

LLS: Local Learning Rule for Deep Neural Networks Inspired by Neural Activity Synchronization

Marco P. E. Apolinario, Arani Roy, Kaushik Roy
 School of Electrical and Computer Engineering, Purdue University
 West Lafayette, IN, USA
 {mapolina, roy173, kaushik}@purdue.edu

Abstract

Training deep neural networks (DNNs) using traditional backpropagation (BP) presents challenges in terms of computational complexity and energy consumption, particularly for on-device learning where computational resources are limited. Various alternatives to BP, including random feedback alignment, forward-forward, and local classifiers, have been explored to address these challenges. These methods have their advantages, but they can encounter difficulties when dealing with intricate visual tasks or demand considerable computational resources. In this paper, we propose a novel Local Learning rule inspired by neural activity Synchronization phenomena (LLS) observed in the brain. LLS utilizes fixed periodic basis vectors to synchronize neuron activity within each layer, enabling efficient training without the need for additional trainable parameters. We demonstrate the effectiveness of LLS and its variations, LLS-M and LLS-MxM, on multiple image classification datasets, achieving accuracy comparable to BP with reduced computational complexity and minimal additional parameters. Specifically, LLS achieves comparable performance with up to $300\times$ fewer multiply-accumulate (MAC) operations and half the memory requirements of BP. Furthermore, the performance of LLS on the Visual Wake Word (VWW) dataset highlights its suitability for on-device learning tasks, making it a promising candidate for edge hardware implementations.¹

1. Introduction

Currently, stochastic gradient-based optimization schemes serve as the default method for training deep neural network (DNN) models. These schemes leverage the backpropagation (BP) algorithm, enabling the computation of gradients of the loss function with respect to the trainable parameters (weights) in the hidden layers. However, BP

is associated with high time and memory complexities, leading to significant energy consumption. For instance, in a model with L layers and n neurons per layer, BP exhibits time and memory complexities of $O(Ln^2)$ and $O(Ln)$, respectively. While suitable for offline training in environments with ample computational resources (such as the cloud), these computational demands render BP inefficient for on-device learning on low-power edge devices, where computation resources are severely constrained [1, 24, 30]. Studies such as [1] and [24] highlight the large energy consumption associated with extensive external memory accesses and gradient computations in BP. Consequently, there is a need for hardware-friendly algorithms to facilitate efficient on-device learning on low-power edge devices.

With this consideration in mind, numerous works have explored alternatives to backpropagation (BP), trying to eliminate the need of computationally expensive gradient calculations associated with BP. Methods like feedback alignment (FA) and its variant, direct feedback alignment (DFA), utilize random matrices to propagate error signals or directly project errors to each layer, offering some reduction in dependency across layers but still requiring similar memory demands [5, 18, 27]. An alternative to this approach is proposed by [8], which uses random matrices to project targets instead of errors, thereby enabling each layer to be updated independently. Although promising, these methods do not scale well for deep neural networks (DNNs). In contrast, [23] proposes a local learning rule that matches BP performance in large models at the cost of significantly increasing the number of trainable parameters and computational complexity. Recent research works have attempted to replace BP's backward pass with an additional forward pass, aiming to enhance biological plausibility, though they suffer from slow convergence and have not yet proven effective for deep networks [6, 10]. Additionally, [13] proposes a biologically inspired method using a soft winner-take-all mechanism to facilitate unsupervised learning in simpler DNN models. In contrast, [2, 22] and [28] proposed to use auxiliary networks as local classifiers. These

¹Code available at <https://github.com/mapolinario94/LLS-DNN>

methods [2, 22, 28] avoid using end-to-end BP by breaking the problem into smaller pieces and generating error signals with the aid of such local classifiers per layer or group of layers. Since these methods necessitate additional layers to generate the learning signal, we categorize them as hybrids between local learning and BP.

The aforementioned learning methods often struggle to scale to complex vision tasks without high computational costs [6, 8, 13, 18, 23, 27]. Hybrid approaches using local classifiers [2, 22] offer a better balance for on-device learning but at the cost of increasing trainable parameters, thus increasing memory and energy demands. To address this, we propose a Local Learning rule inspired by brain-like neural activity Synchronization (LLS). This rule bypasses intensive gradient calculations of BP and scales to complex vision tasks and deep networks.

Neuronal activity synchronization in the brain reflects the correlation of brain signals. Studies in [3, 9, 11, 14, 21], have demonstrated that neuronal ensembles in the brain synchronize their activity during cognitive learning processes or in response to visual stimuli. Inspired from this biological process, LLS utilizes fixed periodic basis vectors to synchronize neuron activity within same layers of the model. Our experiments show that simple periodic functions like cosine and square enable effective learning in complex image classification tasks. These functions are computationally lightweight, allowing on-the-fly generation on low-power devices without additional trainable parameters. Furthermore, we explore variations of LLS, such as LLS-M and LLS-MxM, to enhance performance on more complex tasks. LLS-M learns to modulate the amplitude of the fixed basis, while LLS-MxM learns to construct an improved basis through a linear combination of the fixed basis. Both variants require minimal trainable parameters, on the order of $O(C)$ and $O(C^2)$, where C represents the number of classes. Evaluation on public image classification datasets, including CIFAR10, CIFAR100, IMAGENETTE, TinyIMAGENET, and Visual Wake Words (VWW), demonstrates that our method achieves high accuracy comparable to BP, with significant reductions in MAC operations, memory usage, and minimal additional parameters. Notably, our method's performance on the VWW dataset underscores its suitability for on-device learning hardware implementations.

The main contributions of the paper are as follows:

- A novel local learning rule that utilizes fixed periodic basis vectors to synchronize neural activity per layer, achieving high accuracy with reduced MAC operations, memory usage, and minimal additional trainable parameters.
- Evaluation of the effectiveness of our method on various image classification datasets, demonstrating accu-

racy comparable to BP.

- Demonstration of the suitability of our method for on-device learning tasks by evaluating its performance on the Visual Wake Word (VWW) dataset, achieving high performance with low computational complexity.

2. Background

2.1. Backpropagation (BP)

As noted earlier, the backpropagation (BP) algorithm is central to deep learning. We explore its mechanics here and introduce key notations used in this work. A neural network model can be represented as a parameterized function $F(\mathbf{x}; \mathbb{W})$, where \mathbf{x} is the input data and \mathbb{W} is the set of parameters. For an L -layer model, the parameters are $\mathbb{W} = \{(\mathbf{W}^{(l)})_{l=1}^L\}$, with $\mathbf{W}^{(l)}$ representing the weights of the l -th layer. Each layer produces an output, $\mathbf{h}^{(l)}$, obtained by applying a linear transformation over the input $\mathbf{h}^{(l-1)}$ based on the parameters $\mathbf{W}^{(l)}$, resulting in an intermediate representation $\mathbf{z}^{(l)}$, followed by a non-linear element-wise activation function $\mathbf{h}^{(l)} = f(\mathbf{z}^{(l)})$. Given a loss function \mathcal{L} and a labeled dataset $\mathbb{D} = \{(\mathbf{x}_i, \mathbf{y}_i)_{i=1}^N\}$, where \mathbf{x} are the inputs, \mathbf{y} the labels, and N the number of samples. The objective is to find the parameters \mathbb{W} that minimize the loss, i.e., $\mathbb{W} := \arg \min_{\mathbb{W}} \frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathbf{y}_i, F(\mathbf{x}_i; \mathbb{W}))$. For this purpose, the conventional approach is to use mini-batch stochastic gradient descent (SGD), which randomly samples a mini-batch of data (\mathbf{X}, \mathbf{Y}) from the dataset to estimate the gradient of the loss function. Such a learning algorithm, with a learning rate (η), has the following update rule for the parameters:

$$\mathbf{W}^{(l)} := \mathbf{W}^{(l)} - \eta \nabla_{\mathbf{W}^{(l)}} \mathcal{L} \quad (1)$$

The gradient $\nabla_{\mathbf{W}^{(l)}} \mathcal{L}$ is computed based on the BP algorithm. BP operates in two phases: the forward pass and the backward pass. During the forward pass, a mini-batch input \mathbf{X} is propagated layer by layer through the model to obtain a model prediction $\mathbf{H}^{(L)} = F(\mathbf{X}; \mathbb{W})$, and the loss $\mathcal{L}(\mathbf{Y}, \mathbf{H}^{(L)})$ is computed. In this process, all intermediate representations $\mathbf{Z}^{(l)}$ are saved. Then, in the backward pass, the chain rule is used to compute the gradients as follows:

$$\begin{aligned} \nabla_{\mathbf{W}^{(l)}} \mathcal{L} &= \frac{\partial \mathcal{L}}{\partial \mathbf{H}^{(l)}} \frac{\partial \mathbf{H}^{(l)}}{\partial \mathbf{Z}^{(l)}} \frac{\partial \mathbf{Z}^{(l)}}{\partial \mathbf{W}^{(l)}} \\ &= \frac{\partial \mathcal{L}}{\partial \mathbf{H}^{(L)}} \prod_{i=l+1}^L \frac{\partial \mathbf{H}^{(i)}}{\partial \mathbf{H}^{(i-1)}} \frac{\partial \mathbf{H}^{(l)}}{\partial \mathbf{Z}^{(l)}} \frac{\partial \mathbf{Z}^{(l)}}{\partial \mathbf{W}^{(l)}} \end{aligned} \quad (2)$$

Here, $\frac{\partial \mathcal{L}}{\partial \mathbf{H}^{(l)}}$ is the learning signal obtained by propagating errors from the last layer (L) to layer l . Additionally, $\frac{\partial \mathbf{H}^{(l)}}{\partial \mathbf{Z}^{(l)}}$ corresponds to the derivative of the activation function $f'(\mathbf{Z}^{(l)})$, and $\frac{\partial \mathbf{Z}^{(l)}}{\partial \mathbf{W}^{(l)}}$ is equivalent to the input of the l -th

layer, i.e., $\mathbf{H}^{(l-1)}$. From (2), it can be observed that while the latter two factors on the right-hand side of (2) depend only on the inputs and outputs of layer l , the learning signal depends on all successive layers. Therefore, the weight updates must be sequential (i.e., update-locking problem). Moreover, the computational and memory complexity of BP are $O(Ln^2)$ and $O(Ln)$, respectively, with n representing the average number of neurons per layer.

2.2. Local learning for DNN

The non-locality and update-locking features of BP, among others, have been argued as reasons that make BP unlikely as the learning rule used by the brain [19]. Different local learning mechanisms that may not rely on the propagation of errors using symmetric weights have been explored in many works [6, 8, 10, 13, 27]. Here, we refer to local learning as learning rules that compute weight updates ($\Delta \mathbf{W}^{(l)}$) based only on inputs ($\mathbf{H}^{(l-1)}$), outputs ($\mathbf{Z}^{(l)}$), and some other global factors. An example is the DFA method [27], which uses random feedback weights ($\mathbf{R}^{(l)}$) to produce the learning signal. In this method, $\frac{\partial \mathcal{L}}{\partial \mathbf{H}^{(l)}}$ in (2) is replaced by $\frac{\partial \mathcal{L}}{\partial \mathbf{H}^{(l)}} \mathbf{R}^{(l)}$. A similar method is proposed by [8], denoted as DRTP, which uses fixed random learning signals produced by propagating the labels instead of error. In other words, the learning signals are $\mathbf{Y} \mathbf{R}^{(l)}$. Other approaches, such as those by [6, 10], use two forward passes to produce the learning signal, or produce a learning signal based on a soft competition mechanism as proposed by [13].

2.3. Neural activity synchronization in the brain

Neural activity synchronization refers to the correlated neuronal signals across different regions of the brain. Groups of neurons that co-activate in response to sensory stimuli or during spontaneous activity are often referred to as ensembles. These ensembles play a crucial role in various cognitive functions, including the processing of visual stimuli in the cortex [21], memory formation [11], and behavior regulation [3]. In addition to these roles, modulations in oscillatory neuronal activity are commonly observed when humans engage in cognitive tasks. For instance, as highlighted by [9], the complex, high-dimensional dynamics of neuronal activity can collapse into low-dimensional oscillatory modes, which in turn facilitates memory enhancement and learning. This synchronization not only simplifies the representation of neuronal dynamics but also captures both linear and non-linear aspects of neuronal interactions. Drawing inspiration from these biological processes, we propose a local learning rule (LLS) that employs fixed periodic vectors for each class to synchronize neural activity within the same layer of a neural network. This approach is intended to enhance the efficiency of learning in artificial systems. By using periodic vectors, the LLS encourages groups of neurons, distributed periodically within

the same layer, to exhibit high activity in response to specific visual stimuli (such as images of a particular class). This design is inspired in the concept of neuronal ensembles within artificial neural networks.

3. LLS: Local Learning Rule inspired by Neural Activity Synchronization

LLS aims to synchronize neural activity within the same layer while minimizing computational complexity and additional trainable parameters. We emphasize three core aspects of LLS: (1) locality, (2) update-unlocking, and (3) minimal parameter requirements.

First, LLS operates locally within each layer, updating synaptic connections ($\mathbf{W}^{(l)}$) based on local inputs ($\mathbf{H}^{(l-1)}$), outputs ($\mathbf{H}^{(l)}$), and generated learning signals. The locally generated learning signals are obtained by projecting $\mathbf{H}^{(l)}$ onto a set of fixed periodic basis vectors $\mathbf{B}^{(l)}$, which align with specific classes to optimize layer performance. Local operation reduces computational overhead of computing the weight gradients.

Second, LLS's update-unlocking feature is a by-product of locality and enables independent weight updates per layer, eliminating the need to save the output activations of all the layers in the model during training. This results in a memory complexity of $O(n_{max})$, where n_{max} represents the maximum number of neurons in a layer. Unlike methods employing auxiliary local classifiers, LLS requires no additional trainable parameters, utilizing fixed periodic vectors for alignment. However, for tasks with numerous classes, relying solely on fixed vectors may present challenges, as discussed in Section 4. To address these limitations, we also propose LLS-M and LLS-MxM as variations of LLS. LLS-M enables learning of optimal modulation for fixed basis vectors, while LLS-MxM learns to form a superior basis via a linear combination of fixed vectors. Both variations entail minimal additional trainable parameters on the order of $O(C)$ and $O(C^2)$, respectively, where C denotes the number of classes in a task.

3.1. Technical details

The hidden layers are trained based on the alignment of their output activations ($\mathbf{H}^{(l)}$) with predefined set of fixed basis vectors ($\mathbf{B}^{(l)}$), as shown in Fig. 1. Alignment is measured as the inner product of a layer's output activations and the basis. To encourage synchronicity in neural responses among neurons, the fixed basis vectors are constructed using periodic functions $g(f, t) = g(f, t + 1/f)$, where f represents spatial frequency.

For a classification problem with C classes, each class c has its own vector $\mathbf{B}_{c,:}^{(l)} = g(f_c, \mathbf{t}^{(l)})$ where $\mathbf{t}^{(l)} = [1, 2, 3, \dots, T^{(l)}]$, $T^{(l)}$ is the length of l -th layer's output ($\mathbf{H}_{i,:}^{(l)}$) and f_c is a fixed frequency for class c . Note that

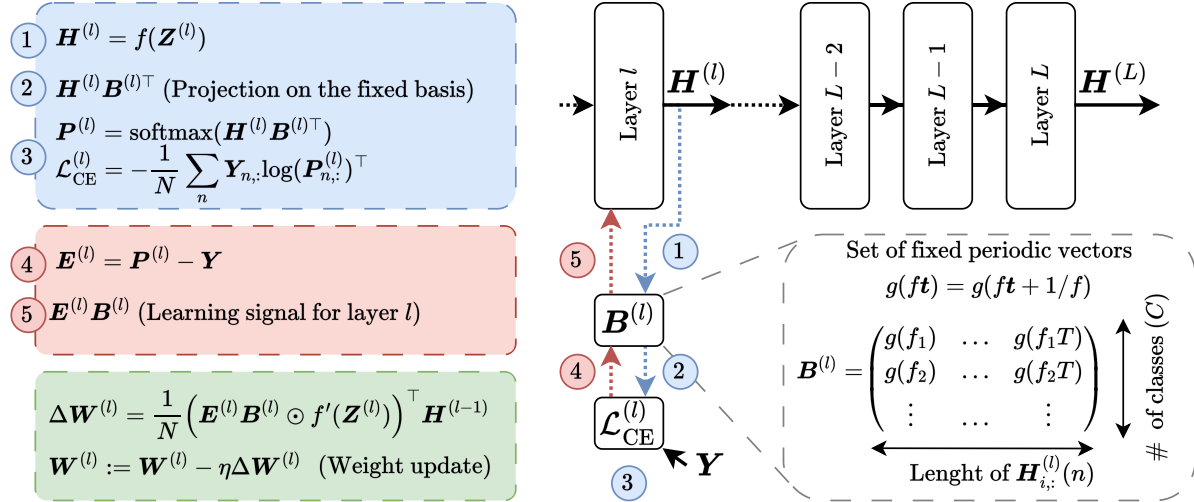


Figure 1. Overview of LLS. Weight updates for the l -th hidden layer within an L -layer neural network are derived via per-layer minimization of cross-entropy loss ($\mathcal{L}_{CE}^{(l)}$) on the projection of output activations ($\mathbf{H}^{(l)}$) over a fixed basis of periodic vectors $\mathbf{B}^{(l)}$, i.e., $\mathbf{H}^{(l)} \mathbf{B}^{(l)\top}$. This produces a local error signal as the difference between the softmax of the projection ($\mathbf{P}^{(l)}$) and the one-hot encoded labels \mathbf{Y} . Subsequently, this error signal is multiplied with the fixed basis to generate the learning signal. Weight updates are then determined by multiplying the locally generated learning signal with the layer's inputs and outputs. Consequently, LLS enables independent layer updates based on local information, resulting in low time and memory complexities of $O(LCn)$ and $O(n_{max})$, respectively. It is noteworthy that the fixed basis $\mathbf{B}^{(l)}$ comprises C vectors, where C represents the number of classes for the classification task. Furthermore, the fixed basis vectors are constructed using periodic functions $g(f_c, t) = g(f_c, t + 1/f_c)$, where f_c denotes the spatial frequency associated with class c .

these basis vectors have the same frequencies for all layers but with different lengths, also $\mathbf{B}_{c,:}^{(l)}$ is the c -th row of $\mathbf{B}^{(l)} \in \mathbb{R}^{C \times T^{(l)}}$. The weight updates can be derived as a per-layer minimization of cross-entropy loss ($\mathcal{L}^{(l)}$) on the projection of the activations over the fixed basis ($\mathbf{H}^{(l)} \mathbf{B}^{(l)\top}$), as illustrated in Fig. 1. Specifically, the per-layer cross-entropy loss is described as follows:

$$\begin{aligned} \mathcal{L}^{(l)} &= -\frac{1}{N} \sum_{n=1}^N \mathbf{Y}_{n,:} \log(\mathbf{P}_{n,:}^\top) \\ &= -\frac{1}{N} \sum_{n=1}^N \log \frac{\exp(\mathbf{H}_{n,:}^{(l)} \mathbf{B}_{c_n^*}^{(l)\top})}{\sum_{c=1}^C \exp(\mathbf{H}_{n,:}^{(l)} \mathbf{B}_{c,:}^{(l)\top})} \end{aligned} \quad (3)$$

Here, N is the number of samples in the mini-batch, c_n^* is the class index for the n -th sample in the mini-batch, and $\mathbf{P}_{n,:}^{(l)}$ is probability vector obtained of applying the softmax function over the projection vector $\mathbf{H}_{n,:}^{(l)} \mathbf{B}^{(l)\top}$. Solving the per-layer minimization problem, $\min_{\mathbf{W}^{(l)}} \mathcal{L}^{(l)}(\mathbf{H}^{(l)}, \mathbf{Y})$, results in the following expression for weight updates on the l -th layer:

$$\begin{aligned} \Delta \mathbf{W}^{(l)} &= \frac{1}{N} \left((\mathbf{P}^{(l)} - \mathbf{Y}) \mathbf{B}^{(l)} \odot f'(\mathbf{Z}^{(l)}) \right)^\top \mathbf{H}^{(l-1)} \\ &= \frac{1}{N} \left(\mathbf{E}^{(l)} \mathbf{B}^{(l)} \odot f'(\mathbf{Z}^{(l)}) \right)^\top \mathbf{H}^{(l-1)} \end{aligned} \quad (4)$$

Here, \odot is the element-wise product. From Equation (4), it is evident that the weight updates for each layer l depend solely on the local variables of that layer, including its inputs, outputs, and the set of fixed basis vectors. Consequently, all layers can be updated independently of the rest of the model. These independent updates are the reason why the memory complexity of LLS depends only on the largest layer (the layer with the highest number of neurons), in contrast with end-to-end training methods that require memory proportional to the number of neurons in the entire model. Moreover, since LLS's learning signals are generated locally, the time complexity to generate them for all the layers is proportional to the number of neurons per layer and the number of classes, that is $O(LCn)$.

The selection of frequencies (f_c) for each class is done to maintain sufficient distance among frequencies of different classes to avoid interference. The range of available frequencies is defined by the length of $\mathbf{H}_{i,:}^{(l)}$. Hence, frequencies can be assigned to be equally distributed in that range or randomly as long as they do not overlap. In practice, we reduce the dimensions of $\mathbf{H}_{i,:}^{(l)}$ of convolutional layers by using average pooling, followed by flattening before projecting it onto the basis $\mathbf{B}^{(l)}$. This helps both in faster convergence of the method and in reducing the number of MAC operations.

3.2. Variations of LLS

So far, we have discussed LLS based on utilizing a basis of periodic vectors $B_{i,:}^{(l)}$, generated from a fixed periodic function $g(\cdot)$. However, such a base may not always be optimal for a given task. For instance, the amplitude of the vectors could be too large making it difficult for the algorithm to converge. Additionally, in problems with a large number of classes, the restriction to fixed periodic vectors may impede the model’s ability to learn semantics in the data, such as grouping similar classes.

To address these concerns, we propose two variations of LLS: LLS-M for learning the appropriate modulation of the fixed basis ($B^{(l)}$), and LLS-MxM for learning to construct a new basis as a linear combination of the original fixed basis.

LLS-M: In this variation, the new basis is simply a modulation of the original fixed basis, defined as $D^{(l)} = M^{(l)} B^{(l)}$, where $M^{(l)}$ is a diagonal matrix of trainable parameters with dimensions equal to the number of classes, i.e., $M^{(l)} \in \mathbb{R}^{C \times C}$. Weight updates for LLS-M follow (4), with $B^{(l)}$ replaced by $D^{(l)}$. The updates for $M^{(l)}$ are computed as follows:

$$\Delta M^{(l)} = I \odot \frac{1}{N} E^{(l)\top} (H^{(l)} B^{(l)\top}) \quad (5)$$

Here, $I \in \mathbb{R}^{C \times C}$ is the identity matrix and \odot is the element-wise product.

LLS-MxM: Here, the new basis vectors ($D^{(l)}$) are obtained as a linear combination of the original fixed periodic vectors: $D^{(l)} = M^{(l)} B^{(l)}$ where $M^{(l)} \in \mathbb{R}^{C \times C}$ with all the $C \times C$ are trainable parameters. Weight updates are obtained following (4), with the basis replaced by $D^{(l)}$. Similar to LLS-M, updates for the matrix $M^{(l)}$ are computed as follows:

$$\Delta M^{(l)} = \frac{1}{N} E^{(l)\top} (H^{(l)} B^{(l)\top}) \quad (6)$$

4. Experimental evaluation

In this section, we assess the efficacy of LLS and its variations across several image classification datasets, which include MNIST [17], FashionMNIST [29], CIFAR10 [15], CIFAR100 [15], IMAGENETTE [7], TinyIMAGENET [16], and Visual Wake Words (VWW) [4].

We primarily evaluate the proposed learning rules using three models: a 5-layer CNN (SmallConv), a VGG8 [22], and MobileNets-V1 (MBNet) [12]. Detailed descriptions of each model are provided in Supplemental Material 1.1. Additionally, information regarding hyperparameters, data pre-processing, and optimizer settings is provided in Supplemental Material 1.2.

4.1. Effect of different basis in learning

First, we compare the effect of different functions $g(\cdot)$ for generating the basis $B^{(l)}$. We consider two simple periodic functions: cosine ($g = \cos(f_c t)$) and square wave ($g = \text{sign}(\cos(f_c t))$). Both functions offer the advantage of being easily generated on-the-fly or require storage with minimal memory overhead due to their periodicity. Additionally, we investigate the scenario where $g(\cdot)$ is a pseudo-random number generator, resulting in a random fixed vector $B_{i,:}^{(l)}$. The results are evaluated on two models, SmallConv and VGG8, across four image classification datasets of increasing complexity. Each model undergoes five training iterations with different random seeds, and the results are reported in Table 1.

We observe that employing any of the three fixed vector bases with LLS yields high accuracy across all four vision tasks. Notably, for the SmallConv model, using LLS with a square basis function present the best accuracy results, followed by cosine basis. In contrast, for the VGG8 model, the random basis exhibits better performance than the periodic basis, with square still performing better than cosine. This discrepancy may be attributed to the increased complexity of per-layer feature representations in deeper models, where a random vector offers more degrees of freedom for such representations. However, it is important to note that a random vector is less hardware-friendly, as it requires specialized pseudo-random number generators, leading to energy and memory overhead, as discussed in [5]. Therefore, in the subsequent sections, we primarily focus on LLS using a square $g(\cdot)$ function (LLS_{square}).

Moreover, employing a periodic function, such as a square function, induces layer neurons to synchronize with the frequency of the basis function. This synchronization is demonstrated in Fig. 2, where the activations for different classes align with the spatial frequencies of the basis function. Here, the spectral decomposition is obtained by applying Fourier transform in the spatial dimension to both basis vectors ($B_{i,:}^{(l)}$) and layer output activations ($h^{(l)}$). As shown in Table 1, synchronization has a beneficial effect on accuracy for small models, such as SmallConv. A reason for this is that such models need to discriminate between classes by transforming inputs through only a few layers. Thus, aligning the layers’ outputs to periodic vectors might be easier than aligning random vectors.

4.2. Comparison with local learning algorithms

In this section, we compare LLS_{square} with other local learning methods that exhibit similar time and memory complexities. These methods include DFA [27], DRTP [8], and PEPITA [6]. For this comparison, we use the MNIST, CIFAR10 and CIFAR100 datasets, with results shown in Table 2.




Figure 2. Neural activity synchronization induced by learning rule $\text{LLS}_{\text{square}}$ on the VGG8 model’s 4th layer output ($\mathbf{h}^{(4)}$) for classes 0 and 1 from the IMAGENETTE dataset. The layer’s response exhibits spatial periodicity coinciding with the periodic function selected as a basis ($\mathbf{B}^{(4)}$).

Table 1. LLS’s performance comparison with different function $g(\cdot)$ to generate the basis $\mathbf{b}^{(l)}$. Test accuracy mean and std are reported over five trials.

Function $g(\cdot)$	Model	MNIST	FashionMNIST	CIFAR10	IMAGENETTE
Cosine	SmallConv	99.50 ± 0.02	89.57 ± 0.22	75.82 ± 0.39	78.03 ± 0.35
Square		99.50 ± 0.02	90.54 ± 0.23	77.79 ± 0.31	79.02 ± 0.76
Random		99.38 ± 0.03	87.30 ± 0.18	74.19 ± 0.57	71.70 ± 1.45
Cosine	VGG8	99.52 ± 0.02	93.04 ± 0.17	86.92 ± 0.27	84.85 ± 0.11
Square		99.54 ± 0.01	93.54 ± 0.06	88.64 ± 0.12	85.62 ± 0.24
Random		99.70 ± 0.02	93.77 ± 0.08	90.45 ± 0.09	87.09 ± 0.28

We observe that training the SmallConv model with DFA, DRTP or PEPITA resulted in low performance or did not converge at all. For DFA, performance improved by increasing the number of channels threefold (SmallConvL). Consequently, we used SmallConvL for reporting results with BP and LLS. However, for DRTP and PEPITA, increasing number of channels did not yield satisfactory results, and hence, we opted for reporting accuracy of each task as reported in the original papers.

As shown in Table 2, LLS demonstrates the best performance among the three local learning methods under consideration. In terms of accuracy, LLS achieves results close to BP, while maintaining significantly lower time and memory complexities compared to BP. In fact, among all the methods in Table 2, only DRTP exhibit a time and memory complexities comparable to LLS. Furthermore, it is worth noting that while DFA, DRTP, and PEPITA do not scale well for deeper models and in many cases require wide DNNs to converge [26], LLS performs well on deeper models, as demonstrated in Section 4.1.

4.3. Performance comparison on deeper models

In this section, we conduct a performance comparison of LLS and its variations on five image classification datasets: CIFAR10, CIFAR100, IMAGENETTE, TinyIMAGENET, and VWV. These datasets cover a wide range of classification tasks, including low to high-resolution images and tasks with few to multiple classes. Notably, we emphasize the experiments conducted on the VWV dataset, as it holds significance for edge vision applications and serves as a rel-

evant use case for on-device learning [4]. The comparison considers four metrics: accuracy, the number of MAC operations required to compute the learning signal, the peak memory usage, and the number of additional trainable parameters needed by each method. We compare our method against BP and the local losses method [22]. Note, local losses method employs a linear classifier per layer.

CIFAR10 and IMAGENETTE First, we examine tasks with a few number of classes and different image resolutions, such as CIFAR10 and IMAGENETTE. As depicted in Table 3, LLS achieves high accuracy, closely following BP and Local Losses. Note, that LLS achieves such high accuracy with approximately $300\times$ fewer MAC operations and half the memory usage compared to BP, and without requiring additional trainable parameters. To further narrow the accuracy gap, we explore variations of LLS, such as LLS-M and LLS-MxM. Both variations improve the accuracy to be closer to BP with almost no increase in MACs and memory usage. Note, however, the accuracy improvement comes at the cost of employing some additional trainable parameters. It is important to note that LLS-MxM still requires approximately $100\times$ fewer trainable parameters than Local Losses.

CIFAR100 and TinyIMAGENET For tasks with hundreds of classes such as CIFAR100 and TinyIMAGENET, LLS exhibits significant accuracy drop compared to BP. This is attributed to the orthogonal nature of the periodic vectors, which compels the model to represent each class

Table 2. Comparison with local learning algorithms (Test accuracy mean and std are reported)

Method	Time	Memory	Model	MNIST	CIFAR10	CIFAR100
BP (baseline)	$O(Ln^2)$	$O(Ln)$	SmallConvL	99.62 ± 0.020	87.57 ± 0.13	62.25 ± 0.29
DFA	$O(LCn)$	$O(Ln)$	SmallConvL	97.90 ± 0.17	71.53 ± 0.38	44.93 ± 0.52
DRTP	$O(LCn)$	$O(n_{max})$	[27]	98.98 ± 0.02	73.10 ± 0.50	41.00 ± 0.30
PEPITA	$O(Ln^2)$	$O(n_{max})$	[8]	98.52 ± 0.15	68.96 ± 0.45	—
LLS _{square} (Ours)	$O(LCn)$	$O(n_{max})$	[6]	98.29 ± 0.13	56.33 ± 1.35	27.56 ± 0.60
LLS _{square} (Ours)	$O(LCn)$	$O(n_{max})$	SmallConvL	99.57 ± 0.03	84.10 ± 0.27	55.32 ± 0.38

Table 3. Performance comparison on image classification datasets. Accuracy mean and std are reported over five trials, the additional params refers to additional trainable parameters, and #MAC is estimated for the number of ops required to generate the learning signal ($\frac{\partial \mathcal{L}}{\partial \mathbf{h}^{(l)}}$).

Method	Model	Accuracy (mean±std)	# MAC ¹ ($\times 10^6$)	Memory ¹ (MB)	Additional params
CIFAR10					
BP	VGG8	94.12 ± 0.12	719.33	1082	—
Local Losses	VGG8	91.93 ± 0.07	2.56	576	1.02×10^5
LLS _{square} (Ours)	VGG8	88.64 ± 0.12	2.46	574	0
LLS-M _{square} (Ours)	VGG8	90.43 ± 0.24	2.46	574	70
LLS-MxM _{square} (Ours)	VGG8	90.89 ± 0.09	2.46	574	700
IMAGENETTE					
BP	VGG8	90.92 ± 0.27	11477.81	15858	—
Local Losses	VGG8	88.06 ± 0.12	36.48	7319	1.02×10^5
LLS _{square} (Ours)	VGG8	85.62 ± 0.24	36.38	7318	0
LLS-M _{square} (Ours)	VGG8	86.60 ± 0.37	36.38	7318	70
LLS-MxM _{square} (Ours)	VGG8	87.29 ± 0.29	36.38	7319	700
CIFAR100					
BP	VGG8	73.69 ± 0.39	719.40	1083	—
Local Losses	VGG8	69.26 ± 0.36	5.33	598	1.02×10^6
LLS _{square} (Ours)	VGG8	58.84 ± 0.33	4.30	577	0
LLS-M _{square} (Ours)	VGG8	62.55 ± 0.24	4.31	577	700
LLS-MxM _{square} (Ours)	VGG8	68.81 ± 0.19	4.51	578	0.70×10^5
TinyIMAGENET					
BP	VGG8	61.10 ± 0.25	2871.20	4048	—
Local Losses	VGG8	54.00 ± 0.11	15.18	1971	2.04×10^6
LLS _{square} (Ours)	VGG8	35.99 ± 0.38	13.13	1928	0
LLS-M _{square} (Ours)	VGG8	41.89 ± 0.20	13.14	1928	1400
LLS-MxM _{square} (Ours)	VGG8	51.41 ± 0.48	13.97	1932	0.28×10^6
Visual Wake Words (VWW)					
BP	MBNet	88.49 ± 0.28	181.83	3036	—
Local Losses	MBNet	82.49 ± 0.17	178.28	730	0.28×10^5
LLS _{square} (Ours)	MBNet	81.91 ± 0.16	178.23	729	0
LLS-M _{square} (Ours)	MBNet	82.71 ± 0.42	178.23	729	28
LLS-MxM _{square} (Ours)	MBNet	83.66 ± 0.21	178.25	729	560

¹: # MAC is estimated for a batch size of 1 and GPU memory is measured for a batch size of 128.

orthogonally, even when semantically some classes have similar representations. Essentially, the basic form of LLS may not effectively capture semantics. Additionally, increasing the number of classes also increases the number of frequencies used to generate the fixed basis, leading to overlapping frequencies. We applied LLS-M learning for the above problems. LLS-M improves the accuracy, but only marginally, as the problems associated with orthogonality of the bases could not be completely solved by simply modulating the bases. In contrast, LLS-MxM learns to

create a better basis as a linear combination of the original basis, offering a larger improvement and bringing the accuracy closer to BP, as show in Table 3. To further verify that LLS-MxM can actually learn semantics, we analyze the learned linear combination matrix ($\mathbf{M}^{(l)}$) used to create the new basis. For instance, for a VGG8 model trained on CIFAR100, we project the $\mathbf{M}^{(l)}$ matrix into a 2D space using t-SNE [20] using the twenty super-classes provided in the dataset as ground truth. The results of this projection are illustrated in Fig. 3, wherein vectors representing similar




Figure 3. Projection of the linear combination matrix $M^{(l)}$ of the fixed basis $B^{(l)}$ using t-SNE. $M^{(l)}$ is obtained after training a VGG8 model with LLS-MxM on CIFAR100. The results provide evidence that our learning rule can learn better basis (as a linear combination of a fixed basis) and can encode semantics within it. Points are colored using the twenty super-class labels provided in CIFAR100.



Figure 4. Visual explanations, obtained with the Grad-CAM method, for predictions of the MBNet model trained with LLS-MxM on the VWW dataset. It can be observed that our method allows the model to learn high level image features to discern about the presence of a person or not in an image.

classes are grouped together. The accuracy improvements shown in Table 3 and the clustering of similar classes illustrated in Fig. 3 demonstrate the ability of LLS-MxM to encode semantic knowledge in the formation of the new basis. Furthermore, it is worth noting that LLS-MxM requires approximately $200\times$ fewer MACs and half memory compared to BP, and approximately $10\times$ fewer trainable parameters than Local Losses.

Visual Wake Words (VWW) Since our learning rule targets on-device learning scenarios, we tested the method on the VWW dataset using a MobileNetsV1 model. Note, the task and the model are suitable for on-device learning. The results are shown in Table 3. For this task, LLS-M and LLS-MxM outperforms the Local Losses method in all metrics (accuracy, MACs, memory, and trainable parameters). Compared to BP, LLS, LLS-M and LLS-MxM show competitive accuracy with fewer MACs and $4\times$ lower memory usage. Moreover, to understand the model’s learning ability, we used the Grad-CAM method [25] to obtain visual expla-

nations of the parts of the image most relevant for a particular prediction. As shown in Fig. 4, the MBNet model trained with LLS-MxM successfully learns high-level image features indicative of the presence of people in a given frame. This provides evidence that our method allows the model to learn complex representations.

5. Conclusions

In this work, we introduced a novel local learning rule, LLS, inspired by the synchronization of neural activity observed in biological systems, which is associated with memory formation and cognitive learning. LLS utilizes fixed periodic basis vectors to synchronize the activity of neurons within the same layer. Moreover, the deliberate choice of simple periodic functions, such as cosine and square functions, enables the generation of such basis easily and on-the-fly on low-power devices without imposing significant hardware overhead. Experimental validation demonstrates that LLS and its variations (LLS-M and LLS-MxM) achieve high accuracy comparable to BP across various image classification datasets, including CIFAR10, CIFAR100, IMAGENETTE, TinyIMAGENET, and VWW. Remarkably, this high accuracy is attained with significantly fewer MAC operations, reduced memory usage, and a minimal number of additional trainable parameters. Furthermore, employing the Grad-CAM method for visual explanations reveals that LLS and its variants can capture high-level information relevant to predictions. In summary, the demonstrated high accuracy and efficiency of LLS make it well-suited for on-device learning applications, particularly in scenarios where computational resources are severely constrained.

Acknowledgments

This work was supported in part by the Center for Co-design of Cognitive Systems (CoCoSys), one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program, and in part by the Department of Energy (DoE).

References

- [1] Aayush Ankit, Izzat El Hajj, Sai Rahul Chalamalasetti, Sapan Agarwal, Matthew Marinella, Martin Foltin, John Paul Strachan, Dejan Milojicic, Wen Mei Hwu, and Kaushik Roy. PANTHER: A Programmable Architecture for Neural Network Training Harnessing Energy-Efficient ReRAM. *IEEE Transactions on Computers*, 69(8):1128–1142, 8 2020. [1](#)
- [2] Eugene Belilovsky, Michael Eickenberg, and Edouard Oyallon. Greedy Layerwise Learning Can Scale to ImageNet. In *International Conference on Machine Learning*, 2018. [1](#), [2](#)
- [3] Luis Carrillo-Reid, Shuting Han, Weijian Yang, Alejandro Akrouh, and Rafael Yuste. Controlling Visually Guided Behavior by Holographic Recalling of Cortical Ensembles. *Cell*, 178(2):447–457, 7 2019. [2](#), [3](#)
- [4] Aakanksha Chowdhery, Pete Warden, Jonathon Shlens, Andrew Howard, and Rocky Rhodes. Visual Wake Words Dataset. *arXiv: 1906.05721*, 6 2019. [5](#), [6](#)
- [5] Brian Crafton, Abhinav Parihar, Evan Gebhardt, and Arijit Raychowdhury. Direct feedback alignment with sparse connections for local learning. *Frontiers in Neuroscience*, 13(MAY), 2019. [1](#), [5](#)
- [6] Giorgia Dellaferriera, Gabriel Kreiman, and Gabriel Kreiman. Error-driven Input Modulation: Solving the Credit Assignment Problem without a Backward Pass. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, pages 4937–4955. PMLR, 7 2022. [1](#), [2](#), [3](#), [5](#), [7](#)
- [7] fast.ai. fastai/imagenette: A smaller subset of 10 easily classified classes from Imagenet, and a little more French, 2021. [5](#)
- [8] Charlotte Frenkel, Martin Lefebvre, and David Bol. Learning Without Feedback: Fixed Random Learning Signals Allow for Feedforward Training of Deep Neural Networks. *Frontiers in Neuroscience*, 15:629892, 2 2021. [1](#), [2](#), [3](#), [5](#), [7](#)
- [9] Ramon Guevara Erra, Jose L Perez Velazquez, and Michael Rosenblum. Neural synchronization from the perspective of non-linear dynamics. *Frontiers in computational neuroscience*, 11:98, 2017. [2](#), [3](#)
- [10] Geoffrey Hinton. The Forward-Forward Algorithm: Some Preliminary Investigations. Technical report, 2022. [1](#), [3](#)
- [11] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79(8):2554, 1982. [2](#), [3](#)
- [12] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *arXiv preprint arXiv:1704.04861*, 4 2017. [5](#)
- [13] Adrien Journé, Hector Garcia Rodriguez, Qinghai Guo, and Timoleon Moraitis. Hebbian Deep Learning Without Feedback. In *2023 International Conference on Learning Representations*, 2023. [1](#), [2](#), [3](#)
- [14] Michael J Jutras and Elizabeth A Buffalo. Synchronous neural activity and memory formation. *Current opinion in neurobiology*, 20(2):150–155, 2010. [2](#)
- [15] Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. 2009. [5](#)
- [16] Ya Le and Xuan S Yang. Tiny ImageNet Visual Recognition Challenge. 2015. [5](#)
- [17] Yann LeCun, Corinna Cortes, and C J Burges. MNIST handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010. [5](#)
- [18] Timothy P. Lillicrap, Daniel Cownden, Douglas B. Tweed, and Colin J. Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature Communications*, 7, 11 2016. [1](#), [2](#)
- [19] Timothy P. Lillicrap, Adam Santoro, Luke Marris, Colin J. Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, 21(6):335–346, 6 2020. [3](#)
- [20] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008. [7](#)
- [21] Jae Eun Kang Miller, Inbal Ayzenshtat, Luis Carrillo-Reid, and Rafael Yuste. Visual stimuli recruit intrinsically generated cortical ensembles. *Proceedings of the National Academy of Sciences of the United States of America*, 111(38):E4053–E4061, 9 2014. [2](#), [3](#)
- [22] Arild Nøkland and Lars H Eidnes. Training Neural Networks with Local Error Signals. In *Proceedings of the 36th International Conference on Machine Learning*, 2019. [1](#), [2](#), [5](#), [6](#)
- [23] Alexander G. Ororbia, Ankur Mali, Daniel Kifer, and C. Lee Giles. Backpropagation-Free Deep Learning with Recursive Local Representation Alignment. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(8):9327–9335, 6 2023. [1](#), [2](#)
- [24] Xiaochen Peng, Shanshi Huang, Hongwu Jiang, Anni Lu, and Shimeng Yu. DNN+NeuroSim V2.0: An End-to-End Benchmarking Framework for Compute-in-Memory Accelerators for On-Chip Training. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 40(11):2306–2319, 11 2021. [1](#)
- [25] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *Proceedings of the IEEE International Conference on Computer Vision*, 2017-October:618–626, 12 2017. [8](#)
- [26] Ganlin Song, Ruitu Xu, and John Lafferty. Convergence and Alignment of Gradient Descent with Random Backpropagation Weights. In *35th Conference on Neural Information Processing Systems (NeurIPS 2021)*, 2021. [6](#)
- [27] Arild Nøkland Trondheim. Direct Feedback Alignment Provides Learning in Deep Neural Networks. *Advances in Neural Information Processing Systems*, 29, 2016. [1](#), [2](#), [3](#), [5](#), [7](#)
- [28] Yulin Wang, Zanlin Ni, Shiji Song, Le Yang, and Gao Huang. Revisiting Locally Supervised Learning: an Alternative to

- End-to-end Training. In *International Conference on Learning Representations*, 2021. [1](#), [2](#)
- [29] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. [5](#)
- [30] Qingtian Zhang, Huaqiang Wu, Peng Yao, Wenqiang Zhang, Bin Gao, Ning Deng, and He Qian. Sign backpropagation: An on-chip learning algorithm for analog RRAM neuromorphic computing systems. *Neural Networks*, 108:217–223, 12 2018. [1](#)