



Direct Feedback Alignment With Sparse Connections for Local Learning

Brian Crafton*, Abhinav Parihar, Evan Gebhardt and Arijit Raychowdhury*

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, United States

OPEN ACCESS

Edited by:

Arindam Basu,
Nanyang Technological University,
Singapore

Reviewed by:

Timothy P. Lillicrap,
Google, United States
Hesham Mostafa,
University of California, San Diego,
United States

*Correspondence:

Brian Crafton
brian.crafton@gatech.edu
Arijit Raychowdhury
arijit.raychowdhury@ece.gatech.edu

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 29 January 2019

Accepted: 07 May 2019

Published: 24 May 2019

Citation:

Crafton B, Parihar A, Gebhardt E and
Raychowdhury A (2019) Direct
Feedback Alignment With Sparse
Connections for Local Learning.
Front. Neurosci. 13:525.
doi: 10.3389/fnins.2019.00525

Recent advances in deep neural networks (DNNs) owe their success to training algorithms that use backpropagation and gradient-descent. Backpropagation, while highly effective on von Neumann architectures, becomes inefficient when scaling to large networks. Commonly referred to as the weight transport problem, each neuron's dependence on the weights and errors located deeper in the network require exhaustive data movement which presents a key problem in enhancing the performance and energy-efficiency of machine-learning hardware. In this work, we propose a bio-plausible alternative to backpropagation drawing from advances in feedback alignment algorithms in which the error computation at a single synapse reduces to the product of three scalar values. Using a sparse feedback matrix, we show that a neuron needs only a fraction of the information previously used by the feedback alignment algorithms. Consequently, memory and compute can be partitioned and distributed whichever way produces the most efficient forward pass so long as a single error can be delivered to each neuron. We evaluate our algorithm using standard datasets, including ImageNet, to address the concern of scaling to challenging problems. Our results show orders of magnitude improvement in data movement and 2× improvement in multiply-and-accumulate operations over backpropagation. Like previous work, we observe that any variant of feedback alignment suffers significant losses in classification accuracy on deep convolutional neural networks. By transferring trained convolutional layers and training the fully connected layers using direct feedback alignment, we demonstrate that direct feedback alignment can obtain results competitive with backpropagation. Furthermore, we observe that using an extremely sparse feedback matrix, rather than a dense one, results in a small accuracy drop while yielding hardware advantages. All the code and results are available under <https://github.com/bcrafton/ssdfa>.

Keywords: bio-plausible algorithms, feedback alignment, local learning, backpropagation, sparse neural networks, hardware acceleration

1. INTRODUCTION

The demise of Dennard scaling (Dennard et al., 1974) and decline of Moore's Law (Moore, 1965) have exposed the fundamental scaling limitations of the von Neumann models of computing. This transition is accompanied by the realization that in a fast evolving, socially interconnected world, we are observing a seismic shift in the amount of unstructured data that need to be processed in

real-time (Najafabadi et al., 2015) which has heralded the third wave of Artificial Intelligence and the exponential growth of Machine Learning in data-analytics, real-time control, computer vision, robotics, and so on. We expect that intelligent systems of the future will be limited by the energy growth of data movement rather than compute. Therefore, we need fundamentally new approaches to sustain the exponential growth in performance beyond the end of the current road-map. In particular, we observe that new computing models that deal with “data analytics” have compute and storage interleaved in a fine grained manner—not separated as in the von Neumann world. Moving forward, computing technology will heavily penalize separation of data and compute and we need to marry them in better ways to handle emergent applications.

The idea of computing locally on data finds its inspiration from the human brain where local processing and updates is preferred to global movement of data. Hence, neuromorphic computing seeks to fundamentally improve the power efficiency of cognitive systems by bringing ideas inspired from biology to electronic hardware, while maintaining the high accuracy and performance that statistical methods have provided. In particular, hardware implementations of deep neural networks (DNNs) and their many variants, either in complementary metal oxide semiconductor (CMOS) (Chen et al., 2016) or emerging technologies (Li et al., 2018), rely on arrays of spatial processors where near-memory (Merolla et al., 2014; Chen et al., 2016; Bankman et al., 2019) or in-memory (Chi et al., 2016) logic computes inference from layers of neurons connected via dense or sparse synaptic connections. This is schematically shown in **Figure 1**. As opposed to a von Neumann architecture (**Figure 1A**) where all the synaptic weights for a particular layer must be loaded in the memory to compute the activations of the successive layers, in a distributed implementation (**Figure 1B**) the weights and logic reside locally and avoid the memory bottleneck. The data movement is minimized, with each neuron computing its activation and sending that information to the next layer. In spite of the success of such spatial processing in the inference mode, such an architecture fails to deliver high efficiency when functions that require global information or weights of multiple layers are implemented. This is particularly evident during the training of DNNs where back-propagation (BP) and stochastic gradient-descent (SGD) have found widespread adoption (LeCun et al., 2015). In BP with SGD, the transpose of the weights of the deeper layers in a network are needed to compute error gradients and the weight updates of the shallower layers, thus requiring global movement of data. This problem is commonly referred to as the weight transport problem (Grossberg, 1987; Lillicrap et al., 2016). As the networks become deeper to keep up with the complexity of the applications, the weight transport problem becomes exacerbated.

While neural networks require many expensive multiply-and-accumulate (MAC) operations, the cost of data movement is higher (Chen et al., 2016). To make the problem worse, for every MAC operation multiple reads and potentially a write are required. Furthermore, the cost of loading data from off chip DRAM is orders of magnitude more expensive than loading from spatially local on chip memories. Kwon et al. (2018) show

through simulation the breakdown of energy consumption in the different layers of a CNN. The majority of the energy consumption comes from memory reads and writes, rather than MAC operations. The strategy these accelerators employ is data reuse. Because the cost of data movement is so high, it is important that each word of data is reused as many times as possible before being flushed from the cache. Optimizing for data reuse, these accelerators can achieve several times better efficiency over data flows that do not use local reuse. Neuromorphic engineering is another research vector which attempts to minimize data movement by borrowing learning rules from the human brain. For example, in spiking neural networks (SNNs), spike timing dependent plasticity (STDP) has gained popularity because of its local update rule. STDP uses local information, available at a synapse, and has been shown to perform well in unsupervised learning (Diehl and Cook, 2015) and supervised learning using Feedback Alignment (Nefcici et al., 2017; Nefcici, 2018). Davies et al. (2018) present a new SNN implementation with tools to perform supervised learning. Rather than seeking to optimize current neural network architectures using data reuse, Davies et al. (2018) uses biological constraints on data movement and calls for new approaches to learning. The fundamental constraint is that a weight can only be accessed and modified by its corresponding destination neuron. This is further illustrated in **Figure 1**, where each neuron is its own module containing compute and memory. To promote local learning, the weights local to the neuron should not be sent to or from the neuron, only activations and error signals. This constraint promotes local learning since the only data movement that occurs are activation signals between adjacent neurons and error computed by the system. With this constraint in place, we define data movement as information a neuron must send or receive for each weight update.

One such promising recent work is Feedback Alignment (FA) (Lillicrap et al., 2016), which has shown that we might be able to bypass the weight transport problem while achieving the same accuracy that BP achieves. FA uses fixed random feedback weights to propagate the errors back through the layers of a DNN rather than using the actual current network weights to compute the partial error. Consequently, the weights in the shallow layers of the network no longer need information about the weights of all the deeper layers. Building on top of this, Nøkland (2016) proposes Direct Feedback Alignment (DFA), where it was shown that the feedback to shallow layers need not be propagated through all the layers. Instead the error signal can be fed back to the shallower layers through completely random linear transformations. This further reduces the amount of information required to update the weights in the network. To further describe the weight transport problem and its relationship with local learning, Baldi et al. (2018) describes the concept of a learning channel. The learning channel is a physical way in which information about targets and deep weights are transported in the network. Backpropagation uses the forward channel in the backward direction. Using the targets and deep weights we compute partial errors to update the rest of the network. Feedback Alignment instead uses a separate channel to transmit the weights and in doing so avoids the weight transport problem.

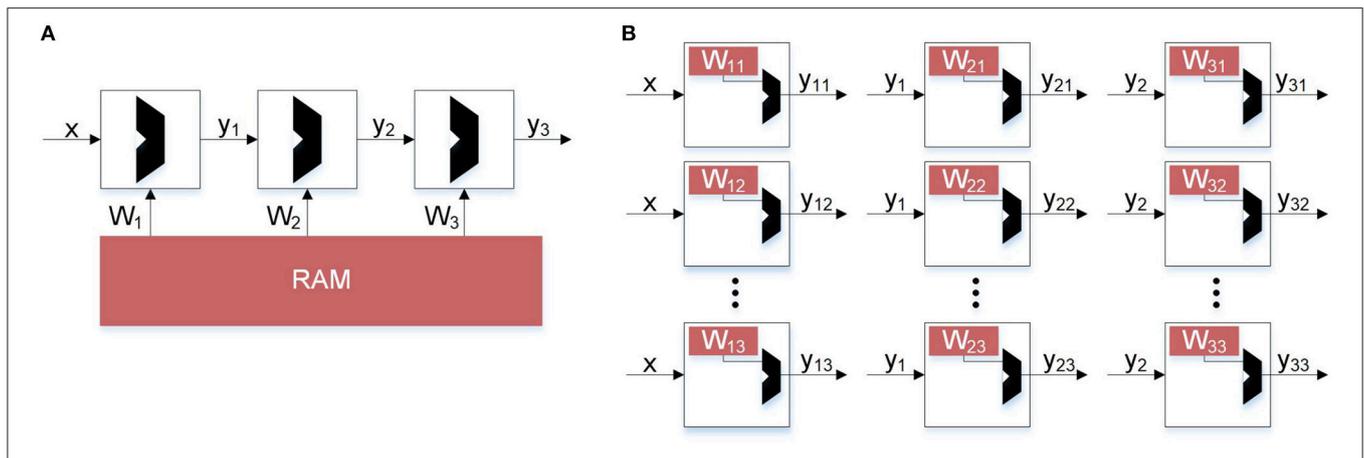


FIGURE 1 | Hardware implementations of inference. Comparing von Neumann architecture with distributed memory architecture to avoid bottleneck. **(A)** Inference constrained by von Neumann architectures. In traditional backpropagation, weight updates not only depend on error but also the other weights. This prevents a distributed architecture. **(B)** Unconstrained inference. Using Direct Feedback Alignment, inference can be distributed and parallel because weight updates depend only on the error and random feedback values.

Other promising algorithms include target propagation (Lee et al., 2015) and local error learning (Mostafa et al., 2018). Target propagation bypasses the weight transport problem while still solving the credit assignment problem like the Feedback Alignment algorithms. Instead of computing a loss gradient, a target value is assigned to each feed forward layer using auto encoders. Local error learning generates local errors at each layer using linear classifiers with fixed random weights. In doing so, errors at the output of the network do not need to be sent back and instead local objective functions are solved rather than a global one. Local error learning bypasses the weight transport problem, solves the credit assignment problem, and also does not even need to send errors to the hidden layers. All these algorithms have been shown to perform similarly on benchmarks. While their performance is impressive, they still suffer the same problems when scaled to larger networks. Each of the algorithms fails to match the performance of BP as both the complexity of the network and difficulty of the benchmark increases. Bartunov et al. (2018) highlighted this problem showing that when applied to problems like CIFAR100 and ImageNet, the biological algorithms failed to come close to backpropagation. For a couple benchmarks we do observe considerable degradation due to fully connected layers, however we show this problem comes primarily from convolutional layers. In fact, when trained layers are transferred from backpropagation and the fully connected layers are trained using feedback alignment the performance is similar to that of backpropagation.

Although DFA does not require the error signals to be transformed by the weights of the deeper layers of the network, each neuron requires feedback weights for each error in the network. This requires a significant number of computations, memory, and data movement that compromises the locality of the algorithm. While shallow weights are no longer dependent on deep weights, the amount of data movement and memory required to compute the error at each neuron prohibits a near memory architecture. In this paper, we show a modified version

of DFA, sparse direct feedback alignment (SDFA), where we propose that sparse feedback of the error signals can result in small drop in the network’s performance but significantly reduces the computational complexity during learning. In an extreme version of SDFA, we demonstrate that even a single error feedback signal can enable the network to learn with a small performance loss. We call this single connection SDFA (SSDFA). We systematically study, through empirical demonstrations, the role of sparsity and rank of the feedback matrix on the network’s performance. Our work demonstrates that SSDFA inherits the computational advantages of local learning similar to bio-mimetic networks, while maintaining the high accuracy of BP with SGD (hitherto simply referred to as BP).

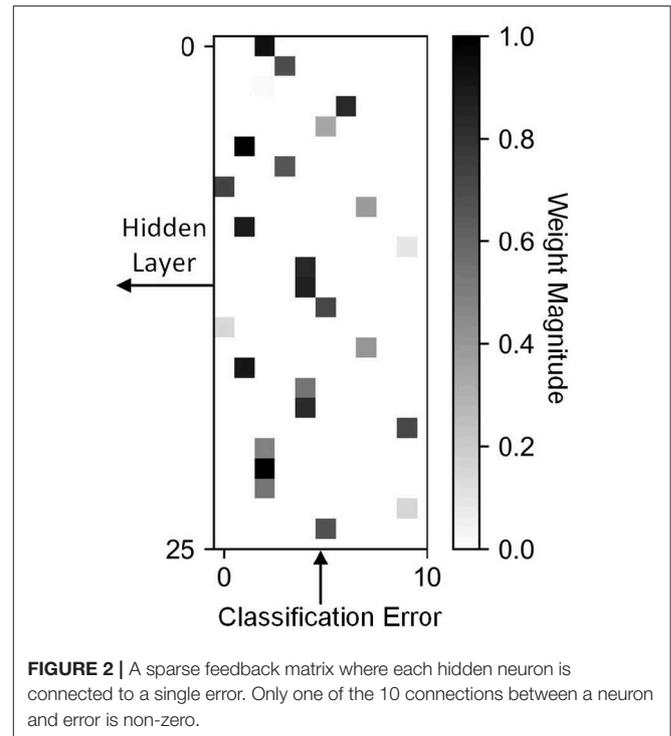
2. SPARSE DIRECT FEEDBACK ALIGNMENT

Direct Feedback Alignment was a remarkable step forward in training fully connected networks. By replacing the backward propagation from deeper layers with a single random matrix, we can avoid the weight transport problem and enable new hardware for training neural networks. In small networks, DFA seems feasible since each neuron requires connections to only a few error signals. However, as the size of the network increases, the size of the feedback matrix also increases and in effect, each weight update needs more information. As an illustrative example, consider a three layered network with 100 hidden neurons which can be trained for classifying the MNIST handwritten digit dataset. Because each image size is 28 by 28 pixels and there are 10 classes, our network size will be $784 (28 \times 28) - 100 - 10$. In this example, each of the 100 hidden neurons require connections to each of the 10 error signals at the output. While this may seem plausible for MNIST, for the AlexNet (Krizhevsky et al., 2012) or VGG16 (Simonyan and Zisserman, 2014) networks that are used to classify the

ImageNet dataset, (Deng et al., 2009) the number of connections becomes much larger. The ImageNet dataset classifies 1,000 different classes of images which are re-sized to 224 by 224 pixels (Simonyan and Zisserman, 2014). In the case of DFA, each of the 4,096 neurons in the first fully connected layer requires a connection to each of the 1000 errors. This quadratic increase in connections prevents DFA from scaling to larger problems without incurring significant computational penalty compared to a more local learning rule.

To relax this problem, we introduce S DFA where the error signal is fed back to all the hidden layer neurons through a highly sparse feedback matrix. S DFA with a sparse feedback matrix enables each neuron to compute its error using a fewer number of error signals. Consequently, as we will demonstrate, a hardware design that implements S DFA based learning requires significantly less data movement in the form of error signals, rendering it more efficient both in terms of throughput and power. We empirically demonstrate that sparsity plays a negligible role on the network's accuracy as long as the feedback matrix is full-rank or near-full-rank. This leads to SS DFA, which can enable local learning requiring only a single global error to be transferred per neuron, while incurring a small loss of accuracy.

We define a feedback as sparse when the values in the feedback matrix that model the connections between the errors and the hidden neurons are mostly zero. The implication on hardware is that most of the connections do not exist, and therefore, do not require data movement or computation. Biological networks share similar properties, where each neuron updates its weights using local values (O'Reilly and Munakata, 2000; Diehl and Cook, 2015). In this work, we use the percentage of zero valued connections to quantify sparsity. In **Figure 2** we present an example of a sparse feedback matrix where each of the 25 hidden neurons is connected to one of 10 errors. In this case the sparsity is 90% because only one of 10 errors are used. We also demonstrate that even an extremely sparse feedback matrix, with 99.9% sparsity, can be used to achieve high accuracy on the ImageNet dataset. In **Figure 3**, we schematically compare BP, DFA, and S DFA for a prototypical fully connected network consisting of four layers of neurons. **Figure 3** shows the number of connections the pre-synaptic neuron in the first hidden layer needs to compute a weight update. In BP, the neuron is dependent on each and every neuron it has a direct or indirect connection to. Hence, the weight update of the synapse is dependent on all model weights deeper in the network. The number of connections grows as the network gets deeper, leading to the weight transport problem. On the other hand, in DFA, the neuron is dependent only on the number of errors of the network. While this decouples the forward and backward pass and relieves the weight transport problem, complex networks with many output neurons will still have a large number of feedback connections. In the proposed S DFA, a weight update is dependent only on the few errors it is connected to. Consequently, even if the network scales in both depth and complexity, the number of connections for updating a synaptic weight is low. In the proposed extreme scenario for S DFA, we name SS DFA, each neuron receives only one error signal and can successfully update its weights with a single error signal.



Using only a single error per neuron, we are able to greatly reduce the amount of data movement in the backwards pass. While there will be significant savings in memory accesses and MAC operations, the key improvement is in data movement. However, in order to take advantage of this, a von Neumann architecture is not an ideal choice for bio-mimetic algorithms. In particular, recent advances in spatial array processors with near-memory computing can achieve significant advantages in both performance and energy-efficiency (Hsu, 2014; Chen et al., 2017). In **Figure 4** we show a comparison of the different algorithms and the data movement required in the backwards pass. We illustrate in **Figures 4A–C** how the data movement greatly reduces. With backpropagation on a von Neumann architecture (**Figure 4A**) we must read the weights and activations to and from the main memory. DFA on a neuromorphic architecture (**Figure 4B**) relaxes the weight transport problem and keeps the read, write, and MAC operations local to the neuron itself with the only data movement occurring when the error is sent backwards. Lastly, in **Figure 4C** the error sent backwards is only a single scalar, reducing data movement in the backwards pass to its minimal form. In **Table 1**, we compare the number of reads, writes, MACs, and data movement across the different algorithms as a function of the neurons and weights in the network. $|A|$, $|W|$, $|B|$, and $|E|$ refer to the number of neurons, weights, feedback weights, and errors, respectively. We also use $|b|$ and $|e|$ to refer to the reduced number of weights and errors from the sparse feedback matrix.

3. MATHEMATICAL FORMULATION

The primary contribution of this work is to show the benefits of using sparse feedback that can eventually enable local learning

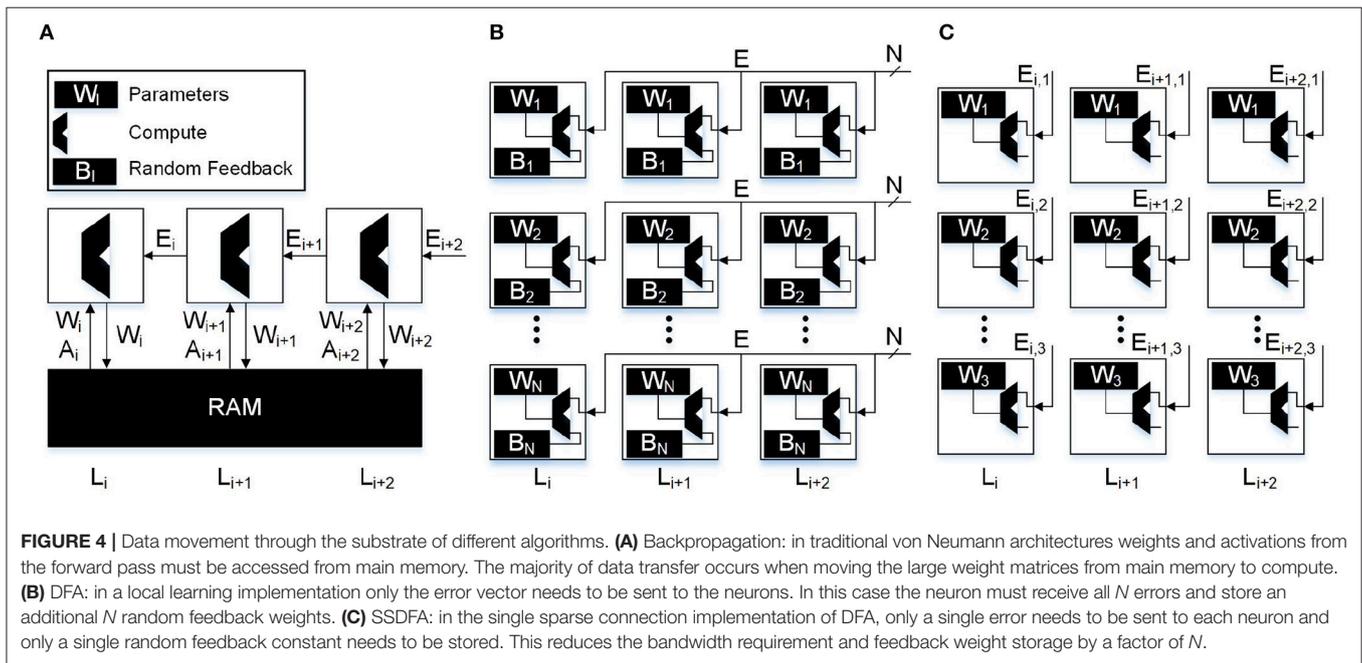
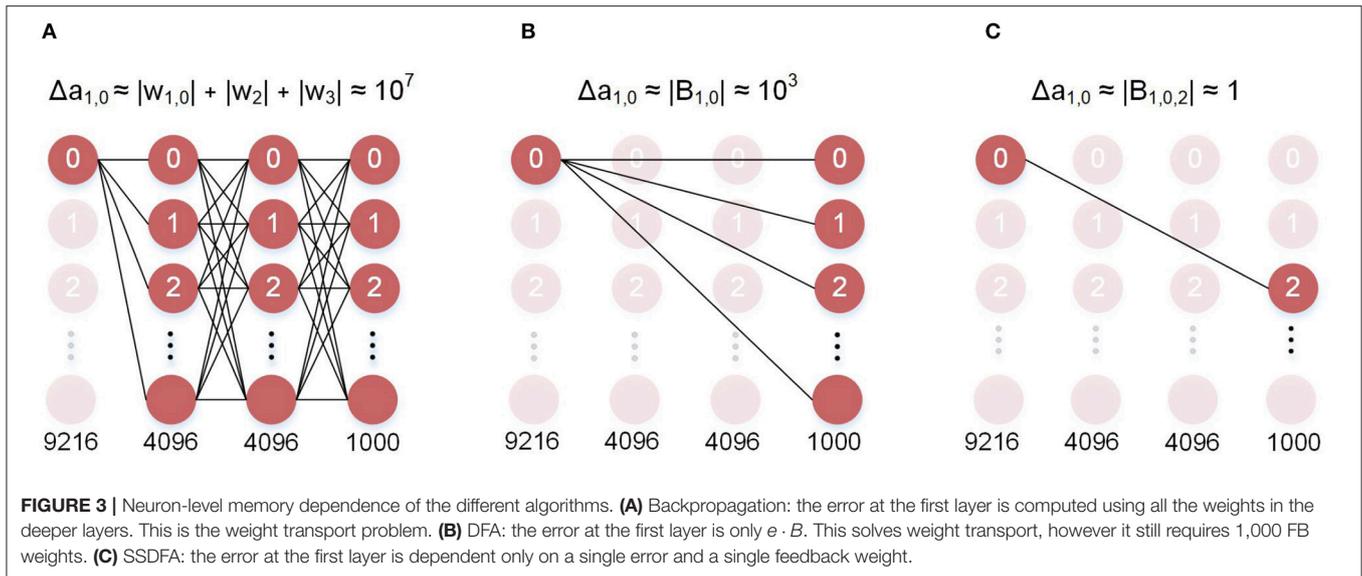


TABLE 1 | Data movement comparison of error assignment algorithms.

Method	Reads	Writes	MACs	Movement
BP	$ W + A $	$ W $	$2 W $	$ W + A + E_{i+1} $
DFA	$ W + B $	$ W $	$ W + B $	$ E $
SDFA	$ W + b $	$ W $	$ W + b $	$ e $

in neuromorphic hardware implementations. Empirically, we show on standard networks and datasets that sparsity, even extreme-sparsity (SSDFA), results in negligible loss of accuracy while reducing data movement during training by orders of magnitude.

We investigate BP, DFA, and SDFA for the fully connected network architecture. The feed forward computation can be written as

$$y_1 = W_1 \cdot x, a_1 = f(y_1) \tag{1}$$

$$y_2 = W_2 \cdot a_1, a_2 = f(y_2) \tag{2}$$

$$y_n = W_n \cdot a_{n-1}, a_n = f(y_n) \tag{3}$$

where x is the feature vector and W_i is the weight matrix connecting layer $i - 1$ to layer i ($y_0 = x$). The dot product of x and W_i yields y_i , and applying the non-linear activation function f results in the activation at layer i , a_i . Each of these algorithms computes the error at a specific layer. The error at the last layer of the network, n , is the classification error e . BP computes the error

at each hidden layer l , δa_l , by transposing the weight matrices W and multiplying by the gradient of the activation function. These layerwise computations for BP can be written as

$$\delta a_n = e \tag{4}$$

$$\delta a_2 = W_2^T \cdot \delta a_3 \odot f'(a_2) \tag{5}$$

$$\delta a_1 = W_1^T \cdot \delta a_2 \odot f'(a_1) \tag{6}$$

where \odot is the element-wise multiplication operator. BP requires all the deeper weights in the network in order to compute the error at a layer earlier in the network. DFA bypasses this dependence, only requiring a random matrix B to be multiplied with the error vector. The layerwise error computations for DFA are

$$\delta a_n = e \tag{7}$$

$$\delta a_2 = B_2 \cdot e \odot f'(a_2) \tag{8}$$

$$\delta a_1 = B_1 \cdot e \odot f'(a_1) \tag{9}$$

where we observe that the error at each layer does not depend on the error at any other layer. In the proposed S DFA, the error computation is identical to Equations (7–9), except with an added constraint that B is sparse. At the neuron level this represents an important detail in the physical implementation of neuromorphic hardware. In DFA the error computation for a neuron i in layer l is

$$\delta a_{l,i} = B_{l,i} \cdot e \cdot f'(a_{l,i}) \tag{10}$$

which is the inner product of two vectors multiplied by a scalar. This is computationally challenging for complex networks. For example, when we use VGG16 network on the ImageNet dataset, the number of MAC operations for each neuron is 1,000 corresponding to the number of output classes. On the other hand, in S DFA, when B is sparse, the number of MACs required for each update will be significantly less, depending on the number of non-zero entries in the corresponding row of B . In an extreme scenario, when each row of B has only 1 non-zero entry, the error computation reduces to

$$\delta a_{l,i} = B_{l,i,j} \cdot e_j \cdot f'(a_{l,i}) \tag{11}$$

which is the product of three scalars.

The significance of this comes from past work in local learning and bio-plausible 3-factor learning rules. Locality is a constraint on the learning rule, and for a learning rule to be local each of the variables used in the learning rule must also be local. Baldi and Sadowski (2016) simplifies learning rules to the following forms

$$\delta W_{ij} = f(o_i, o_j, W_{ij}) \tag{12}$$

$$\delta W_{ij} = f(o_i, o_j, W_{ij}, t_{ij}) \tag{13}$$

where o_i and o_j are the output values of neurons i and j connected by weight W_{ij} . The variable t_{ij} is the target value computed for weight W_{ij} . An example of Equation (12) is Hebbian learning (Hebb, 1949) using Oja’s rule (Oja, 1982). In Oja’s rule the synapse uses only the activations of the neurons that it is connected to,

and its own state variable. Commonly referred to as the three factor learning rule (Baldi and Sadowski, 2016), Equation (13) requires a target value. Each of the algorithms discussed in this work (BP, DFA, SS DFA), can be simplified to Equation (13). In a supervised learning problem, computing the target value requires information about the network error. By the definition we gave earlier, we claim that data movement is information a neuron must send or receive. Therefore, any information sent or received by the neuron is non-local, which implies that error information in the backwards pass is non-local. Furthermore, we realize that any supervised learning algorithm would be non-local since all error information is not local to the neuron. As a result, none of the learning rules discussed in this paper exhibit pure locality, which would only be possible using a variant of Oja’s rule given this definition. Although some recent work (Mostafa et al., 2018) has shown a method of local supervised learning, for practical problems the learning rule fails to remain local. Local Error Learning (Mostafa et al., 2018) locally stores the labels at each layer and computes the prediction error using a random matrix at that layer. Since the prediction error is computed locally, it no longer needs to be transported from the end of the network to the given layer. In theory this prevents data movement, but for practical problems it is not possible to store data labels in this way. Two examples of systems where this technique fails to scale are real time systems and very large datasets. For real time systems like an autonomous drone, we do not have the labels to store locally in each layer. These labels would need to be transported from main memory, which is more expensive than moving errors from the last layer of the network. A similar problem arises for very large datasets. For very large datasets like ImageNet, we would need to store over one million labels locally in each layer or neuron. Storing this amount of information is not possible to do while maintaining locality.

The algorithms discussed in this paper have different levels of locality, which depend on the amount of non-local information required to compute the target value. To compute the target value for BP, we require all the downstream weights and errors. By computing the errors layer by layer, we are able to avoid redundant computation. However, this still requires the most data movement. In the case of DFA, we can store the random feedback weights local to the neurons themselves so the only data movement required are the downstream errors. As we scale DFA to larger datasets with more errors and therefore larger feedback matrices, we further compromise the locality of the algorithm because we cannot store large feedback matrices and route hundreds of errors to a single neuron. When the feedback matrix is made sparse (S DFA), we can scale the size of the network without requiring dense feedback connections. When only a single feedback connection is used (SS DFA), the target value is a function of only a single error and a single feedback weight. Although SS DFA uses non-local information, it is minimal in the sense that each neuron needs to receive only a single error.

4. RESULTS

To evaluate the performance of S DFA and SS DFA, we benchmark the proposed algorithms on standard vision datasets vis-a-vis BP

and DFA. We empirically evaluate the accuracy of the models with respect to the desired characteristics of the feedback matrix.

4.1. Benchmarks and Network Architectures

In our simulations we use four standard benchmarks of varying complexity that have been used in prior work. These benchmarks are MNIST (LeCun et al., 1990), CIFAR10, CIFAR100 (Krizhevsky and Hinton, 2009), and ImageNet (Deng et al., 2009). ImageNet is by far the most exhaustive dataset, and like (Bartunov et al., 2018), we find that using this benchmark reveals scaling issues related to the Feedback Alignment algorithms. For each of these benchmarks we use two networks. For MNIST, CIFAR10, and CIFAR100 we use one fully connected network and one convolutional network which are further described in Table 2. For ImageNet we use two standard convolutional networks: AlexNet (Krizhevsky et al., 2012) and VGG16 (Simonyan and Zisserman, 2014), which have produced high accuracy with BP. For each dataset, the corresponding model architectures are summarized in Table 2.

4.2. S DFA and SSDFA for Fully Connected Network

We start by reporting our findings on fully connected networks. The three properties of the feedback matrix that we study are:

1. **Rank:** The number of linearly independent vectors in the feedback matrix.
2. **Sparsity:** The percentage of zero-valued weights in the feedback matrix connecting each hidden layer to the classification error. We are particularly interested in the case

- where each hidden neuron is connected to a single error (SSDFA).
3. **Angle:** The angle between the vectorized weight matrix (W) and the vectorized feedback matrix (B).

The evaluation platform is setup in TensorFlow (Abadi et al., 2016). For MNIST and CIFAR10 there are 10 output neurons and hence 10 error signals. For each combination of rank and sparsity we simulate the models 10 times and statistically observe the accuracy of the network. To ensure that the network has feedback from all ten errors, the product of rank and connectivity ($1 - \text{sparsity}$) must be equal to or >1 . In order to create a $N \times M$ feedback matrix connecting N hidden neurons and M errors with rank R and sparsity S , we need R linearly independent vectors with sparsity S . When the rank of the matrix is not full, some of the vectors must be linearly dependent. As a result, they must have the same zero valued indices. When the rank and connectivity product is <1 , some columns in our matrix will be completely zero. This implies these errors will never be propagated backwards and the error at the layer will not depend on these classes. Consequently, we are unable to generate results for cases where the product is <1 . If the network does not have feedback from all ten errors, then the network's performance is impacted, not because of lower rank or higher sparsity, but because only some of the error signals are being propagated back. **Rank:** Empirically, we observe that the rank of a feedback matrix has the largest impact on the resulting accuracy of the network. As we increase the rank, the accuracy of the network increases and finally saturates. In Figure 5 we show our results for the MNIST and CIFAR10 datasets. In Figures 5A,B, we show the accuracy and angle vs. the rank of the feedback matrix, for varying sparsity. For the MNIST dataset, we observe that the test accuracy saturates, as expected at 97.5%, and is maximum for the full-rank matrix. However, for CIFAR10, the accuracy continues to increase as a function of rank without saturation. This shows that the rank of the feedback matrix is a critical design parameter and the feedback needs to be a full-rank matrix to maximize the network's accuracy.

Sparsity: Our results show that sparsity of the feedback matrix has very little impact on the resulting accuracy of the network in both the MNIST and CIFAR10 networks (Figure 5). We also observe that sparsity has a negligible impact on the angle between the feedback matrix used and the resulting feed forward weights (Figure 6). For many resource constrained systems, a small difference in accuracy for large improvements in power and performance is an excellent trade-off. We observe that training with highly sparse feedback matrices, even with just a single feedback error, performs very well—while significantly reducing the computational demand, as we describe later.

Angle: Lillicrap et al. (2016) shows that at each hidden layer, the angle between error gradient computed by FA and BP decreases as the network is trained. Instead we look at the angle between the vectorized weight matrix and feedback matrix after training. In the case of propagating errors to shallow layers of the network, the corresponding weight matrix used for the angle calculation is the product of all the weights matrices following this layer. Hence, for the angle calculation for a layer l in a n layer network,

TABLE 2 | Network architectures.

Dataset	Fully connected network	Convolutional network
MNIST	FC 400 Softmax 10	Conv (3×3 1×1 32) Conv (3×3 1×1 64) Pool (2×2 2×2) FC 128 Softmax 10
CIFAR (10, 100)	FC 1000 FC 1000 FC 1000 Softmax (10, 100)	Conv (5×5 1×1 96) Pool (3×3 2×2) Conv (5×5 1×1 128) Pool (3×3 2×2) Conv (5×5 1×1 256) Pool (3×3 2×2) FC 2048 FC 2048 Softmax (10, 100)
ImageNet	–	AlexNet (Krizhevsky et al., 2012)
ImageNet	–	VGG16 (Simonyan and Zisserman, 2014)

The format for convolutional layers is as follows: Conv (kernel size, number output channels, stride size). The format for pool layers is as follows: Pool (kernel size, stride size). The networks used for the CIFAR10 and CIFAR100 datasets were both derived from Nøklund (2016).

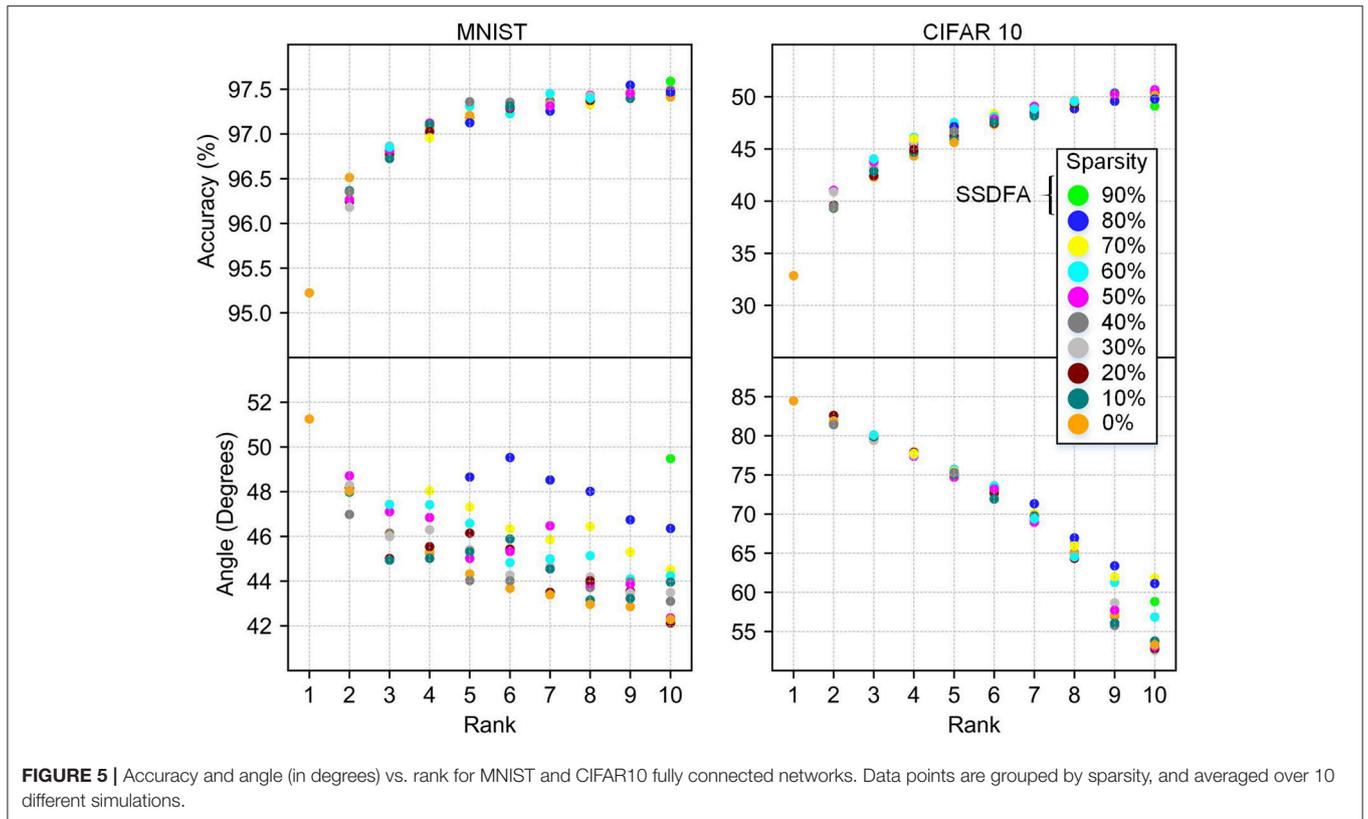


FIGURE 5 | Accuracy and angle (in degrees) vs. rank for MNIST and CIFAR10 fully connected networks. Data points are grouped by sparsity, and averaged over 10 different simulations.

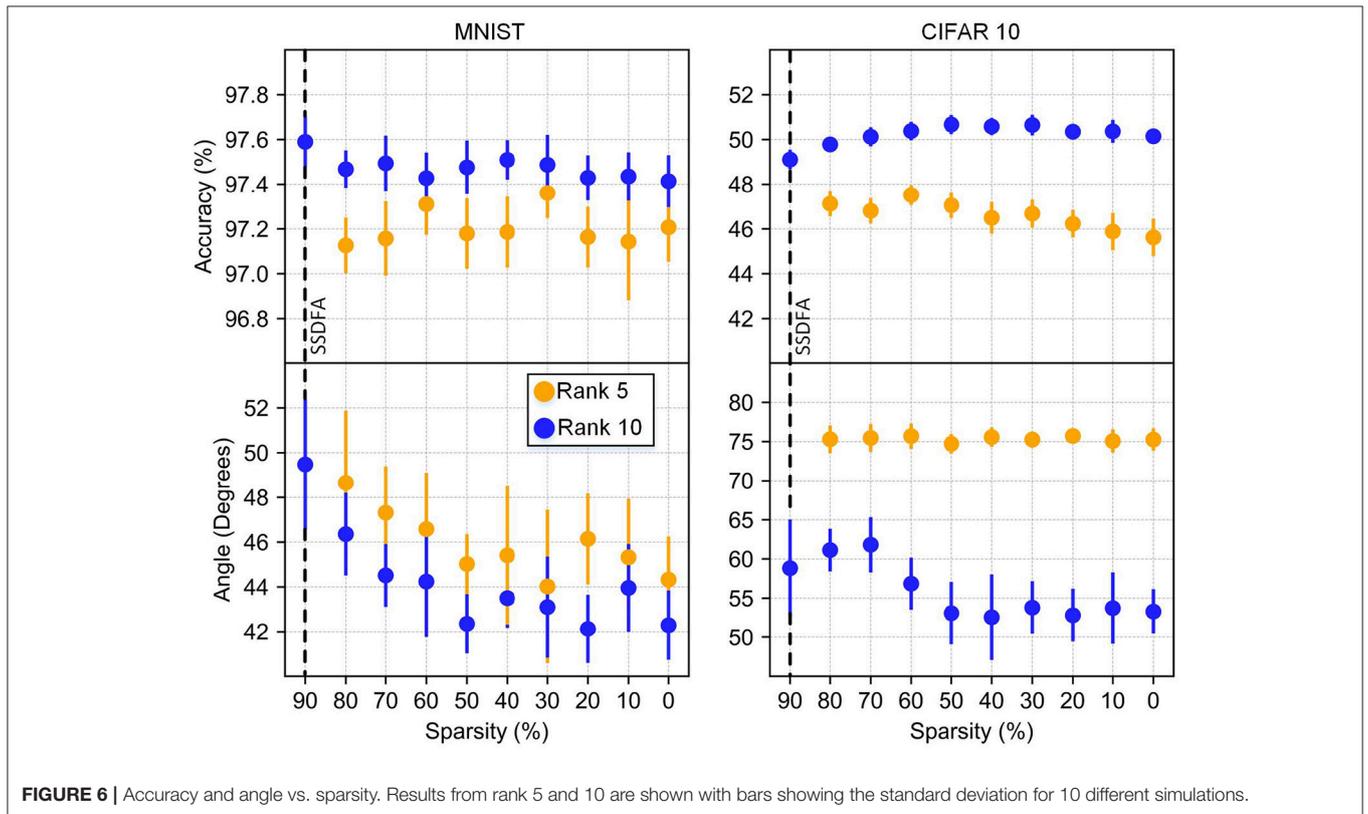


FIGURE 6 | Accuracy and angle vs. sparsity. Results from rank 5 and 10 are shown with bars showing the standard deviation for 10 different simulations.

the angle is given by

$$\angle(B_l, W) = \angle(B_l, \prod_{i=l+1}^n W_i) \quad (14)$$

In **Figure 5** we show a strong correlation of accuracy and angle between the feedback matrix and the weight matrix. This correlation is more apparent for CIFAR10 because it does not plateau as it approaches a full-rank matrix. This correlation is also illustrated in **Figure 6** because neither the angle nor the accuracy changes significantly as a function of sparsity.

In **Table 3** we summarize the results for fully connected networks. We show the results collected from running BP, DFA,

TABLE 3 | Test accuracy (in %) for fully connected networks.

Benchmark	BP	DFA	SSDFA
MNIST	98.2	97.8	97.5
CIFAR10	59.9	58.9	58.6
CIFAR100	33.1	29.4	28.8

TABLE 4 | Test accuracy (in %) for the convolutional neural networks.

Benchmark	BP	DFA	SSDFA
MNIST	99.1	98.9	98.8
CIFAR10	79.6	72.3	73.1
CIFAR100	51.0	43.7	41.8
ImageNet (AlexNet)	53.6	6.2	2.8

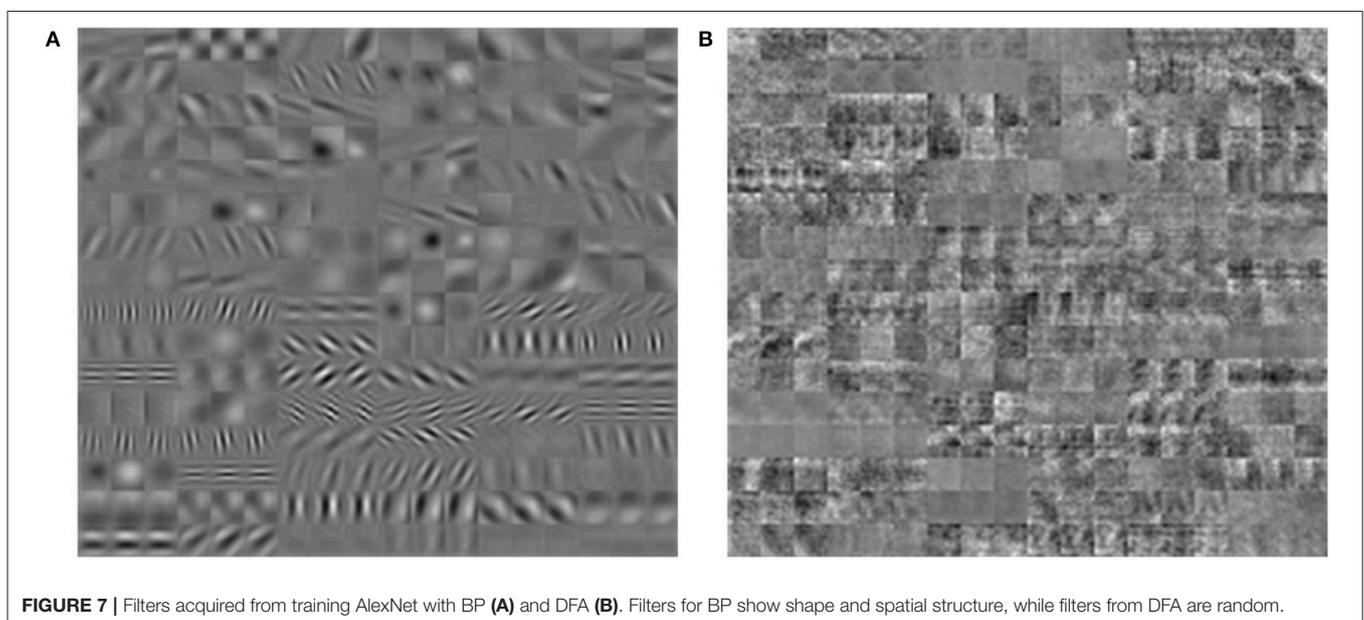
Accuracy results are slightly less for ImageNet than originals because minimal pre-processing and data augmentation techniques were used. All results are reported as Top-1 accuracy.

and SSDFA on MNIST, CIFAR10, and CIFAR100 with fully connected networks of different sizes summarized in **Table 2**. All the network parameters and code to run experiments can be found under: <https://github.com/bcrafton/ssdfa>. The results yielded by BP and DFA are similar to previous work (Nøkland, 2016), which shows that DFA performs similarly to BP for the fully connected networks. Also like previous work (Nøkland, 2016), we observe a non-negligible drop in accuracy for the CIFAR100 benchmark using both DFA and SSDFA. Since both DFA and SSDFA fail to match backpropagation for this benchmark, we infer that the issue is likely caused by the direct feedback and not the sparsity.

4.3. SDFA and SSDFA for Convolutional Neural Networks

Earlier work (Nøkland, 2016) show that DFA can be used in convolutional neural networks (CNNs). However, we note, similar to Bartunov et al. (2018), as we increase the complexity of our network and dataset, DFA fails to match the accuracy of BP. For the convolutional network benchmarks, we show our results in **Table 4**. Consistent with previous work, the CNNs we use for MNIST show similar performance for BP, DFA, and SSDFA. However, as the problem and network complexities scale to CIFAR10, CIFAR100, and ImageNet, the gap between BP and DFA grows to where their performance is no long comparable.

We attribute this problem to the convolutional layers in the network. The convolutional architecture introduces extra constraints on the weight space, and as such, results in stronger constraints on the feedback matrix which can be used. This can be seen in **Figure 7**, where we show the filter weights for the first convolutional layer from training AlexNet on ImageNet using BP and DFA. The resulting filters show completely different patterns. The BP filters show a well-defined spatial structure, while the filters from DFA are noisy with much less structure. From this,



we infer that DFA cannot learn convolutional filters as efficiently as BP for large and complex networks.

One way to alleviate this problem is Transfer Learning (Pan and Yang, 2010). Transfer Learning is a technique where a model from a prior task is used to initialize a model for another task. This is of particular importance in mobile platforms where pre-trained models are transferred to the mobile platform and the models are further refined via on-device learning. By borrowing the weights from the convolutional layers of CNNs which yield good results, we can bypass training these layers with DFA or SSDFA. We demonstrate that by reusing these convolutional

filters and training the fully connected layers using DFA and SSDFA, we can achieve similar performance to BP. In **Table 5**, we show the results for five different benchmarks of varying complexity. In each of these benchmarks, we transfer the weights from CNNs trained with BP and train the fully connected layers of each network. In all of these benchmarks, the performance of DFA and SSDFA is competitive with BP. The largest performance degradation we observe is 2.7% on the ImageNet benchmark using the AlexNet network.

4.4. Computational Advantage

The primary advantage of using SSDFA is that it greatly reduces data movement in the backwards pass. Furthermore, SSDFA also reduces the number of multiply-and-accumulate operations (MACs) and memory-reads. Computationally, this is motivated by biology where memory and compute are interleaved and global movement of data is minimal. New neuromorphic architectures for deep learning (Shin et al., 2017; Lee et al., 2018, 2019) and reinforcement learning (Amaravati et al., 2018a,b; Cao et al., 2019; Kim et al., 2019) seek to apply this constraint to avoid the communication overhead. However, as we have noted earlier, BP violates this constraint. Rather than requiring local information, it requires information from the weights of the deeper layers in the network. By decoupling the forward and backward weights, DFA and SSDFA do not require information about the weights deeper in the network.

TABLE 5 | Test accuracy (in %) for after transfer learning of the convolutional filters followed by SSDFA for the fully connected layers.

Benchmark	BP	DFA	SSDFA
MNIST	99.1	99.1	99.0
CIFAR10	77.1	77.8	78.0
CIFAR100	48.2	49.0	48.2
ImageNet (AlexNet)	49.0	48.8	46.3
ImageNet (VGG16)	65.8	65.3	64.5

Accuracy results are slightly less for ImageNet than originals because minimal pre-processing and data augmentation techniques were used. All results are reported as Top-1 accuracy.

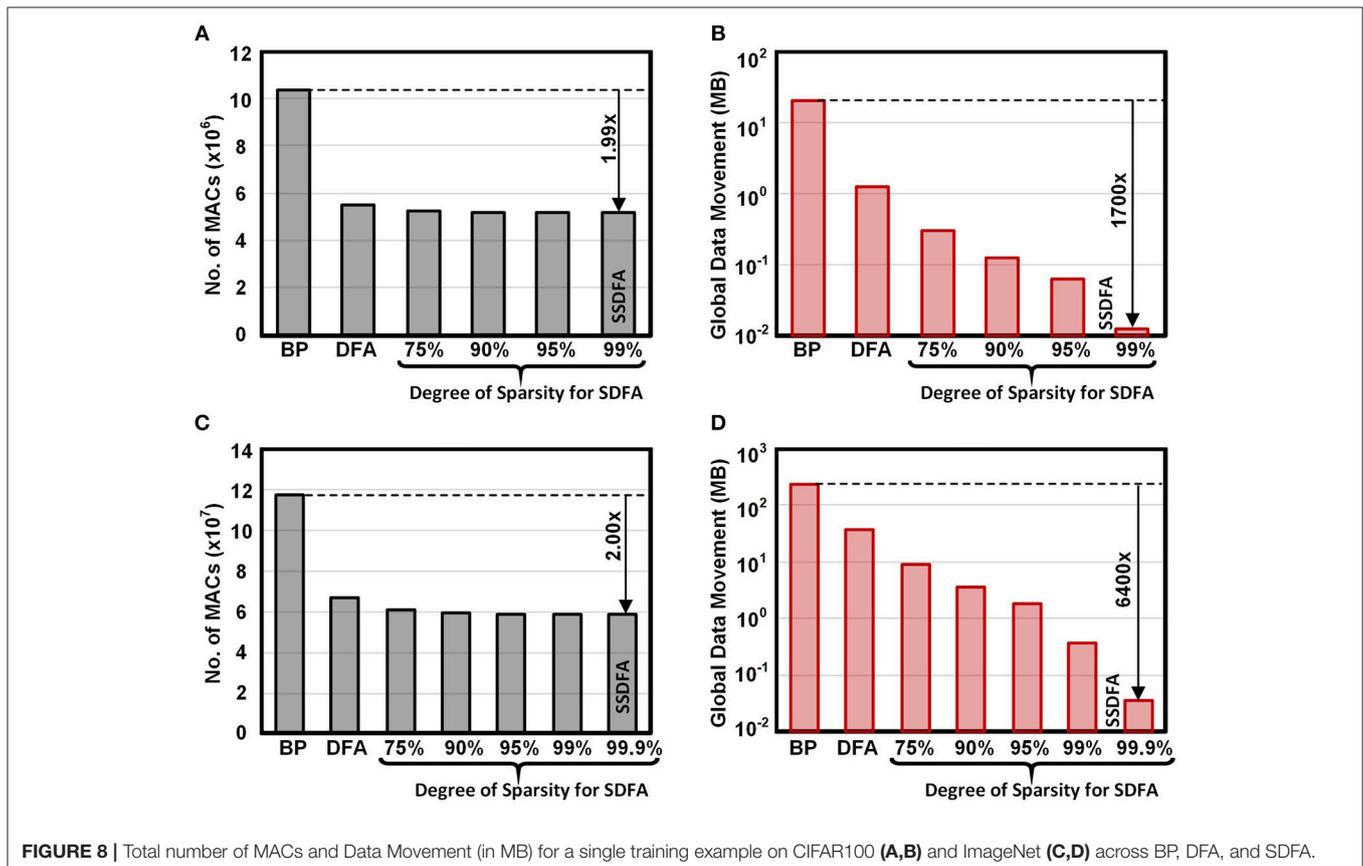


FIGURE 8 | Total number of MACs and Data Movement (in MB) for a single training example on CIFAR100 (A,B) and ImageNet (C,D) across BP, DFA, and SSDFA.

In **Figure 8**, we show the MAC and data movement savings when implemented in a near memory architecture on the CIFAR100 and ImageNet datasets. In section 2, we defined data movement as any data that must be sent outside of the neuron. In the backwards pass this is the error information sent to each neuron in SSDFA and DFA. For BP implemented on a von Neumann architecture, this is equivalent to all the data that must be transported to and from main memory. In **Figures 8B,D** we show the data movement results across SSDFA, DFA, and BP for CIFAR100 and ImageNet. On the ImageNet dataset we observe a $6,400\times$ reduction in data movement from BP to SSDFA. This number reflects our approximations in **Table 1**, where BP requires $|W| + |A| + |E_{i+1}|$ words of data and SSDFA requires just a single error, e . From DFA to SSDFA, we observe a $1,000\times$ reduction in data movement since each neuron receives only a single error rather than 1,000. We show the data movement requirements for the CIFAR100 benchmark in **Figures 8C,D**. We observe that the advantages of using SSDFA are less for smaller benchmarks because the number of errors and size of the hidden layers are less.

In **Figures 8A,C** we show the number of MACs across SSDFA, DFA, and BP for CIFAR100 and ImageNet. The reduction in the number of MACs from BP to SSDFA is nearly a factor of two. This reflects our approximations in **Table 1**, where we show that the number of MACs reduces from $2|W|$ to $|W| + |b|$. The number of MACs required to compute the error at a hidden layer reduces from $|W|$ to $|b|$, but the number of MACs to compute the partial error at each weight remains $|W|$. Since we must compute the error at each feedforward weight, the reduction in MACs will always be bounded. If the feedforward matrix, W , was also sparse, then this bottleneck could be reduced and allow for a larger reduction in MAC operations.

5. DISCUSSION

Feedback Alignment and Direct Feedback Alignment have been proposed to address the weight transport problem in backpropagation. We propose Sparse DFA where the fixed feedback matrix used is constrained to be sparse. Such a sparse feedback matrix helps in a simpler physical implementation of communicating errors. To justify our arguments, we studied the training performance of networks MNIST, CIFAR10, CIFAR100, and ImageNet using feedback matrices with different constraints, such as rank and sparsity. We observe that rank of the feedback matrix has much stronger impact on accuracy, and making the feedback connections sparse has negligible effect on

performance. Furthermore, using an extremely sparse version of SDFA where only a single error is fed back for weight update, we observe comparable performance while minimizing data movement.

As was claimed in Bartunov et al. (2018), Feedback Alignment, Direct Feedback Alignment, and as a result, the proposed Sparse Direct Feedback Alignment, do not scale to large networks. While our results are similar, we show that this is specifically true for convolutional networks due to the additional architectural constraints of repeated filters which they incorporate. We show that by fixing the convolutional filters of the network to those trained using BP, and learning only the fully connected layers using DFA, we observe that the network's performance is close to that of BP. Our approach greatly simplifies the task of propagating the errors from the deeper end of the network to the shallow layers for learning the model weights.

6. METHODS

We construct our networks in TensorFlow and create a new layer type so that we can feedback error directly to each hidden layer. To get optimal results we perform a hyper parameter search and also test different gradient descent optimizers and activation functions. We sweep learning rate and learning rate decay to find the optimal set. For weight initialization we used a uniform distribution in the range $[-1/\sqrt{fanout}, 1/\sqrt{fanout}]$. For feedback matrix initialization we used $[-1/\sqrt{fanin}, 1/\sqrt{fanin}]$ where the input dimension was layer size. For sparse matrices we found $[-1/\sqrt{\frac{fanin \cdot N}{fanout}}, 1/\sqrt{\frac{fanin \cdot N}{fanout}}]$ (where N is the number of sparse connections) to work well.

AUTHOR CONTRIBUTIONS

BC developed the main ideas, worked on the simulation experiments, and wrote the paper. AP worked on the simulation experiments and wrote the paper. EG worked on the simulation experiments. AR developed the main ideas and wrote the paper.

FUNDING

This work was funded by the U.S. Department of Defense's Multidisciplinary University Research Initiatives (MURI) Program under grant number FOA: N00014-16-R-FO05 and the Semiconductor Research Corporation under the Center for Brain Inspired Computing (C-BRIC).

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). "Tensorflow: a system for large-scale machine learning," in *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)* (Savannah, GA), 265–283.
- Amaravati, A., Nasir, S. B., Thangadurai, S., Yoon, I., and Raychowdhury, A. (2018b). "A 55 nm time-domain mixed-signal neuromorphic accelerator with stochastic synapses and embedded reinforcement learning for autonomous micro-robots," in *Solid-State Circuits Conference-(ISSCC), 2018 IEEE International* (San Francisco, CA: IEEE), 124–126.
- Amaravati, A., Nasir, S. B., Ting, J., Yoon, I., and Raychowdhury, A. (2018a). "A 55-nm, 1.0–0.4 v, 1.25-pj/mac time-domain mixed-signal neuromorphic accelerator with stochastic synapses for reinforcement learning in autonomous mobile robots. *IEEE J. Solid State Circuits*. 54, 75–87. doi: 10.1109/JSSC.2018.2881288

- Baldi, P., and Sadowski, P. (2016). A theory of local learning, the learning channel, and the optimality of backpropagation. *Neural Netw.* 83, 51–74. doi: 10.1016/j.neunet.2016.07.006
- Baldi, P., Sadowski, P., and Lu, Z. (2018). Learning in the machine: random backpropagation and the deep learning channel. *Artif. Intell.* 260, 1–35. doi: 10.1016/j.artint.2018.03.003
- Bankman, D., Yang, L., Moons, B., Verhelst, M., and Murmann, B. (2019). An always-on 3.8 μ J 86% cifar-10 mixed-signal binary cnn processor with all memory on chip in 28-nm cmos. *IEEE J. Solid State Circuits* 54, 158–172. doi: 10.1109/JSSC.2018.2869150
- Bartunov, S., Santoro, A., Richards, B. A., Hinton, G. E., and Lillicrap, T. (2018). Assessing the scalability of biologically-motivated deep learning algorithms and architectures. *arXiv preprint arXiv:1807.04587*.
- Cao, N., Chang, M., and Raychowdhury, A. (2019). “14.1 A 65 nm 1.1-to-9.1 tops/w hybrid-digital-mixed-signal computing platform for accelerating model-based and model-free swarm robotics,” in *2019 IEEE International Solid-State Circuits Conference-ISSCC* (San Francisco, CA: IEEE), 222–224.
- Chen, Y.-H., Krishna, T., and Emer, J. S., and Sze, V. (2016). Eyeriss: an energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE J. Solid-State Circuits* 52, 127–138. doi: 10.1109/JSSCC.2016.7418007
- Chen, Y.-H., Krishna, T., Emer, J. S., and Sze, V. (2017). Eyeriss: an energy-efficient reconfigurable accelerator for deep convolutional neural networks. *IEEE J. Solid-State Circuits* 52, 127–138. doi: 10.1109/JSSC.2016.2616357
- Chi, P., Li, S., Xu, C., Zhang, T., Zhao, J., Liu, Y., et al. (2016). “PRIME: a novel processing-in-memory architecture for neural network computation in ReRAM-based main memory,” in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)* (Seoul), 27–39. doi: 10.1109/ISCA.2016.13
- Davies, M., Srinivasa, N., Lin, T.-H., Chinya, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro*. 38, 82–99. doi: 10.1109/MM.2018.112130359
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). “Imagenet: a large-scale hierarchical image database,” in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (Miami, FL: IEEE), 248–255.
- Dennard, R. H., Gaensslen, F. H., Rideout, V. L., Bassous, E., and LeBlanc, A. R. (1974). Design of ion-implanted mosfet’s with very small physical dimensions. *IEEE J. Solid State Circuits* 9, 256–268.
- Diehl, P. U., and Cook, M. (2015). Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* 9:99. doi: 10.3389/fncom.2015.00099
- Grossberg, S. (1987). Competitive learning: from interactive activation to adaptive resonance. *Cogn. Sci.* 11, 23–63.
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. New York, NY: Wiley.
- Hsu, J. (2014). Ibm’s new brain [news]. *IEEE Spectr.* 51, 17–19. doi: 10.1109/MSPEC.2014.6905473
- Kim, C., Kang, S., Shin, D., Choi, S., Kim, Y., and Yoo, H.-J. (2019). “A 2.1 tflops/w mobile deep rl accelerator with transposable pe array and experience compression,” in *2019 IEEE International Solid-State Circuits Conference-ISSCC* (San Francisco, CA: IEEE), 136–138.
- Krizhevsky, A., and Hinton, G. (2009). *Learning Multiple Layers of Features From Tiny Images*. Technical report, Citeseer.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems* (Tahoe, CA: IEEE), 1097–1105.
- Kwon, H., Pellauer, M., and Krishna, T. (2018). Maestro: an open-source infrastructure for modeling dataflows within deep learning accelerators. *arXiv preprint arXiv:1805.02566*.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539
- LeCun, Y., Boser, B. E., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W. E., et al. (1990). “Handwritten digit recognition with a back-propagation network,” in *Advances in Neural Information Processing Systems* (Denver, CO: IEEE), 396–404.
- Lee, D.-H., Zhang, S., Fischer, A., and Bengio, Y. (2015). “Difference target propagation,” in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (Porto: Springer), 498–515.
- Lee, J., Kim, C., Kang, S., Shin, D., Kim, S., and Yoo, H.-J. (2018). “UNPU: a 50.6 tops/w unified deep neural network accelerator with 1b-to-16b fully-variable weight bit-precision,” in *2018 IEEE International Solid-State Circuits Conference-ISSCC* (San Francisco, CA: IEEE), 218–220.
- Lee, J., Lee, J., Han, D., Lee, J., Park, G., and Yoo, H.-J. (2019). “7.7 Inpu: A 25.3 tflops/w sparse deep-neural-network learning processor with fine-grained mixed precision of fp8-fp16,” in *2019 IEEE International Solid-State Circuits Conference-ISSCC* (San Francisco, CA: IEEE), 142–144.
- Li, C., Belkin, D., Li, Y., Yan, P., Hu, M., Ge, N., et al. (2018). Efficient and self-adaptive *in-situ* learning in multilayer memristor neural networks. *Nat. Commun.* 9:2385. doi: 10.1038/s41467-018-04484-2
- Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. (2016). Random synaptic feedback weights support error backpropagation for deep learning. *Nat. Commun.* 7:13276. doi: 10.1038/ncomms13276
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Moore, G. E. (1965). *Cramming More Components Onto Integrated Circuits*. New York, NY: McGraw-Hill.
- Mostafa, H., Ramesh, V., and Cauwenberghs, G. (2018). Deep supervised learning using local errors. *Front. Neurosci.* 12:608. doi: 10.3389/fnins.2018.00608
- Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., and Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. *J. Big Data* 2:1. doi: 10.1186/s40537-014-0007-7
- Neftci, E. O. (2018). Data and power efficient intelligence with neuromorphic learning machines. *iScience* 5:52. doi: 10.1016/j.isci.2018.06.010
- Neftci, E. O., Augustine, C., Paul, S., and Detorakis, G. (2017). Event-driven random back-propagation: enabling neuromorphic deep learning machines. *Front. Neurosci.* 11:324. doi: 10.3389/fnins.2017.00324
- Nøkland, A. (2016). “Direct feedback alignment provides learning in deep neural networks,” in *Advances in Neural Information Processing Systems* (Barcelona: IEEE), 1037–1045.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *J. Math. Biol.* 15, 267–273.
- O’Reilly, R. C., and Munakata, Y. (2000). *Computational Explorations in Cognitive Neuroscience: Understanding the Mind by Simulating the Brain*. Cambridge, MA: MIT Press.
- Pan, S. J., and Yang, Q. (2010). A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22, 1345–1359. doi: 10.1109/TKDE.2009.191
- Shin, D., Lee, J., Lee, J., and Yoo, H.-J. (2017). “14.2 dnpu: an 8.1 tops/w reconfigurable cnn-rnn processor for general-purpose deep neural networks,” in *2017 IEEE International Solid-State Circuits Conference (ISSCC)* (San Francisco, CA: IEEE), 240–241.
- Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2019 Crafton, Parihar, Gebhardt and Raychowdhury. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.