

FastHebb: Scaling Hebbian Training of Deep Neural Networks to ImageNet Level [★]

Gabriele Lagani^{1,2}, Claudio Gennaro², Hannes Fassold³, and Giuseppe Amato²

¹ Dept. of Computer Science, University of Pisa, 56127, Pisa, Italy

`gabriele.lagani@phd.unipi.it`

² ISTI-CNR, 56124, Pisa, Italy

`{gabriele.lagani,giuseppe.amato,claudio.gennaro}@isti.cnr.it`

³ Joanneum Research, 8010, Graz, Austria

`hannes.fassold@joanneum.at`

Abstract. Learning algorithms for Deep Neural Networks are typically based on supervised end-to-end Stochastic Gradient Descent (SGD) training with error backpropagation (backprop). Backprop algorithms require a large number of labelled training samples to achieve high performance. However, in many realistic applications, even if there is plenty of image samples, very few of them are labelled, and semi-supervised sample-efficient training strategies have to be used. Hebbian learning represents a possible approach towards sample efficient training; however, in current solutions, it does not scale well to large datasets. In this paper, we present *FastHebb*, an efficient and scalable solution for Hebbian learning which achieves higher efficiency by 1) merging together update computation and aggregation over a batch of inputs, and 2) leveraging efficient matrix multiplication algorithms on GPU. We validate our approach on different computer vision benchmarks, in a semi-supervised learning scenario. FastHebb outperforms previous solutions by up to 50 times in terms of training speed, and notably, for the first time, we are able to bring Hebbian algorithms to ImageNet scale.

Keywords: Hebbian Learning · Deep Learning · Neural Networks · Semi-Supervised · Sample Efficiency · Content-Based Image Retrieval

1 Introduction

In the past few years, Deep Neural Networks (DNNs) have emerged as a powerful technology in the domain of computer vision [10, 19]. DNNs started gaining popularity also in the domain of large scale multimedia Content-Based Image Retrieval (CBIR), replacing handcrafted feature extractors [2, 36] and using activations of internal layers as feature vectors for similarity search. Learning algorithms for DNNs are typically based on supervised end-to-end Stochastic Gradient Descent (SGD) training with error backpropagation (*backprop*). This approach is considered biologically implausible by neuroscientists [32]. Instead,

[★] This work was partially supported by the H2020 project AI4Media (GA 951911).

they propose *Hebbian* learning as a biological alternative to model synaptic plasticity [9].

Moreover, backprop-based algorithms need a large number of labeled training samples in order to achieve high results, which are expensive to gather, as opposed to unlabeled samples. Therefore, researchers started to investigate *semi-supervised* learning strategies, which aim to exploit large amounts of unlabeled data, in addition to the fewer labeled data, for sample efficient learning [5, 28]. In this context, a possible direction that has been proposed is to perform an unsupervised pre-training stage on all the available samples, which is then followed by a supervised fine-tuning stage on the few labeled samples only [16, 38].

In recent work, Hebbian learning has begun to gain attention from the computer science community as an effective method for unsupervised pre-training, since Hebbian algorithms do not require supervision, achieving promising results in scenarios with scarce labeled data [22, 24]. However, current solutions for Hebbian training (such as [3, 20, 23, 25, 26, 35]) are still limited in terms of computational efficiency, making it difficult to scale to large datasets such as ImageNet [6].

In order to address this issue, we present *FastHebb*, a novel solution for Hebbian training that achieves enhanced efficiency by leveraging two observations. First, when a mini-batch of inputs has to be processed, the weight update corresponding to each input is first computed, and then the various updates are aggregated over the mini-batch; however, update computation and aggregation can be merged together with a significant speedup. Second, Hebbian learning rules can be reformulated in terms of matrix multiplications, which enables to exploit efficient matrix multiplication algorithms on GPU.

We validate our method on various computer vision benchmarks. Since Hebbian algorithms are unsupervised, we consider a semi-supervised training scenario, in which Hebbian learning is used to perform unsupervised network pre-training, followed by fine-tuning with traditional backprop-based supervised learning. We also consider sample efficiency scenarios, in which we assume that only a small fraction of the training data is labeled, in order to study the effectiveness of Hebbian pre-training in scenarios with scarce data. In order to make comparisons with backprop-based methods, we consider Variational Auto-Encoder (VAE) [15, 16] pre-training as a baseline for comparisons. We show that our approach achieves comparable results, but with a significant speed-up both in terms of number of epochs, as well as total training time. In particular, our method achieves up to 50x speed-up w.r.t. previous Hebbian learning solutions, allowing to scale up our experiments to ImageNet level. To the best of our knowledge, this is the first time that Hebbian algorithms are applied at such scale.

In summary, our contribution is twofold:

1. We propose a novel efficient solution to Hebbian learning algorithms, with code available online ⁴;
2. We performed extensive experimental evaluation of the performance of our solution on various computer vision benchmarks. In particular, for the first

⁴ github.com/GabrieleLagani/HebbianLearning/tree/fasthebb

time (to the best of our knowledge) results of Hebbian algorithms on ImageNet are provided.

The remainder of this paper is structured as follows: Section 2 introduces some background and related work on Hebbian learning; Section 3 presents our FastHebb method; Section 4 provides the details of our experimental setup; Section 5 presents the results of our experiments; Finally, Section 6 outlines some concluding remarks and hints for future work.

2 Background and related work

In this section, we illustrate some of the Hebbian learning rules from literature that recently provided promising results, and we describe some related work focusing on the application of such rules on computer vision tasks, in particular in semi-supervised training scenarios. Since a thorough explanation of the Hebbian rules would be outside the scope of this paper, here we just give the update equations of interest, referring the interested reader to the vast literature on the topic [7, 9, 25–27].

Let us start by considering a neuron, identified by an index i , with weight vector \mathbf{w}_i , which receives as input a vector \mathbf{x} , and produces a corresponding output y_i . One of the Hebbian approaches that we focus on is the *soft Winner-Takes-All* (SWTA) competitive learning rule [8, 25, 29], which can be expressed as follows:

$$\Delta \mathbf{w}_i = \eta r_i (\mathbf{x} - \mathbf{w}_i) \quad (1)$$

where η is the learning rate, and the coefficient r_i is a score computed as the *softmax* of the neural activations: $r_i = \frac{e^{y_i/T}}{\sum_j e^{y_j/T}}$. Here, T is the *temperature* parameter of the softmax, which serves to cope with the variance of the activations (the name comes from statistical mechanics, where this operation was first defined). The effect of such a defined score is to allow each neuron to specialize on a different cluster of input patterns.

The other learning rule that we consider is Hebbian Principal Component Analysis (HPCA) [4, 14, 24, 26]:

$$\Delta \mathbf{w}_i = \eta y_i \left(\mathbf{x} - \sum_{j=1}^i y_j \mathbf{w}_j \right) \quad (2)$$

WTA competition was studied in past work as a possible approach for training relatively shallow neural networks [20, 35] (with up to 2-3 hidden layers). The investigation was further extended to deeper networks, and to hybrid architectures where some layers were trained by backprop and others by Hebbian learning [1, 26]. Experimental results on CNNs showed promises of HPCA-like learning mechanisms initially with shallow networks [3], and then with deeper networks as well [23, 24, 26].

Since the HPCA and SWTA learning rules are unsupervised, they have found application in the context of semi-supervised neural network training, in order

to perform an unsupervised pre-training stage [22–25]. In particular, they were found to be particularly useful in sample efficient learning scenarios, i.e. situations with scarce availability of labeled data. Related approaches for unsupervised pre-training are based on autoencoding architectures [5, 16, 28, 38]. Results on various computer vision benchmarks suggest that Hebbian pre-training allows to significantly improve performance on such scenarios compared to other unsupervised pre-training methods such as Variational Auto-Encoder (VAE) pre-training [15, 16]. Application of Hebbian learning to semi-supervised settings seems a promising direction. Other approaches to semi-supervised learning are based on pseudo-labeling/consistency-based methods [13, 34]. However, these methods are not in contrast with unsupervised pre-training, and they could actually be integrated together. This possible future direction will also be highlighted in Section 6.

The problem with current Hebbian learning solutions is that they do not scale well to large datasets. Note that, the learning rules mentioned above describe the weight update for a single input \mathbf{x} . When there is a batch of inputs to be processed, the weight updates are aggregated over the batch dimension, typically by averaging (or weighted averaging, for SWTA, the weights being the competition scores r_i , check [1, 21] for details). Similarly, in a convolutional layer, \mathbf{x} would correspond to a patch extracted from an input at a given offset, and weight updates computed at different offsets need to be aggregated over all the extracted patches. In this contribution, we notice that these two phases (update computation and aggregation) can be merged together, which allows to reformulate Hebbian learning rules more efficiently in terms of matrix multiplications, which are particularly suitable for GPU computation. We show that our solution is able to scale well to large datasets such as ImageNet.

3 Efficient Hebbian learning with FastHebb method

Let us start by introducing some preliminary information about the multi-dimensional tensor data that we need to work with, and the notation that will be used in the following.




Fig. 1: Types of tensor objects involved in our scenario.

We define a *tensor* simply as a multi-dimensional array of data. In particular, our tensors are three-dimensional. We denote such tensors with capital letters,

followed by as many indices as dimensions (three in our case). A dimension of size 1, also known as a *singleton* dimension, is denoted with the symbol 1 as index. Moreover, we adopt the following convention: index $b = 1 \dots B$ denotes the *batch* dimension, index $n = 1 \dots N$ denotes the *neuron* dimension, and index $s = 1 \dots S$ denotes the *size* dimension. Note that the meaning of an index is inferred by the corresponding letter and not by its position. With reference to Fig. 1, the first tensor (from left to right) is a typical input tensor, consisting of a mini-batch of B inputs, each being a vector of size S . The second tensor represents a typical weight matrix, consisting of one weight vector for each of the N neurons, each of size S . The third is a typical output tensor, with each output being a vector on N elements, one for each neuron, and there is one such vector for each of the B elements in the batch. The last is a typical reconstruction tensor, which extends over all the dimensions.

Finally, in order to make the use of matrix multiplication explicit in our formulas, we will use the notation $\text{matmul}(\cdot, \cdot)$ as follows:

$$C_{d,e,g} = \sum_f A_{d,e,f} B_{d,g,f} = \sum_f A_{d,e,f} B_{d,f,g} := \text{matmul}(A_{d,e,f}, B_{d,f,g}) \quad (3)$$

Note that we are taking the tensor product between tensors A and B , identifying index d and contracting index f . This corresponds to a batch matrix multiplication over index d , i.e. mapping d pairs of matrices with indices (e, f) and (f, g) , to d matrices with indices (e, g) : $(e, f) \times (f, g) \rightarrow (e, g)$. If more than three dimensions are present, then the last two denote height and width of the matrices, and all the previous dimensions are considered as batch dimensions (and thus identified). If a batch dimension of one of the multiplied tensors happens to be a singleton, then it undergoes *broadcasting* to match the other tensor dimension, as done in common mathematical frameworks. In all the other cases the corresponding batch dimensions of the two tensors must have the same size (as well as the contracted f dimension). Sums, subtractions, and multiplications by constants over tensor are performed component-wise, but all dimensions must match. Also in this case, a singleton dimension of one tensor undergoes broadcast to match the corresponding dimension of the other tensor (in case of singleton dimensions, and only in this case, correspondence is inferred from the position of the indices).

Using the notation introduced above, we can express the Hebbian rules discussed in this paper, including the aggregation step, as follows:

$$\Delta W_{1,n,s} = \sum_b C_{b,n,1} \Delta W_{b,n,s} = \text{matmul}(C_{n,1,b}, \Delta W_{n,b,s}) \quad (4)$$

Tensor C represents the coefficients for (weighted) averaging during the update aggregation step. With our notation, we consider the batch index b to run over all the patches extracted from the inputs, and also over all the inputs in the mini-batch. In other words, all the patches extracted from all the images in the mini-batch are considered as a unique larger mini-batch over which aggregation is performed.

Notice that, at this point, update computation and aggregation phases are considered together. In fact, merging these two phases is an essential step towards achieving the performance improvement addressed in this work, as described below. In particular, as the dimension associated with index b is very large, since it runs over all the patches extracted from all the inputs, it would be beneficial to contract this index as soon as possible in our computations, possibly before larger tensors such as $\Delta W_{b,n,s}$ are obtained. We proceed differently depending on the Hebbian rule under consideration.

Hebbian Winner-Takes-All. The (soft-)WTA learning rule can be rewritten as follows:

$$\begin{aligned}\Delta W_{1,n,s} &= \eta \sum_b C_{b,n,1} R_{b,n,1} (X_{b,1,s} - W_{1,n,s}) \\ &= \eta \sum_b (C R)_{b,n,1} (X - W)_{b,n,s} \\ &= \eta \text{matmul}((C R)_{n,1,b}, (X - W)_{n,b,s})\end{aligned}\tag{5}$$

Where $C_{b,n,1} = \frac{R_{b,n,1}}{\sum_b R_{b,n,1}}$.

Note that this formulation requires $O(B N S)$ complexity both in time and space. In particular, it needs to store a $B \times N \times S$ tensor. All the elements are stored simultaneously in order to parallelize operations over each dimension through vectorized or GPU hardware. If the amount of memory required is prohibitive, it is possible to serialize computations over one or more dimensions. However, computational performance can be improved by rewriting:

$$\begin{aligned}\Delta W_{1,n,s} &= \eta \sum_b C_{b,n,1} R_{b,n,1} (X_{b,1,s} - W_{1,n,s}) = \\ &= \eta \sum_b (C R)_{b,n,1} X_{b,1,s} - \eta \sum_b (C R)_{b,n,1} W_{1,n,s} = \\ &= \eta \text{matmul}((C R)_{1,n,b}, X_{1,b,s}) - \eta \sum_b (C R)_{b,n,1} W_{1,n,s} = \\ &= \eta \text{matmul}((C R)_{1,n,b}, X_{1,b,s}) - \eta Q_{1,n,1} W_{1,n,s}\end{aligned}\tag{6}$$

Where $Q_{1,n,1} = \sum_b (C R)_{b,n,1}$.

By contracting index b early, we have obtained a new formulation that requires only $O(N(B + S))$ space. The time complexity depends on the algorithm employed for matrix multiplication, which can be made lower than $O(BNS)$.

Hebbian Principal Component Analysis. The Hebbian PCA learning rule can be rewritten as follows:

$$\begin{aligned}
\Delta W_{1,n,s} &= \eta \frac{1}{B} \sum_b Y_{b,n,1} \left(X_{b,1,s} - \sum_{n'=1}^n Y_{b,n',1} W_{1,n',s} \right) \\
&= \eta \frac{1}{B} \sum_b Y_{b,n,1} \left(X_{b,1,s} - \sum_{n'=1}^N L_{n,n'} Y_{b,n',1} W_{1,n',s} \right) \\
&= \eta \frac{1}{B} \sum_b Y_{b,n,1} E_{b,n,s} \\
&= \eta \frac{1}{B} \text{matmul}(Y_{n,1,b}, E_{n,b,s})
\end{aligned} \tag{7}$$

Where $E_{b,n,s} = \left(X_{b,1,s} - \sum_{n'=1}^N L_{n,n'} Y_{b,n',1} W_{1,n',s} \right)$, and $L_{n,n'}$ is simply a lower-triangular matrix with all ones on and below the main diagonal and all zeros above.

In this case, the computation requires $O(BN^2S)$ space and time, but this can be improved by rewriting:

$$\begin{aligned}
\Delta W_{1,n,s} &= \eta \frac{1}{B} \sum_b Y_{b,n,1} \left(X_{b,1,s} - \sum_{n'=1}^N L_{n,n'} Y_{b,n',1} W_{1,n',s} \right) \\
&= \eta \frac{1}{B} \sum_b Y_{b,n,1} X_{b,1,s} - \eta \frac{1}{B} \sum_b Y_{b,n,1} \sum_{n'=1}^N L_{n,n'} Y_{b,n',1} W_{1,n',s} \\
&= \eta \frac{1}{B} \text{matmul}(Y_{1,n,b}, X_{1,b,s}) - \eta \frac{1}{B} \sum_{n'=1}^N \sum_b Y_{b,n,1} Y_{b,n',1} L_{n,n'} W_{1,n',s} \\
&= \eta \frac{1}{B} \text{matmul}(Y_{1,n,b}, X_{1,b,s}) - \eta \frac{1}{B} \sum_{n'=1}^N \text{matmul}(Y_{1,n,b}, Y_{1,b,n'}) L_{n,n'} W_{1,n',s} \\
&= \eta \frac{1}{B} \text{matmul}(Y_{1,n,b}, X_{1,b,s}) - \eta \frac{1}{B} \sum_{n'=1}^N P_{1,n,n'} W_{1,n',s} \\
&= \eta \frac{1}{B} \text{matmul}(Y_{1,n,b}, X_{1,b,s}) - \eta \frac{1}{B} \text{matmul}(P_{1,n,n'}, W_{1,n',s})
\end{aligned} \tag{8}$$

Here, $P_{1,n,n'} = \text{matmul}(Y_{1,n,b}, Y_{1,b,n'}) L_{n,n'}$.

This computation requires $O(N^2 + NS)$ space, and at most $O(BNS + BN^2 + N^2S)$ time.

4 Experimental setup

In order to validate our method, we performed experiments on various datasets in the computer vision domain. We evaluated both the computing time required by

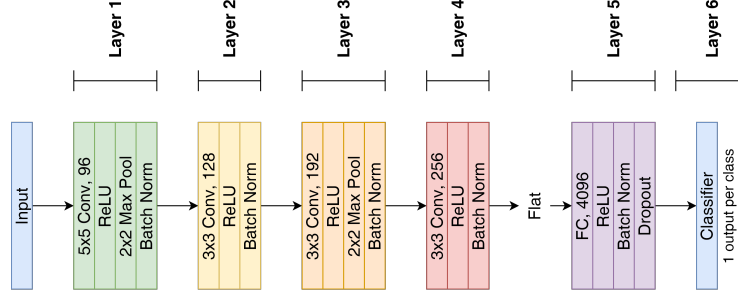


Fig. 2: The neural network used for the experiments.

Hebbian algorithms, with and without the FastHebb optimization, and their performance in sample efficiency scenarios, also making comparisons with backprop-based learning. In the following, we describe the details of our experiments and comparisons, discussing the network architecture and the training procedure.

4.1 Datasets and sample efficiency regimes

The datasets that we considered for our experiments are CIFAR10 [18], CIFAR100, Tiny ImageNet [37], and ImageNet [6]. We performed our experiments in various regimes of label scarcity. We define an $s\%$ *sample-efficiency regime* as a scenario in which on $s\%$ of the training set elements is assumed to be labeled. We considered 1%, 2%, 3%, 4%, 5%, 10%, 25%, and 100% sample efficiency regimes.

For each of the above regimes, we run our experiments in a semi-supervised training fashion: first, an unsupervised pre-training stage was performed, exploiting the Hebbian learning rules, using all the available training samples; this was followed by a supervised backprop-based fine-tuning stage on the labeled samples only.

4.2 Network architecture and training

We considered a six layer neural network as shown in Fig. 2: five deep layers plus a final linear classifier. The various layers were interleaved with other processing stages (such as ReLU nonlinearities, max-pooling, etc.), and the overall architecture was inspired by AlexNet [19].

A similar, but bigger model was used for ImageNet classification, which is shown in Fig. 3.

For each sample efficiency regime, we trained the network with our semi-supervised approach in a classification task. First, we used Hebbian unsupervised pre-training rules in the internal layers. This was followed by the fine tuning stage with SGD training, involving the final classifier as well as the previous layers, in an end-to-end fashion.




Fig. 3: The bigger neural network used for the experiments on ImageNet.

For each configuration we also created a baseline for comparison. In this case, we used another popular unsupervised method, namely the Variational Auto-Encoder (VAE) [15], for the unsupervised pre-training stage. This was again followed by the supervised end-to-end fine tuning based on SGD. VAE-based semi-supervised learning was also the approach considered in [16].

Both classification accuracy and training time were evaluated and used as metrics for comparisons.

4.3 Details of training

We implemented our experiments using PyTorch. All the hyperparameters mentioned below resulted from a parameter search aimed at maximizing the validation accuracy on the respective datasets, following the Coordinate Descent (CD) approach [17].

Training was performed in 20 epochs using mini-batches of size 64. No more epochs were necessary, since the models had already reached convergence at that point. Networks were fed input images of size 32x32 pixels, except for the case of ImageNet, where images of size 210x210 were used.

During Hebbian training, the learning rate was set to 10^{-3} (10^{-4} for ImageNet). No L2 regularization or dropout was used, since the learning method did not present overfitting issues.

For VAE training, the network backbone without the classifier acted as encoder, with an extra layer mapping the output to 256 gaussian latent variables, while a specular network branch acted as decoder. VAE training was performed without supervision, in an end-to-end encoding-decoding task, optimizing the β -VAE Variational Lower Bound [11], with coefficient $\beta = 0.5$.

Both for VAE training and for the supervised training stage, based on SGD, the initial learning rate was set to 10^{-3} and kept constant for the first ten epochs, while it was halved every two epochs for the remaining ten epochs. We also used momentum coefficient 0.9, and Nesterov correction. During supervised training, we also used dropout rate 0.5, L2 weight decay penalty coefficient set to $5 \cdot 10^{-2}$ for CIFAR10, 10^{-2} for CIFAR100, $5 \cdot 10^{-3}$ for Tiny ImageNet, and $1 \cdot 10^{-3}$ for ImageNet. Cross-entropy loss was used as optimization metric.

Table 1: Training times on each dataset, for VAE, Hebbian PCA (HPCA), Hebbian PCA with FastHebb (HPCA-FH), soft-WTA (SWTA), and soft-WTA with FastHebb (SWTA-FH) methods.

Dataset	Method	Epoch Duration	Num. Epochs	Total Duration
CIFAR10	VAE	14s	17	3m 58s
	SWTA	4m 14s	1	4m 14s
	SWTA-FH	18s	1	18s
	HPCA	6m 23s	12	1h 16m 36s
	HPCA-FH	19s	12	3m 48s
CIFAR100	VAE	15s	15	3m 45s
	SWTA	4m 16s	1	4m 16s
	SWTA-FH	18s	1	18s
	HPCA	6m 25s	7	44m 55s
	HPCA-FH	19s	7	2m 13s
Tiny ImageNet	VAE	33s	20	11m
	SWTA	9m 41s	1	9m 41s
	SWTA-FH	41s	1	41s
	HPCA	14m 20s	14	3h 20m 40s
	HPCA-FH	43s	14	10m 2s
ImageNet	VAE	2h 59m 19s	16	47h 49m 4s
	SWTA	105h 13m 24s	3	315h 40m 12s
	SWTA-FH	3h 38m 6s	3	10h 54m 18s
	HPCA	155h 41m 39s	3	467h 4m 57s
	HPCA-FH	3h 39m 18s	3	10h 57m 54s

To obtain the best possible generalization, *early stopping* was used in each training session, i.e. we chose as final trained model the state of the network at the epoch when the highest validation accuracy was recorded.

Experiments were performed on an Ubuntu 20.4 machine, with Intel Core I7 10700K Processor, 32GB Ram, and NVidia Geforce 3060 GPU with 12GB dedicated memory. The experiments were implemented using the Pytorch package, version 1.8, and Python 3.7.

5 Results and discussion

In this section, the experimental results obtained with each dataset are presented and analyzed. We report the training times on each dataset, for all the approaches explored. Moreover, we report the classification accuracy in the semi-supervised task, in the various sample efficiency regimes. Experiment results from five independent iterations were averaged, and the differences between methods were tested for statistical significance with a p value of 0.05.

5.1 Training time performance evaluation

Table 1 shows the training time measured on the various datasets, for each of the considered approaches. We measured the average epoch duration, the total

Table 2: Accuracy results on each dataset (top-1 for CIFAR10, and top-5 for the other datasets, since they have many more classes), for the various approaches explored.

Regime	Method	CIFAR10	CIFAR100	Tiny	ImageNet
1%	VAE	22.54	12.28	5.55	2.72
	SWTA	30.23	15.30	6.20	6.69
	HPCA	39.75	22.63	11.38	8.65
2%	VAE	26.78	15.25	6.74	6.14
	SWTA	36.59	20.76	8.56	11.52
	HPCA	45.51	30.83	15.71	13.64
3%	VAE	29.00	16.44	7.74	15.35
	SWTA	41.54	23.69	10.26	15.67
	HPCA	48.80	35.04	18.23	17.28
4%	VAE	31.15	17.89	8.45	23.97
	SWTA	45.31	26.91	11.52	19.95
	HPCA	51.28	38.89	20.55	20.39
5%	VAE	32.75	18.48	9.29	29.04
	SWTA	48.35	29.57	12.55	24.87
	HPCA	52.20	41.42	22.46	23.28
10%	VAE	45.67	23.80	13.51	43.73
	SWTA	58.00	38.26	16.70	41.54
	HPCA	57.35	48.93	28.13	34.27
25%	VAE	68.70	52.59	37.89	61.33
	SWTA	69.85	56.26	24.96	59.34
	HPCA	64.77	58.70	37.10	56.92
100%	VAE	85.23	79.97	60.23	76.84
	SWTA	85.37	79.80	54.94	76.10
	HPCA	84.38	74.42	53.96	77.28

number of training epochs required by each method, and the total training time. The number of epochs is counted by considering the training over when the network performance stops improving. The reported number of epochs refers to the pre-training phase only, and not to the successive fine-tuning, as we observed no statistically significant difference in the duration of the latter phase for different pre-training methods. Training time of FastHebb methods are compared to the previous respective best known solutions for Hebbian learning, that were also based on GPU [26].

We can see that, in terms of total training time, Hebbian methods are almost five times faster than VAE on ImageNet. Among the Hebbian approaches, soft-WTA is faster, thanks to its lower time complexity. Most importantly, as shown from the ImageNet performance results, thanks to the novel optimization, FastHebb algorithms scale gracefully also to large scale datasets.

5.2 Semi-supervised, sample efficiency scenario

Table 2 shows the classification accuracy results obtained on the various dataset, for each of the considered approaches. Top-1 accuracy was used for CIFAR10, and top-5 for all the other datasets, since they have many more classes. Note that, in this case, we show the results for HPCA and soft-WTA, but these are the same with or without the FastHebb optimization. In fact, the optimization does not change the update rule itself.

We can observe that Hebbian approaches perform better than VAE in sample efficiency regimes with very scarce label availability, below 4-5%. In particular, we can observe performance improvements of HPCA of almost 20% in the 5% regime for the CIFAR10 dataset. On the other hand VAE-based pre-training only improves when the available number of labeled training samples for the successive supervised fine-tuning phase becomes larger. When scaling up to ImageNet dataset, we still have a slight advantage of Hebbian methods in scarce data regimes (from 2 to 6%, depending on the regime). However, when higher regimes are considered, the performance of Hebbian pre-training is slightly lower than VAE, but this is compensated, as shown before, by a significant advantage in terms of training time.

6 Conclusions and future work

We have shown how the FastHebb approach can be leveraged to optimize running times of Hebbian learning algorithms for DNN training. Thanks to this optimization, we were able to scale Hebbian learning experiments to ImageNet level. To the best of our knowledge, this is the first solution able to bring Hebbian learning to such scale. Experiments in semi-supervised scenarios show the efficacy of Hebbian approaches for unsupervised network pre-training, compared to backprop-based VAE pre-training, both in terms of classification accuracy and training time, especially in sample efficiency scenarios where the labeled data for supervised fine tuning are scarce (less than 4-5% of the overall available data).

As possible future work directions, we suggest to perform further studies of FastHebb on other large-scale application scenarios, such as Content Based Image Retrieval (CBIR) to evaluate the quality of deep features extracted by this method. Preliminary work in this direction is promising [22]. Moreover, further Hebbian rules can also be derived, for example from Independent Component Analysis (ICA) [12] and sparse coding [30,31,33]. Finally, in the context of semi-supervised learning, Hebbian approaches can also be combined with pseudo-labeling and consistency-based methods mentioned in Section 2 [13,34].

References

1. Amato, G., Carrara, F., Falchi, F., Gennaro, C., Lagani, G.: Hebbian learning meets deep convolutional neural networks. In: International Conference on Image Analysis and Processing. pp. 324–334. Springer (2019)

2. Babenko, A., Slesarev, A., Chigorin, A., Lempitsky, V.: Neural codes for image retrieval. In: European conference on computer vision. pp. 584–599. Springer (2014)
3. Bahroun, Y., Soltoggio, A.: Online representation learning with single and multi-layer hebbian networks for image classification. In: International Conference on Artificial Neural Networks. pp. 354–363. Springer (2017)
4. Becker, S., Plumbley, M.: Unsupervised neural network learning procedures for feature extraction and classification. *Applied Intelligence* **6**(3), 185–203 (1996)
5. Bengio, Y., Lamblin, P., Popovici, D., Larochelle, H.: Greedy layer-wise training of deep networks. In: Advances in neural information processing systems. pp. 153–160 (2007)
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
7. Gerstner, W., Kistler, W.M.: Spiking neuron models: Single neurons, populations, plasticity. Cambridge university press (2002)
8. Grossberg, S.: Adaptive pattern classification and universal recoding: I. parallel development and coding of neural feature detectors. *Biological cybernetics* **23**(3), 121–134 (1976)
9. Haykin, S.: Neural networks and learning machines. Pearson, 3 edn. (2009)
10. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
11. Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., Lerchner, A.: beta-vae: Learning basic visual concepts with a constrained variational framework (2016)
12. Hyvarinen, A., Karhunen, J., Oja, E.: Independent component analysis. *Studies in informatics and control* **11**(2), 205–207 (2002)
13. Iscen, A., Tolias, G., Avrithis, Y., Chum, O.: Label propagation for deep semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5070–5079 (2019)
14. Karhunen, J., Joutsensalo, J.: Generalizations of principal component analysis, optimization problems, and neural networks. *Neural Networks* **8**(4), 549–562 (1995)
15. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
16. Kingma, D.P., Mohamed, S., Jimenez Rezende, D., Welling, M.: Semi-supervised learning with deep generative models. *Advances in neural information processing systems* **27**, 3581–3589 (2014)
17. Kolda, T.G., Lewis, R.M., Torczon, V.: Optimization by direct search: New perspectives on some classical and modern methods. *SIAM review* **45**(3), 385–482 (2003)
18. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
19. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25**, 1097–1105 (2012)
20. Krotov, D., Hopfield, J.J.: Unsupervised learning by competing hidden units. *Proceedings of the National Academy of Sciences* **116**(16), 7723–7731 (2019)
21. Lagani, G.: Hebbian learning algorithms for training convolutional neural networks. Master’s thesis, School of Engineering, University of Pisa, Italy (2019), <https://etd.adm.unipi.it/theses/available/etd-03292019-220853/>

22. Lagani, G., Bacciu, D., Gallicchio, C., Falchi, F., Gennaro, C., Amato, G.: Deep features for cbir with scarce data using hebbian learning. *arXiv preprint arXiv:2205.08935* (2022)
23. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Evaluating hebbian learning in a semi-supervised setting. In: *International Conference on Machine Learning, Optimization, and Data Science*. pp. 365–379. Springer (2021)
24. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Hebbian semi-supervised learning in a sample efficiency setting. *Neural Networks* **143**, 719–731 (2021)
25. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Training convolutional neural networks with competitive hebbian learning approaches. In: *International Conference on Machine Learning, Optimization, and Data Science*. pp. 25–40. Springer (2021)
26. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Comparing the performance of hebbian against backpropagation learning using convolutional neural networks. *Neural Computing and Applications* **34**(8), 6503–6519 (2022)
27. Lagani, G., Mazziotti, R., Falchi, F., Gennaro, C., Cicchini, G.M., Pizzorusso, T., Cremisi, F., Amato, G.: Assessing pattern recognition performance of neuronal cultures through accurate simulation. In: *2021 10th International IEEE/EMBS Conference on Neural Engineering (NER)*. pp. 726–729. IEEE (2021)
28. Larochelle, H., Bengio, Y., Louradour, J., Lamblin, P.: Exploring strategies for training deep neural networks. *Journal of machine learning research* **10**(1) (2009)
29. Nowlan, S.J.: Maximum likelihood competitive learning. In: *Advances in neural information processing systems*. pp. 574–582 (1990)
30. Olshausen, B.A.: Learning linear, sparse, factorial codes. *Massachusetts Institute of Technology, AIM-1580* (1996)
31. Olshausen, B.A., Field, D.J.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**(6583), 607 (1996)
32. O’Reilly, R.C., Munakata, Y.: *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. MIT press (2000)
33. Rozell, C.J., Johnson, D.H., Baraniuk, R.G., Olshausen, B.A.: Sparse coding via thresholding and local competition in neural circuits. *Neural computation* **20**(10), 2526–2563 (2008)
34. Sellars, P., Aviles-Rivero, A.I., Schönlieb, C.B.: Laplacenet: A hybrid energy-neural model for deep semi-supervised classification. *arXiv preprint arXiv:2106.04527* (2021)
35. Wadhwa, A., Madhow, U.: Bottom-up deep learning using the hebbian principle (2016)
36. Wan, J., Wang, D., Hoi, S.C.H., Wu, P., Zhu, J., Zhang, Y., Li, J.: Deep learning for content-based image retrieval: A comprehensive study. In: *Proceedings of the 22nd ACM international conference on Multimedia*. pp. 157–166 (2014)
37. Wu, J., Zhang, Q., Xu, G.: Tiny imagenet challenge. *Tech. rep., Stanford University* (2017)
38. Zhang, Y., Lee, K., Lee, H.: Augmenting supervised neural networks with unsupervised objectives for large-scale image classification. In: *International conference on machine learning*. pp. 612–621 (2016)