} \sum_{t=1}^T \eta_t \end{aligned} \quad (12)$$

If  $\eta_t = 1/\sqrt{t}$ , then  $\sum_{t=1}^T \eta_t \leq 2\sqrt{T} - 1$  (Zinkevich, 2003), so from (12),

$$R(T) \leq \frac{D^2}{2}\sqrt{T} + (G + \varepsilon)^2 \left(\sqrt{T} - 1/2\right) \quad (13)$$

□

This is a sublinear regret and therefore, average regret  $R(T)/T$  is bounded above by 0 in the limit as  $T \rightarrow \infty$ . To achieve this result, we constrained  $\mathbf{g}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) - \frac{c}{2}\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 \leq \tilde{\mathbf{g}}^t \cdot (\mathbf{w}^t - \mathbf{w}^*)$ . We now examine sufficient conditions for this inequality to be satisfied.

$$\begin{aligned} \mathbf{g}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) - \frac{c}{2}\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 &\leq \tilde{\mathbf{g}}^t \cdot (\mathbf{w}^t - \mathbf{w}^*) \\ \varepsilon^t \cdot (\mathbf{w}^* - \mathbf{w}^t) &\leq \frac{c}{2}\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 \end{aligned} \quad (14)$$

Since  $\varepsilon^t \cdot (\mathbf{w}^* - \mathbf{w}^t) \leq \|\varepsilon^t\| \cdot \|\mathbf{w}^* - \mathbf{w}^t\|$  by Cauchy-Schwarz, it is sufficient for:

$$\begin{aligned} \|\varepsilon^t\| \cdot \|\mathbf{w}^* - \mathbf{w}^t\| &\leq \frac{c}{2}\|\mathbf{w}^t - \mathbf{w}^*\|_2^2 \\ \|\varepsilon^t\| &\leq \frac{c}{2}\|\mathbf{w}^t - \mathbf{w}^*\| \end{aligned} \quad (15)$$

## B.1 CONSIDERATIONS FOR RANK DEFICIENT HESSIANS

In the preceding proof, we assumed  $c > 0$ . However, it is common for this to not hold. For example, in linear regression, where  $c = \lambda_{\min}(\mathbf{X}\mathbf{X}^\top)$  for sample input  $\mathbf{X} \in \mathbb{R}^{(n_i \times B)}$  (Haykin, 2014, Chapter 4.3), if  $B < n_i$  then  $c = 0$ . We can modify (15) to handle this case. Let  $\tilde{c}$  be the minimum non-zero Eigenvalue of  $\mathbf{X}\mathbf{X}^\top$  and let  $\tilde{\mathbf{w}} \in \mathbb{R}^B$  represent  $\mathbf{w} \in \mathbb{R}^N$  in the Eigenbasis of  $\mathbf{X}\mathbf{X}^\top$ . Then (15) becomes:

$$\|\varepsilon^t\| \leq \frac{\tilde{c}}{2}\|\tilde{\mathbf{w}}^t - \tilde{\mathbf{w}}^*\| \quad (16)$$

## B.2 ESTIMATES FOR THE SKS ERROR

We can estimate the SKS error  $\|\varepsilon^t\|$  in both the biased and unbiased cases. For the **biased, zero-variance case**, we get rid of the lowest singular value (out of  $q$  singular values) as we see each sample. Thus, at a given sample  $i$ , the error is  $\sigma_q^{(i)} \cdot \mathbf{Q}_{L,q}^{(i)} \cdot \mathbf{Q}_{R,q}^{(i)}$  and the average squared error is  $(1/N)\sigma_q^{(i)2}$ . We can treat this as a per-element variance. If these smallest singular components are uncorrelated from sample to sample, then the variances add:

$$\sigma_\varepsilon^2 \approx \frac{1}{N} \sum_{i=1}^B \left(\sigma_q^{(i)}\right)^2 \quad (17)$$

For the **unbiased, minimum-variance case**, Theorem A.4 from Benzing et al. (2019) states that the minimum variance is  $s_1^2/k + s_2$  where  $s_1 = \sum_{i=m}^q \sigma_i$ ,  $s_2 = \sum_{i=m}^q \sigma_i^2$ , and  $k, m$  are as defined in Section 4.1.2. Since  $m$  is chosen to minimize variance, we can upper bound the variance by choosing  $m = r$  and therefore  $k = 1$ ,  $s_1 = \sigma_r + \sigma_q$ , and  $s_2 = \sigma_r^2 + \sigma_q^2$ . Empirically, this tends to be a good approximation. Then, the average per-element variance added at sample  $i$  is approximately  $(2/N) \left(\sigma_r^{(i)} \sigma_q^{(i)}\right)$ . Assuming errors between samples are uncorrelated, this leads to a total variance:

$$\sigma_\varepsilon^2 \approx \frac{2}{N} \sum_{i=1}^B \sigma_r^{(i)} \sigma_q^{(i)} \quad (18)$$

For either case,  $\|\varepsilon\|^2 \approx N\sigma_\varepsilon^2$ . For the  $t$ -th batch and  $i$ -th sample, we denote  $\sigma_q^{(t,i)}$  as the  $q$ -th singular value. For simplicity, we focus on the biased, zero-variance case (the unbiased case is similar). From (15), an approximately sufficient condition for sublinear-regret convergence is:

$$\sum_{i=1}^B \left(\sigma_q^{(t,i)}\right)^2 \leq \frac{c^2}{4} \|\mathbf{w}^t - \mathbf{w}^*\|^2 \quad (19)$$

### B.3 DISCUSSION ON CONVERGENCE

Equation (19) suggests that as  $\mathbf{w}^t \rightarrow \mathbf{w}^*$ , the constraints for achieving sublinear-regret convergence become more difficult to maintain. However, in practice this may be highly problem-dependent as the  $\sigma_q$  will also tend to decrease near optimal solutions. To get a better sense of the behavior of the left-hand side of (19), suppose that:

$$\sum_{i=1}^B \left(\sigma_q^{(t,i)}\right)^2 \approx \sum_{i=q}^B \left(\sigma_i(\mathbf{G}^t)\right)^2 \leq \sum_{i=1}^B \left(\sigma_i(\mathbf{G}^t)\right)^2 = \|\mathbf{G}^t\|_F^2$$

where  $\mathbf{G}^t = \nabla_{\mathbf{W}^t} f^t(\mathbf{W}^t) \in \mathbb{R}^{(n_o \times n_i)}$  are the matrix weight  $\mathbf{W}^t$  gradients at batch  $t$  and  $\|\cdot\|_F$  is a Frobenius norm. We therefore expect both the left (proportional to  $\|\mathbf{G}^t\|_F^2$ ) and the right (proportional to  $\|\mathbf{w}^t - \mathbf{w}^*\|^2$ ) of (19) to decrease during training as  $\mathbf{w}^t \rightarrow \mathbf{w}^*$ . This behavior is in fact what is seen in Figure 1(b). If achieving convergence is found to be difficult, (19) provides some insight for convergence improvement methods.

One solution is to reduce batch size  $B$  to satisfy the inequality as necessary. This minimizes the weight updates during more repetitive parts of training while allowing dense weight updates (possibly approaching standard SGD with small batch sizes) during more challenging parts of training.

Another solution is to reduce  $\sigma_q$ . One way to do this is to increase the rank  $r$  so that the spectral energy of the updates are spread across more singular components. There may be alternate approaches based on conditioning the inputs to shape the distribution of singular values in a beneficial way.

A third method is to focus on  $c$ , the lower bound on curvature of the convex loss functions. Perhaps a technique such as weight regularization can increase  $c$  by adding constant curvature in all Eigen-directions of the loss function Hessian (although this may also increase the LHS of (19)). Alternatively, perhaps low-curvature Eigen-directions are less important for loss minimization, allowing us to raise the  $c$  that we effectively care about. This latter approach requires no particular action on our part, except the recognition that fast convergence may only be guaranteed for high-curvature directions. This is exemplified in Figure 1(b), where we can see SVD track the curve for  $C$  more so than  $c$ .

Finally, we note that this analysis focuses solely on the errors introduced by a floating-point version of SKS. Quantization noise can add additional error into the  $\varepsilon^t$  term. We expect this to add a constant offset to the LHS of (19). For a weight LSB  $\Delta$ , quantization noise has variance  $\Delta^2/12$ , so we desire:

$$N \frac{\Delta^2}{12} + \sum_{i=1}^B \left(\sigma_q^{(t,i)}\right)^2 \leq \frac{c^2}{4} \|\mathbf{w}^t - \mathbf{w}^*\|^2 \quad (20)$$

## C KRONECKER SUMS IN NEURAL NETWORK LAYERS

### C.1 DENSE LAYER

A dense or fully-connected layer transforms an input  $\mathbf{a} \in \mathbb{R}^{n_i \times 1}$  to an intermediate  $\mathbf{z} = \mathbf{W} \cdot \mathbf{a} + \mathbf{b}$  to an output  $\mathbf{y} = \sigma(\mathbf{z}) \in \mathbb{R}^{n_o \times 1}$  where  $\sigma$  is a non-linear activation function. Gradients of the loss function with respect to the weight parameters can be found as:

$$\nabla_{\mathbf{W}} \mathcal{L} = \underbrace{(\nabla_{\mathbf{z}} \mathcal{L})}_{d\mathbf{z}} \odot \underbrace{(\nabla_{\mathbf{W}} \mathbf{z})}_{\mathbf{a}^\top} = d\mathbf{z} \otimes \mathbf{a} \quad (21)$$

which is exactly the per-sample Kronecker sum update we saw in linear regression. Thus, at every training sample, we can add  $(d\mathbf{z}^{(i)} \otimes \mathbf{a}^{(i)})$  to our low rank estimate with SKS.

## C.2 CONVOLUTIONAL LAYER

A convolutional layer transforms an input feature map  $\mathbf{A} \in \mathbb{R}^{h_{in} \times w_{in} \times c_{in}}$  to an intermediate feature map  $\mathbf{Z} = \mathbf{W}_{kernel} * \mathbf{A} + \mathbf{b} \in \mathbb{R}^{h_{out} \times w_{out} \times c_{out}}$  through a 2D convolution  $*$  with weight kernel  $\mathbf{W}_{kernel} \in \mathbb{R}^{c_{out} \times k_h \times k_w \times c_{in}}$ . Then it computes an output feature map  $y = \sigma(z)$  where  $\sigma$  is a non-linear activation function.

Convolutions can be interpreted as matrix multiplications through the `im2col` operation which converts the input feature map  $\mathbf{A}$  into a matrix  $\mathbf{A}_{col} \in \mathbb{R}^{(h_{out}w_{out}) \times (k_h k_w c_{in})}$  where the  $i^{\text{th}}$  row is a flattened version of the sub-tensor of  $\mathbf{A}$  which is dotted with  $\mathbf{W}_{kernel}$  to produce the  $i^{\text{th}}$  pixel of the output feature map (Ren & Xu, 2015). We can multiply  $\mathbf{A}_{col}$  by a flattened version of the kernel,  $\mathbf{W} \in \mathbb{R}^{c_{out} \times (k_h k_w c_{in})}$  to perform the  $\mathbf{W}_{kernel} * \mathbf{A}$  convolution operation with a matrix multiplication. Under the matrix multiplication interpretation, weight gradients can be represented as:

$$\nabla_{\mathbf{W}} \mathcal{L} = \underbrace{(\nabla_{\mathbf{Z}_{col}} \mathcal{L})}_{d\mathbf{Z}_{col}^{\top}} \odot \underbrace{(\nabla_{\mathbf{W}} \mathbf{Z})}_{\mathbf{A}_{col}} = \sum_{i=1}^{h_{out}w_{out}} d\mathbf{Z}_{col,i}^{\top} \otimes \mathbf{A}_{col,i}^{\top} \quad (22)$$

which is the same as  $h_{out}w_{out}$  Kronecker sum updates. Thus, at every output pixel  $j$  of every training sample  $i$ , we can add  $(d\mathbf{Z}_{col,j}^{(i)\top} \otimes \mathbf{A}_{col,j}^{(i)\top})$  to our low rank estimate with SKS.

Note that while we already save an impressive factor of  $B/q$  in memory when computing gradients for the dense layer, we save a much larger factor of  $Bh_{out}w_{out}/q$  in memory when computing gradients for the convolution layers, making the low rank training technique even more crucial here.

However, some care must be taken when considering activation memory for convolutions. For compute-constrained edge devices, image dimensions may be small and result in minimal intermediate feature map memory requirements. However, if image dimensions grow substantially, activation memory could dominate compared to weight storage. Clever dataflow strategies may provide a way to reduce intermediate activation storage even when performing backpropagation<sup>9</sup>.

## D HARDWARE QUANTIZATION MODEL

In a real device, operations are expected to be performed in fixed point arithmetic. Therefore, all of our training experiments are conducted with quantization in the loop. Our model for quantization is shown in Figure 3. The green arrows describe the forward computation. Ignoring quantization for a moment, we would have  $\mathbf{a}^{\ell} = \text{ReLU}(\alpha^{\ell} \mathbf{W}^{\ell} * \mathbf{a}^{\ell-1} + \mathbf{b}^{\ell})$ , where  $*$  can represent either a convolution or a matrix multiply depending on the layer type and  $\alpha^{\ell}$  is the closest power-of-2 to He initialization (He et al., 2015). For quantization, we rely on four basic quantizers:  $Q_w, Q_b, Q_a, Q_g$ , which describe weight quantization, bias and intermediate accumulator quantization, activation quantization, and gradient quantization, respectively. All quantizers use fixed clipping ranges as depicted and quantize uniformly within those ranges to the specified bitwidths.

In the backward pass, follow the orange arrows from  $\delta^{\ell}$ . Backpropagation follows standard backpropagation rules including using the straight-through estimator (Bengio et al., 2013) for quantizer gradients. However, because we want to perform training on edge devices, these gradients must themselves be quantized. The first place this happens is after passing backward through the ReLU derivative. The other two places are before feeding back into the network parameters  $\mathbf{W}^{\ell}, \mathbf{b}^{\ell}$ , so that  $\mathbf{W}^{\ell}, \mathbf{b}^{\ell}$  cannot be used to accumulate values smaller than their LSB. Finally, instead of deriving  $\Delta \mathbf{W}^{\ell}$  from a backward pass through the  $*$  operator, the SKS method is used.

<sup>9</sup>For example, one could compute just a sliding window of rows of every feature map, discarding earlier rows as later rows are computed, resulting in a square-root reduction of activation memory. To incorporate backpropagation, compute the forward pass once fully, then compute the forward pass again, as well as the backward pass using the sliding window approach in both directions.



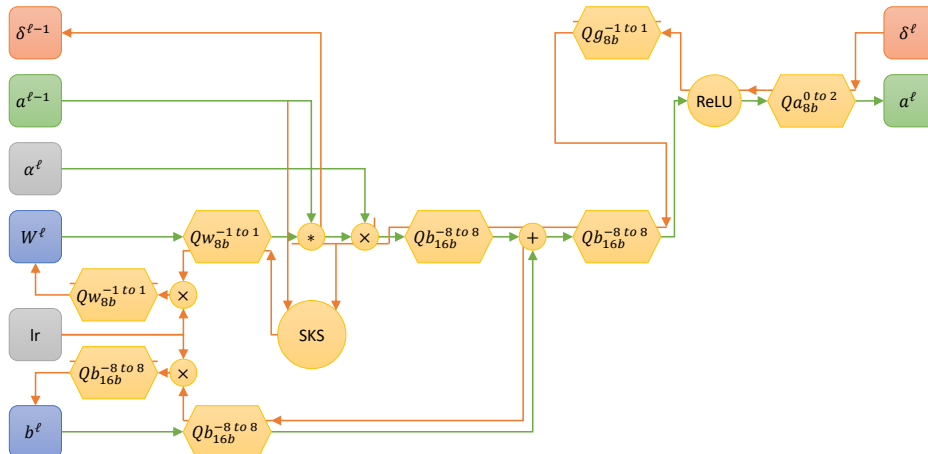


Figure 3: Signal flow graph for a forward and backward quantized convolutional or dense layer.

SKS collects  $\alpha^{l-1}$ ,  $\mathbf{dz}^l$  for many samples before computing the approximate  $\Delta\tilde{\mathbf{W}}^l$ . It accumulates information in two low rank matrices  $\mathbf{L}$ ,  $\mathbf{R}$  which are themselves quantized to 16 bits with clipping ranges determined dynamically by the max absolute value of elements in each matrix. While SKS accumulates for  $B$  samples, leading to a factor of  $B$  reduction in the rate of updates to  $\mathbf{W}^l$ ,  $\mathbf{b}^l$  is updated at every sample. This is feasible in hardware because  $\mathbf{b}^l$  is small enough to be stored in more expensive forms of memory that have superior endurance and write power performance.

Because of the coarse weight LSB size, weight gradients may be consistently quantized to 0, preventing them from accumulating. To combat this, we only apply an update if a minimum update density  $\rho_{\min} = 0.01$  would be achieved, otherwise we continue accumulating samples in  $\mathbf{L}$  and  $\mathbf{R}$ , which have much higher bitwidths. When an update does finally happen, the “effective batch size” will be a multiple of  $B$  and we increase the learning rate correspondingly. In the literature, a linear scaling rule is suggested (see Goyal et al. (2017)), however we empirically find square-root scaling works better (see Appendix H).

### E GRADIENT MAX-NORMING




Figure 4: Maximum magnitude of weight gradients versus training step for standard SGD on a CNN trained on MNIST.

Figure 4 plots the magnitude of gradients seen in a weight tensor over training steps. One apparent property of these gradients is that they have a large dynamic range, making them difficult to quantize. Even when looking at just the spikes, they assume a wide range of magnitudes. One potential method of dealing with this dynamic range is to scale tensors so that their max absolute element is 1 (similar to a per-tensor AdaMax (Kingma & Ba, 2014) or Range Batch-Norm (Banner et al., 2018) applied to gradients). Optimizers such as Adam, which normalize by gradient variance, provide a justification for why this sort of scaling might work well, although they work at a per-element

rather than per-tensor level. We choose max-norming rather than variance-based norming because the former is easier computational and potentially more ammenable to quantization. However, a problem with the approach of normalizing tensors independently at each sample is that noise might be magnified during regions of quiet as seen in the Figure. What we therefore propose is normalization by the maximum of both the current max element and a moving average of the max element.

Explicitly, max-norm takes two parameters - a decay factor  $\beta = 0.999$  and a gradient floor  $\varepsilon = 10^{-4}$  and keeps two state variables - the number of evaluations  $k := 0$  and the current maximum moving average  $x_{mv} := \varepsilon$ . Then for a given input  $x$ , max-norm modifies its internal state and returns  $x_{norm}$ :

$$\begin{aligned} k &:= k + 1 \\ x_{max} &:= \max(|\mathbf{x}|) + \varepsilon \\ x_{mv} &:= \beta \cdot x_{mv} + (1 - \beta) \cdot x_{max} \\ \tilde{x}_{mv} &:= \frac{x_{mv}}{1 - \beta^k} \\ \mathbf{x}_{norm} &:= \frac{\mathbf{x}}{\max(x_{max}, \tilde{x}_{mv})} \end{aligned}$$

## F STREAMING BATCH NORMALIZATION

Standard batch normalization (Ioffe & Szegedy, 2015) normalizes a tensor  $\mathbf{X}$  along some axes, then applies a trainable affine transformation. For each slice  $\mathbf{X}$  of  $\mathbf{X}$  that is normalized independently:

$$\mathbf{Y} = \gamma \cdot \frac{\mathbf{X} - \mu_b}{\sqrt{\sigma_b^2 + \varepsilon}} + \beta$$

where  $\mu_b, \sigma_b$  are mean and standard deviation statistics of a minibatch and  $\gamma, \beta$  are trainable affine transformation parameters.

In our case, we do not have the memory to hold a batch of samples at a time and must compute  $\mu_b, \sigma_b$  in an online fashion. To see how this works, suppose we knew the statistics of each sample  $\mu_i, \sigma_i$  for  $i = 1 \dots B$  in a batch of  $B$  samples. For simplicity, assume the  $i^{\text{th}}$  sample is a vector  $\mathbf{X}_{i,:} \in \mathbb{R}^n$  containing elements  $X_{i,j}$ . Then:

$$\mu_b = \frac{1}{B} \sum_{i=1}^B \mu_i \tag{23}$$

$$\sigma_b^2 = \frac{1}{B} \sum_{i=1}^B \frac{1}{n} \sum_{j=1}^n X_{i,j}^2 - \mu_b^2 = \frac{1}{B} \sum_{i=1}^B (\sigma_i^2 + \mu_i^2) - \mu_b^2 \neq \frac{1}{B} \sum_{i=1}^B \sigma_i^2 \tag{24}$$

In other words, the batch variance is not equal to the average of the sample variances. However, if we keep track of the sum-of-square values of samples  $\sigma_i^2 + \mu_i^2$ , then we can compute  $\sigma_b^2$  as in (24). We keep track of two state variables:  $\mu_s, sq_s$  which we update as  $\mu_s := \mu_s + \mu_i$  and  $sq_s := sq_s + \sigma_i^2 + \mu_i^2$  for each sample  $i$ . After  $B$  samples, we divide both state variables by  $B$  and apply (23, 24) to get the desired batch statistics. Unfortunately, in an online setting, all samples prior to the last one in a given batch will only see statistics generated from a portion of the batch, resulting in noisier estimates of  $\mu_b, \sigma_b$ .

In *streaming* batch norm, we alter the above formula slightly. Notice that in online training, only the most recently viewed sample is used for training, so there is no reason to weight different samples of a given batch equally. Therefore we can use an exponential moving average instead of a true average to track  $\mu_s, sq_s$ . Specifically, let:

$$\begin{aligned} \mu_s &:= \eta \cdot \mu_s + (1 - \eta) \cdot \mu_i \\ sq_s &:= \eta \cdot sq_s + (1 - \eta) \cdot (\sigma_i^2 + \mu_i^2) \end{aligned}$$

If we set  $\eta = 1 - 1/B$ , a weighting of  $1/B$  is seen on the current sample, just as in standard averages with a batch of size  $B$ , but now all samples receive similarly clean batch statistic estimates, not just the last few samples in a batch.

### G ONLINE DATASET

For our experiments, we construct a dataset comprising an offline training, validation, and test set, as well as an online training set. Specifically, we start with the standard MNIST dataset of LeCun et al. (1998) and split the 60k training images into partitions of size 9k, 1k, and 50k. Elastic transforms (Simard et al., 2003; Ernestus, 2016) are used to augment each of these partitions to 50k offline training samples, 10k offline validation samples, and 100k online training samples, respectively. Elastic transforms are also applied to the 10k MNIST test images to generate the offline test samples.

The source images for the 100k online training samples are randomly drawn with replacement, so there is a certain amount of data leakage in that an online algorithm may be graded on an image that has been generated from the same image a previous sample it has trained on has been generated from. This is intentional and is meant to mimic a real-life scenario where a deployed device is likely to see a restrictive and repetitive set of training samples. Our experiments include comparisons to standard SGD to show that SKS’s improvement is not merely due to overfitting the source images.

From the online training set, we also generate a “distribution shift” dataset by applying unique additional augmentations to every contiguous 10k samples of the 100k online training samples. Four types of augmentations are explored. Class distribution clustering biases training samples belonging to similar classes to have similar indices. For example, the first thousand images may be primarily “0”s and “3”s, whereas the next thousand might have many “5”s. Spatial transforms rotate, scale, and shift images by random amounts. Background gradients both scale the contrast of the images and apply black-white gradients across the image. Finally, white noise is random Gaussian noise added to each pixel. Figure 5 shows some representative examples of what these augmentations look like. The augmentations are meant to mimic different external environments an edge devices might need to adapt to.




Figure 5: Samples of different types of distribution shift augmentations.

In addition to distribution shift for testing adaptation, we also look at internal statistical shift of weights in two ways - analog and digital. For analog weight drift, we apply independent additive Gaussian noise to each weight every  $d = 10$  steps with  $\sigma = \sigma_0/\sqrt{1M/d}$  where  $\sigma_0 = 10$  and re-clip the weights between -1 and 1. This can be interpreted as each cell having a Gaussian cumulative error with  $\sigma = \sigma_0$  after  $1M$  steps. For digital weight drift, we apply independent binary random flips to the weight matrix bits every  $d$  steps with probability  $p = p_0/(1M/d)$  where  $p_0 = 10$ . This can be interpreted as each cell flipping an average of  $p_0$  times over  $1M$  steps. Note that in real life,  $\sigma_0, p_0$  depend on a host of issues such as the environmental conditions of the device (temperature, humidity, etc), as well as the rate of seeing training samples.

## H HYPERPARAMETER SELECTION

In order to compare standard SGD with the SKS approach, we sweep the learning rates of both to optimize accuracy. In Figure 6, we compare accuracies across a range of learning rates for four different cases: SGD or SKS with or without max-norming gradients. Optimal accuracies are found when learning rate is around 0.01 for all cases. For most experiments, 8b weights, activations, and gradients, and 16b biases are used. Experiments similar to those in Section I are used to select some of the hyperparameters related to the SKS method in particular. In most experiments, rank-4 SKS with batch sizes of 10 (for convolution layers) or 100 (for fully-connected layers) are used. Additional details can be found in the supplemental code.




Figure 6: The left two heat maps are used to select the base / standard SGD learning rate. The right two heat maps are used to select the SKS learning rate using the optimal SGD learning rate for bias training from the previous sweeps. For the SKS sweeps, the learning rate is scaled proportional to the square-root of the batch size  $B$ . This results in an approximately constant optimal learning rate across batch size, especially for the max-norm case. Accuracy is reported averaged over the last 500 samples from a 10k portion of the online training set, trained from scratch.

## I ADDITIONAL STUDIES

In Figure 7, rank and weight bitwidth is swept for SKS with gradient max-norming. As expected, training accuracy improves with both higher SKS rank and bitwidth. In dense NVM applications, higher bitwidths may be achievable, allowing for corresponding reductions in the SKS rank and therefore, reductions in the auxiliary memory requirements.

In Table 2, biased (zero-variance) and unbiased (low-variance) versions of SKS are compared. Accuracy improvements are generally seen moving from biased to unbiased SKS although the pattern differs between the no-norm and max-norm cases. In the no-norm case, a significant improvement is seen favoring unbiased SKS for fully-connected layers. In the max-norm case, the choice of biased or unbiased SKS has only a minor impact on accuracy. It might be expected that as the number of accumulated samples for a given pseudobatch increases, lower variance would be increasingly important at the expense of bias. For our network, this implies convolutions, which receive updates at every pixel of an output feature map, would preferentially have biased SKS, while the fully-connected layer would preferentially be unbiased. This hypothesis is supported by the no-norm experiments, but not by the max-norm experiments.

In Table 3, several ablations are performed on SKS with max-norm. Most notably, weight training is found to be extremely important for accuracy as bias-only training shows a  $\approx 15 - 30\%$  accuracy hit depending on whether max-norming is used. Streaming batch norm is also found to be quite helpful, especially in the no-norm case.

Now, we explain the  $\kappa_{th}$  ablation. In Section 4.1.1, we found the SVD of a small matrix  $C$  and its singular values  $\sigma_1, \dots, \sigma_q$ . This allows us to easily find the condition number of  $C$  as  $\kappa(C) = \sigma_1/\sigma_q$ . We suspect high condition numbers provide relatively useless update information akin to noise, especially in the presence of  $L, R$  quantization. Therefore, we prefer not to update  $L, R$  on samples whose condition number exceeds threshold  $\kappa_{th}$ . We can avoid performing an actual SVD (saving computation) by noting that  $C$  is often nearly diagonal, leading to the approximation

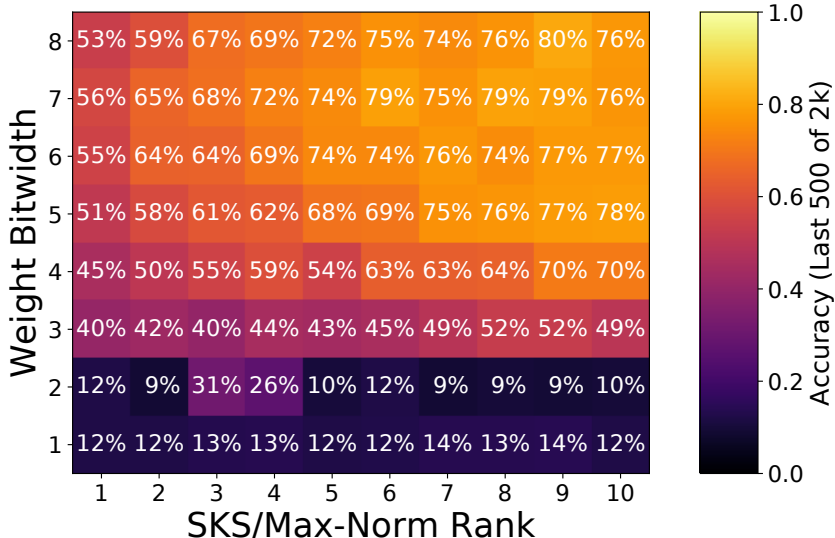


Figure 7: Accuracy across a variety of SKS ranks and weight bitwidths, showing the expected trends of increasing accuracy with rank and bitwidth. Accuracy is calculated by averaging the accuracy on the last 500 samples from a 2k portion of the training data. For bitwidths of 1 and 2, mid-rise quantization is used (e.g., 1 bit quantizes values to -0.5 and 0.5 instead of -1 and 0).

Table 2: Importance of unbiased SVD. Accuracy is calculated from the last 500 samples of 10k samples trained from scratch. Mean and unbiased standard deviation are calculated from five runs of different random seeds.

Conv SKS	FC SKS	Accuracy (no-norm)	Accuracy (max-norm)
Biased	Biased	79.7% $\pm$ 1.1%	82.7% $\pm$ 1.3%
Biased	Unbiased	83.0% $\pm$ 0.9%	82.4% $\pm$ 1.2%
Unbiased	Biased	77.7% $\pm$ 1.5%	84.6% $\pm$ 2.0%
Unbiased	Unbiased	81.0% $\pm$ 0.9%	83.6% $\pm$ 2.5%

$\kappa(\mathbf{C}) \approx C_{1,1}/C_{q,q}$ . Empirically, this rough heuristic works well to reduce computation load while having minor impact on accuracy. In Table 3,  $\kappa_{th} = 10^8$  does not appear to ubiquitously improve on the default  $\kappa_{th} = 100$ , despite being  $\approx 2\times$  slower to compute.

Table 3: Miscellaneous selected ablations. Accuracy is calculated from the last 500 samples of 10k samples trained from scratch. Mean and unbiased standard deviation are calculated from five runs of different random seeds.

Modified Condition	Accuracy (no-norm)	Accuracy (max-norm)
baseline (no modifications)	80.2% $\pm$ 1.0%	83.0% $\pm$ 1.1%
bias-only training	51.8% $\pm$ 3.2%	68.6% $\pm$ 1.4%
no streaming batch norm	68.2% $\pm$ 1.9%	81.8% $\pm$ 1.3%
no bias training	81.3% $\pm$ 1.0%	83.0% $\pm$ 1.4%
$\kappa_{th} = 10^8$ instead of 100	79.8% $\pm$ 1.4%	84.2% $\pm$ 1.4%