

Evaluating the Feasibility of Sampling-Based Techniques for Training Multilayer Perceptrons

Sana Ebrahimi
University of Illinois Chicago
Department of Computer Science
Chicago, Illinois, USA
sebrah7@uic.edu

Rishi Advani
University of Illinois Chicago
Department of Computer Science
Chicago, Illinois, USA
radvani2@uic.edu

Abolfazl Asudeh
University of Illinois Chicago
Department of Computer Science
Chicago, Illinois, USA
asudeh@uic.edu

ABSTRACT

The training process of neural networks is known to be time-consuming, and having a deep architecture only aggravates the issue. This process consists mostly of matrix operations, among which matrix multiplication is the bottleneck. Several sampling-based techniques have been proposed for speeding up the training time of deep neural networks by approximating the matrix products. These techniques fall under two categories: (i) sampling a subset of nodes in every hidden layer as active at every iteration and (ii) sampling a subset of nodes from the previous layer to approximate the current layer’s activations using the edges from the sampled nodes. In both cases, the matrix products are computed using only the selected samples. In this paper, we evaluate the feasibility of these approaches on CPU machines with limited computational resources. Making a connection between the two research directions as special cases of approximating matrix multiplications in the context of neural networks, we provide a negative theoretical analysis that shows feedforward approximation is an obstacle against scalability. We conduct comprehensive experimental evaluations that demonstrate the most pressing challenges and limitations associated with the studied approaches. We observe that the hashing-based node selection method is not scalable to a large number of layers, confirming our theoretical analysis. Finally, we identify directions for future research.

1 INTRODUCTION

The database community has played a significant role in tackling the complexities of big data, developing advanced technologies for addressing scalability challenges in very large data. Many such technologies have been used to address big-data challenges across various domains. For example, Asudeh et al. [3, 4] leverage sampling-based similarity joins [10] and selectivity estimation queries [7, 24] for signal reconstruction at scale. Among such examples, data management for machine learning (ML) [32, 35, 62] shines as an interdisciplinary domain that led to many key advances, including SystemML [8, 9, 20], join optimization for feature selection [33], federated learning [41], ML model materialization [26, 61], etc. Following the theme of *Extended Database Technologies for Machine Learning*, this paper studies the extension of sampling-based techniques for efficient training of deep neural networks (DNNs) on CPU machines with limited resources.

One of the crucial factors in developing DNN models with high performance is the network architecture. The task and dataset at hand determine the most appropriate architecture to use. In addition to a large number of layers, DNNs often have a high

number of nodes per layer. While large models can better generalize, training them can be computationally expensive, requiring extensive amounts of data and powerful hardware, including expensive GPUs. On the other hand, “the ubiquity of CPUs provides a workaround to the GPU’s dominance” [49], motivating “democratiz[ing] AI with CPUs” (Shrivastava as quoted by Smith [49]). Nevertheless, limited resources on personal computers with ordinary CPUs or mobile devices leads to difficulties in training DNNs to a sufficient level of accuracy. DNNs need to compute “activation values” for every layer in a forward pass and calculate gradients to update weights in the backpropagation. This requires performing expensive matrix multiplications that make the training process inefficient. Furthermore, large matrices often do not fit in the cache, and storing them in main memory necessitates constant communication between the processor and memory, which is even more time consuming.

In this work, we explore the scalability of two directions in sampling-based approaches based on locality-sensitive hashing (LSH) [21, 28] and Monte-carlo (MC) estimations [46] for efficient training of DNNs, which can be applied on memory- and computation-constrained devices. LSH was introduced in VLDB’99 [21] for near-neighbor queries. Since then, there have been consistent contributions to this line of research, including recent publications [5, 54, 58]. Similarly, sampling-based Monte-carlo estimations and inner-product estimation (and search) have been two core techniques for approximate query processing [2, 44] and Vector-DBs [43]. We hope that evaluating sampling-based techniques for efficient training of DNNs in this paper inspires (DB) researchers to address the open problems identified.

Our contributions can be summarized as follows.

- We make a connection between two separate sampling-based research directions for training DNNs by showing that both techniques can be viewed as special cases of matrix approximation, where one samples rows of the weight matrix while the other sample its columns. To the best of our knowledge, there is no previous work in the literature to make this observation.
- After careful exploration of different techniques, we provide negative theoretical results that show estimation errors during the feedforward step propagate across layers. In particular, for the model of [50], we prove that estimation error increases *exponentially* with the number of hidden layers.
- We provide extensive experiments using five training approaches and six benchmark datasets to evaluate the scalability of sampling-based approaches. Our experimental results confirm our theoretical analysis that feedforward approximation is an obstacle against scalability. In addition to other findings, our experiments reveal that while the model of [1] is scalable for mini-batch gradient decent when the batch size is relatively large, there is a research gap when it comes to designing scalable sampling-based approaches for stochastic gradient decent.

The rest of our paper is organized as follows. We first discuss

© 2025 Copyright held by the owner/author(s). Published in Proceedings of the 28th International Conference on Extending Database Technology (EDBT), 25th March-28th March, 2025, ISBN 978-3-98318-097-4 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

some of the potential benefits of training DNNs on CPU machines in §2, followed by related work in §3.

We define the problem formulation and provide a taxonomy of sampling-based approaches for efficient training in §4. We discuss two of these approaches in further detail in §5 and §6. We present our theoretical analysis in §7.

We discuss the extensions to convolutional neural networks in our technical report [17].

Experiment details and takeaways are discussed in §8–10, and we offer concluding remarks in §11.

2 BENEFITS OF TRAINING NEURAL NETWORKS ON CPUS

The pursuit of advancing DNN training on CPUs unveils a compelling avenue replete with practical advantages. Below, we briefly explain some of these potential benefits:

Abundance of CPU Machines. CPU-equipped personal computing devices, including PCs and smartphones, enjoy widespread availability and accessibility among a vast segment of the population. Remarkably, the computational potential of these devices often remains underutilized. Leveraging such resources for DNN training introduces the opportunity to conduct this computational-intensive task at no additional hardware cost for personal endeavors. Furthermore, while individual devices possess limited capacity, their collective potential can effectively address a multitude of moderate-sized artificial intelligence (AI) challenges. Recognizing this collective capability, recent endeavors have emerged to design client-side AI frameworks, exemplified by JavaScript packages like Tensorflow.js, facilitating machine learning on the client-side. Advancements in DNN training on CPU machines directly benefit these platforms.

Independence from Backend Servers. Personalized AI necessitates the training, or at the very least, fine-tuning of machine learning models with user-specific data. Opting for DNN training on CPU machines renders this process independent of backend GPU servers. Instead of transmitting data to the server, each personal device can locally fine-tune the models using its own data. This approach instantly confers several additional advantages:

Privacy By refraining from transmitting data to a server, concerns regarding data privacy are substantially alleviated.

Reduced Backend Computation The computational burden is shifted to the client side, at no expense to the server.

Offline Availability By localizing computations, the need for communication with a server is eliminated. This is especially useful for users with limited/unreliable internet access.

Democratizing DNN Training. GPU-equipped machines, while gradually becoming more affordable, still pose a considerable financial barrier. These costs manifest in the form of GPU access or enterprise APIs, especially in the context of large models like ChatGPT. Consequently, such resources remain inaccessible to a significant portion of the population. The facilitation of DNN training on CPU machines effectively dismantles this accessibility barrier.

Incorporating these considerations into the discourse of DNN training on CPU machines not only enriches the academic discussion but also underscores the profound implications of this research direction in addressing pressing real-world challenges and democratizing AI accessibility. It is evident that significant research efforts have been judiciously directed towards the princi-

pal trajectory of DNN training on GPU systems. Conversely, the avenue of training DNNs on CPU machines remains relatively under-explored within the research landscape.

3 RELATED WORK

The increasing importance of DNN applications opened the door to a variety of challenges associated with training these models. While there are numerous works on techniques for scaling DNNs, many of them have expensive hardware requirements and use GPUs to accelerate training [18]. Unlike GPUs, CPUs are available on any device, so optimizing training performance on CPUs is beneficial. There have been studies in which distributed, concurrent, or parallel programming on CPUs has been used to accelerate training [13, 25, 27, 29, 50, 55], but these methods are not always applicable due to variation in hardware requirements. Thus, algorithms focused on algorithmic optimization of feedforward and backpropagation are essential. Usually, methods with little to no special hardware requirements are preferred. Several algorithms apply a variety of sampling-based [1, 6, 19, 36, 37, 50] or non-sampling-based approximations [25, 38, 39, 63] to improve training for DNNs.

Several studies have shown that one way to scale up DNNs is to simplify the matrix-matrix or vector-matrix operations involved [1, 6, 50, 52, 57]. The complexity of matrix multiplication dominates the computational complexity of training a neural network. The multiplication of large matrices is known to be the main bottleneck in training DNNs. Often, we try to sparsify the matrices, which can minimize the communication between the memory and the processors [60]. Pruning the network and limiting the calculations in both directions to a subset of nodes per layer is one solution to train DNNs efficiently. This is what dropout-type algorithms suggest [6, 19, 52]. Dropout-type methods either use a data-dependent sampling distribution [1, 6, 57] or a predetermined sampling probability [50, 52]. Note that these techniques are able to provide a good approximation only if used in the context of neural networks; they are not necessarily applicable to general matrix multiplication.

4 PRELIMINARIES

In this section, we describe the neural network model, feedforward step, and backpropagation in the form of matrix operations. Finally, we discuss state-of-the-art sampling-based algorithms.

4.1 Problem Description

Many neural network architectures have been studied over the past decade. In this paper we focus on the standard multi-layer perceptron (MLP) model and analyze two major directions of sampling-based techniques for training neural networks.

Consider a feedforward neural network with m_i inputs, m_o outputs, and ℓ hidden layers. In general, every hidden layer k contains n_k hidden nodes while the nodes in the $(k - 1)$ -th and k -th layers are fully connected. Without loss of generality, for ease of explanation, we assume all hidden layers have exactly n hidden nodes (Figure 1). For each layer k , we denote the vector of outputs by $a^k \in \mathbb{R}^{1 \times n}$. Similarly, $W^k \in \mathbb{R}^{n \times n}$ and $b^k \in \mathbb{R}^{1 \times n}$ are the weights and the biases of layer k ,¹ respectively.

Let f be the activation function (e.g., sigmoid or ReLU). With input vector $x \in \mathbb{R}^{1 \times m_i}$, the feedforward step is a chain of matrix products and activation functions that maps input vector x to

¹The first and last layer are exceptions: $W^1 \in \mathbb{R}^{m_i \times n}$ and $W^\ell \in \mathbb{R}^{n \times m_o}$.

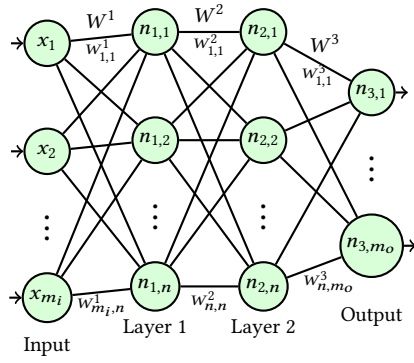


Figure 1: Neural network with $\ell = 3$ layers and n nodes per hidden layer.

Table 1: Table of Notations

Notation	Description
a^k	output vector of k -th layer
z^k	input vector of the k -th layer
W^k	weight matrix of the k -th layer
W_i^k	row i of W^k
$W_{:,j}^k$	column j of W^k
$\langle x_1, x_2 \rangle$	inner product of x_1 and x_2

output vector y ; it can be represented as follows (with $a^0 = x$):

$$z^k = a^{k-1} W^k + b^k \quad a^k = f(z^k)$$

In the setting described above, matrix-vector multiplication can be done in $\Theta(n^2)$ time and applying the element-wise activation function takes $\Theta(n)$ time for each layer. Thus, the entire feedforward process for the whole network is in the order of $\Theta(\ell n^2)$.

The final aspect in the training of neural networks is backpropagation, an efficient method of computing gradients for the weights to move the network towards an optimal solution via stochastic gradient descent² (SGD) [22]. The weight gradients for the backpropagation step can be computed recursively using Equation 1, where \mathcal{L} is the loss function and \odot is the Hadamard product.

$$\begin{aligned} \delta^\ell &= \nabla_{z^\ell} \mathcal{L} = f'(z^\ell) \odot \nabla_{a^\ell} \mathcal{L} & \nabla_{W^k} \mathcal{L} &= a^k \delta^{k+1} \\ \delta^k &= \nabla_{z^k} \mathcal{L} = f'(z^k) \odot W \delta^{k+1} & \nabla_{b^k} \mathcal{L} &= \delta^{k+1} \end{aligned} \quad (1)$$

With gradient $a^k \delta_{k+1}$ and learning rate η , the weight matrix W^k will be updated to $W^k - \eta a^k \delta_{k+1}$. The gradient computation and update operations are also in form of vector-matrix operations that take $\Theta(n^2)$ time for each layer. As a result, the backpropagation step in SGD also requires $\Theta(\ell n^2)$ time.

²While SGD uses only one data point to compute the gradients, an alternative approach is mini-batch gradient descent (MGD), where a small sample set (mini-batch) of the training set is used for estimations. Note that SGD can be viewed as a special case of MGD where the batch size is 1. Following our scope in this paper, SGD is considered for problem formulation, explaining the learning algorithm, and analysis. Nevertheless, as we shall later explain in §6.1, one of the evaluated approaches, MC-APPROX, is based on MGD, the generalization of SGD. While SGD operations are in form of vector to matrix multiplication, MGD operations are in form of matrix (vectors of samples in the mini-batch) to matrix multiplication.

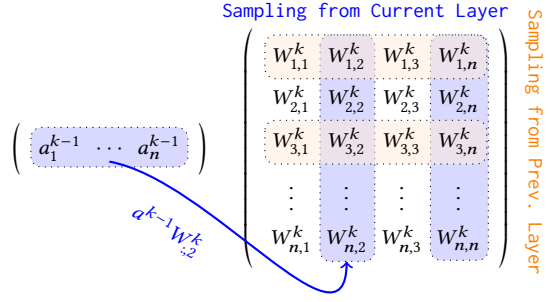
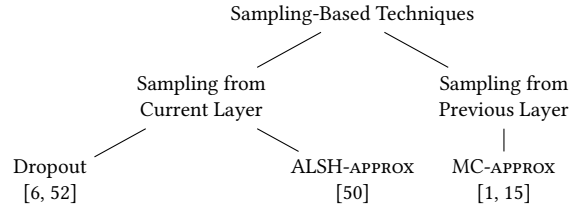


Figure 2: High-level idea of the sampling-based techniques.

4.2 Taxonomy of Sampling-Based Techniques

The computation bottleneck in the training of a DNN is matrix multiplication, in form of a $(1 \times n)$ to $(n \times n)$ vector-matrix product for SGD. Sampling-based approaches seek to speed up this operation by skipping a large portion of the scalar operations. SGD is a noisy algorithm by nature. As such, it is more tolerant of small amounts of noise [40], allowing for approximation. At a high level, these approaches fall in two categories, as shown below:



Sampling from Current Layer. The approaches in this category select a small subset of nodes in each layer during each feedforward-backpropagation step, and update the values of only those nodes. In Figure 2, each column $W_{:,j}^k$ corresponds to the node n_j^k in layer k , while each cell $W_{i,j}^k$ in that column represents the weight of the edge from n_i^{k-1} to n_j^k . As a result, these approaches can be viewed as *selecting a small subset of the columns* of W^k (e.g., the highlighted columns) and conducting the inner product only for those.

Sampling from Previous Layer. Instead of selecting a subset of columns and computing the exact inner-product for them, the alternative is to select all columns but compute the inner-product approximately for them by *selecting a small subset of rows* of W^k (e.g., highlighted rows in Figure 2). That is, instead of computing the sum for all n scalars in an inner-product, to *estimate* the sum by sampling a small number of scalars.

5 EFFICIENT TRAINING BY SAMPLING FROM CURRENT LAYER

5.1 Dropout

Srivastava et al. [52] introduced DROPOUT, a computationally efficient training approach that reduces the risk of overfitting. During each feedforward step, the algorithm picks a subset of nodes uniformly at random in each hidden layer and drops the remaining nodes temporarily. The sampled nodes are then used for feedforward evaluation and backpropagation.

While DROPOUT was originally introduced to fix overfitting, it introduced a computation reduction to the training process.

In many cases, DROPOUT improved the runtime efficiency compared to the standard training process on the same architecture. However, there are scenarios in which training under DROPOUT requires more training iterations and eventually hurts the runtime. One can observe that due to the randomness in sampling with a fixed probability (usually $p = 1/2$), there is a risk of dropping nodes that significantly affect the output values. Ba and Frey [6] addressed this issue by proposing ADAPTIVE-DROPOUT, which uses a data-dependent distribution that is an approximation of the Bayesian posterior distribution over the model architecture and updates the sampling ratio adaptively w.r.t the current network. This method avoids randomly dropping significant nodes in the model.

5.2 Asymmetric Locality-Sensitive Hashing for Training Approximation

Unlike in DROPOUT, one might want to intelligently select a small subset of so-called *active nodes* for each layer for computing the inner products. In particular, given the vector a^{k-1} , the goal is to find a small portion of nodes j in layer k for which the value of $a^{k-1}W_{:,j}^k$ is maximized in order to avoid computing inner products for small values (estimating them as zero). Given a set S of vectors (in this case, the set of columns in W^k) and a query vector a , the problem of finding a vector $w^* \in S$ with maximum inner product $\langle a, w^* \rangle$ is called maximum inner-product search (MIPS). To solve MIPS, Shrivastava and Li [48] employ asymmetric locality-sensitive hashing (ALSH).

Definition 5.1 (Asymmetric Locality-Sensitive Hashing [48]). Given a similarity threshold S_0 and similarity function $\text{sim}(\cdot)$, a family \mathcal{H} of hash functions are (S_0, cS_0, p_1, p_2) -sensitive for c -NNS³ with $a \in \mathbb{R}^n$ as query and a set of $w \in \mathbb{R}^n$ vectors if for all $h \in \mathcal{H}$ chosen uniformly, the following conditions are satisfied:

$$\begin{aligned} \text{sim}(w, a) \geq S_0 &\implies \Pr[h(Q(a)) = h(P(w))] \geq p_1 \\ \text{sim}(w, a) \leq cS_0 &\implies \Pr[h(Q(a)) = h(P(w))] \leq p_2 \end{aligned}$$

For $w, a \in \mathbb{R}^n$ with $\|w\| \leq C$, where C is a constant less than 1, and $\|a\| = 1$, they define the transformations P and Q for the inner product as follows.

$$\begin{aligned} P: \mathbb{R}^n &\rightarrow \mathbb{R}^{n+m}, & w &\mapsto [w; \|w\|^{2^1}, \dots, \|w\|^{2^m}] \\ Q: \mathbb{R}^n &\rightarrow \mathbb{R}^{n+m}, & a &\mapsto [a; 1/2, \dots, 1/2] \end{aligned} \quad (2)$$

In other words, to generate P , w is padded with m terms, where term i is the ℓ_2 norm of w to the power of 2^i . Q is generated by padding a with m copies of the constant $1/2$. Shrivastava and Li [48] prove that NNS in the transformed space is equivalent to the maximum inner product in the original space:

$$\arg \max_w \langle w, a \rangle \approx \arg \min_w \|Q(a) - P(w)\|. \quad (3)$$

Equation 3 motivates using MIPS for efficient training of DNNs. Spring and Shrivastava [50] build their algorithm (referred to as **ALSH-APPROX** in this paper) upon Equation 3. As explained in §4, the feedforward step and backpropagation consist of many matrix multiplications, each of which involve a set of inner products as large as each hidden layer. ALSH-APPROX uses ALSH to prune each layer by finding active nodes, in this case, nodes with maximum activation values. This is equivalent to solving MIPS in each layer.

³c-approximation of nearest neighbor search [28].

Essentially, ALSH-APPROX uses ALSH to find active nodes j whose weight vector $W_{:,j}^k$ collides with an input vector a^{k-1} under the same hash function. The probability of collision captures the similarity of vectors in each hidden layer. To do so, it sets the query vector as $q = a^{k-1}$ and the set of vectors using the columns of W^k as $S = \{W_{:,1}^k, \dots, W_{:,n}^k\}$. Then, after constructing Q and P based on Equation 2, we have

$$\arg \max_j \langle W_{:,j}^k, a^{k-1} \rangle \approx \arg \min_j \|Q(a^{k-1}) - P(W_{:,j}^k)\|. \quad (4)$$

ALSH-APPROX constructs L independent hash tables with 2^K hash buckets and assigns a K -bit randomized hash function to every table. Each layer has been assigned L hash tables and a meta hash function to compute a hash signature for the weight vectors and fill all the hash tables before training. In this setting, K and L are tunable hyperparameters that affect the active set's size and quality.

During training, ALSH-APPROX computes the hash signature of each incoming input using the existing hash functions. Then, a set of weight vectors will be returned using the hash values corresponding to the hash bucket. The active nodes in a layer are the union of their corresponding weight vectors from probing L hash tables. Then, the model only performs the exact inner product on these active nodes and skips the rest. Finally, the gradient will only backpropagate through the active nodes and update the corresponding weights. In other words, ALSH is used to sample a subset of nodes with probability $1 - (1 - p^K)^L$ if p is the probability of collision.

Updating the hash tables ensures that the modified weight vectors are recognized. Based on the results reported by Spring and Shrivastava [50], the number of active nodes for each input can be as small as 5% of the total nodes per layer. Thus, ALSH-APPROX performs a significantly smaller set of inner products in each iteration. Moreover, due to the sparsity of the active sets belonging to different data inputs, the overlap between them throughout the dataset is small, so the weight gradient updates corresponding to these inputs are sparse as well. Thus, the hash table updates are executed after processing a batch of inputs and can be executed in parallel. The main advantage of ALSH-APPROX is that, unlike DROPOUT, it finds the active nodes *before* computing the inner products.

6 EFFICIENT TRAINING BY SAMPLING FROM PREVIOUS LAYER

While techniques discussed in §5 reduce the vector-matrix multiplication time by selecting a subset of columns (nodes) from each weight matrix W^k and computing the inner product exactly for them, an alternative approach is to select all columns but to compute inner products approximately. This idea has been proposed by Adelman et al. [1]. This paper is built on the MC method by Drineas et al. [15] for fast approximation of matrix multiplication. We first review the work of Drineas et al. [15] in §6.1 and then in §6.2 we explain how Adelman et al. [1] adapt the method to develop an algorithm for efficient training of DNNs.

6.1 Fast Approximation of Matrix Multiplication

For many applications, a fast estimation of the matrix product is good enough. In addition to hardware/software oriented optimizations such as cache management [18, 23] or half-precision computations [40, 55], MC methods [46] have been used for such

estimations. At a high level, MC methods use repeated sampling and the law of large numbers to estimate aggregate values.

Recall that given two matrices $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$, the product AB is an $m \times p$ matrix, where every element $AB_{i,j}$ is the inner product of i -th row of A with the j -th column of B :

$$AB_{i,j} = \langle A_{i,:}^T, B_{:,j} \rangle = \sum_{t=1}^n A_{i,t} B_{t,j} \quad (5)$$

In an MC estimation of $AB_{i,j}$, instead of computing the sum over all $t \in [1, n]$, only a small sample of elements $\sigma \subset \{1, \dots, n\}$, where $c = |\sigma| \ll n$, are considered. Arguing that uniform sampling would add a high error in estimating AB , Drineas et al. [15] introduce a nonuniform sampling method to generate σ with a probability proportional to the magnitude of data. Specifically, they develop a randomized algorithm that samples each column i of A and row i of B with probability

$$p_i = \frac{\|A_{:,i}\| \cdot \|B_{i,:}\|}{\sum_{t=1}^n \|A_{:,t}\| \cdot \|B_{t,:}\|}, \quad (6)$$

where $\|\cdot\|$ is the ℓ_2 norm. They define the matrices $C \in \mathbb{R}^{m \times c}$ and $R \in \mathbb{R}^{c \times p}$ as $C = ASD$ and $R = (SD)^T B$, respectively. S is then defined as an $n \times c$ sampling matrix, with $S_{ij} = 1$ if the i th row of A is the j th sample. D is a $c \times c$ diagonal scaling matrix with $D_{jj} = \frac{1}{\sqrt{cp_j}}$. The authors prove that defining p_i as in Equation 6 minimizes the expected estimation error, $\mathbb{E}[\|AB - CR\|_F]$, where $\|\cdot\|_F$ is the Frobenius norm. Each element $AB_{i,j}$ is estimated as $\sum_{t=1}^c C_{i,t} R_{t,j} = \sum_{t=1}^c \frac{1}{cp_i} A_{i,t} B_{t,j} \approx AB_{i,j}$. Sampling row-column pairs w.r.t p_i reduces the complexity of matrix multiplication to $O(mcp)$.

6.2 MC-APPROX

Training DNNs involves computationally expensive matrix multiplication operations. However, the gradients computed during backpropagation only approximate directions towards the optimal solution, so the training process has a high tolerance to small amounts of noise. This makes approximation of matrix multiplication a reasonable choice to speed up training of DNNs. Following this idea, Adelman et al. [1] propose a MC approximation method for fast training of DNNs (in this paper, referred to as **MC-APPROX_M** for the mini-batch setting and **MC-APPROX_S** for the stochastic setting) based on the MC estimation of matrix multiplication explained in §6.1. Despite the fact that Drineas et al. [15] provide an unbiased estimate for the matrix multiplication AB (i.e., $E[CR] = AB$), Adelman et al. [1] prove that the sampling distribution is not able to provide an unbiased estimation of the weight gradient if it is used for both the forward step and backward pass simultaneously.

One way to eliminate the bias is to use MC approximation only in the forward pass, propagate the gradient through the entire network, and perform the exact computations. However, experiments show this approach results in lower accuracy in practice. Therefore, Adelman et al. [1] propose a new sampling distribution that yields an unbiased estimate of the weight gradient $\nabla_W \mathcal{L}$ when it is used only during the feedforward step. Specifically, they sample column-row pairs independently from $A \in \mathbb{R}^{m \times n}$ and $B \in \mathbb{R}^{n \times p}$.

Let k be the number of samples for estimation, let $V \in \mathbb{R}^{n \times n}$ be a diagonal sampling matrix with $V_{i,i} = Z_i \sim \text{Bernoulli}(p_i)$, where $\sum_{i=0}^n p_i = k$, and let $D \in \mathbb{R}^{n \times n}$ be a diagonal scaling matrix with $D_{i,i} = \frac{1}{\sqrt{p_i}}$. Then, the multiplication of matrices A and

B can be estimated as $AB \approx \sum_{i=0}^n \frac{Z_i}{p_i} A_{:,i} B_{i,:} = AVDD^T V^T B = A' B'$. The estimation error in this case is $E[\|AB - A' B'\|_F^2] = \sum_{i=0}^n \frac{1-p_i}{p_i} \|A_{:,i}\|^2 \|B_{i,:}\|^2$. Hence, the following assignment of probabilities minimizes the estimation error under the constraint $\sum_{i=0}^n p_i = k$.

$$p_i = \min \left\{ \frac{k \|A_{:,i}\| \|B_{i,:}\|}{\sum_{t=0}^n \|A_{:,t}\| \|B_{t,:}\|}, 1 \right\} \quad (7)$$

The authors prove that training a neural network by approximating matrix products in backpropagation converges with the same rate as standard SGD and results in an unbiased estimator when nonlinearities are not considered. When accounting for nonlinearities, the results hold as long as the MC approximation of Wx is unbiased and the activation and loss functions are β -Lipschitz.

7 THEORETICAL ANALYSIS

As explained in §5.2 and §6.1, sampling-based approaches seek to speed up the training of DNNs by skipping a large number of computations and approximating matrix multiplications. In this section we provide negative theoretical results for scalability against the feedforward approximation. At a high level, we show that small estimation errors in the initial layers get propagated and compounded in subsequent layers. Adelman et al. [1] already observed the low performance of MC-APPROX when the feedforward step is approximated and therefore only applied approximation during backpropagation for MLPs. As such, we focus on ALSH-APPROX for our analysis. First, let us introduce the following notation, which we will use throughout this section.

- \bar{a}_j^k : the estimation of a_j^k by ALSH-APPROX.
- $e_j^k = a_j^k - \bar{a}_j^k$: the activation value estimation error.
- \uparrow_j^k : the set of active nodes for n_j^k .

LEMMA 7.1. *Let f be a linear activation function such that $a = f(z) = z$. Assuming the active nodes are detected exactly, the estimation error for the node n_j^k by ALSH-APPROX is as follows.*

$$e_j^k = \begin{cases} \sum_{i \in \uparrow_j^1} x_i W_{i,j}^1 & \text{if } k = 1 \\ e^{k-1} W_{:,j}^k + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^1 & \text{otherwise} \end{cases}$$

PROOF. We want to show that the estimation error for the node n_j^k by ALSH-APPROX is as follows.

$$e_j^k = \begin{cases} \sum_{i \in \uparrow_j^1} x_i W_{i,j}^1 & \text{if } k = 1 \\ e^{k-1} W_{:,j}^k + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^1 & \text{otherwise} \end{cases}$$

First, for $k = 1$:

$$\begin{aligned} a_j^1 &= \sum_{i=1}^n x_i W_{i,j}^1 = \sum_{l \in \uparrow_j^1} x_l W_{l,j}^1 + \sum_{i \notin \uparrow_j^1} x_i W_{i,j}^1 = \bar{a}_j^1 + \sum_{i \notin \uparrow_j^1} x_i W_{i,j}^1 \\ \implies e_j^1 &= a_j^1 - \bar{a}_j^1 = \sum_{i \notin \uparrow_j^1} x_i W_{i,j}^1 \end{aligned}$$

Analogously, when $k > 1$:

$$\begin{aligned}
a_j^k &= \sum_{i=1}^n a_i^{k-1} W_{i,j}^{k-1} = \sum_{i=1}^n (a_i^{k-1} + e_i^{k-1}) W_{i,j}^{k-1} \\
&= \sum_{l \in \uparrow_j^k} \bar{a}_l^{k-1} W_{l,j}^{k-1} + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^{k-1} + \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1} \\
&= \bar{a}_j^k + e^{k-1} W_{:,j}^k + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^{k-1} \\
\Rightarrow e_j^1 &= a_j^1 - \bar{a}_j^1 = e^{k-1} W_{:,j}^k + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^{k-1}
\end{aligned}$$

□

Lemma 7.1 provides a recursive formula for the activation value estimation error in terms of the weighted summation over active nodes versus inactive nodes. To provide a non-recursive and easier to understand formula, in Theorem 7.2 we assume that the weighted summation over the active nodes is always c times that of the inactive nodes.

THEOREM 7.2. *Let f be a linear activation function such that $a = f(z) = z$. Suppose for any node n_p^l ,*

$$\sum_{i \in \uparrow_p^l} a_i^{l-1} W_{i,p} = c \sum_{i \notin \uparrow_p^l} a_i^{l-1} W_{i,p}.$$

Then, $a_j^k = \bar{a}_j^k \left(\frac{c+1}{c} \right)^k$. That is, $e_j^k = \bar{a}_j^k \left(\left(\frac{c+1}{c} \right)^k - 1 \right)$.

PROOF. For any node n_p^l , we have

$$\sum_{i \in \uparrow_p^l} a_i^{l-1} W_{i,p} = c \sum_{i \notin \uparrow_p^l} a_i^{l-1} W_{i,p}.$$

We then use induction to prove $a_j^k = \bar{a}_j^k \left(\frac{c+1}{c} \right)^k$.

Base case. When $k = 1$:

$$a_j^1 = \sum_{i=1}^n x_i W_{i,j}^1 = \sum_{l \in \uparrow_j^1} x_l W_{l,j}^1 + \sum_{i \notin \uparrow_j^1} x_i W_{i,j}^1 = \bar{a}_j^1 + \frac{1}{c} \bar{a}_j^1 = \bar{a}_j^1 \frac{c+1}{c}$$

Inductive step. Assuming $a_j^{k-1} = \bar{a}_j^{k-1} \left(\frac{c+1}{c} \right)^{k-1}$:

$$\begin{aligned}
a_j^k &= \sum_{i=1}^n a_i^{k-1} W_{i,j}^{k-1} = \sum_{i=1}^n (\bar{a}_i^{k-1} + e_i^{k-1}) W_{i,j}^{k-1} \\
&= \sum_{l \in \uparrow_j^k} \bar{a}_l^{k-1} W_{l,j}^{k-1} + \sum_{i \notin \uparrow_j^k} \bar{a}_i^{k-1} W_{i,j}^{k-1} + \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1} \\
&= \bar{a}_j^k + \frac{1}{c} \bar{a}_j^k + \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1} = \frac{c+1}{c} \bar{a}_j^k + \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1}
\end{aligned} \tag{8}$$

Let $A = \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1}$. We have

$$\begin{aligned}
e_i^{k-1} &= a_i^{k-1} - \bar{a}_i^{k-1} = a_i^{k-1} - a_i^{k-1} \left(\frac{c}{c+1} \right)^{k-1} \\
&= a_i^{k-1} \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right).
\end{aligned}$$

Thus,

$$\begin{aligned}
A &= \sum_{i=1}^n e_i^{k-1} W_{i,j}^{k-1} = \sum_{i=1}^n a_i^{k-1} \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right) W_{i,j}^{k-1} \\
&= \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right) \sum_{i=1}^n a_i^{k-1} W_{i,j}^{k-1} = \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right) a_j^k.
\end{aligned}$$

Now, plugging A back into Equation 8, we get:

$$\begin{aligned}
a_j^k &= \frac{c+1}{c} \bar{a}_j^k + A = \frac{c+1}{c} \bar{a}_j^k + \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right) a_j^k \\
\Rightarrow a_j^k &\left(1 - \left(1 - \left(\frac{c}{c+1} \right)^{k-1} \right) \right) = \frac{c+1}{c} \bar{a}_j^k \Rightarrow a_j^k = \bar{a}_j^k \left(\frac{c+1}{c} \right)^k
\end{aligned}$$

□

Theorem 7.2 proves that the estimation error increases **exponentially** with the number of layers. As a result, due to the sharp increase in the estimation error, ALSH-APPROX does not scale to DNNs. To better observe this, suppose $c = 5$ (i.e., the weighted sum for the active nodes is five times that of the inactive nodes). Then, using Theorem 7.2, the error-to-estimate ratios for different numbers of layers are as follows.

k	1	2	3	4	5	6
e_j^k / \bar{a}_j^k	0.2	0.44	0.72	1.07	1.48	1.98

From the above table, it is evident that as soon as the depth of the network gets larger than 3, the estimation error dominates the estimation value. This is consistent with our experiment results, where ALSH-APPROX failed to scale with more than 3 hidden layers.

8 EXPERIMENT SETUP

8.1 Hardware

This paper aims to evaluate sampling-based approaches for training DNNs on regularly available machines; thus, we ran all experiments on a single-CPU machine (Intel Core i9-9920X machine with 128 GB of memory) without a GPU.

8.2 Datasets

We used the following six benchmark datasets for our experiments.

MNIST [14] 70,000 handwritten digits, each in the form of a 28×28 grayscale image, and 10 classes (digits zero to nine).

Kuzushiji-MNIST [11] 70,000 cursive Japanese characters, each in the form of a 28×28 grayscale image, and 10 classes.

Fashion-MNIST [59] 70,000 fashion products, each in the form of a 28×28 grayscale image, and 10 classes.

EMNIST-Letters [12] 145,600 handwritten letters, each in the form of a 28×28 grayscale image, and 26 classes.

NORB [34] 48,600 photographs of 50 toys from different angles, each in the form of a 96×96 grayscale image, and 5 classes.

CIFAR-10 [31] 60,000 color images, each of dimensions 32×32 , and 10 classes.

We randomly partition the datasets as shown below:

Dataset	Train	Test	Validation
(Kuzushiji/Fashion-)MNIST	55000	10000	5000
EMNIST-Letters	104800	20800	20000
NORB	22300	24300	2000
CIFAR-10	45000	10000	5000

We obtained similar results across different datasets. For brevity, we provide a detailed discussion only using the results on the MNIST dataset, but extensive results for other datasets are provided in our technical report [17].

8.3 Methods Evaluated

We evaluated four sampling-based approaches for training DNNs discussed in §5 and §6, namely DROPOUT [52], ADAPTIVE-DROPOUT [6], MC-APPROX [1], and ALSH-APPROX [50], on fully connected DNNs. In addition, the regular training approach, referred to as STANDARD, has been implemented for comparison purposes. All implementations are in Python 3.9 using the PyTorch library. For MC-APPROX,⁴ ALSH-APPROX,^{5,6} DROPOUT,⁷ and ADAPTIVE-DROPOUT⁸ we used the publicly available code.

8.4 Default Values

To train our models, we use SGD. The activation function used for hidden layers is ReLU due to its simplicity, ease of computation, and the fact that it helps with the vanishing gradients problem [22]. The output layer activation function is log softmax, and the loss function used throughout experiments is the negative log-likelihood. The learning rate is always either 10^{-4} or 10^{-3} depending on the setting, and the models are trained for 50 epochs. In particular, we set the learning rate to 10^{-4} for MC-APPROX_S.⁹ The hyperparameters of our implementation are the best values reported for each approach. For MC-APPROX the batch size is set to 20 and $k = 10$. For ALSH-APPROX, $K = 6$, $L = 5$, and $m = 3$ (Equation 2) as specified in [50]. In order to have a fair comparison with ALSH-APPROX, we set the probability of picking nodes for DROPOUT and ADAPTIVE-DROPOUT to $p = 0.05$, and we use a network with 3 hidden layers and 1000 hidden units per layer across algorithms. The implementation of ALSH-APPROX provided in [51] performs better with the Adam optimizer [30] than with Adagrad [16], which is used in the original implementation in [50]. Hence, we use Adam in our experiments.

For the convolutional setting, we used ResNet-18 with two fully-connected layers as a classifier to run our experiments. We limit the approximation to the classifier and keep the convoluted operations exact. Also, for CIFAR-10, we use pure SGD instead of Adam.

8.5 Experiment Plan

We are mainly interested in evaluating the following.

Accuracy How do the algorithms perform when training networks with different depths?

Time How scalable are the evaluated algorithms (in particular, ALSH-APPROX and MC-APPROX) w.r.t training time?

Hyperparameters How do hyper-parameter choices (e.g., batch size) affect training time and accuracy?

Accuracy here refers to the percentage of correct predictions on the entire dataset. Since we focus on multi-class classification, we also provide confusion matrices.

⁴github.com/acsl-technion/approx

⁵github.com/rdspring1/LSH-Mutual-Information

⁶github.com/rdspring1/LSH_DeepLearning

⁷github.com/gngdb/adaptive-standout

⁸see footnote 7.

⁹Across all tables and plots we use the subscripts “S” and “M” to refer to the SGD and mini-batch SGD (with default batch size 20) settings, respectively. When there is no subscript, the default is MGD for MC-APPROX and SGD for all other methods.

9 EXPERIMENT RESULTS

9.1 Accuracy

We begin our experiments by addressing the first question in §8.5. To do so, we generate models with different numbers of hidden layers (1 to 7) and evaluate each method on all six datasets discussed in §8.2 for both stochastic¹⁰ and mini-batch settings. The confusion matrices for all algorithms are provided in Figure 3. Every row in the figure shows the performance of an algorithm, while different columns represent networks with different numbers of hidden layers. In all plots contained within the figure, the x-axis shows the model prediction and the y-axis shows the true labels. Consequently, the diagonal cells show correct predictions, while all other cells are incorrect predictions. Ideally, the models should have (near-)zero values on non-diagonal cells.

Baselines. STANDARD,¹¹ DROPOUT, and ADAPTIVE-DROPOUT (the first three rows in Figure 3) are our baselines for comparisons. For DROPOUT, the nodes are sampled randomly with probability p , and for ADAPTIVE-DROPOUT, p is updated w.r.t the Bayesian posterior distribution of data input. In standard feedforward training, we expect to observe a decrease in generalization error over complex datasets as we add layers to the network and an increase in the ability to learn nonlinear functions. Clearly, this does not include the cases in which STANDARD overfits.

ALSH-APPROX. The experiment results on ALSH-APPROX (Row 4 in Figure 3) confirm a decrease in accuracy as the number of layers increases. In particular, Figures 3o and 3p show a sharp decrease in performance on 5 to 7 layers. This is confirmed in Figure 7, where the accuracy of ALSH-APPROX drops from 70.07% to 25.14% from 5 to 7 layers. Comparing ALSH-APPROX with STANDARD_S in Figure 3, though initially the two algorithms perform similarly on a small number of layers (Columns 1 and 2), the performance gap quickly increases with the number of layers — confirming the *lack of scalability* of ALSH-APPROX for DNNs. This is also observed in Figure 7. We can see from Table 2 that, compared with DROPOUT_S, ALSH-APPROX performs better, but it does not outperform STANDARD_S and ADAPTIVE-DROPOUT_S.

MC-APPROX. MC-APPROX is designed for use with mini-batch SGD. As we shall further investigate in §9.3, even though MC-APPROX_S outperforms other methods evaluated (Table 2), the runtime for large numbers of layers and datasets is so high that it is infeasible for computation-limited systems. This is reflected in Table 3. Therefore, as indicated in §8.4, we use mini-batch SGD (with batch size 20) as the default setting in our experiments. The experiment results on MC-APPROX_M are provided in the last row of Figure 3. MC-APPROX_M shows equally good performance across different numbers of layers, confirming its scalability for DNNs. In particular, when varying the number of hidden layers (Figure 7), the minimum accuracy obtained by MC-APPROX_M is 92.71% for one hidden layer. Comparing the confusion matrices of MC-APPROX with ADAPTIVE-DROPOUT and STANDARD in Figure 3, we can see that performance is consistent across the three algorithms. As shown in Figure 7, in most cases, MC-APPROX_M outperformed STANDARD_M with 2–4% difference in accuracy. This is also evident in Table 2, where MC-APPROX_M and MC-APPROX_S outperformed other algorithms on the MNIST and Fashion-MNIST datasets with 3 hidden layers. To further assess the scalability of MC-APPROX_M in deeper networks, we

¹⁰when the batch size is equal to 1.

¹¹training the neural network without sampling.

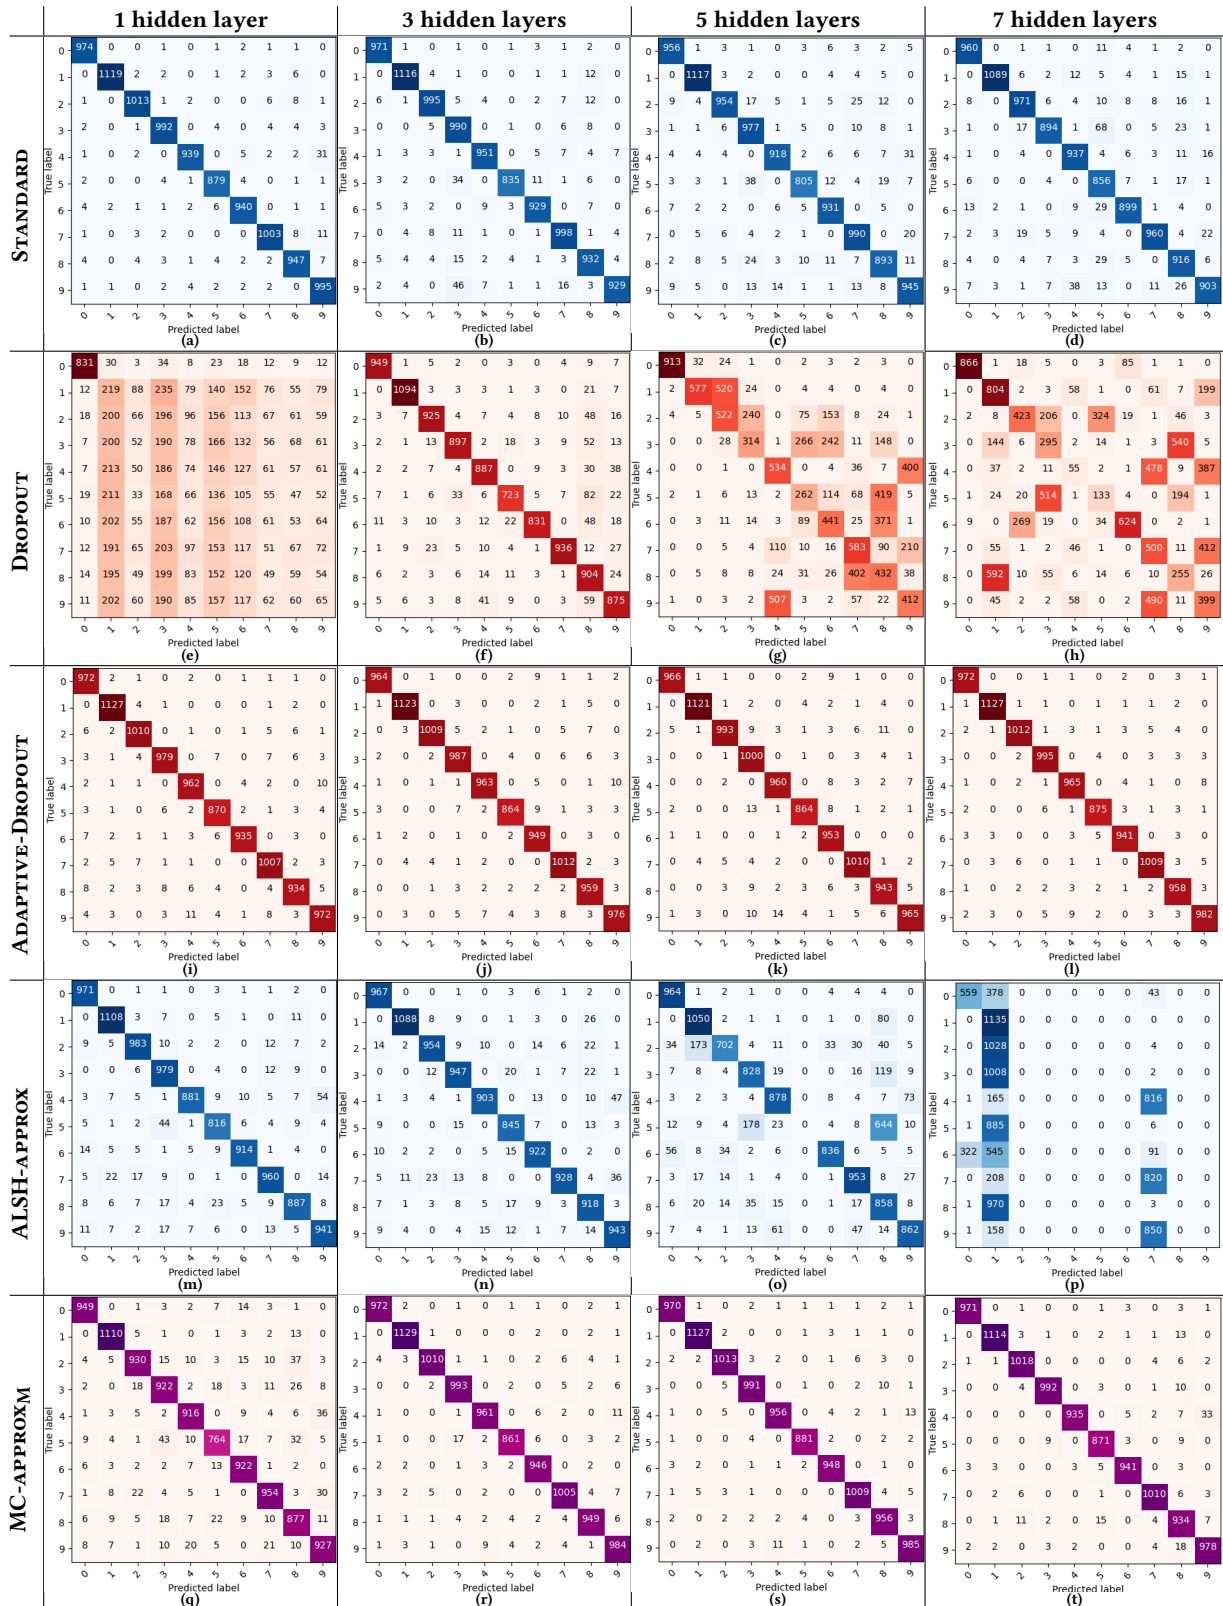


Figure 3: Confusion matrices of different algorithms for different numbers of hidden layers. In all plots, x-axis and y-axis are the predicted and true labels (0 to 9), respectively.

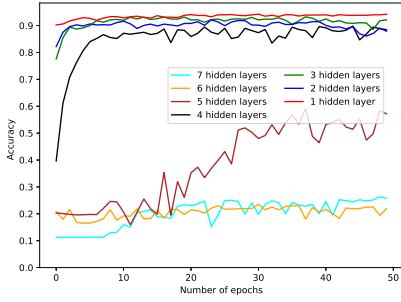


Figure 4: ALSH-APPROX: validation accuracy for different numbers of layers.

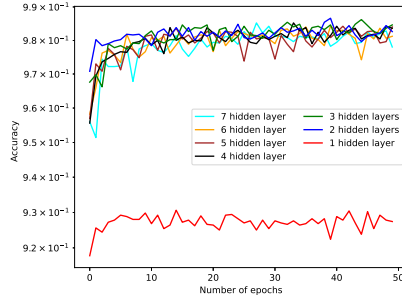


Figure 5: MC-APPROX: validation accuracy for different numbers of layers.

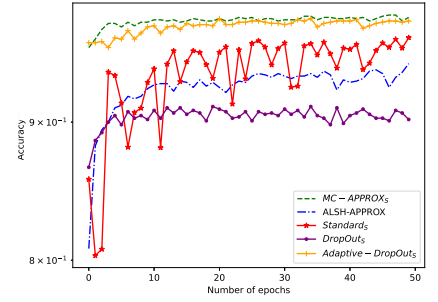


Figure 6: Validation accuracies with 3 hidden layers.

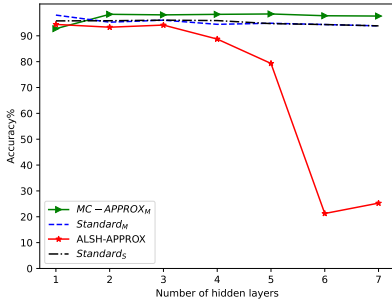


Figure 7: Test accuracy for different numbers of hidden layers.

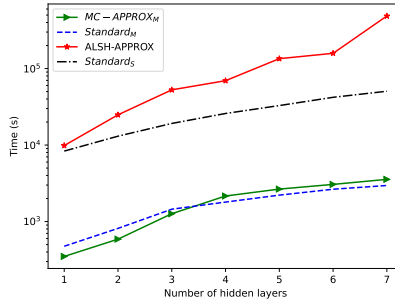


Figure 8: Training time for different numbers of hidden layers.

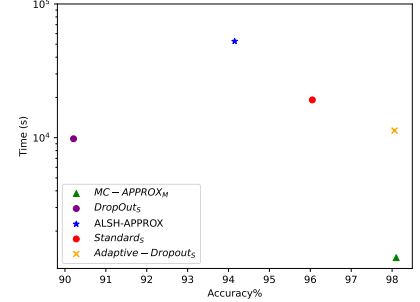


Figure 9: Accuracy vs. training time (3 hidden layers).

conducted additional experiments using 10- and 20-layer networks. The obtained accuracy rates of 97.32% for 10 layers and 95.71% for 20 layers validate our earlier findings. Finally, looking at Figure 7, the only case in which MC-APPROX_M fails to obtain the highest accuracy compared with ALSH-APPROX is when the model has only 1 hidden layer. ALSH-APPROX performs (94.4%) slightly better than MC-APPROX_M (92.71%).

9.2 Training Time

After studying the impact of network depth on accuracy, we next turn our attention to efficiency (i.e., training time). The results from all five methods on three hidden layers, on one CPU and without parallelization, are summarized in Table 3 and Table 4. Even though from Table 4 it is evident that MC-APPROX_M significantly outperforms other approaches with batch size 20, MC-APPROX_S is slower than ADAPTIVE-DROPOUT_S, STANDARD_S, and DROPOUT_S. Similarly, Figure 8 shows the runtime comparison of MC-APPROX_M and ALSH-APPROX with STANDARD_S and STANDARD_M (baseline) for different numbers of layers. The results confirm the superiority of MC-APPROX_M over the other algorithms up to 3 layers. Note that the observed increase in the training time of ADAPTIVE-DROPOUT per epoch in comparison to STANDARD can be attributed to the additional computational overhead of the construction of dropout masks and their subsequent multiplication with the weight matrices in each layer (Table 4).

ALSH-APPROX. Before discussing our efficiency results from ALSH-APPROX, let us emphasize that ALSH-APPROX is a scalable algorithm that significantly *benefits from parallelization* [50]. During training, the hash table construction, computing hash

signature, querying hash tables, and updating weight vectors by sparse weight gradients are parallelized, which makes the algorithm fast with parallel processing using multiple processing units. We refer interested readers to the details and results of Spring and Shrivastava [50].

ALSH-APPROX needs to reconstruct the hash tables after a set of weight updates. Following the original implementation of ALSH-APPROX, in our experiments, for the first 10000 training data points, we reconstruct hash tables every 100 images. Then gradually, we expand the set to avoid time-consuming table reconstructions and update the tables every 1000 images. This helps with directing the gradient and decreasing the redundancy in the dataset. Table 3 shows that ALSH-APPROX exhibits slower performance compared to all other methods in the statistic setting.

Also, in models with additional hidden layers, we can see an increase in training time as shown in Figure 8 that is larger than other methods on the same network structure. This is consistent with the results presented by Spring and Shrivastava [50], where it is shown that the runtime significantly drops with parallelization. Evidently, as shown by Spring and Shrivastava [50], parallelization has no impact on the accuracy of ALSH-APPROX. In other words, the accuracy scalability results of ALSH-APPROX discussed in the previous section are independent of parallelization.

MC-APPROX. Due to the sampling ratio of MC-APPROX ($p \approx 0.1$), MC-APPROX performs more atomic scalar operations than ALSH-APPROX with 5% of the nodes. However, based on the results from Figure 8 and Table 3, MC-APPROX_M and MC-APPROX_S are around 20 and 2 times faster than ALSH-APPROX, respectively. This is

Table 2: Test accuracy (%) for a network with 3 hidden layers.

Dataset	ALSH-APPROX	MC-APPROX _M	MC-APPROX _S	DROPOUT _S	ADAPTIVE-DROPOUT _S	STANDARD _S
MNIST	94.15	98.10	98.38	90.21	98.06	96.46
Kuzushiji-MNIST	72.87	91.78	96.50	9.84	90.73	83.86
Fashion-MNIST	78.11	87.85	88.58	76.28	86.12	73.64
EMNIST-Letters	64.97	89.84	90.75	4.96	89.50	85.34
NORB	78.57	92.05	97.52	32.73	96.60	51.61
CIFAR-10	10.31	73.26	62.11	67.85	75.55	93.02

Table 3: Training time per epoch (sec.) with 3 hidden layers and batch size 1 on MNIST.

Method	ALSH-APPROX	MC-APPROX _S	DROPOUT _S	ADAPTIVE-DROPOUT _S	STANDARD _S
Total	807.50 ± 22.92	422.23 ± 3.36	196.15 ± 0.55	225.85 ± 1.91	361.51 ± 5.13
Feedforward	168.02 ± 3.34	28.44 ± 0.076	32.32 ± 0.04	59.03 ± 0.10	28.976 ± 0.04
Backpropagation	356.16 ± 7.28	110.98 ± 0.69	131.02 ± 0.57	132.68 ± 1.89	61.88 ± 1.02

Table 4: Training time per epoch (sec.) with 3 hidden layers and batch size 20 on MNIST.

Method	MC-APPROX _M	DROPOUT _M	ADAPTIVE-DROPOUT _M	STANDARD _M
Total	35.141 ± 0.385	42.818 ± 0.108	53.412 ± 0.301	39.310 ± 1.087
Feedforward	4.733 ± 0.054	4.789 ± 0.010	10.539 ± 0.024	4.213 ± 0.052
Backpropagation	11.323 ± 0.071	7.957 ± 0.029	8.493 ± 0.271	23.143 ± 1.080

because of the significantly lower overhead of MC-APPROX compared to ALSH-APPROX. Figure 8 shows that MC-APPROX_M outperforms other methods while maintaining a training time comparable to STANDARD_M. Notably, despite the additional computational load in the backpropagation process with MC-APPROX_M, it achieves backpropagation times that are twice as fast as those of STANDARD on the same architecture (Table 4). For networks with fewer than 4 hidden layers, MC-APPROX_M is slightly faster than STANDARD_M, and for deeper networks, the opposite is true. Nevertheless, Figure 5 confirms the higher accuracy of the MC-APPROX_M for various numbers of layers on MNIST. From Table 3, it is evident that STANDARD_S is faster than MC-APPROX_S. The reason is that, in order to estimate probabilities based on Equation 7 for each mini-batch, MC-APPROX makes a pass over the mini-batch and the matrix W . As a result, in SGD, where mini-batch size is one, the overhead time and the time to approximate the matrix multiplication exceeds the required time for exact multiplication (STANDARD_S). In §9.4, we provide a thorough analysis of the runtime comparison between DROPOUT, ADAPTIVE-DROPOUT, and MC-APPROX_M.

Finally, Figure 9 shows that MC-APPROX performs better in both speed and accuracy compared to other methods.

9.3 Hyperparameters

MC-APPROX is designed for mini-batch stochastic gradient descent, as it uses the set of samples in the mini-batch for error estimation in Equation 7 to identify which rows in W^k to select. In SGD, the estimations would be made using only one sample, and hence are not reliable. As a result, while MC-APPROX performs well for mini-batch SGD with a large-enough batch size (20), its efficiency drops significantly for SGD. In the stochastic setting, MC-APPROX_S is slower than STANDARD_S (Table 3) and is prone to overfitting.

To evaluate this, we run experiments on stochastic setting where batch size is set to 1 (MC-APPROX_S) along with different

batch sizes. The results are provided in Figures 6, 10, 11, and 12 and Tables 3 and 4. The results in Figure 10 show the decrease in accuracy for small mini-batches with the same learning rate: the accuracy drops from 98% to 84%. The optimal learning rate to use is smaller for smaller batch sizes [47], so to resolve the overfitting in MC-APPROX_S, we decreased the learning rate from 10^{-3} to 10^{-4} . As shown in Table 2 and Figure 6, MC-APPROX_S performs well in terms of accuracy. Moreover, as the model gets more complex by adding hidden layers (Figure 12) and the variance increases, the model is unable to avoid overfitting even with decreasing the learning rate. We discussed in §6.2 that MC-APPROX chooses row-column pairs from matrices such that the columns are from the first input matrix $X \in \mathbb{R}^{m \times n}$ and the corresponding rows are sampled from the second matrix $W \in \mathbb{R}^{n \times n}$. Figure 12 provides evidence of the lack of scalability of MC-APPROX_S for deep networks. This can be attributed to the specific sampling procedure employed by MC-APPROX_S. In the stochastic setting, $X_{:,j}$ is reduced to a singleton set. As a result, the time overhead increases, while the reliability of probability estimation for row-column selection decreases.

9.4 Memory Analysis

Our computing architecture is equipped with a hierarchical memory layout that includes a 384 KB Level 1 (L1) cache, 12 MB Level 2 (L2) cache, and a 19.3 MB Level 3 (L3) cache. We find that the ALSH-APPROX algorithm initially requires 24 MB of memory for table setup and subsequently expands by 3,731.9 MB by the end of training. This substantial increase in memory usage indicates that the data that is not cache resident contributes to significant system overload. This plays a crucial role in the execution time of the ALSH-APPROX algorithm, as it necessitates frequent data retrieval from slower, non-cache memory sources. MC-APPROX demonstrates a 45 MB increment in memory usage by the end of training. Notably, memory usage decreases to 16.4 and 15.7 MB for DROPOUT and ADAPTIVE-DROPOUT, respectively. However,

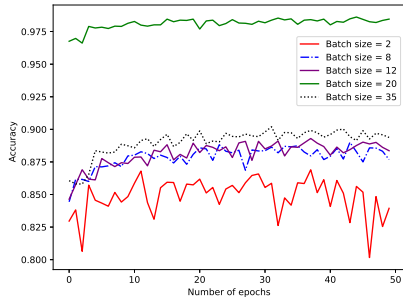


Figure 10: MC-APPROX: validation accuracy for different batch sizes (learning rate = 0.001).

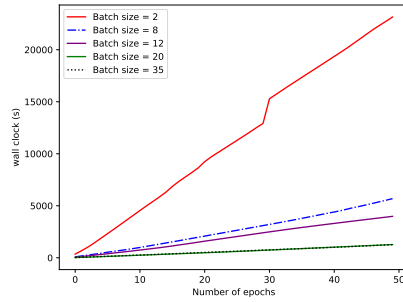


Figure 11: MC-APPROX: training time for different batch sizes (learning rate = 0.001).

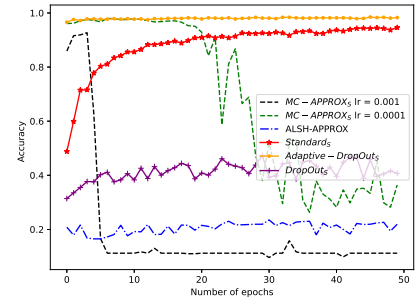


Figure 12: Validation accuracies with 7 hidden layers.

the I/O traffic remains high, and the runtime still increases similarly to MC-APPROX. This phenomenon is attributable to cache misses during training with DROPOUT and ADAPTIVE-DROPOUT. As Mohan et al. [42] illustrate, minimizing storage I/O per epoch is more critical than caching data. If data is evicted from the cache before use or is never cached, accessing it from memory increases the I/O overhead. Cache profiling results indicate a roughly 24% increase in cache misses with DROPOUT and a 27% increase with ADAPTIVE-DROPOUT compared to MC-APPROX, explaining the rise in runtime.

10 LESSONS AND DISCUSSION

It is important to recognize that no single method serves as the optimal solution for optimizing training across all architectures and datasets. On the positive side though, our theoretical analysis and experimental evaluations confirm the effectiveness of existing approaches under various settings. Specifically, our experimental findings discussed in §9 confirm the superiority of MC-APPROX. When used on MLPs with an appropriately sized mini-batch larger than one, it offers enhanced accuracy, speed and memory efficiency. MC-APPROX effectively approximates the inner product at each layer by sampling nodes from the previous layer, and utilizing a batch of data point vectors significantly improves the sampling quality. Another notable observation is that MC-APPROX decreases the number of cache misses which significantly contributes to its increased speed. Moreover, our theoretical analysis in §7 demonstrate a small estimation error for ALSH-APPROX, when the number of hidden layers is small. This was observed in our experiment in Figures 3m, Figure 3n, and Table 2. This confirm the effectiveness of ALSH-APPROX, especially with parallel computing [50].

10.1 Feedforward Approximation Scalability

A major takeaway in this paper is the negative impact of approximation during the feedforward process. First, in §7, we theoretically analyzed the error propagation effect from layer to layer. In particular, Theorem 7.2 highlights the *exponential* increase of gradient estimation error in ALSH-APPROX as the number of hidden layers increases. As a result, for neural networks with more than 3 hidden layers, the error can become even larger than the estimation value. Consequently, the gradient estimation can become utterly arbitrary, resulting in completely inaccurate weight updates during the backpropagation step, which leads to an inaccurate model. For MC-APPROX, Adelman et al. [1] did not observe consistent behavior across different models in their

experiments. Interestingly, the authors provide the theoretical result that (i) approximating both feedforward and backpropagation operations leads to biased estimates, and (ii) approximating only feedforward operations is unbiased. However, their method for feedforward approximation *failed* in experiments [1]. As a result, MC-APPROX (the algorithm evaluated in this paper) only adds approximation during backpropagation. We observed similar results for ALSH-APPROX in §9. It is evident in Figures 3, 7, 12, and 4 that ALSH-APPROX failed to scale for DNNs, confirming our theoretical analysis in Theorem 7.2.

We observed that backpropagation generally took significantly longer than the feedforward step (Tables 3 and 4). Fortunately, backpropagation optimization can significantly improve training time [22, 53]; introducing approximation only during the backpropagation step has the potential to significantly reduce training time. Nevertheless, designing scalable sampling-based algorithms that introduce approximation on both feedforward and backpropagation in DNNs on CPU machines remains an open research direction.

10.2 DNNs and Small Batch Size

As observed in our theoretical analysis and experiment results, ALSH-APPROX does not scale to DNNs with more than a few hidden layers. MC-APPROX, on the other hand, scales for DNNs, but it is designed based on mini-batch gradient descent and performs well when the batch size is reasonably large (greater than 10). However, the performance of MC-APPROX quickly drops for small batch sizes under the same setting. In particular, we observed a swift drop in time efficiency (Figure 11) under SGD. While MC-APPROX₅ demonstrated a high accuracy in some cases (Table 2), this comes at a cost of a significant increase in training time (even compared to STANDARDS) and a high risk of overfitting, especially for deep networks (Figure 12). Thus, designing scalable sampling-based algorithms for SGD on CPU remains an open research direction.

10.3 ALSH-APPROX Prediction in DNNs

We would like to conclude this section with an interesting observation on ALSH-APPROX. Let us consider the confusion matrices of ALSH-APPROX in Figure 3 (Row 4) once again. From Figure 3m, one can confirm that (i) there is no class imbalance in the test set (approximately same number of samples in each class), and (ii) having high accuracy, the model predictions are uniformly distributed across different classes (approximately same number of samples predicted to be in each class). On the other hand, in

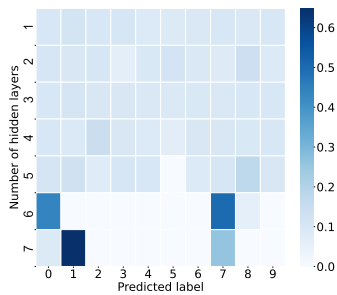
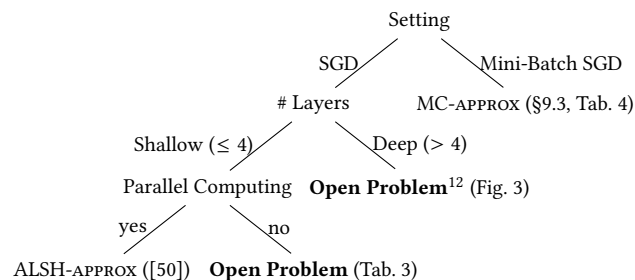


Figure 13: ALSH-APPROX: impact of network depth on label prediction ratio distribution.

Figure 3p, it seems not only the model is inaccurate, but interestingly only a few labels from the class labels are generated in predictions (all samples are labeled as either 0, 1, or 7). We repeated the experiment multiple times and observed consistent behavior. Furthermore, comparing Row 2 (DROPOUT) with Row 4 (ALSH-APPROX) in Figure 3, we note that, while both methods failed to scale with the number of hidden layers, DROPOUT maintains the label diversity in its prediction, which demonstrates randomness. To better present this, in Figure 13, we provide the ratio of test-set samples predicted for each label (column) for networks with various numbers of hidden layers. It is clear that, while initially the label prediction distribution is uniform, as the number of layers increases, the predictions concentrate around a few arbitrary classes. The reason is that, while training the model using ALSH-APPROX, as the gradient estimation error increases for deeper networks, a small subset of nodes remains active in deeper layers, regardless of the input sample. As a result, the set of edges for which the weights get updated remains almost the same. Therefore, when predicting the label of an input sample, the same set of nodes is “activated”, resulting in a small set of predictions generated.

10.4 Optimal Choice of Training Method

Building on the insights and lessons from our study, we present a decision tree to guide users in selecting the most effective method for training DNNs on CPU machines:



Our primary objective is to identify methods that deliver performance nearly comparable to standard training in terms of accuracy while ensuring faster execution on our CPU system across various settings. Our results, presented in Section 9.3 and Table 4, indicate that MC-APPROX surpasses other methods in mini-batch settings. In stochastic environments, the effective method varies with the network depth. The experimental evalua-

¹²“Open problem” refers to settings where existing sampling-based algorithms failed; further research is needed to design algorithms for those settings.

tion by Spring and Shrivastava [50] confirms that ALSH-APPROX scales effectively using parallel computing with multi-processors up to 2^6 processors, for up to four layers. Therefore, we believe ALSH-APPROX performs optimally up to four layers under such conditions.

11 CONCLUSION

Many of the advanced technologies originally developed for addressing big data challenges have been extended to solve scalability complexities across various domains. In this paper, we evaluated one of these settings, where sampling-based techniques were proposed for training DNN on CPU machines with limited resources. To this end, we made connection between two sampling-based research directions that can be viewed as matrix multiplication approximations. We provided theoretical analyses, followed by extensive empirical evaluations.

Our results demonstrate a correlation between the number of hidden layers and approximation error in DNNs under hashing-based methods. In addition, we provided insights into the performance of fast training methods in different settings and highlight areas for further research.

As a final note, energy consumption during DNN training raises environmental concerns. Recent studies have explored the significant carbon footprint associated with large-scale neural networks, primarily due to their energy consumption [45, 56]. One interesting direction for future work is to study the impact of sampling-based techniques on energy efficiency.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation, Grant No. 2107290, and the UIC University Fellowship. The authors would like to thank the anonymous reviewers and the meta-reviewer for their invaluable feedback.

REFERENCES

- [1] Menachem Adelman, Kfir Levy, Ido Hakimi, and Mark Silberstein. 2021. Faster Neural Network Training with Approximate Tensor Operations. In *Advances in Neural Information Processing Systems*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (Eds.), Vol. 34. Curran Associates, Inc., 27877–27889.
- [2] Thomas Dybdahl Ahle, Rasmus Pagh, Ilya Razenshteyn, and Francesco Silvestri. 2016. On the complexity of inner product similarity join. In *Proceedings of the 35th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*. 151–164.
- [3] Abolfazl Asudeh, Jeess Augustine, Saravanan Thirumuruganathan, Azade Nazi, Nan Zhang, Gautam Das, and Divesh Srivastava. 2021. Scalable signal reconstruction for a broad range of applications. *Commun. ACM* 64, 2 (2021), 106–115.
- [4] Abolfazl Asudeh, Azade Nazi, Jeess Augustine, Saravanan Thirumuruganathan, Nan Zhang, Gautam Das, and Divesh Srivastava. 2018. Leveraging similarity joins for signal reconstruction. *Proceedings of the VLDB Endowment* 11, 10 (2018), 1276–1288.
- [5] Martin Aumüller and Matteo Ceccarello. 2022. Implementing Distributed Approximate Similarity Joins using Locality Sensitive Hashing. In *EDBT: International Conference on Extending Database Technology*.
- [6] Lei Jimmy Ba and Brendan Frey. 2013. Adaptive Dropout for Training Deep Neural Networks. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (NIPS '13)*. Curran Associates Inc., Red Hook, NY, USA, 3084–3092.
- [7] Kevin Beyer, Rainer Gemulla, Peter J Haas, Berthold Reinwald, and Yann Sismanis. 2009. Distinct-value synopses for multiset operations. *Commun. ACM* 52, 10 (2009), 87–95.
- [8] Matthias Boehm, Michael W Dusenberry, Deron Eriksson, Alexandre V Evfimievski, Faraz Makari Manshadi, Niketan Pansare, Berthold Reinwald, Frederick R Reiss, Prithviraj Sen, Arvind C Surve, et al. 2016. SystemML: Declarative machine learning on spark. *PVLDB* 9, 13 (2016), 1425–1436.
- [9] Matthias Boehm, Shirish Tatikonda, Berthold Reinwald, Prithviraj Sen, Yuanyuan Tian, Douglas R Burdick, and Shivakumar Vaithyanathan. 2014. Hybrid parallelization strategies for large-scale machine learning in systemml. *Proceedings of the VLDB Endowment* 7, 7 (2014), 553–564.

- [10] Surajit Chaudhuri, Venkatesh Ganti, and Raghav Kaushik. 2006. A primitive operator for similarity joins in data cleaning. In *ICDE*. IEEE.
- [11] Tarin Clanuwat, Mikel Bober-Irizar, Asanobu Kitamoto, Alex Lamb, Kazuaki Yamamoto, and David Ha. 2018. Deep Learning for Classical Japanese Literature. (Dec. 2018). <https://doi.org/10.48550/arXiv.1812.01718>
- [12] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and Andre van Schaik. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 International Joint Conference on Neural Networks (IJCNN)*. Institute of Electrical and Electronics Engineers (IEEE), 2921–2926. <https://doi.org/10.1109/ijcnn.2017.7966217>
- [13] Jeffrey Dean, Greg Corrado, Rajat Monga, Kai Chen, Matthieu Devin, Mark Mao, Marc'auelio Ranzato, Andrew Senior, Paul Tucker, Ke Yang, Quoc Le, and Andrew Ng. 2012. Large Scale Distributed Deep Networks. In *Advances in Neural Information Processing Systems*, F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (Eds.), Vol. 25. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2012/file/6aca97005c68f1206823815f66102863-Paper.pdf
- [14] Li Deng. 2012. The MNIST Database of Handwritten Digit Images for Machine Learning Research. *IEEE Signal Processing Magazine* 29, 6 (Nov. 2012), 141–142. <https://doi.org/10.1109/msp.2012.2211477>
- [15] Petros Drineas, Ravi Kannan, and Michael W. Mahoney. 2006. Fast Monte Carlo Algorithms for Matrices I: Approximating Matrix Multiplication. *SIAM J. Comput.* 36, 1 (Jan. 2006), 132–157. <https://doi.org/10.1137/s0097539704442684>
- [16] John Duchi, Elad Hazan, and Yoram Singer. 2011. Adaptive Subgradient Methods for Online Learning and Stochastic Optimization. *Journal of Machine Learning Research (JMLR)* 12, 61 (2011), 2121–2159. <http://jmlr.org/papers/v12/duchi11a.html>
- [17] Sana Ebrahimi, Rishi Advani, and Abolfazl Asudeh. 2023. Sampling-Based Techniques for Training Deep Neural Networks with Limited Computational Resources: A Scalability Evaluation. *arXiv preprint arXiv:2306.09293* (2023).
- [18] Kayvon Fatahalian, Jeremy Sugerman, and Pat Hanrahan. 2004. Understanding the Efficiency of GPU Algorithms for Matrix-Matrix Multiplication. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware (HWS '04)*. Association for Computing Machinery, New York, NY, USA, 133–137.
- [19] Trevor Gale, Erich Elsen, and Sara Hooker. 2019. The state of sparsity in deep neural networks. (Feb. 2019).
- [20] Amol Ghoting, Rajasekar Krishnamurthy, Edwin Pednault, Berthold Reinwald, Vikas Sindhwani, Shirish Tatikonda, Yuanyan Tian, and Shivakumar Vaithyanathan. 2011. SystemML: Declarative machine learning on MapReduce. In *2011 IEEE 27th International conference on data engineering*. IEEE, 231–242.
- [21] Aristides Gionis, Piotr Indyk, and Rajeev Motwani. 1999. Similarity Search in High Dimensions via Hashing. In *Proceedings of the 25th International Conference on Very Large Data Bases (VLDB '99)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 518–529.
- [22] Ian J. Goodfellow, Yoshua Bengio, and Aaron Courville. 2016. *Deep Learning*. MIT Press, Cambridge, MA, USA. <http://www.deeplearningbook.org>.
- [23] Kazushige Goto and Robert A. van de Geijn. 2008. Anatomy of high-performance matrix multiplication. *ACM Transactions on Mathematical Software (TOMS)* 34, 3 (2008), 1–25.
- [24] Marios Hadjieleftheriou, Xiaohui Yu, Nick Koudas, and Divesh Srivastava. 2008. Hashed samples: selectivity estimators for set similarity selection queries. *PVLDB* 1, 1 (2008), 201–212.
- [25] Song Han, Huizi Mao, and William J. Dally. 2016. Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding. *ICLR*.
- [26] Sona Hasani, Saravanan Thirumuruganathan, Abolfazl Asudeh, Nick Koudas, and Gautam Das. 2018. Efficient construction of approximate ad-hoc ML models through materialization and reuse. *Proceedings of the VLDB Endowment* 11, 11 (2018), 1468–1481.
- [27] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. 2018. AMC: AutoML for Model Compression and Acceleration on Mobile Devices. In *Computer Vision – ECCV 2018*. Springer International Publishing, 815–832. https://doi.org/10.1007/978-3-030-01234-2_48
- [28] Piotr Indyk and Rajeev Motwani. 1998. Approximate Nearest Neighbors: Towards Removing the Curse of Dimensionality. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing (STOC '98)*. Association for Computing Machinery, 604–613.
- [29] Dhiraj Kalamkar, Evangelos Georganas, Sudarshan Srinivasan, Jianping Chen, Mikhail Shiryayev, and Alexander Heinecke. 2020. Optimizing Deep Learning Recommender Systems Training on CPU Cluster Architectures. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–15. <https://doi.org/10.1109/SC41405.2020.00047>
- [30] Diederik P. Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. (Dec. 2014). <https://doi.org/10.48550/arXiv.1412.6980>
- [31] Alex Krizhevsky. 2009. *Learning Multiple Layers of Features from Tiny Images*. Technical Report.
- [32] Arun Kumar, Matthias Boehm, and Jun Yang. 2017. Data management in machine learning: Challenges, techniques, and systems. In *Proceedings of the 2017 ACM International Conference on Management of Data*. 1717–1722.
- [33] Arun Kumar, Jeffrey Naughton, Jignesh M Patel, and Xiaojin Zhu. 2016. To join or not to join?: Thinking twice about joins before feature selection. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 19–34.
- [34] Yann LeCun, Fu Jie Huang, and Léon Bottou. 2004. Learning methods for generic object recognition with invariance to pose and lighting. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004. (CVPR '04)*, Vol. 2. Institute of Electrical and Electronics Engineers (IEEE). <https://doi.org/10.1109/CVPR.2004.1315150>
- [35] Guoliang Li, Xuanhe Zhou, and Lei Cao. 2021. AI meets database: AI4DB and DB4AI. In *Proceedings of the 2021 International Conference on Management of Data*. 2859–2866.
- [36] Zhe Li, Boqing Gong, and Tianbao Yang. 2016. Improved Dropout for Shallow and Deep Learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NIPS '16)*. Curran Associates Inc., Red Hook, NY, USA, 2531–2539.
- [37] Yuzhe Ma, Ran Chen, Wei Li, Fanhua Shang, Wenjian Yu, Minsik Cho, and Bei Yu. 2019. A Unified Approximation Framework for Compressing and Accelerating Deep Neural Networks. In *2019 IEEE 31st International Conference on Tools with Artificial Intelligence (ICTAI)*. Institute of Electrical and Electronics Engineers (IEEE), 376–383. <https://doi.org/10.1109/ictai.2019.00060>
- [38] Alireza Makhzani and Brendan Frey. 2015. Winner-Take-All Autoencoders. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2 (NIPS '15)*. MIT Press, Cambridge, MA, USA, 2791–2799.
- [39] Giosuè Cataldo Marinò, Flavio Furia, Dario Malchiodi, and Marco Frasca. 2023. Efficient and Compact Representations of Deep Neural Networks via Entropy Coding. *IEEE Access* 11 (2023), 106103–106125. <https://doi.org/10.1109/ACCESS.2023.3317293>
- [40] Stefano Markidis, Steven W. D. Chien, Erwin Laure, Ivy Bo Peng, and Jeffrey S. Vetter. 2018. NVIDIA Tensor Core Programmability, Performance & Precision. *2018 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* (2018), 522–531.
- [41] B McMahan and Daniel Ramage. 2017. *Federated learning: Collaborative machine learning without centralized training data*. Technical Report. Technical report, Google.
- [42] Jayashree Mohan, Amar Phanishayee, Ashish Raniwala, and Vijay Chidambaram. 2020. Analyzing and mitigating data stalls in DNN training. *arXiv preprint arXiv:2007.06775* (2020).
- [43] James Jie Pan, Jianguo Wang, and Guoliang Li. 2023. Survey of vector database management systems. *arXiv preprint arXiv:2310.14021* (2023).
- [44] Yongjoo Park, Barzan Mozafari, Joseph Sorenson, and Junhao Wang. 2018. Verdictdb: Universalizing approximate query processing. In *Proceedings of the 2018 International Conference on Management of Data*. 1461–1476.
- [45] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R So, Maud Texier, and Jeff Dean. 2022. The carbon footprint of machine learning training will plateau, then shrink. *Computer* 55, 7 (2022), 18–28.
- [46] Christian P. Robert. 2016. *Monte Carlo Methods*. Wiley, 1–13 pages. <https://doi.org/10.1002/9781118445112.stat03876.pub2>
- [47] Christopher J. Shallow, Jaehoon Lee, Joseph Antognini, Jascha Sohl-Dickstein, Roy Frostig, and George E. Dahl. 2019. Measuring the Effects of Data Parallelism on Neural Network Training. *Journal of Machine Learning Research (JMLR)* 20, 112 (2019), 1–49.
- [48] Anshumali Shrivastava and Ping Li. 2014. Asymmetric LSH (ALSH) for Sublinear Time Maximum Inner Product Search (MIPS). In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc.
- [49] Matthew S. Smith. 2023. The Case for Running AI on CPUs Isn't Dead Yet. *IEEE Spectrum* (1 June 2023).
- [50] Ryan Spring and Anshumali Shrivastava. 2017. Scalable and Sustainable Deep Learning via Randomized Hashing. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 445–454.
- [51] Ryan Spring and Anshumali Shrivastava. 2020. Mutual Information Estimation using LSH Sampling. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, Christian Bessiere (Ed.). International Joint Conferences on Artificial Intelligence Organization, 2807–2815. <https://doi.org/10.24963/ijcai.2020/389>
- [52] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research (JMLR)* 15, 1 (Jan. 2014), 1929–1958.
- [53] Xu Sun, Xuancheng Ren, Shuming Ma, Bingzhen Wei, Wei Li, and Houfeng Wang. 2017. Training Simplification and Model Simplification for Deep Learning: A Minimal Effort Back Propagation Method. *IEEE Transactions on Knowledge and Data Engineering* 32 (2017), 374–387.
- [54] Yao Tian, Xi Zhao, and Xiaofang Zhou. 2023. DB-LSH 2.0: Locality-Sensitive Hashing With Query-Based Dynamic Bucketing. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [55] Vincent Vanhoucke, Andrew Senior, and Mark Z. Mao. 2011. Improving the speed of neural networks on CPUs. (2011).
- [56] Roberto Verdecchia, June Sallou, and Luís Cruz. 2023. A systematic review of Green AI. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 13, 4 (2023), e1507.
- [57] Sida Wang and Christopher Manning. 2013. Fast dropout training. In *Proceedings of the 30th International Conference on Machine Learning (Proceedings of Machine Learning Research)*, Sanjoy Dasgupta and David McAllester (Eds.), Vol. 28. PMLR, Atlanta, Georgia, USA, 118–126. <https://doi.org/10.5555/3042817.3042907>

- [58] Ruidi Wei and Florian Kerschbaum. 2023. Cryptographically Secure Private Record Linkage using Locality-Sensitive Hashing. *Proceedings of the VLDB Endowment* 17, 2 (2023), 79–91.
- [59] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. (Sept. 2017). <https://doi.org/10.48550/arXiv.1708.07747>
- [60] Tiechui Yao, Jue Wang, Junyu Gu, Yumeng Shi, Fang Liu, Xiaoguang Wang, Yangang Wang, and Xuebin Chi. 2023. A Sparse Matrix Optimization Method for Graph Neural Networks Training. In *International Conference on Knowledge Science, Engineering and Management*. Springer, 114–123.
- [61] Ce Zhang, Arun Kumar, and Christopher Ré. 2016. Materialization optimizations for feature selection workloads. *ACM Transactions on Database Systems (TODS)* 41, 1 (2016), 2.
- [62] Xuanhe Zhou, Chengliang Chai, Guoliang Li, and Ji Sun. 2020. Database meets artificial intelligence: A survey. *IEEE Transactions on Knowledge and Data Engineering* 34, 3 (2020), 1096–1116.
- [63] Xiaotian Zhu, Wengang Zhou, and Houqiang Li. 2018. Adaptive layerwise quantization for deep neural network compression. In *2018 IEEE International Conference on Multimedia and Expo (ICME)*. Institute of Electrical and Electronics Engineers (IEEE), 1–6.