# Faster Neural Network Training with Approximate Tensor Operations

Menachem Adelman Intel & Technion adelman.menachem@gmail.com

> Ido Hakimi Technion idohakimi@gmail.com

Kfir Y. Levy \* Technion kfirylevy@technion.ac.il

Mark Silberstein Technion mark@ee.technion.ac.il

# Abstract

We propose a novel technique for faster deep neural network training which systematically applies sample-based approximation to the constituent tensor operations, i.e., matrix multiplications and convolutions. We introduce new sampling techniques, study their theoretical properties, and prove that they provide the same convergence guarantees when applied to SGD training. We apply approximate tensor operations to single and multi-node training of MLP and CNN networks on MNIST, CIFAR-10 and ImageNet datasets. We demonstrate up to 66% reduction in the amount of computations and communication, and up to 1.37x faster training time while maintaining negligible or no impact on the final test accuracy.

# 1 Introduction

Approximation techniques for faster inference and training of deep neural networks (DNNs) have received considerable attention. Examples include quantization [1–5], low-rank and structured-sparse models [6–9], weight extrapolations [10], and partial/asynchronous gradient updates in the context of distributed training [11, 12]. Sampling-based approximations were used to accelerate inference [13, 14], but using them in training [15–17] has not been systematically studied nor demonstrated end-to-end GPU performance benefits in practice.

We propose a novel approach to accelerating *DNN training* by systematically approximating tensor operations via sampling. At a high level, the original matrix products and convolutions are replaced with their faster approximate versions. The approximation is applied separately to each tensor operation, keeping the network architecture and tensor dimensions intact, thereby facilitating the adoption of this technique in existing DNN training frameworks, potentially in combination with other approximation techniques. Furthermore, when combined with distributed training, our technique allows for seamless reduction in the communication bandwidth and increased performance gains.

We begin by reviewing the plethora of existing methods for approximating matrix multiplication. We compare several known algorithms [18–25], and find *column-row sampling (CRS)* [20] to be the most suitable for approximating matrix multiplications in training. In order to compute the product of two matrices  $A^{\top}B$ , the CRS algorithm samples the columns of  $A^{\top}$  and the corresponding rows of *B* thus constructing smaller matrices which are then multiplied as usual. This method incurs low sampling overheads and lends itself to an efficient implementation using existing dense matrix product routines. CRS minimizes the approximation error for the Frobenius norm of the resulting matrix while keeping the approximation unbiased.

<sup>\*</sup>A Viterbi fellow

<sup>35</sup>th Conference on Neural Information Processing Systems (NeurIPS 2021).

Sampling-based approximations can be interpreted as a form of Dropout [26], and we discuss the similarities and differences between the two. While Dropout aims to prevent overfitting, we focus on approximations as means to accelerate training by reducing the amount of computation.

In this work we aim to answer two main questions. First, can neural networks be trained while using approximate tensor operations? Second, what are the relations between using exact or approximate operations during training?

We start by analyzing the simpler case of linear regression, where we can derive the effects of approximations in closed form. We define a new loss function that takes the sampling into account, and observe that the resulting gradients differ from the exact training due to the dependency between sampled features. To this end, we propose a new *Bernoulli-CRS* variant which achieves statistical independence of samples, study its properties, and show that in linear regression it is equivalent to dynamic  $L_2$  weight regularization of the original, non-approximate loss.

We then turn to the more general case of non-linear deep neural networks. We show that using sampling-based approximations in the backward pass provides the same convergence guarantees as the exact SGD for bounded weights. The convergence result holds for unbounded weights as well if approximation is applied only to the weight gradients and if the activation functions are bounded.

We also study a new *TopK-CRS* algorithm which deterministically selects the top-k column-row pairs with the highest norms. We show that this algorithm is equivalent to the minimal mean square error estimator (MMSE) in case column-row pairs are pairwise independent with zero mean.

Last, we generalize matrix product approximation to convolutions and analyze the approximation error to derive the optimal sampling policy. This allows us to apply approximations to training of convolutional neural networks (CNNs).

We implement our techniques in PyTorch [27]<sup>2</sup> and evaluate them on several DNN topologies, including MLP and CNN networks on MNIST [28], Wide ResNet 28-10 [29] on CIFAR-10 [30], and ResNet-50 and ResNet-152 [31] on ImageNet [32]. We demonstrate up to 66% reduction in the number of FLOPs and up to 1.33x faster training time with little or no degradation in model accuracy.

We develop another flavor of *TopK-CRS* which samples according to the weight norms only. When sampling the same subset of weights for different workers in a data-parallel setting, our sampling technique enables reducing the amount of gradient communication between workers. Notably, our algorithm is compatible with the standard AllReduce approach used in distributed deep learning. We implement an AllReduce scheme that takes advantage of the smaller gradient footprint and demonstrate 1.09x-1.37x speedup in multi-node training.

Our contributions are as follows:

- We derive general convergence guarantees for training with approximate tensor operations.
- We develop novel sampling algorithms and analyze their theoretical properties.
- We extend sampling-based algorithms to fast approximation of multi-channel convolutions.
- We show that our approach can reduce the computation, communication and total training time on several popular neural network architectures with little or no accuracy degradation.

# 2 Related work

To the best of our knowledge, we are the first to study the application of sample-based approximations of tensor operations to speed up DNN training. However, there have been several prior efforts to accelerate DNN computations via approximation which we survey below.

Several works accelerate inference through model compression [33–39]. A large body of work is devoted to quantization and low-precision datatypes (see for example [1, 2, 5]). Approximation was used to extrapolate weight values [10]. Another approach enforces low-rank or structured-sparse structure on the layers, resulting in lower computational cost both for training and inference [6–9]. Other works accelerate inference by approximating large matrices as products of lower-ranked ones [40, 41] or through locality-sensitive hashing [42].

<sup>&</sup>lt;sup>2</sup>https://github.com/acsl-technion/approx

In the context of distributed training, several works targeted communication bottlenecks by gradient quantization [3, 4], delayed weight updates [11, 12] and low-rank approximation of the gradient matrix [43]. These methods are complementary and compatible with ours.

Sampling-based approximations were used to accelerate inference [13, 14], but using them for training[15–17] has not been systematically studied nor shown to speed up training on GPUs without accuracy degradation. Sub-sampling whole layers was shown to enable training of very deep CNNs [44].

#### 2.1 Approximate matrix multiplication

There are several known algorithms for approximating matrix product. However, only those that meet the following requirements will be effective for DNN training. First, the algorithm should apply to dense matrices of arbitrary dimensions. Second, to reduce training time, the overall multiplication including input transformation should be faster than the original product. Last, the algorithm should be amenable to efficient implementation on commodity hardware.

Using these criteria, we consider the following algorithms:

**Random walk** [18] This algorithm performs random walks on a graph representation of the input matrices, but is applicable to non-negative matrices only.

**Random projections** [21–23] The two matrices to be multiplied are first projected into a lowerdimensional subspace by a scaled random sign matrix. These algorithms require both input matrices to be roughly square, otherwise the cost of projection will be similar to the cost of original product. In DNNs, however, it is common for one dimension to be smaller than the other.

**FFT** [24, 25] These algorithms represent each column-row outer product as a polynomial multiplication and then calculate it using Fast Fourier Transform. The complexity depends on the sparsity of the input matrices, decreasing as the sparsity increases. Therefore, these algorithms might not be effective for dense matrices.

**SVD** [19, 33, 36] Several algorithms replace one input matrix with its low-rank approximation using truncated SVD. These algorithms are suitable for inference where the weight matrix factorization can be pre-computed offline, but are not applicable to training since the high cost of factorization is incurred in every matrix product.

**Column-row sampling (CRS)** [19, 20] The sampling algorithm approximates matrix product  $A^{\top}B$  by sampling k columns of  $A^{\top}$  and respective rows of B to form smaller matrices, which are then multiplied as usual.

We choose CRS as the basis for our current work because it meets all the criteria above: It is applicable to fully-connected layers of any size, its effectiveness does not depend on the matrix contents, its sampling is computationally lightweight, and may use regular matrix multiplication algorithms since the sampled sub-matrices remain dense.

## 2.2 CRS

Let  $A \in \mathbb{R}^{n \times m}$ ,  $B \in \mathbb{R}^{n \times p}$ . Their product  $A^{\top}B$  is approximated as a weighted sum of outer products between sampled columns of  $A^{\top}$  and corresponding rows of B:

$$A^{\top}B \approx \sum_{t=1}^{k} \frac{1}{kp_{i_t}} A^{\top(i_t)} B_{(i_t)}$$
(1)

where  $A^{\top(i)}$ ,  $B_{(i)}$  denote the matrix *i*'th column and row respectively, *k* is the number of samples (satisfying  $1 \le k \le n$ ),  $\{p_i\}_{i=1}^n$  is a probability distribution over the column-row pairs of  $A^{\top}$ , *B* and  $i_t \in \{1, ..., n\}$ . This algorithm allows linear reduction in complexity from O(mnp) to O(mkp).

(1) can also be expressed as  $A^{\top}DS^{\top}SDB$ , where  $D \in \mathbb{R}^{n \times n}$  is a diagonal scaling matrix with:

$$(D)_{j,j} = \frac{1}{\sqrt{kp_j}} \tag{2}$$

and  $S \in \mathbb{R}^{k \times n}$  is a sampling matrix that selects k features, possibly with replacement. S is a random matrix, where each row has 1 in one entry and zeros in others. In each row, the probability of having the non-zero entry in column j is  $p_j$ .

Drineas et al. [20] show that CRS is unbiased:

$$\mathbb{E}\left[A^{\top}DS^{\top}SDB\right] = A^{\top}B \tag{3}$$

and also derive upper bounds for the expected Frobenius and spectral norms of the error matrix  $||A^{\top}B - A^{\top}DS^{\top}SDB||$ . They show that the error is minimized when the sampling probabilities are proportional to the product of the column-row Euclidean norms:

$$p_{i} = \frac{|A_{(i)}||B_{(i)}|}{\sum_{j=1}^{n} |A_{(j)}||B_{(j)}|}$$
(4)

In which case the expected Frobenius error is:

$$\frac{1}{k} \left( \sum_{t=1}^{k} |A_{(i_t)}| |B_{(i_t)}| \right)^2 - \frac{1}{k} \left| \left| A^\top B \right| \right|_F^2$$
(5)

#### 2.3 Approximate Tensor Operations and Dropout

Sampling-based approximations can be interpreted as a flavor of Dropout [26], a popular technique to prevent overfitting by randomly zeroing individual activations during training. However, the sparsity pattern resulting from Dropout is unstructured and therefore cannot be exploited efficiently by GPUs despite recent advances in structured sparsity support[8]. Prior works on fast Dropout training [45, 46] are different than ours and do not demonstrate acceleration of large networks while maintaining accuracy.

Some works proposed non-uniform Dropout probabilities for individual activations [47] or full channels [48]. Their sampling heuristics are different from ours which are derived from optimal approximation. Furthermore, they use Dropout only for preventing overfitting and do not leverage it to speed up training. In our experiments we demonstrate the utility of sampling-based approximations for DNNs with and without Dropout. Conversely, we did not observe improved accuracy from approximations which could have been attributed to reduced overfitting.

## **3** Approximate Linear Regression

We now analyze CRS in the simpler context of linear regression, where we can derive the effects of approximations in closed form. We show that this leads to biased gradient estimates.

Let  $X \subset \mathbb{R}^{n \times M}$  a dataset containing M examples, each a vector in  $\mathbb{R}^n$ . Every  $x^i \in X$  is associated with a "ground truth" value  $y^i \in \mathbb{R}$ .

Let  $w \in \mathbb{R}^n$  be parameters in a linear model that predicts a value  $\bar{y}^i \in \mathbb{R}$  for every  $x^i \in X$ :

$$\bar{y}^i = w^\top x^i \tag{6}$$

To simplify the notation we do not include an explicit bias term. We do so without loss of generality since we can always add another entry of 1 to the features.

Let us define the MSE (Mean Square Error) loss function:

$$\ell = \sum_{i=1}^{M} (\bar{y}^i - y^i)^2 \tag{7}$$

When using SGD (Stochastic Gradient Descent), we are given a single example  $x^i \in \mathbb{R}^n$  in each step, and update w using the gradients  $\frac{\partial \ell}{\partial w}$ . For notation simplicity we omit the superscript i from  $x^i, \bar{y}^i, y^i$ .

The gradients are given by the chain rule as:

$$\frac{\partial \ell}{\partial w} = 2x(w^{\top}x - y) \tag{8}$$

Now, let us assume that the multiplication  $w^{\top}x$  is approximated using CRS. The linear regression model now becomes:

$$\hat{y} = w^{\top} D S^{\top} S D x \tag{9}$$

and for the MSE loss the gradients will be:

$$\frac{\partial \ell}{\partial w} = 2DS^{\top}SDx(w^{\top}DS^{\top}SDx - y)$$
(10)

Where  $\frac{\widehat{\partial \ell}}{\partial w}$  denotes the CRS weight gradients.

Note that  $DS^{\top}SD$  appears twice in (10). In  $w^{\top}DS^{\top}SDx$  it represents sampling in the forward pass, while in  $DS^{\top}SDx$  it results in passing gradients only to the elements of w that were sampled in the forward pass.

It should be emphasized that (9) and (10) in fact describe gradients with respect to a different loss function compared to (8): one loss function uses  $\hat{y}$  while the other uses  $\bar{y}$ . If the approximate gradients are unbiased estimates of the non-approximate gradients, we could relate the approximate training process to the original one. However, the weight gradients do not satisfy this unbiasedness property:

$$\mathbb{E}\left[\left(\frac{\widehat{\partial \ell}}{\partial w}\right)_{j}\right] = 2x_{j}\left(\sum_{t=1}^{n} w_{t} \mathbb{E}\left[\left(\tilde{S}\right)_{j,j}(\tilde{S})_{t,t}\right] x_{t} - y\right)$$
(11)

where we denote  $\tilde{S} \triangleq DS^{\top}SD$ , and use the fact that  $\tilde{S}$  is diagonal.

For the expression in (11) to be equal to that in (8) we need that  $\mathbb{E}\left[(\tilde{S})_{j,j}(\tilde{S})_{t,t}\right] = 1$ . However, this is not the case because  $(\tilde{S})_{j,j}$  and  $(\tilde{S})_{t,t}$  are not independent random variables: an entry in the diagonal of  $\tilde{S}$  is the (scaled) number of times a column-row pair was selected out of the k total samples in CRS. The event of selecting a particular pair therefore affects selecting others.

We note that if instead of changing the loss function we treated the approximate multiplication as a "black box" that replaces the original product, we could use (8) and only replace the forward pass product  $w^{\top}x$  with the "black box" substitute of  $w^{\top}DS^{\top}SDx$ . This would yield:

$$\frac{\partial \ell}{\partial w} = 2x(w^{\top}DS^{\top}SDx - y)$$
(12)

which satisfies:

$$\mathbb{E}\left[\frac{\widetilde{\partial\ell}}{\partial w}\right] = \frac{\partial\ell}{\partial w}.$$
(13)

(12) is equivalent to applying the approximate computation in the forward pass, but propagating the gradients to all weight entries in the same way as if the computation were exact. In practice we find that this approach leads to significantly lower accuracy in deep neural networks compared to sampling the same entries in the forward and backward pass, or to applying approximations in the backward pass only.

# 4 Bernoulli-CRS

We now turn to develop Bernoulli-CRS, a new variant of sampling approximation that enables to sample column-row pairs independently and without replacement. Applied to linear regression, we show that using Bernoulli-CRS is equivalent to employing unbiased gradient estimates in addition to a bias term which can be interpreted as scale dependent weight regularization.

**Bernoulli-CRS:** These aforementioned properties can be achieved by assigning a separate Bernoulli sampling probability  $p_i$  for each column-row pair *i*, and sampling pairs independently of each other. To control the amount of sampling, we add another constraint that all the probabilities will sum up to an integer *k*:

$$\sum_{i=1}^{n} p_i = k \tag{14}$$

Let us define  $K \in \mathbb{R}^{n \times n}$  a random diagonal sampling matrix where  $K_{j,j} \sim \text{Bernoulli}(p_j)$  for  $1 \le j \le n$ . Furthermore, let us define another diagonal scaling matrix  $P \in \mathbb{R}^{n \times n}$  where  $P_{j,j} = \frac{1}{\sqrt{p_j}}$  for  $1 \le j \le n$ .

Using the K and P matrices we may now define our new Bernoulli-CRS algorithm. Let  $A \in \mathbb{R}^{n \times m}$ and  $B \in \mathbb{R}^{n \times p}$ . The product  $A^{\top}B$  can be approximated with  $\tilde{A}^{\top}\tilde{B}$  defined as follows:

$$\tilde{A}^{\top}\tilde{B} := \sum_{i=1}^{n} \frac{Z_i}{p_i} A^{\top(i)} B_{(i)} = A^{\top} P K K P B$$
(15)

where  $\{Z_i \sim \text{Bernoulli}(p_i)\}_{i=1}^n$  are independent random variables. We denote  $\tilde{A} \triangleq KPA$  and  $\tilde{B} \triangleq KPB$ .

In the appendix we develop the properties of Bernoulli-CRS. We show it is unbiased and derive bounds on the error variance both in expectation and in high probability. We derive the optimal sampling probabilities minimizing the expected variance, and show that under certain conditions they are given by the simpler expression:

$$p_{i} = \min\left\{\frac{k|A_{(i)}||B_{(i)}|}{\sum_{j=1}^{n}|A_{(j)}||B_{(j)}|}, 1\right\}$$
(16)

In the appendix we show that applying Bernoulli-CRS in linear regression leads to unbiased estimate of the original gradients with an additional regularization term  $\mathcal{R}(w)$ , which we define as:

$$\mathcal{R}(w) = \mathbf{E}\left[\sum_{j=1}^{n} \frac{1-p_j}{p_j} x_j^2 w_j^2\right]$$
(17)

and the expectation is with respect to the distribution of the data samples.

The term  $\mathcal{R}(w)$  can be interpreted as input-dependent  $L_2$  regularization. The regularization is higher as  $x_j$  grows in magnitude and as  $p_j$  decreases. Both serve to reduce the impact on the weights if they were chosen with small probabilities or mostly because of the input size. We note that Wager et al. [49] conduct a similar analysis for Dropout in the particular case where the same sampling probability is used for all features.

To summarize, sampling in the simpler case of linear regression minimizes the original loss function with an added regularization term.

## 5 Approximate Backpropagation in Non-Linear Deep Networks

The analysis of approximate linear regression cannot simply generalize to deep non-linear networks: non-linearity leads to biased network output even if the approximate multiplication is itself unbiased. Still, we are able to obtain strong theoretical results on the relations between exact and approximate training if the approximations are limited to the backward pass: the forward pass is calculated as usual, and the matrix products in the backward pass are performed using approximate matrix multiplication.

#### We prove the following theorem:

**Theorem 1.** Let f(x, W, b) be a multi-layer neural network with  $\beta$ -Lipschitz activation functions  $\sigma$ . Let  $\ell$  be a  $\beta$ -Lipschitz loss function, and let the network be trained with SGD using properly

decreasing learning rate. Assume that the weights are bounded; and further assume that the matrix products in the backward pass are approximated using an unbiased approximation scheme, i.e.,

$$\mathbb{E}\left[A^{\top}B - \operatorname{approx}(A^{\top}B)\right] = 0$$

and that there exists a constant C and a norm  $||\cdot||$  such that:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|^{2}\right] \leq C \left|\left|A\right|\right|^{2} \left|\left|B\right|\right|^{2}.$$

Then the approximated NN gradient estimates are unbiased, and their second moments are bounded.

**Corollary.** Based on recent works on non-convex optimization (see e.g. [50]), the unbiasedness and bounded second moments ensured by Theorem 1 imply that approximate backpropagation enjoys the same convergence guarantees as regular SGD training.

In the appendix we show that CRS and other sampling algorithms satisfy the property

$$\mathbb{E}\left[\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|^{2}\right] \leq C \left|\left|A\right|\right|^{2} \left|\left|B\right|\right|^{2}$$

Note that for Theorem 1 we required that weights will be bounded during the training process. This is a strong assumption which could be justified if weight regularization or clipping is used. In the appendix we prove the same results without relying on these assumptions, if only the weight gradients are approximated and if the activation function is bounded (such as sigmoid).

# 6 Sampling Without Scaling and Top-k Selection

We now consider a different sampling scheme where k column-row pairs are selected deterministically without scaling. This can be viewed as a special case of Bernoulli-CRS, where the sampling probabilities are either 0 or 1. We now show that under certain assumptions on the distribution of the input matrices, this scheme can lead to the optimal estimation:

**Theorem 2.** Let A be an  $n \times m$  random matrix and B be an  $n \times p$  random matrix, such that

$$\mathbb{E}\left[A^{\top(i)}B_{(i)}\right] = 0$$

for  $1 \le i \le n$ . Assume k column-row pairs with indices  $\{j\}_1^n$  are sampled from A and B. Then, the MMSE estimator of the product  $A^{\top}B$  is  $\tilde{A}^{\top}\tilde{B}$  where  $\tilde{A}, \tilde{B}$  are constructed from the sampled column-row pairs without scaling.

Furthermore, if  $A^{\top(i)}B_{(i)}$  and  $A^{\top(j)}B_{(j)}$  are independent for different i, j then the MSE is minimized when sampling k pairs with the maximum norm multiplication  $|A_{(i)}||B_{(i)}|$ .

The assumptions in Theorem 2 can hold in practice if weights are initialized with a zero-centered distribution[51], if the distribution of weights remains centered around zero during training [51–54], and if different weights can be considered pairwise-independent [55].

We study the approximation quality of CRS, Bernoulli-CRS and top-k selection on synthetic matrix multiplication. We generate  $100 \times 100$  random matrices and compute the error metric:

$$\frac{\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|_{F}}{\left|\left|A\right|\right|_{F}\left|\left|B\right|\right|_{F}}$$
(18)

Figures 1(a),1(b) show the approximation error for different algorithms and sampling ratios, averaged over 1000 runs. We observe that Bernoulli-CRS outperforms CRS in higher sampling ratios. Also, when one matrix has i.i.d entries with zero mean, Bernoulli-CRS outperforms CRS and top-k selection performs the best as expected from Theorem 2.

We also consider a different flavor of top-k selection, which we refer to as "top-k-weights": sampling k column-row pairs corresponding to rows of B with the highest norms. While not providing the theoretical guarantees of Theorem 2, the new variant has a desirable property for data parallel distributed training, where weights are identical between different workers. A deterministic selection algorithm that only depends on the weights will sample the same weights for all workers, allowing to reduce the gradient communication between nodes to the selected weights only.



(a) Matrix product: both matrix entries drawn from  $\mathcal{N}(1,1)$ 

(b) Matrix product: one matrix entries drawn from  $\mathcal{N}(0, 1)$ , the other from  $\mathcal{N}(1, 1)$ 

Figure 1: Approximation error depending on the amount of performed computations. Lower is better.

# 7 Approximating Convolutions

We extend the basic CRS algorithm to the approximation of multi-channel convolutions. In matrix multiplication sampling is performed over the common dimension. The analogue for multi-channel convolution is to sample over the input channels dimension, illustrated in Figure 2. As in the matrix case, the output dimensions remain the same.



Figure 2: Sampling one input channel out of three

In the appendix we derive the optimal sampling probabilities and scaling factors. Bernoulli-CRS and top-*k* algorithms can be developed for convolutions as well in an analogous way.

# 8 Experimental Results

We implement CRS, Bernoulli-CRS and top-k selection approximation algorithms in PyTorch both for matrix multiplication and convolution. Our implementation allows to control the sampling degree and the application of approximation in the forward or backward passes.

We replace exact tensor operations with their sampling-based approximations, *without changing training hyper-parameters*. Only column-row pairs sampled in the forward pass are used during backpropagation as the rest do not affect the loss. Hence, sampling in the forward pass also reduces the amount of backward pass computations by the same ratio. We apply approximations only during training, and use exact computations for validation/test evaluation.

We evaluate our approximate training technique on several network architectures and datasets: MLP and CNN on MNIST, Wide ResNet 28-10 on CIFAR-10 and ResNet-50 and ResNet-152 on ImageNet. We train the networks on a single node using NVidia V100 GPUs (two GPUs for ResNet-152, one for the rest), and measure the reduction in multiply-accumulate operations due to sampling as well as the overall speedup in total training time versus the non-approximate baseline. The appendix includes additional details on the models and the training process.

| Network                  |             | COMPUTE<br>REDUCTION | Communication<br>Reduction | ACCURACY<br>(BASELINE) | Training<br>Speedup |
|--------------------------|-------------|----------------------|----------------------------|------------------------|---------------------|
| MLP (MNIST)              |             | 50%                  | -                          | 98.22% (98.22%)        | -                   |
| CNN (MNIST)              |             | 66%                  | -                          | 99.25% (99.35%)        | -                   |
| WRN-28-10 (CIFAR-10)     |             | 50%                  | -                          | 95.89% (96.17%)        | 1.33x               |
| ResNet-50 (ImageNet)     |             | 6.5%                 | -                          | 75.63% (75.6%)         | 1.04x               |
| ResNet-152<br>(ImageNet) | SINGLE NODE | 40%                  | -                          | 76.44% (77.65%)        | 1.16x               |
|                          |             | 9%                   | -                          | 77.66% (77.65%)        | 1.04x               |
|                          | 8 Nodes     | 40%                  | 48%                        | 76.44% (77.65%)        | 1.37x               |
|                          |             | 12%                  | 23%                        | 77.48% (77.65%)        | 1.13x               |
|                          |             | 9%                   | 13%                        | 77.8% (77.65%)         | 1.09x               |

Table 1: Compute reduction, communication reduction and wall-clock speedup of training with approximate tensor operations.







Figure 3: CIFAR-10 (approximate forward and backward)

mate backward)

100

95

90

85

80

Test Accuracy

WRN-28-10 on Figure 4: 3-layer MLP on Figure 5: AllReduce with top-k-MNIST (exact forward, approxi- weights sampling (% fewer gradients sent).

Our results using top-k sampling are summarized in Table 1. We see a reduction of up to 66% in the computational cost with little or no degradation in model accuracy, and up to 1.33x faster training time. We believe the gap between compute reduction and end-to-end speedup can be reduced by fusing the sampling with the matrix multiplication routines, or running on a different HW architecture that allows fast sampling and norm computation. We note that the small MNIST models do not exhibit training time speedup since they are not compute-intensive enough to saturate the GPU. The ratio between compute reduction and actual speedup is smaller in ResNet-152 compared to ResNet-50 and WRN-28-10 because the batch size per GPU is lower due to the limited GPU memory capacity.

Sampling Algorithms We compare CRS, Bernoulli-CRS and top-k selection on MNIST and CIFAR-10 and find empirically that top-k results in higher accuracy and faster training time (Figure 3). This result is consistent with that of approximate  $\mathcal{N}(0,1)$  matrix product (Figure 1(b)). This is not surprising given Theorem 2 and our empirical observation that the weight distribution is close to symmetrical around zero throughout training.

Approximations in Forward Pass and Backpropagation For the small MNIST models we are able to perform as low as 10% of the computations in the backward pass without harming accuracy (Fig. 4). However, in the larger models (WRN-28-10) we find empirically that accuracy drops when approximating only the backward pass. Therefore, in Table 1 we report results for consistent sampling in the forward and backward passes.

**Sampling Ratio Sensitivity** We find that the achievable compute reduction is not consistent across networks and datasets. For MNIST and CIFAR-10 we maintain good accuracy while reducing 50%-66% of the computations. However, ImageNet proved to be more sensitive and we kept the accuracy intact when applying 50% sampling to the ResNet layers with 1024 or more channels only. Figure 6 shows the learning curves under different sampling ratios compared to the non-approximate baseline.



Figure 6: Learning curves for validation accuracy under different top-k sampling ratios

**Distributed Training** We implement an AllReduce "top-*k*-weights" scheme in PyTorch. This scheme performs reduction only for the gradients of the sampled weights, reducing inter-node communications. Table 1 shows the accuracy-speedup trade-off for ResNet-152 distributed training. Figure 5 shows the respective scaling behavior of these schemes relative to the exact baseline. We note that compute savings did not lead to significant single-node speedup since in this experiment the V100 GPUs (from Amazon AWS) had lower memory capacity, which led to smaller batch size per GPU. The multi-node training speedup is therefore mostly due to the communication savings.

# 9 Conclusion

In this work we have demonstrated the utility of sample-based approximation of tensor operations for neural network training, both theoretically and empirically. We believe that further acceleration could be achieved through dedicated GPU primitives fusing sampling and matrix multiplication/convolution, as well as varying and adaptive sampling rates for different layers and iterations. Studying other approximation algorithms, applications in resource-constrained environments and bridging the gaps between our theoretical results and what worked best in practice are all promising directions for future research. Overall, we believe that sample-based approximations and fast approximations in general are valuable additions to the toolbox of techniques for deep learning acceleration.

## Acknowledgments and Disclosure of Funding

K.Y. Levy acknowledges support from the Israel Science Foundation (grant No. 447/20). I. Hakimi acknowledges support from the Hasso Plattner Institute at the Technion. M. Silberstein acknowledges support from the Israel Science Foundation (grant No. 1027/18).

## References

- [1] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *arXiv preprint arXiv:1609.07061*, 2016.
- [2] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaev, Ganesh Venkatesh, et al. Mixed precision training. arXiv preprint arXiv:1710.03740, 2017.
- [3] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [4] Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In Advances in Neural Information Processing Systems, pages 1508–1518, 2017.
- [5] Ayan Chakrabarti and Benjamin Moseley. Backprop with approximate activations for memory-efficient network training. *Advances in Neural Information Processing Systems*, 32:2429–2438, 2019.

- [6] Franck Mamalet and Christophe Garcia. Simplifying convnets for fast learning. In *International Conference* on Artificial Neural Networks, pages 58–65. Springer, 2012.
- [7] Oleksii Kuchaiev and Boris Ginsburg. Factorization tricks for lstm networks. *arXiv preprint arXiv:1703.10722*, 2017.
- [8] Asit Mishra, Jorge Albericio Latorre, Jeff Pool, Darko Stosic, Dusan Stosic, Ganesh Venkatesh, Chong Yu, and Paulius Micikevicius. Accelerating sparse deep neural networks. arXiv preprint arXiv:2104.08378, 2021.
- [9] Itay Hubara, Brian Chmiel, Moshe Island, Ron Banner, Seffi Naor, and Daniel Soudry. Accelerated sparse neural training: A provable and efficient method to find n: m transposable masks. *arXiv preprint arXiv:2102.08124*, 2021.
- [10] Sagar V Kamarthi and Stefan Pittner. Accelerating neural network training using weight extrapolations. *Neural networks*, 12(9):1285–1299, 1999.
- [11] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In Advances in neural information processing systems, pages 693–701, 2011.
- [12] Nikko Strom. Scalable distributed dnn training using commodity gpu cloud computing. In *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [13] Weizhe Hua, Yuan Zhou, Christopher M De Sa, Zhiru Zhang, and G Edward Suh. Channel gating neural networks. In *Advances in Neural Information Processing Systems*, pages 1886–1896, 2019.
- [14] Xitong Gao, Yiren Zhao, Łukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331*, 2018.
- [15] Xu Sun, Xuancheng Ren, Shuming Ma, and Houfeng Wang. meprop: Sparsified back propagation for accelerated deep learning with reduced overfitting. In *International Conference on Machine Learning*, pages 3299–3308, 2017.
- [16] Beidi Chen, Tharun Medini, James Farwell, Sameh Gobriel, Charlie Tai, and Anshumali Shrivastava. Slide: In defense of smart algorithms over hardware acceleration for large-scale deep learning systems. arXiv preprint arXiv:1903.03129, 2019.
- [17] Negar Goli and Tor M Aamodt. Resprop: Reuse sparsified backpropagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1548–1558, 2020.
- [18] Edith Cohen and David D Lewis. Approximating matrix multiplication for pattern recognition tasks. *Journal of Algorithms*, 30(2):211–252, 1999.
- [19] Petros Drineas and Ravi Kannan. Fast monte-carlo algorithms for approximate matrix multiplication. In *FoCS*, volume 1, pages 452–459, 2001.
- [20] Petros Drineas, Ravi Kannan, and Michael W Mahoney. Fast monte carlo algorithms for matrices i: Approximating matrix multiplication. *SIAM Journal on Computing*, 36(1):132–157, 2006.
- [21] Avner Magen and Anastasios Zouzias. Low rank matrix-valued chernoff bounds and approximate matrix multiplication. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1422–1436. SIAM, 2011.
- [22] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on, pages 143–152. IEEE, 2006.
- [23] Kenneth L Clarkson and David P Woodruff. Numerical linear algebra in the streaming model. In Proceedings of the forty-first annual ACM symposium on Theory of computing, pages 205–214. ACM, 2009.
- [24] Rasmus Pagh. Compressed matrix multiplication. *ACM Transactions on Computation Theory (TOCT)*, 5 (3):9, 2013.
- [25] Konstantin Kutzkov. Deterministic algorithms for skewed matrix products. In 30th International Symposium on Theoretical Aspects of Computer Science, page 466, 2013.

- [26] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15 (1):1929–1958, 2014.
- [27] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In Advances in neural information processing systems, pages 8026–8037, 2019.
- [28] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [29] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.
- [30] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- [31] 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*, pages 770–778, 2016.
- [32] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015.
- [33] Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In Advances in neural information processing systems, pages 1269–1277, 2014.
- [34] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *Proceedings of the British Machine Vision Conference. BMVA Press*, 2014.
- [35] Vadim Lebedev, Yaroslav Ganin, Maksim Rakhuba, Ivan Oseledets, and Victor Lempitsky. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. arXiv preprint arXiv:1412.6553, 2014.
- [36] Kazuki Osawa, Akira Sekiya, Hiroki Naganuma, and Rio Yokota. Accelerating matrix multiplication in deep learning by using low-rank approximation. In *High Performance Computing & Simulation (HPCS)*, 2017 International Conference on, pages 186–192. IEEE, 2017.
- [37] Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014.
- [38] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [39] Xu Sun, Xuancheng Ren, Shuming Ma, Bingzhen Wei, Wei Li, and Houfeng Wang. Training simplification and model simplification for deep learning: A minimal effort back propagation method. arXiv preprint arXiv:1711.06528, 2017.
- [40] Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. arXiv preprint arXiv:2009.14794, 2020.
- [41] Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. arXiv preprint arXiv:2006.04768, 2020.
- [42] Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. Reformer: The efficient transformer. *arXiv preprint arXiv:2001.04451*, 2020.
- [43] Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. Powersgd: Practical low-rank gradient compression for distributed optimization. arXiv preprint arXiv:1905.13727, 2019.
- [44] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016.
- [45] Sida Wang and Christopher Manning. Fast dropout training. In *international conference on machine learning*, pages 118–126. PMLR, 2013.
- [46] Ben Graham, Jeremy Reizenstein, and Leigh Robinson. Efficient batchwise dropout training using submatrices. *arXiv preprint arXiv:1502.02478*, 2015.

- [47] Zhe Li, Boqing Gong, and Tianbao Yang. Improved dropout for shallow and deep learning. In Advances in Neural Information Processing Systems, pages 2523–2531, 2016.
- [48] Saihui Hou and Zilei Wang. Weighted channel dropout for regularization of deep convolutional neural network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 8425–8432, 2019.
- [49] Stefan Wager, Sida Wang, and Percy S Liang. Dropout training as adaptive regularization. In Advances in neural information processing systems, pages 351–359, 2013.
- [50] Rong Ge, Furong Huang, Chi Jin, and Yang Yuan. Escaping from saddle points—online stochastic gradient for tensor decomposition. In *Conference on Learning Theory*, pages 797–842, 2015.
- [51] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.
- [52] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pages 1613–1622. PMLR, 2015.
- [53] Song Han, Jeff Pool, John Tran, and William J Dally. Learning both weights and connections for efficient neural networks. arXiv preprint arXiv:1506.02626, 2015.
- [54] Martin Thoma. Analysis and optimization of convolutional neural network architectures. *arXiv preprint arXiv:1707.09725*, 2017.
- [55] Zhongzhan Huang, Wenqi Shao, Xinjiang Wang, and Ping Luo. Convolution-weight-distribution assumption: Rethinking the criteria of channel pruning. arXiv preprint arXiv:2004.11627, 2020.
- [56] Dimitri P Bertsekas and John N Tsitsiklis. Neuro-dynamic programming. Athena Scientific, 1996.
- [57] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [58] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

# Appendices

## **A Bernoulli-CRS Properties**

Let us define  $K \in \mathbb{R}^{n \times n}$  a random diagonal sampling matrix where  $K_{j,j} \sim \text{Bernoulli}(p_j)$  for  $1 \le j \le n$ .

Let us define another diagonal scaling matrix  $P \in \mathbb{R}^{n \times n}$  where  $P_{j,j} = \frac{1}{\sqrt{p_j}}$  for  $1 \le j \le n$ .

Using the K and P matrices we may now define our new Bernoulli-CRS algorithm. Let  $A \in \mathbb{R}^{n \times m}$ and  $B \in \mathbb{R}^{n \times p}$ . The product  $A^{\top}B$  can be approximated with  $\tilde{A}^{\top}\tilde{B}$  defined as follows:

$$\tilde{A}^{\top}\tilde{B} := \sum_{i=1}^{k} \frac{Z_i}{p_i} A^{\top(i)} B_{(i)} = A^{\top} P K K P B$$
(19)

where  $\{Z_i \sim \text{Bernoulli}(p_i)\}_{i=1}^n$  are independent random variables. We denote  $\tilde{A} \triangleq KPA$  and  $\tilde{B} \triangleq KPB$ .

First, we show that the above holds in expectation:

**Proposition 1.**  $\mathbb{E}\left[\tilde{A}^{\top}\tilde{B}\right] = A^{\top}B.$ 

Let T = Trace(K) the number of non-zero diagonal elements in K. We note that to perform the actual computation it is enough to sample the T column-row pair with the corresponding element in K being non-zero. Unlike CRS, the lower rank of the sampled matrices is not constant and depends on the random matrix K. Its expectation is controlled through the parameter k:

**Proposition 2.**  $\mathbb{E}[T] = k$ .

Therefore, Bernoulli-CRS will perform on average the same amount of computations as in the fixed-rank CRS.

Let us further derive the properties of the proposed sampling algorithm. Specifically, what are the optimal values for the probabilities  $p_i$  under the constraint  $\sum_{i=1}^{n} p_i = k$ ?

First, let us calculate the variance of  $\tilde{A}^{\top}\tilde{B}$ :

**Proposition 3.** 

$$\mathbf{Var}\left[(\tilde{A}^{\top}\tilde{B})_{i,j}\right] = \sum_{t=1}^{n} \frac{1-p_t}{p_t} A_{t,i}^2 B_{t,j}^2$$

We will be interested in the Frobenius norm of the error matrix  $\left\| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right\|_{F}^{2}$ , which can be derived from the following theorem:

**Theorem 3.** The expected Frobenius norm of the error matrix  $\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right]$  is  $\sum_{t=1}^{n} \frac{1-p_{t}}{p_{t}} |A_{(t)}|^{2} |B_{(t)}|^{2}.$ 

Furthermore, under the constraint  $\sum_{i=1}^{n} p_i = k$  it is minimized for the probabilities:

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbf{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbf{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}}$$

where  $\mu$  is a root of the following function:

$$G(\mu) := \sum_{i=1}^{n} \left( \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbb{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbb{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}} \right) - k$$

Corollary. The sampling probabilities

$$p_i = \min\left\{\frac{k|A_{(i)}||B_{(i)}|}{\sum_{j=1}^n |A_{(j)}||B_{(j)}|}, 1\right\}$$

are optimal if  $k \leq \frac{\sum_{i=1}^{n} |A_{(i)}| |B_{(i)}|}{\max_i |A_{(i)}| |B_{(i)}|}$ 

From Theorem 3 it follows that for the probabilities:

$$p_{i} = \min\left\{\frac{k|A_{(i)}||B_{(i)}|}{\sum_{j=1}^{n}|A_{(j)}||B_{(j)}|}, 1\right\}$$
(20)

the expected Frobenius error is:

$$\frac{1}{k} \left( \sum_{i=1}^{n} e_i |A_{(i)}| |B_{(i)}| \right)^2 - \sum_{i=1}^{n} e_i |A_{(i)}|^2 |B_{(i)}|^2$$
(21)

where we denote:

$$e_{i} \triangleq \begin{cases} 1 & |A_{(i)}||B_{(i)}| \le \frac{\sum_{j=1}^{n} |A_{(j)}||B_{(j)}|}{k} \\ 0 & \text{else} \end{cases}$$
(22)

Comparing that with the bound in (5), we can see that different values of A, B determine which algorithm performs better.

Knowing the expected Frobenius error also implies a bound on the spectral norm of the error matrix, since the spectral and Frobenius norms are related by:

$$|A|| \le ||A||_F \le \sqrt{r} \, ||A|| \tag{23}$$

where r is the rank of A and ||A|| denotes its spectral norm.

The following theorem yields high probability bounds for the Frobenius and spectral norms for the Bernoulli-CRS algorithm:

**Theorem 4.** Let  $A \in \mathbb{R}^{n \times m}$  and  $B \in \mathbb{R}^{n \times p}$ . Let  $\tilde{A}, \tilde{B}$  be the sampled matrices according to the Bernoulli-CRS algorithm described above. Denote

$$R \triangleq \max_{i} \left| \left| A^{\top(i)} B_{(i)} \right| \right|$$

and

$$\sigma^{2} \triangleq \frac{1}{k} \left( \sum_{i=1}^{n} e_{i} |A_{(i)}| |B_{(i)}| \right)^{2} - \sum_{i=1}^{n} e_{i} |A_{(i)}|^{2} |B_{(i)}|^{2}$$

then, for all  $t \ge 0$ :

$$\mathbb{P}\left\{ \left| \left| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right| \right| \ge t \right\} \le (m+p) \cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right)$$
$$\mathbb{P}\left\{ \left| \left| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right| \right|_F \ge t \right\} \le (m+p)^{3/2} \cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right)$$

#### A.1 Bernoulli-CRS in linear regression

We now show that applying Bernoulli-CRS in linear regression leads to unbiased estimate of the original gradients with an additional term that can be interpreted as regularization. The analysis for linear regression using Bernoulli-CRS is the same as in Section 3, with the sampling and scaling matrices  $DS^{T}SD$  replaced with PKKP

The expression for the weight gradient (similar to (11)) now becomes:

$$\mathbb{E}\left[\left(\frac{\partial \tilde{\ell}}{\partial w}\right)_{j}\right] = 2x_{j}\left(\sum_{t=1}^{n} w_{t}\left(\mathbb{E}\left[(\tilde{K})_{j,j}(\tilde{K})_{t,t}\right]x_{t} - y\right)\right)$$
(24)

$$=2x_j(w^{\top}x-y+w_j\left(\mathbb{E}\left[(\tilde{K})_{j,j}^2\right]-1\right)x_j)$$
(25)

$$=2x_{j}(w^{\top}x - y + \frac{1 - p_{j}}{p_{j}}w_{j}x_{j})$$
(26)

where we denote  $\tilde{K} \triangleq PKKP$ .

When comparing (26) and (8) we see that using Bernoulli-CRS yields unbiased estimates of the original gradients with an additional bias term that is related to a scale-dependent regularization  $\mathcal{R}(w)$ , which we define as:

$$\mathcal{R}(w) = \mathbf{E}\left[\sum_{j=1}^{n} \frac{1-p_j}{p_j} x_j^2 w_j^2\right]$$
(27)

and the expectation is with respect to the distribution of the data samples.

This term can be interpreted as input-dependent  $L_2$  regularization. The regularization is higher as  $x_j$  grows in magnitude and as  $p_j$  decreases. Both serve to reduce the impact on the weights if they were chosen with small probabilities or mostly because of the input size.

# **B** Approximating Convolutions - Details

Formally, let  $I \in \mathbb{R}_{B \times IH}^{IW \times IC}$  be the input tensor, where *B* is the batch size, *IH*, *IW* are the input height and width, and *IC* are the input channels. Let  $K \in \mathbb{R}_{KH \times KW}^{IC \times OC}$  be the kernels tensor, where *KH*, *KW* are the kernel height and width, and *IC*, *OC* are the input and output channels respectively. Let  $O \in \mathbb{R}_{B \times OH}^{OW \times OC}$  be the output tensor, where *OH*, *OW* are the output height and width.

The multi-channel convolution operation is defined as:

$$O_{b,oh}^{ow,oc} = I * K = \sum_{i=1}^{IC} \sum_{kh=1}^{KH} \sum_{kw=1}^{KW} I_{b,oh+kh-1}^{ow+kw-1,i} \cdot K_{kh,kw}^{i,oc}$$
(28)

For notation simplicity, we assume zero padding. The inner sums in (28) can be written as 1-channel convolutions:

$$O_{b,oh}^{ow,oc} = \sum_{i=1}^{IC} I^{[i]} * K_{[i]}$$
<sup>(29)</sup>

where  $I^{[i]} \in \mathbb{R}_{B \times IH}^{IW \times 1}$  denotes a tensor with one input channel that corresponds to the *i*'th input channel of *I*, i.e  $I^{[i]}_{b,ih}^{iw,1} = I^{iw,i}_{b,ih}$ . Similarly,  $K_{[i]} \in \mathbb{R}_{KH \times KW}^{1 \times OC}$  corresponds to the *i*'th input channel of *K*.

This formulation immediately hints at the possibility to sample over the input channel dimension, similarly to sampling column-row pairs in matrices. We propose to approximate convolutions by sampling lower-rank tensors:

$$\tilde{O} = \sum_{t=1}^{k} \frac{1}{k p_{i_t}} I^{[i_t]} * K_{[i_t]} \triangleq \tilde{I} * \tilde{K}$$
(30)

where  $\{i_t\}_{t=1}^k$  are such that  $i_t \in \{1, ..., IC\}$  and  $\{p_i\}_{i=1}^{IC}$  is a probability distribution over the input channels,  $\tilde{I}$  is a tensor composed of sampled input channels of I scaled by  $\sqrt{\frac{1}{kp_i}}$ , and  $\tilde{K}$  is a tensor composed of corresponding sampled input channels of K scaled by the same factor.

Computing the convolution of the smaller tensors  $\tilde{I} * \tilde{K}$  can be done using standard efficient convolution implementations. Figure 2 illustrates the sampling operation.

The properties of the approximation in (30) can be derived similarly to the CRS derivations for matrix multiplication. In particular, we prove the approximation is unbiased, and similar to matrix CRS, we use sampling probabilities proportional to the tensor Euclidean norms:

$$p_{i} = \frac{\left\|I^{[i]}\right\|_{F} \cdot \left\|K_{[i]}\right\|_{F}}{\sum_{j=1}^{IC} \left\|I^{[j]}\right\|_{F} \cdot \left\|K_{[j]}\right\|_{F}}$$
(31)

In section C.4 we show that the optimal sampling probabilities are significantly more complicated to calculate, but under certain conditions they reduce to (31).

Bernoulli-CRS and top-k algorithms can be developed for convolutions as well in an analogous way.

# C Proofs

#### C.1 Proofs for Section 5 - Approximate Backpropagation

**Theorem 1.** Let f(x, W, b) be a multi-layer neural network with  $\beta$ -Lipschitz activation functions  $\sigma$ . Let  $\ell$  be a  $\beta$ -Lipschitz loss function, and let the network be trained with SGD using properly decreasing learning rate. If the matrix products in the backward pass are approximated using an unbiased approximation scheme satisfying:

$$\mathbb{E}\left[A^{\top}B - \operatorname{approx}(A^{\top}B)\right] = 0$$

and:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|^{2}\right] \leq C \left|\left|A\right|\right|^{2} \left|\left|B\right|\right|^{2}$$

for some finite constant C and some norm  $||\cdot||$ ,

and if the weights are bounded, then the approximated gradients are unbiased with bounded second moments.

**Corollary.** Based on recent works on non-convex optimization [50], Theorem 1 implies that approximate backpropagation enjoys the same convergence guarantees as regular SGD training.

*Proof.* The network *f* can be described by:

$$h_1 = W_1^\top x + b_1$$

$$a_1 = \sigma(h_1)$$

$$h_l = W_l^\top a_{l-1} + b_l$$

$$a_l = \sigma(h_l)$$

$$\hat{y} = W_L^\top a_{L-1}$$

where  $x \in \mathbb{R}^n, W_1 \in \mathbb{R}^{n \times d_1}, W_l \in \mathbb{R}^{d_{l-1} \times d_l}, b_l \in \mathbb{R}^{d_l}, \ell$  is the number of layers and  $\hat{y} \in \mathbb{R}^{d_L}$  is the network output.

Let us denote the weight, bias and activation gradients with respect to a loss function  $\ell$  by  $\nabla W_l, \nabla b_l, \nabla a_l$  respectively. Let us denote and the gradients yielded by the approximation scheme as  $\nabla \tilde{W}_l, \nabla \tilde{b}_l, \nabla \tilde{a}_l$ .

Lemma 1.

$$\mathbb{E}\left[\nabla \tilde{W}_l\right] = \nabla W_l \quad and \quad \mathbb{E}\left[\nabla \tilde{b}_l\right] = \nabla b_l$$

*Proof.* We prove by induction. The last layer satisfies:

$$\nabla W_L = a_{L-1} \nabla \hat{y} \qquad \nabla a_{L-1} = W_L \nabla \hat{y}$$

and its approximation is given by:

$$abla ilde W_L = ext{approx}(a_{L-1} 
abla \hat y) \qquad 
abla ilde a_{L-1} = ext{approx}(W_L 
abla \hat y)$$

Since the approximation methods satisfies:

$$\mathbb{E}\left[A^{\top}B - \texttt{approx}(A^{\top}B)\right] = 0$$

we get:

$$\mathbb{E}\left[\nabla \tilde{W}_{L}\right] = \nabla W_{L} \qquad \mathbb{E}\left[\nabla \tilde{a}_{L-1}\right] = \nabla a_{L-1}$$

for the induction step, we will show that if  $\mathbb{E}[\nabla \tilde{a}_l] = \nabla a_l$  then:

$$\mathbb{E}\left[\nabla \tilde{W}_{l-1}\right] = \nabla W_{l-1}$$
$$\mathbb{E}\left[\nabla \tilde{b}_{l-1}\right] = \nabla b_{l-1}$$
$$\mathbb{E}\left[\nabla \tilde{a}_{l-1}\right] = \nabla a_{l-1}$$

 $\nabla \tilde{W}_{l-1}$  is given by:

$$\nabla \tilde{W}_{l-1} = \operatorname{approx}(a_{l-1} \nabla \tilde{h}_l^{\top}) = \operatorname{approx}(a_{l-1} \Sigma'(h_l) \nabla \tilde{a}_l^{\top})$$

where  $\Sigma'(h_l)$  is a diagonal matrix with the diagonal being the derivative of  $\sigma$  in location  $h_l$ . Taking the expectation we get:

$$\mathbb{E}\left[\nabla \tilde{W}_{l-1}\right] = \mathbb{E}\left[\operatorname{approx}(a_{l-1}\Sigma'(h_l)\nabla \tilde{a}_l^{\top})\right]$$
$$= \mathbb{E}\left[\mathbb{E}\left[\operatorname{approx}(a_{l-1}\Sigma'(h_l)\nabla \tilde{a}_l^{\top})|\nabla \tilde{a}_l^{\top}\right]\right]$$
$$= \mathbb{E}\left[a_{l-1}\Sigma'(h_l)\nabla \tilde{a}_l^{\top}\right]$$
$$= a_{l-1}\Sigma'(h_l)\nabla a_l^{\top}$$
$$= \nabla W_{l=1}$$

where we used the unbiased approximation property of approx and the law of total expectation. Similar arguments for  $\mathbb{E}[\nabla \tilde{a}_{l-1}]$  yield:

$$\begin{split} \mathbb{E}\left[\nabla \tilde{a}_{l-1}\right] &= \mathbb{E}\left[\operatorname{approx}(W_l \nabla \tilde{h}_l)\right] \\ &= \mathbb{E}\left[\operatorname{approx}(W_l \Sigma'(h_l) \nabla \tilde{a}_l)\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\operatorname{approx}(W_l \Sigma'(h_l) \nabla \tilde{a}_l) | \nabla \tilde{a}_l\right]\right] \\ &= \mathbb{E}\left[W_l \Sigma'(h_l) \nabla \tilde{a}_l\right] \\ &= \nabla a_{l-1} \end{split}$$

and for  $\mathbb{E}\left[\nabla \tilde{b}_{l-1}\right]$ :

$$\mathbb{E}\left[\nabla \tilde{b}_{l-1}\right] = \mathbb{E}\left[\nabla \tilde{h}_{l-1}\right]$$
$$= \mathbb{E}\left[\Sigma'(h_l)\nabla \tilde{a}_l\right]$$
$$= \Sigma'(h_l)\nabla a_l$$
$$= \nabla b_{l-1}$$

In other words, the unbiased estimation of the gradients follows from the linearity of backpropagation with respect to the gradients, even for non-linear activation functions.

We can write the training step using SGD and the approximate gradients  $\nabla \tilde{W}_l^t$  for layer l at iteration t as:

$$W_l^{t+1} = W_l^t - \alpha_t (\nabla W_l^t + \omega_t)$$

where  $\omega_t$  is a gradient noise defined as:

$$\omega_t \triangleq \nabla \tilde{W}_l^t - \nabla W_l^t$$

Based on Lemma 1, the gradient noise  $\omega_t$  is a martingale difference sequence satisfying:

$$\mathbb{E}\left[\omega_t | W_{t-1}\right] = \mathbb{E}\left[\nabla \tilde{W}_l^t - \nabla W_l^t | W_{t-1}\right] = 0$$

Lemma 2. Under the assmuptions in Theorem 1:

$$\mathbb{E}\left[\left|\left|\omega_{t}\right|\right|^{2}\left|W_{t-1}\right] < D$$

for some constant D.

*Proof.* We prove by induction. Since  $\ell$  is  $\beta$ -Lipschitz, the gradients  $\nabla y$  are bounded. During backpropagation the gradients are propagated by:

$$\nabla \tilde{a}_{l-1} = \operatorname{approx}(W_l \Sigma'(h_l) \nabla \tilde{a}_l)$$

Let us assume  $\nabla \tilde{a}_l$  is bounded and show that  $\nabla \tilde{a}_{l-1}$  is bounded in expectation as well:

$$\mathbb{E}\left[\left|\left|\nabla\tilde{a}_{l-1}\right|\right|^{2}\right] \leq \mathbb{E}\left[\left|\left|\nabla\tilde{a}_{l-1}-W_{l}\Sigma'(h_{l})\nabla\tilde{a}_{l}\right|\right|^{2}\right] + \mathbb{E}\left[\left|\left|W_{l}\Sigma'(h_{l})\nabla\tilde{a}_{l}\right|\right|^{2}\right] \\ \leq C\left|\left|W_{l}\right|\right|^{2}\left|\left|\Sigma'(h_{l})\nabla\tilde{a}_{l}\right|\right|^{2} \\ < D'$$

for some constant D', where the second inequality follows from the properties of approx and last inequality follows from the  $\beta$ -Lipschitz of  $\Sigma$ , the induction assumption on the boundness of  $\nabla \tilde{a}_l$  and the assumption on the boundness of  $W_l$ .

The gradients  $\nabla \tilde{W}$  are calculated by:

$$\nabla W_{l-1} = \texttt{approx}(a_{l-1}\Sigma'(h_l)\nabla \tilde{a}_l^{+})$$

and therefore:

$$\mathbb{E}\left[\left|\left|\omega_{t}\right|\right|^{2}\left|W_{t-1}\right] = \mathbb{E}\left[\left|\left|\nabla\tilde{W}_{l}^{t}-\nabla W_{l}^{t}\right|\right|^{2}\left|W_{t-1}\right]\right]$$

$$= \mathbb{E}\left[\left|\left|\nabla\tilde{W}_{l}^{t}-\left(a_{l-1}\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right)+\left(a_{l-1}\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right)+\nabla W_{l}^{t}\right|\right|^{2}\left|W_{t-1}\right]$$

$$\leq \mathbb{E}\left[\left|\left|\nabla\tilde{W}_{l}^{t}-a_{l-1}\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right|\right|^{2}\right] + \mathbb{E}\left[\left|\left|a_{l-1}\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right|\right|^{2}\right] + \mathbb{E}\left[\left|\left|\nabla W_{l}^{t}\right|\right|^{2}\right]$$

$$\leq C_{1}\left|\left|a_{l-1}\right|\right|^{2}\mathbb{E}\left[\left|\left|\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right|\right|^{2}\right] + C_{2}\left|\left|a_{l-1}\right|\right|^{2}\mathbb{E}\left[\left|\left|\Sigma'(h_{l})\nabla\tilde{a}_{l}^{\top}\right|\right|^{2}\right] + \mathbb{E}\left[\left|\left|\nabla W_{l}^{t}\right|\right|^{2}\right]$$

$$\leq D$$

In the second inequality we used the properties of approx. In the last inequality we used the boundness of  $\Sigma', \nabla W_l^t$  from the assumptions, the boundness of  $\mathbb{E}\left[||\nabla \tilde{a}_l^\top||^2\right]$  from above. In addition, we assumed boundness of the activations  $a_{l-1}$ . This assumption holds if the activation function  $\sigma$  is bounded (for example, sigmoid) and in the general case it also requires the assumptions on the boundness of weights and inputs.

The same arguments can be made for the bias and the approximate bias gradients.

Based on Lemmas 1 and 2 and using standard analysis of SGD (for example [56] and [50]) the SGD convergence guarantees hold for approximate backpropagation as well.  $\Box$ 

**Remark.** Both CRS and Bernoulli-CRS satisfy the property

$$\mathbb{E}\left[\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|^{2}\right] \leq C \left|\left|A\right|\right|^{2} \left|\left|B\right|\right|^{2}$$

since the expected Frobenius norm for the error matrix satisfies:

$$\begin{split} \mathbb{E} \left[ \left| \left| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right| \right|_{F}^{2} \right] &= \\ \frac{1}{k} \left( \sum_{i=1}^{n} e_{i} |A_{(i)}| |B_{(i)}| \right)^{2} - \sum_{i=1}^{n} |A_{(i)}|^{2} |B_{(i)}|^{2} \\ &\leq \left( \sum_{i=1}^{n} e_{i} |A_{(i)}| |B_{(i)}| \right)^{2} \\ &\leq \left( \sum_{i=1}^{n} |A_{(i)}|^{2} \right) \left( \sum_{i=1}^{n} |B_{(i)}|^{2} \right) \\ &= \left| |A| \right|_{F}^{2} \left| |B| \right|_{F}^{2} \end{split}$$

where we used Theorem 3 and the Cauchy-Schwarz inequality.

**Corollary.** Let f(x, W, b) be a multi-layer neural network with bounded  $\beta$ -Lipschitz activation functions  $\sigma$ . Let  $\ell$  be a  $\beta$ -Lipschitz loss function, and let the network be trained with SGD using properly decreasing learning rate. If the weight gradient matrix products in the backward pass are approximated using an unbiased approximation scheme satisfying:

$$\mathbb{E}\left[A^{\top}B - \operatorname{approx}(A^{\top}B)\right] = 0$$

and:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \operatorname{approx}(A^{\top}B)\right|\right|^{2}\right] \leq C \left|\left|A\right|\right|^{2} \left|\left|B\right|\right|^{2}$$

for some finite constant C and some norm  $||\cdot||$ ,

then then the approximated gradients are unbiased with bounded second moments.

Proof. Lemma 1 under these assumptions holds by the same arguments. We now prove the equivalent of Lemma 2:

$$\mathbb{E}\left[\left|\left|\omega_{t}\right|\right|_{F}^{2}\left|W_{t-1}\right]\right] = \mathbb{E}\left[\left|\left|\nabla\tilde{W}_{l}^{t}-\nabla W_{l}^{t}\right|\right|_{F}^{2}\left|W_{t-1}\right]\right]$$
$$= \mathbb{E}\left[\left|\left|\nabla\tilde{W}_{l}^{t}-a_{l-1}\Sigma'(h_{l})\nabla a_{l}^{\top}\right|\right|_{F}^{2}\right]$$
$$\leq \left|\left|a_{l-1}\right|\right|_{F}^{2}\left|\left|\Sigma'(h_{l})\nabla a_{l}^{t}\right|\right|_{F}^{2}$$
$$\leq D$$

The first inequality follows from the properties of approx. The second inequality follows from the  $\beta$ -Lipschitz property of  $\ell, \Sigma$  bounding the second term, and from the boundness of the activation function  $\sigma$  bounding the first term. П

#### C.2 Proofs for Section 6 - Sampling Without Scaling and Top-k Selection

**Theorem 2.** Let A be a  $n \times m$  random matrix and B be  $n \times p$  random matrix, such that - -

$$\mathbb{E}\left[A^{+}(i)B_{(i)}\right] = 0$$

\_

for  $1 \le i \le n$ . Assume k column-row pairs with indices  $\{j\}_1^n$  are sampled from A and B.

Then, the MMSE estimator for the matrix product  $A^{\top}B$  would be  $\tilde{A}^{\top}\tilde{B}$  where  $\tilde{A}, \tilde{B}$  are constructed from the sampled column-row pairs without scaling.

Furthermore, if  $A^{\top(i)}B_{(i)}$  and  $A^{\top(j)}, B_{(j)}$  are independent for different *i* and *j* then the MSE will be minimized when sampling k pairs with the maximum norm multiplication  $|A_{(i)}||B_{(i)}|$ .

*Proof.* Given sampled pairs  $j_1, ..., j_k$  the MMSE estimator would be:

$$\begin{split} \widehat{A}^{\top}\widehat{B} &= \mathbb{E}\left[A^{\top}B|A_{(j_{1})}, ..., A_{(j_{k})}, B_{(j_{1})}, ..., B_{(j_{k})}\right] \\ &= \mathbb{E}\left[\sum_{i=1}^{k} A_{(j_{i})}^{\top}B_{(j_{i})} + \sum_{i \notin \{j\}_{1}^{k}} A^{\top(i)}B_{(i)}|A_{(j_{1})}, ..., A_{(j_{k})}, B_{(j_{1})}, ..., B_{(j_{k})}\right] \\ &= \sum_{i=1}^{k} A_{(j_{i})}^{\top}B_{(j_{i})} + \sum_{i \notin \{j\}_{1}^{k}} \mathbb{E}\left[A^{\top(i)}B_{(i)}\right] \\ &= \sum_{i=1}^{k} A_{(j_{i})}^{\top}B_{(j_{i})} \\ &= \widetilde{A}^{\top}\widetilde{B} \end{split}$$

The MSE would be:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right] = \mathbb{E}\left[\left|\left|\sum_{i \notin \{j\}_{1}^{k}} A^{\top(i)}B_{(i)}\right|\right|_{F}^{2}\right]$$

if we assume independence between different column-row pairs  $A^{\top(i)}B_{(i)}, A^{\top(j)}B_{(j)}$  then the last expression reduces to:

$$\sum_{i \notin \{j\}_1^k} \mathbb{E}\left[ \left| \left| A^{\top(i)} B_{(i)} \right| \right|_F^2 \right] = \sum_{i \notin \{j\}_1^k} \mathbb{E}\left[ |A_{(i)}|^2 |B_{(i)}|^2 \right]$$

and therefore will be minimized for a top-k selection scheme that samples the pairs with the highest norm.

## C.3 Proofs for Section A - Bernoulli-CRS

**Proposition 1.**  $\mathbb{E}\left[\tilde{A}^{\top}\tilde{B}\right] = A^{\top}B$ 

Proof.

$$\mathbb{E} \left[ A^{\top} P K K P B \right] = A^{\top} P P \mathbb{E} \left[ K K \right] B$$
$$= A^{\top} P P \mathbb{E} \left[ K \right] B$$
$$= A^{\top} B$$

where we used that fact that K is diagonal and that  $K_{i,i} \in \{0,1\}$ .

**Proposition 2.**  $\mathbb{E}[T] = k$ 

Proof.

$$\mathbb{E}[T] = \mathbb{E}\left[\sum_{j=1}^{n} K_{j,j}\right] = \sum_{j=1}^{n} \mathbb{E}[K_{j,j}] = \sum_{j=1}^{n} p_j = k$$

**Proposition 3.** 

$$\mathbf{Var}\left[(\tilde{A}^{\top}\tilde{B})_{i,j}\right] = \sum_{t=1}^{n} \frac{1-p_t}{p_t} A_{t,i}^2 B_{t,j}^2$$

*Proof.* Fix *i*, *j*. From Proposition 1:

$$\mathbb{E}\left[ (\tilde{A}^{\top} \tilde{B})_{i,j} \right] = (A^{\top} B)_{i,j}$$

Calculating the second moment:

 $\mathbb E$ 

$$\begin{split} \left[ (\tilde{A}^{\top} \tilde{B})_{i,j}^{2} \right] &= \mathbb{E} \left[ \left( \sum_{t=1}^{n} A_{t,i} \frac{K_{t,t}}{p_{t}} B_{t,j} \right)^{2} \right] \\ &= \mathbb{E} \left[ \sum_{t=1}^{n} \sum_{u=1}^{n} A_{t,i} \frac{K_{t,t}}{p_{t}} B_{t,j} A_{u,i} \frac{K_{u,u}}{p_{u}} B_{u,j} \right] \\ &= \mathbb{E} \left[ \sum_{t=1}^{n} \sum_{u \neq t}^{n} A_{t,i} \frac{K_{t,t}}{p_{t}} B_{t,j} A_{u,i} \frac{K_{u,u}}{p_{u}} B_{u,j} \right] \\ &+ \mathbb{E} \left[ \sum_{t=1}^{n} A_{t,i}^{2} B_{t,j}^{2} \frac{K_{t,t}}{p_{t}^{2}} \right] \\ &= \sum_{t=1}^{k} \sum_{u \neq t}^{k} A_{t,i} B_{t,j} A_{u,i} B_{u,j} + \sum_{t=1}^{n} \frac{1}{p_{t}} A_{t,i}^{2} B_{t,j}^{2} \\ &= (A^{\top} B)_{i,j}^{2} - \sum_{t=1}^{n} A_{t,i}^{2} B_{t,j}^{2} + \sum_{t=1}^{n} \frac{1-p_{t}}{p_{t}} A_{t,i}^{2} B_{t,j}^{2} \end{split}$$

Therefore:

$$\mathbf{Var}\left[(\tilde{A}^{\top}\tilde{B})_{i,j}\right] = \mathbb{E}\left[(\tilde{A}^{\top}\tilde{B})_{i,j}^{2}\right] - \mathbb{E}\left[(\tilde{A}^{\top}\tilde{B})_{i,j}\right]^{2}$$
$$= \sum_{t=1}^{k} \frac{1 - p_{t}}{p_{t}} A_{t,i}^{2} B_{t,j}^{2}$$

**Theorem 3.** The expected Frobenius norm of the error matrix  $\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right]$  is  $\sum_{t=1}^{n} \frac{1-p_{t}}{p_{t}} |A_{(t)}|^{2} |B_{(t)}|^{2}.$ 

Furthermore, under the constraint  $\sum_{i=1}^{n} p_i = k$  it is minimized for the probabilities:

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbb{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbb{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}}$$

where  $\mu$  is a root of the following function:

$$G(\mu) := \sum_{i=1}^n \left( \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbf{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbf{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}} \right) - k$$

Proof. Note:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right] = \sum_{i=1}^{m} \sum_{j=1}^{p} \mathbb{E}\left[\left(A^{\top}B - \tilde{A}^{\top}\tilde{B}\right)_{i,j}^{2}\right]$$
$$= \sum_{i=1}^{m} \sum_{j=1}^{p} \mathbf{Var}\left[\left(\tilde{A}^{\top}\tilde{B}\right)_{i,j}\right]$$

Therefore, using Proposition 3 we get:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right] = \sum_{i=1}^{m} \sum_{j=1}^{p} \sum_{t=1}^{n} \frac{1 - p_{t}}{p_{t}} A_{t,i}^{2} B_{t,j}^{2}$$
$$= \sum_{t=1}^{n} \frac{1 - p_{t}}{p_{t}} \left(\sum_{i=1}^{m} A_{t,i}^{2}\right) \left(\sum_{j=1}^{p} B_{t,j}^{2}\right)$$
$$= \sum_{t=1}^{n} \frac{1 - p_{t}}{p_{t}} |A_{(t)}|^{2} |B_{(t)}|^{2}$$

Let us now find the optimal sampling probabilities that minimize the Frobenius error. Define the function:

$$f(p_1, p_2, ..., p_n) = \sum_{t=1}^n \frac{1-p_t}{p_t} |A_{(t)}|^2 |B_{(t)}|^2$$

We can now consider the optimization problem:

$$\min_{p_1,\dots,p_n} \quad f(p_1,\dots,p_n) \\ \text{s.t} \quad p_i - 1 \le 0 \\ -p_i \le 0 \\ \sum_{i=1}^n p_i - k = 0$$

We define the Lagrangian as:

$$L(p_1, ..., p_n, \lambda_1, ..., \lambda_n, \nu_1, ..., \nu_n, \mu) \triangleq f(p_1, p_2, ..., p_k) + \sum_{i=1}^n \lambda_i (p_i - 1) - \sum_{i=1}^n \nu_i p_i + \mu \left(\sum_{i=1}^n p_i - k\right)$$

where  $\lambda_i \geq 0$ ,  $\nu_i \geq 0$  and  $\mu \in \mathbb{R}$ .

Applying KKT stationarity condition:

$$0 = \frac{\partial}{\partial p_i} L = -\frac{1}{p_i^2} |A_{(i)}|^2 |B_{i}|^2 + \lambda_i - \nu_i + \mu = 0$$

Therefore:

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\lambda_i - \nu_i + \mu}}$$

Next we divide into 3 cases,

**Case 1:** If  $p_i \in (0,1)$ : In this case due to complementary-slackness we obtain  $\lambda_i = \nu_i = 0$ , and therefore,

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}}$$

Case 2: If  $p_i = 1$ : In this case due to complementary-slackness we obtain  $\nu_i = 0$ , and therefore,

$$1 = p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu + \lambda_i}}$$

**Case 3:** If  $p_i = 0$ : In this case due to complementary-slackness we obtain  $\lambda_i = 0$ , which implies that,

$$0 = p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu - \nu_i}}$$

but this can only happen if  $|A_{(i)}||B_{(i)}| = 0$ .

Combining the above we conclude that given  $\mu$  one can write the solution as follows,

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbf{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbf{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}}$$

Now, in order to satisfy the equality conditions  $\mu$  should satisfy the following equality,

$$\sum_{i=1}^{n} \left( \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbb{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbb{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}} \right) = k$$

Now, one can actually find  $\mu$  using bisection, To see this consider the following function,

$$G(\mu) := \sum_{i=1}^{n} \left( \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbb{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbb{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}} \right) - k$$

And note that  $G(\mu)$  is a one dimensional monotonically decreasing (actually non-increasing) function of  $\mu$ .

Also, if we sorts the  $|A_{(i)}||B_{(i)}|$ 's, i.e.  $|A_{(1)}||B_{(1)}| \le |A_{(2)}||B_{(2)}| \le \dots |A_{(n)}||B_{(n)}|$ , then given j such that  $\mu \in (|A_{(j)}||B_{(j)}|, |A_{(j+1)}||B_{(j+1)}|)$ , then we can find the exact value of  $\mu$  from the equality constraints equation:

$$\sum_{i=1}^{n} \left( \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\mu}} \mathbb{1}_{\{0 < |A_{(i)}||B_{(i)}| < \mu\}} + \mathbb{1}_{\{|A_{(i)}||B_{(i)}| \ge \mu\}} \right) = k$$

Corollary. The sampling probabilities

$$p_{i} = \min\left\{\frac{k|A_{(i)}||B_{(i)}|}{\sum_{j=1}^{n}|A_{(j)}||B_{(j)}|}, 1\right\}$$

are optimal if  $k \leq \frac{\sum_{i=1}^{n} |A_{(i)}| |B_{(i)}|}{\max_i |A_{(i)}| |B_{(i)}|}$ 

*Proof.* As a simpler, sub-optimal solution for the above optimization problem we propose the following relaxation. First, we solve the optimization problem without the inequality conditions:

$$0 \le p_i \le 1$$

Then, for each optimal  $p_i^*$  we clamp the value between the range [0, 1]. This allows us to comply with the inequality conditions that allows to treat  $p_i$  as a parameter to Bernoulli distribution at the expense of relaxing the constraint on the sum of the parameters  $p_i$ , leading to potentially sub-optimal solution.

As the first step, we therefore solve the problem:

$$\min_{p_1,\dots,p_n} \quad f(p_1,\dots,p_n)$$
  
s.t 
$$\sum_{i=1}^n p_i - k$$

To minimize f subject to the constraint  $\sum_{i=1}^{n} p_i = k$  we use the Lagrange multiplier  $\lambda$  and define the function:

$$g(p_1, p_2, ..., p_n) = f(p_1, p_2, ..., p_n) + \lambda \left(\sum_{i=1}^n p_i - k\right)$$

Deriving and equaling to zero we get:

$$0 = \frac{\partial g}{\partial p_i} = -\frac{1}{p_i^2} |A_{(i)}|^2 |B_{ij}|^2 + \lambda$$

Therefore:

$$p_i = \frac{|A_{(i)}||B_{(i)}|}{\sqrt{\lambda}}$$

Substituting in  $\sum_{i=1}^{n} p_i = k$ :

$$\sum_{i=1}^{n} \frac{|A_{(i)}||B_{i}||}{\sqrt{\lambda}} = k$$
$$\sqrt{\lambda} = \frac{\sum_{i=1}^{n} |A_{(i)}||B_{(i)}|}{k}$$

And therefore we get:

$$p_i = \frac{k|A_{(i)}||B_{(i)}|}{\sum_{i=1}^n |A_{(i)}||B_{(i)}|}$$

And the final result after clamping would be:

$$p_i = \min\left\{\frac{k|A_{(i)}||B_{(i)}|}{\sum_{i=1}^n |A_{(i)}||B_{(i)}|}, 1\right\}$$

Note that this solution yields  $p_i \ge 0$ , satisfying one of the original inequality conditions. What about  $p_i \le 1$ ?

If

$$k \le \frac{\sum_{i=1}^{n} |A_{(i)}| |B_{(i)}|}{\max_{i} |A_{(i)}| |B_{(i)}|}$$

then the second inequality conditions holds as well and the solution is indeed the optimal solution to the original problem.

Substituting in the expression for the Frobenius error we get:

$$\mathbb{E}\left[\left|\left|A^{\top}B - \tilde{A}^{\top}\tilde{B}\right|\right|_{F}^{2}\right] = \sum_{t=1}^{n} \frac{1 - p_{t}}{p_{t}} |A_{(t)}|^{2} |B_{(t)}|^{2}$$
$$= \frac{1}{k} \left(\sum_{i=1}^{n} |A_{(i)}| |B_{(i)}|\right)^{2} - \sum_{i=1}^{n} |A_{(i)}|^{2} |B_{(i)}|^{2}$$

The following theorem yields high probability bounds for the Frobenius and spectral norms for the Bernoulli-CRS algorithm:

**Theorem 4.** Let  $A \in \mathbb{R}^{n \times m}$  and  $B \in \mathbb{R}^{n \times p}$ . Let  $\tilde{A}, \tilde{B}$  be the sampled matrices according to the Bernoulli-CRS algorithm described above. Denote

$$R \triangleq \max_{i} \left| \left| A^{\top(i)} B_{(i)} \right| \right|$$

and

$$\sigma^{2} \triangleq \frac{1}{k} \left( \sum_{i=1}^{n} e_{i} |A_{(i)}| |B_{(i)}| \right)^{2} - \sum_{i=1}^{n} e_{i} |A_{(i)}|^{2} |B_{(i)}|^{2}$$

then, for all  $t \ge 0$ :

$$\mathbb{P}\left\{ \left| \left| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right| \right| \ge t \right\} \le (m+p) \cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right)$$
$$\mathbb{P}\left\{ \left| \left| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right| \right|_F \ge t \right\} \le (m+p)^{3/2} \cdot \exp\left(\frac{-t^2/2}{\sigma^2 + Rt/3}\right)$$

Proof. The Matrix Bernstein concentration inequality states:

**Theorem (Matrix Bernstein [57]).** Consider a finite sequence  $\{\mathbf{Z}_k\}$  of independent, random matrices with dimensions  $d_1 \times d_2$ . Assume that each random matrix satisfies

$$\mathbb{E}[\mathbf{Z}_k] = 0$$
 and  $||\mathbf{Z}_k|| \leq R$  almost surely.

Define

$$\sigma^{2} \triangleq \max\left\{ \left\| \sum_{k} \mathbb{E}\left[ \mathbf{Z}_{k} \mathbf{Z}_{k}^{\top} \right] \right\|, \left\| \sum_{k} \mathbb{E}\left[ \mathbf{Z}_{k}^{\top} \mathbf{Z}_{k} \right] \right\| \right\}$$

Then, for all  $t \geq 0$ ,

$$\mathbb{P}\left\{\left\|\sum_{k} \mathbf{Z}_{k}\right\| \ge t\right\} \le (d_{1} + d_{2}) \cdot \exp\left(\frac{-t^{2}/2}{\sigma^{2} + Rt/3}\right)$$

In our sampling algorithm, we can define:

$$\mathbf{Z}_{k} \triangleq A_{(k)}^{\top} B_{(k)} - \frac{1}{p_{k}} K_{k,k} A_{(k)}^{\top} B_{(k)}$$

when  $K_{k,k}$  is a Bernoulli random variable with parameter  $p_k$  as defined above. It is clear that  $\mathbb{E}[\mathbf{Z}_k] = 0$ .

Also, let us define:

$$R \triangleq \max_{k} \left| \left| A_{(k)}^{\top} B_{(k)} \right| \right|$$

so it it also clear that  $||\mathbf{Z}_k|| \leq R$ .

By construction,  $\{\mathbf{Z}_k\}$  are independent.

We can also define:

$$\sigma^{2} \triangleq \max\left\{ \left\| \sum_{k} \mathbb{E} \left[ \mathbf{Z}_{k} \mathbf{Z}_{k}^{\top} \right] \right\|, \left\| \sum_{k} \mathbb{E} \left[ \mathbf{Z}_{k}^{\top} \mathbf{Z}_{k} \right] \right\| \right\}$$
$$= \max\left\{ \left\| \mathbb{E} \left[ \sum_{k} \mathbf{Z}_{k} \mathbf{Z}_{k}^{\top} \right] \right\|, \left\| \mathbb{E} \left[ \sum_{k} \mathbf{Z}_{k}^{\top} \mathbf{Z}_{k} \right] \right\| \right\}$$
$$= \max\left\{ \left\| \mathbb{E} \left[ (A^{\top}B - \tilde{A}^{\top}\tilde{B})(A^{\top}B - \tilde{A}^{\top}\tilde{B})^{\top} \right] \right\|,$$
$$\left\| \mathbb{E} \left[ (A^{\top}B - \tilde{A}^{\top}\tilde{B})^{\top}(A^{\top}B - \tilde{A}^{\top}\tilde{B}) \right] \right\| \right\}$$
$$\leq \max\left\{ \operatorname{Tr} \left( \mathbb{E} \left[ (A^{\top}B - \tilde{A}^{\top}\tilde{B})(A^{\top}B - \tilde{A}^{\top}\tilde{B})^{\top} \right] \right),$$
$$\operatorname{Tr} \left( \mathbb{E} \left[ (A^{\top}B - \tilde{A}^{\top}\tilde{B})^{\top}(A^{\top}B - \tilde{A}^{\top}\tilde{B}) \right] \right) \right\}$$
$$= \mathbb{E} \left[ \left\| A^{\top}B - \tilde{A}^{\top}\tilde{B} \right\|_{F}^{2} \right]$$
$$= \frac{1}{k} \left( \sum_{i=1}^{n} e_{i} |A_{(i)}| |B_{(i)}| \right)^{2} - \sum_{i=1}^{n} e_{i} |A_{(i)}|^{2} |B_{(i)}|^{2}$$

where we used the linearity of expectation and trace, the property  $||A|| \leq \text{Tr}[A]$  for positive semidefinite matrices and the expected Frobenius norm from Theorem 3.

The bound on the spectral norm follows immediately from the Matrix Berenstein inequality.

Using the property:

$$||A||_F \le \sqrt{r} \, ||A||$$

26

we get the similar result for the Frobenius norm, factored by  $\sqrt{m+p}$ .

#### C.4 Proofs for Section B - Approximating Convolutions

The following proofs go along the same lines of [20], generalizing them to multi-channel convolutions (zero-padding assumed).

**Lemma 3.** Suppose  $I \in \mathbb{R}_{B \times IH}^{IW \times IC}$ ,  $K \in \mathbb{R}_{KW \times KW}^{IC \times OC}$ ,  $1 \le k \le IC$ ,  $\{p_i\}_{i=1}^{IC}$  is a probability distribution over  $\{1, ..., IC\}$  and  $\{i_t\}_{t=1}^k$  are such that  $i_t \in \{1, ..., IC\}$ .

Let  $O \in \mathbb{R}_{B \times OH}^{OW \times OC} = I * K$  be the multi-channel convolution of I, K as defined in (28) and let  $\tilde{O}$  be its approximation by sampling k input channels as defined in (30). Then:

$$\mathbb{E}\left[\tilde{O}\right] = O$$

*Proof.* We show that every b, oh, ow, oc satisfies  $\mathbb{E}\left[\tilde{O}_{b,oh}^{ow,oc}\right] = O_{b,oh}^{ow,oc}$ .

For  $t \in \{1, ..., k\}$ , define  $X_t = (\frac{I^{[i_t]} * K_{[i_t]}}{p_{i_t}})^{ow,oc}_{b,oh}$ .

Using (30) we can write  $\tilde{O}_{b,oh,ow,oc} = \sum_{t=1}^{k} \frac{1}{k} X_t$ .

Taking the expectation, we get:

$$(\mathbb{E}\left[\tilde{O}\right])_{b,oh}^{ow,oc} = \mathbb{E}\left[\sum_{t=1}^{k} \frac{1}{k} X_t\right] = \mathbb{E}\left[X_t\right] = \sum_{i=1}^{IC} p_i \cdot \frac{(I^{[i]} * K_{[i]})_{b,oh}^{ow,oc}}{p_i} = O_{b,oh}^{ow,oc}$$
(32)

Lemma 4. Suppose the same as Lemma 3. Then:

$$\begin{aligned} \mathbf{Var}\left[\tilde{O}_{b,oh}^{ow,oc}\right] &= \frac{1}{k} \sum_{i=1}^{IC} \frac{1}{p_i} \sum_{h=1}^{KH} \sum_{w=1}^{KW} (I_{b,oh+h-1}^{ow+w-1,i})^2 (K_{h,w}^{i,oc})^2 \\ &+ \frac{1}{k} \sum_{i=1}^{IC} \frac{1}{p_i} \sum_{\substack{h,h'=1\\h\neq h'}}^{KH} \sum_{w\neq w'}^{KW} I_{b,oh+h-1}^{ow+w-1,i} I_{b,oh+h'-1}^{ow+w'-1,i} K_{h,w}^{i,oc} K_{h',w'}^{i,oc} \\ &- \frac{1}{k} (O_{b,oh}^{ow,oc})^2 \end{aligned}$$

*Proof.* Define  $X_t$  as in Lemma 3. From (30) and the independence of different  $X_t$ :

$$\mathbf{Var}\left[\tilde{O}_{b,oh}^{ow,oc}\right] = \mathbf{Var}\left[\sum_{t=1}^{k} \frac{1}{k} X_{t}\right] = \frac{1}{k} \mathbf{Var}\left[X_{t}\right] = \frac{1}{k} (\mathbb{E}\left[X_{t}^{2}\right] - \mathbb{E}\left[X_{t}\right]^{2})$$
(33)

$$\mathbb{E}\left[X_{t}^{2}\right] = \sum_{i=1}^{IC} p_{i} \cdot \frac{\left(\left(I^{[i]} * K_{[i]}\right)_{b,oh}^{ow,oc}\right)^{2}}{p_{i}^{2}}$$

$$= \sum_{i=1}^{IC} \frac{1}{p_{i}} \left(\sum_{h=1}^{KH} \sum_{w=1}^{KW} I_{b,oh+h-1}^{ow+w-1,i} K_{h,w}^{i,oc}\right)^{2}$$
(34)

From (32) we get  $\mathbb{E}[X_t] = O$ .

Substituting both expressions in (33) and expanding concludes the proof.

**Lemma 5.** *Suppose the same as Lemma 3. Then:* 

$$\mathbb{E}\left[\left\|O - \tilde{O}\right\|_{F}^{2}\right] = \sum_{i=1}^{IC} \frac{\left\|I^{[i]}\right\|_{F}^{2} \cdot \left\|K_{[i]}\right\|_{F}^{2} - E_{IK}^{i} + R_{IK}^{i}}{kp_{i}} - \frac{1}{k}\left\|O\right\|_{F}^{2}$$

where

$$\begin{split} E_{IK}^{i} &= \sum_{b=1}^{B} \sum_{\substack{oh,ow \ s.t. \\ oh < KH \ or \\ ow < KW}} \sum_{\substack{oc=1 \ h=1}}^{KH} \sum_{\substack{h,w \ s.t. \\ h > oh \ or \\ w > ow}}^{OC} (I_{b,oh}^{ow,i})^{2} (K_{h,w}^{i,oc})^{2} \\ R_{IK}^{i} &= \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \sum_{\substack{h,h'=1 \\ h \neq h'}}^{KH} \sum_{\substack{w,w'=1 \\ w \neq w'}}^{KW} I_{b,oh+h-1}^{ow+w'-1,i} I_{b,oh+h'-1}^{ow+w'-1,i} K_{h,w}^{i,oc} K_{h',w'}^{i,oc} \end{split}$$

The expected error is minimized when the sampling probabilities are:

$$p_{i} = \frac{\sqrt{\left\|I^{[i]}\right\|_{F}^{2} \cdot \left\|K_{[i]}\right\|_{F}^{2} - E_{IK}^{i} + R_{IK}^{i}}}{\sum_{j=1}^{IC} \sqrt{\left\|I^{[j]}\right\|_{F}^{2} \cdot \left\|K_{[j]}\right\|_{F}^{2} - E_{IK}^{j} + R_{IK}^{j}}}$$

**Remark.** We use here the Frobenius norm in its generalization for tensors. For a tensor T of rank r:

$$||T||_F = \sqrt{\sum_{j_1, j_2, \dots, j_r} T_{j_1, j_2, \dots, j_r}^2}$$

*Proof.* Note that:

$$\mathbb{E}\left[\left\|O - \tilde{O}\right\|_{F}^{2}\right] = \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \mathbb{E}\left[\left((O - \tilde{O})_{b,oh}^{ow,oc}\right)^{2}\right] = \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \mathbf{Var}\left[\tilde{O}_{b,oh}^{ow,oc}\right]$$

Substituting the result from Lemma 4:

$$\mathbb{E}\left[\left\|O-\tilde{O}\right\|_{F}^{2}\right] = \sum_{i=1}^{IC} \frac{1}{kp_{i}} \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \sum_{h=1}^{KH} \sum_{w=1}^{KW} (I_{b,oh+h-1}^{ow+w-1,i})^{2} (K_{h,w}^{i,oc})^{2} + \sum_{i=1}^{IC} \frac{1}{kp_{i}} \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \sum_{h,h'=1}^{KH} \sum_{w,w'=1}^{KW} I_{b,oh+h-1}^{ow+w-1,i} I_{b,oh+h'-1}^{ow+w'-1,i} K_{h,w}^{i,oc} K_{h',w'}^{i,oc} - \frac{1}{k} \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} (O_{b,oh}^{ow,oc})^{2}$$

$$(35)$$

This expression includes 3 terms. The first involves products between each element of  $I^{[i]}$  and all the corresponding entries in  $K_{[i]}$ , except for the upper and left edges of  $I^{[i]}$ . We therefore add and subtract the correction term  $E^i_{IK}$  to get:

$$\begin{split} &\sum_{i=1}^{IC} \frac{1}{kp_i} \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} \sum_{h=1}^{KH} \sum_{w=1}^{KW} (I_{b,oh+h-1}^{ow+w-1,i})^2 (K_{h,w}^{i,oc})^2 \\ &= \sum_{i=1}^{IC} \frac{1}{kp_i} \left( \left( \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} (I_{b,oh+h-1}^{ow+w-1,i})^2 \right) \left( \sum_{oc=1}^{OC} \sum_{h=1}^{KH} \sum_{w=1}^{KW} (K_{h,w}^{i,oc})^2 \right) - E_{IK}^i \right) \\ &= \sum_{i=1}^{IC} \frac{\left\| I^{[i]} \right\|_F^2 \cdot \left\| K_{[i]} \right\|_F^2 - E_{IK}^i}{kp_i} \end{split}$$

The second term is  $\sum_{i=1}^{IC} \frac{1}{kp_i} R_{IK}^i$ .

The third term can be written as  $\frac{1}{k} \sum_{b=1}^{B} \sum_{oh=1}^{OH} \sum_{ow=1}^{OW} \sum_{oc=1}^{OC} (O_{b,oh,ow,oc})^2 = \frac{1}{k} \|O\|_F^2$ Substituting these terms in (35) yields the result of (5). To find  $\{p_i\}_{i=1}^{IC}$  that minimize the expression in (5) it is enough to minimize the function  $f = \sum_{i=1}^{IC} \frac{\alpha_i^2}{p_i}$  under the constraints  $\sum p_i = 1$  and  $p_i > 0$ . We can write the numerator as  $\alpha_i^2$  because the expression in (34) is non-negative.

This minimization problem has a straightforward solution in Lemma 4 of [20], which is  $p_i = \frac{\alpha_i}{\sum_{j=1}^{IC} \alpha_j}$ .

In our case, 
$$\alpha_i = \sqrt{\|I^{[i]}\|_F^2 \cdot \|K_{[i]}\|_F^2 - E_{IK}^i + R_{IK}^i}$$
, and therefore the optimal probabilities are  

$$p_i = \frac{\sqrt{\|I^{[i]}\|_F^2 \cdot \|K_{[i]}\|_F^2 - E_{IK}^i + R_{IK}^i}}{\sum_{j=1}^{IC} \sqrt{\|I^{[j]}\|_F^2 \cdot \|K_{[j]}\|_F^2 - E_{IK}^j + R_{IK}^j}}$$

The terms  $E_{IK}^i$ ,  $R_{IK}^i$  emerge for convolutions when the kernel spatial dimensions are greater than one. However, computing them is too expensive, precluding efficient implementation of the approximate version. We therefore omit them and verify empirically whether the resulting norm-proportional probabilities:

$$p_{i} = \frac{\left\|I^{[i]}\right\|_{F} \cdot \left\|K_{[i]}\right\|_{F}}{\sum_{j=1}^{IC} \left\|I^{[j]}\right\|_{F} \cdot \left\|K_{[j]}\right\|_{F}}$$

yield better results than the uniform sampling. Intuitively, in some (common) cases these terms are much smaller than  $\|I^{[i]}\|_F^2 \cdot \|K_{[i]}\|_F^2$ , so their omission does not significantly impact the final outcome.  $E_{IK}^i$  amounts to the outer spatial dimensions of the input not being convolved with the entire kernel, so it is likely to be smaller than the Frobenius norm of the whole input.  $R_{IK}^i$  is the sum of products of different input and kernel entries. If different kernels are lowly-correlated and weights are centered around zero, the sum will include terms of similar magnitudes but opposite signs.

## **D** Implementation Details

All single-node results were obtained using 2.2GHz Intel Xeon Silver 4210 CPU with four NVidia V100 GPUs with 32GB of memory. Wall-time speedup were measured when running with a single GPU, except ResNet-152 where 2 GPUs are used due to memory capacity. We used PyTorch version 1.7.0 with CUDA 10.1 and Python version 3.6.9.

#### D.1 MLP for MNIST

The MNIST dataset [28] includes 60K training examples and 10K test examples. We use 5K as validation set. Each example is a  $28 \times 28$  gray-scale image of a handwritten digit.

Our MLP model contains the following layers:

- $784 \times 500$  fully-connected layer with RELU activations.
- $500 \times 500$  fully-connected layer with RELU activations.
- $500 \times 10$  fully-connected layer with RELU activations.
- · Log Softmax

We use the Adam optimizer [58] with default parameters (learning rate= $0.001, \beta_1 = 0.9, \beta_2 = 0.999, \epsilon = 1e - 08$ ). As loss function we use negative log likelihood. We use minibatch size of 50 and train the model for 20 epochs.

We apply sampling to all the fully connected layers. When sampling in the backward pass, we do not reduce the batch dimension below 10 in the weight gradient computation.

Figure 7(a) shows the MNIST test accuracy for different sampling algorithms and sampling ratios in the forward pass. We observe that top-k performs the best. Figure 7(b) shows the same when approximations are applied in the backward pass only. In this case, all sampling algorithms are similar when performing above 30% of the backward pass computations.



Figure 7: MNIST test accuracy for MLP, under different approximating algorithms and different sampling ratios

## D.2 CNN for MNIST

The network is composed of the following layers:

- $5 \times 5 \times 32$  convolution layer with RELU activation, followed by  $2 \times 2$  max pooling.
- $5 \times 5 \times 64$  convolution layer with RELU activation, followed by  $2 \times 2$  max pooling.
- Dropout layer with p = 0.5.
- $3136 \times 1024$  fully connected layer with RELU activation.
- $1024 \times 10$  fully connected layer.
- Dropout layer with p = 0.5.
- Log Softmax

The model is trained using Adam optimizer with default parameters (learning rate=0.001, $\beta_1 = 0.9$ , $\beta_2 = 0.999$ , $\epsilon = 1e - 08$ ) and negative log likelihood loss. We use minibatch size of 50 and train the model for 20 epochs.

We apply sampling to the convolutional layers. When sampling in the backward pass, we do not reduce the batch dimension below 10 in the weight gradient computation.

Figure 8(a) shows the MNIST test accuracy for different sampling algorithms and sampling ratios in the forward pass. We observe that top-k performs the best. Figure 8(b) shows the same when approximations are applied in the backward pass only. In this case, all sampling algorithms are similar when performing above 30% of the backward pass computations.



Figure 8: MNIST test accuracy for CNN, under different approximating algorithms and different sampling ratios

#### D.3 Wide ResNet-28-10 for CIFAR-10

The CIFAR-10 dataset [30] consists of  $32 \times 32$  color images from 10 classes, split into 50K training set and 10K test set.

For WRN-28-10 [29] we use the implementation in https://github.com/meliketoy/wide-resnet.pytorch, avialable under MIT License.

WRN-28-10 includes the following layers:

- conv1  $3 \times 3 \times 16$  input convolution layer
- conv2 eight  $3 \times 3 \times 160$  convolution layers
- conv3 eight  $3 \times 3 \times 320$  convolution layers
- conv4 eight  $3 \times 3 \times 640$  convolution layers
- Batch normalization,  $8 \times 8$  Average pooling, fully connected+softmax layers.

Every two subsequent convolution layers are followed by a residual connection that adds the input to these layers to the result. the first convolution conv3 and conv4 has a stride of 2, halving the spatial dimensions. For additional details see [29].

Image preprocessing includes padding to 36x36 and random crop, horizontal flipping and per-image whitening. The optimizer is Momentum SGD with momentum=0.9 and 5e-4 weight decay. Learning rate is 0.15 for the first 60 epochs, 0.03 until epoch 120, 0.006 until epoch 160 and 0.0012 afterwards. We use batch size of 256, cross-entropy loss and train the model for 200 epochs.

We apply sampling to the convolutional layers except the first layer due to the small number of input channels (3) and the single fully-connected layer which amounts only to 0.01% of the total computations in WRN-28-10. When sampling in the backward pass, we do not reduce the batch dimension below 10 in the weight gradient computation.

Figure 9(a) shows the CIFAR-10 test accuracy for different sampling algorithms and sampling ratios in the forward pass. We observe that top-k performs the best. Figure 9(b) shows the same when approximations are applied in the backward pass only. In this case, Bernoulli-CRS performs the best but is still below 1% of the baseline accuracy until 90% sampling ratio.



Figure 9: CIFAR-10 test accuracy for WRN-28-10, under different approximating algorithms and different sampling ratios

Figure 6(a) shows the CIFAR-10 validation accuracy learning curves for different forward-pass top-k sampling ratios, compared to the non-approximate baseline. We observe that higher sampling ratios lead to slower learning at the early training stages but the gap is decreasing as the training progresses. Figure 10 focuses on the last training epochs to observe the accuracies in more detail. We observe that 50% sampling is slightly lower than the non-approximate baseline, while less aggressive approximations that perform 70% or 90% of the computations achieve identical or slightly higher validation accuracy.



Figure 10: Learning curves for WRN-28-10 CIFAR-10 validation accuracy under different top-k sampling ratios. Focused view of last training epochs

#### D.4 ResNet-50 and ResNet-152 for ImageNet

The ImageNet [32] ILSVRC 2012 dataset contain 1.2 million training images of varying dimensions split into 1000 classes. The validation set includes 50K images and the test set consists of 100K images.

For ResNet-50 [31] we use the implementation in https://github.com/pytorch/examples/ tree/master/imagenet, available under BSD 3-Clause License. See [31] for further details on ResNet-50 architecture.

Image preprocessing includes random 224x224 crop, horizontal flipping and image normalization. The optimizer is Momentum SGD with momentum=0.9 and 1e-4 weight decay. Learning rate is 0.1 and it is decayed by 10 every 30 epochs. We use batch size of 256, cross-entropy loss and train the model for 90 epochs.

We apply sampling to the convolutional layers except the first layer due to the small number of input channels (3) and the fully-connected layer.

Figure 11(a) shows the top-1 accuracy of ResNet-50 for different sampling ratios. The different data points correspond to 50% top-k sampling applied to all the layers, all layers with at least 128 channels, 256 channels, 512 channels and 1024 channels.

Figure 11(b) shows the top-1 accuracy of ResNet-152 for different sampling ratios. The different data points correspond to 50% top-k sampling applied to all the layers, all layers with at least 256 channels, 512 channels and 1024 channels.



Figure 11: ResNet-50 and ResNet-152 ImageNet top-1 test accuracy. The accuracy increases as higher amounts of the computations are performed

Figures 6(b) and 6(c) show the top-1 validation accuracy learning curves for different forward-pass top-k sampling ratios, compared to the non-approximate baseline. We observe that ResNet-50 and ResNet-152 are more sensitive to sampling compared to WRN-28-10 on CIFAR10. Nonetheless, applying 50% sampling in the layers with 1024 channels, corresponding to 93% of the computations in ResNet-50 and 91% of the computations in ResNet-152, follow the non-approximate learning curves almost identically.

#### **D.5** Distributed Training

To evaluate the accuracy of top-k-weights algorithm for ResNet-152 on Imagenet we used the same settings as in the previous section and trained on a single node with 4 GPUs. The accuracy results are shown in figure 12. The different data points correspond to 50% top-k-weights sampling applied to all layers with at least 256 channels, 512 channels and 1024 channels.



Figure 12: ResNet-152 ImageNet top-1 test accuracy, using top-k-weights algorithm.

For the distributed training experiments we used eight AWS EC2 instances equipped with 2.7GHz Intel Xeon E5-2686v4 CPU, one V100 GPU with 16 GB of memory, 10 Gbps networking, PyTorch version 1.7.1, CUDA 11 and Python 3.7.6.

For the distributed measurement we used the same hyper-parameters except the minibatch size which we set to 32 per GPU. We could not increase the batch size since the AWS EC2 GPU we used had 16GB of memory and could not support higher batch size. We note that the eight-node setting has a total global batch size of 256, which matches the batch size used in the accuracy evaluation.