

Online Representation Learning with Single and Multi-layer Hebbian Networks for Image Classification

Yanis Bahroun and Andrea Soltoggio

Loughborough University, Computer Science Department,
Loughborough, Leicestershire, United Kingdom
{y.bahroun, a.soltoggio}@lboro.ac.uk

To appear in: Proceedings of the 26th International Conference on Artificial Neural Networks (ICANN 2017), Alghero, Italy, 11-15 September, 2017.

Abstract. Unsupervised learning permits the development of algorithms that are able to adapt to a variety of different datasets using the same underlying rules thanks to the autonomous discovery of discriminating features during training. Recently, a new class of Hebbian-like and local unsupervised learning rules for neural networks have been developed that minimise a *similarity matching* cost-function. These have been shown to perform sparse representation learning. This study tests the effectiveness of one such learning rule for learning features from images. The rule implemented is derived from a nonnegative classical multidimensional scaling cost-function, and is applied to both single and multi-layer architectures. The features learned by the algorithm are then used as input to an SVM to test their effectiveness in classification on the established CIFAR-10 image dataset. The algorithm performs well in comparison to other unsupervised learning algorithms and multi-layer networks, thus suggesting its validity in the design of a new class of compact, online learning networks.

Keywords: Classification; Competitive learning; Feature learning; Hebbian learning; Online algorithm; Neural networks; Sparse coding; Unsupervised learning.

1 Introduction

Biological synaptic plasticity is hypothesized to be one of the main phenomena responsible for human learning and memory. One mechanism of synaptic plasticity is inspired by the Hebbian learning principle which states that connections between two units, e.g., neurons, are strengthened when they are simultaneously activated. In artificial neural networks, implementations of Hebbian plasticity are known to learn recurring patterns of activations. The use of extensions of this rule, such as Oja's rule [9] or the Generalized Hebbian rule, also called Sanger's rule [15], have permitted the development of algorithms that have proved particularly efficient at tasks such as online dimensionality reduction. Two important properties of brain-inspired models, namely competitive learning [14] and sparse coding [10] can be performed using Hebbian and anti-Hebbian learning rules. Such properties can be achieved with inhibitory connections, which extend the capabilities of Hebbian rules beyond simple extraction of the principal component of input data. The continuous and local update dynamics of Hebbian learning also make it suitable for learning from a continuous stream of data. Such an algorithm can take one image at a time with memory requirements that are independent of the number of samples.

This study employs Hebbian/anti-Hebbian learning rules derived from a similarity matching cost-function [12] and applies it to perform online unsupervised learning of features from multiple image datasets. The rule proposed in [12] is applied here for the first time to online features learning for image classification with single and multi-layer architectures. The quality of the features is assessed visually and by performing classification with a linear classifier working on the learned features. The simulations show that a simple single-layer Hebbian network can outperform more complex models such as Sparse Autoencoders (SAE) and Restricted Boltzmann machines (RBM) for image classifications tasks [3]. When applied to multi-layer architectures, the rule learns additional features. This study is the first of its kind to perform multi-layer sparse dictionary learning based on the similarity matching principle developed in [12] and to apply it to image classification. A preliminary draft of this paper has been made available on arXiv repository [2].

2 Hebbian/anti-Hebbian Network Derived From a Similarity Matching Cost-Function

The rule implemented by the Hebbian/anti-Hebbian network used in this work derives from an adaptation of Classical MultiDimensional Scaling (CMDS). CMDS is a popular embedding technique [4]. Unlike most dimensionality reduction techniques, e.g. PCA, the CMDS uses as input the matrix of similarity between inputs to generate a set of embedding coordinates. The advantage of MDS is that any kind of distance or similarity matrix can be analyzed. However, in its simplest form, CMDS produces dense features maps which are often unsuitable when considered for image classification. Therefore an adaptation of the CMDS introduced recently in [12] is used to overcome this weakness. The model implemented is a nonnegative classical multidimensional scaling that has three properties: it takes a similarity matrix as input, it produces sparse codes, and can be implemented using a new biologically plausible Hebbian model. The Hebbian/anti-Hebbian rule introduced in [12] is given as follows: for a set of inputs $x^t \in \mathbb{R}^n$ for $t \in \{1, \dots, T\}$, the concatenation of the inputs defines an input matrix $X \in \mathbb{R}^{n \times T}$. The output matrix Y of encodings is an element of $\mathbb{R}^{m \times T}$ that corresponds to a sparse overcomplete representation of the input if $m > n$, or to a low-dimensional embedding of the input if $m < n$. The objective function proposed by [12] is:

$$Y^* = \arg \min_{Y \geq 0} \|X'X - Y'Y\|_F^2 \quad (1)$$

where F is the Frobenius norm and $X'X$ is the Gram matrix of the inputs which corresponds to the similarity matrix. Solving Eq.1 directly requires storing $Y \in \mathbb{R}_+^{m \times T}$ which increases with time T making online learning difficult. Thus instead an online learning version of Eq.1 is expressed as:

$$(y^T)^* = \arg \min_{y^T \geq 0} \|X'X - Y'Y\|_F^2 \quad . \quad (2)$$

The components of the solution of Eq.2, found in [12] using coordinate descent, are :

$$(y_i^T)^* = \max \left(W_i^T x^T - M_i^T (y^T)^*, 0 \right) \quad \forall i \in \{1, \dots, m\}, \quad (3)$$

$$\text{where } W_{ij}^T = \frac{\sum_{t=1}^{T-1} y_i^t x_j^t}{\sum_{t=1}^{T-1} (y_i^t)^2} ; \quad M_{ij}^T = \frac{\sum_{t=1}^{T-1} y_i^t y_j^t}{\sum_{t=1}^{T-1} (y_i^t)^2} \mathbf{1}_{i \neq j} . \quad (4)$$

W^T and M^T can be found using the recursive formulations:

$$W_{ij}^T = W_{ij}^{T-1} + \left(y_i^{T-1} (x_j^{T-1} - W_{ij}^{T-1} y_i^{T-1}) / \hat{Y}_i^T \right) \quad (5)$$

$$M_{ij \neq i}^T = M_{ij \neq i}^{T-1} + \left(y_i^{T-1} (y_j^{T-1} - M_{ij \neq i}^{T-1} y_i^{T-1}) / \hat{Y}_i^T \right) \quad (6)$$

$$\hat{Y}_i^T = \hat{Y}_i^{T-1} + (y_i^{T-1})^2 . \quad (7)$$

W^T (green arrows) and M^T (blue arrows) can be interpreted respectively as feed-forward connections between the input and the hidden layer and lateral synaptic inhibitory connections within the hidden layer. The weight matrices are of fixed sizes and updated sequentially, which makes the model suitable for online learning. The architecture of the Hebbian/anti-Hebbian network is represented in Figure 1.




Fig. 1: Hebbian/anti-Hebbian network with lateral connections derived from Eq.2

3 A Model to Learn Features From Images

In the new model presented in this study, the input data vectors (x^1, \dots, x^T) are composed of patches taken randomly from a training dataset of images. For every new input x^t presented, the model first computes a sparse post-synaptic activity y^t . Second, the synaptic weights are modified based on local Hebbian/anti-Hebbian learning rules requiring only the current pre- post-synaptic neuronal activities. The model can be seen as a sparse encoding followed by a recursive updating scheme, which are both well suited to solve large-scale online problems.

A multi-class SVM classifies the pictures using output vectors obtained by a simple pooling of the feature vectors, Y^* , obtained for the input images from the trained network. In particular, given an input image, each neuron in the output layer produces a new image, called a feature map, which is pooled in quadrants [3] to form 4 terms of the input vector for the SVM.

3.1 Multi-layer Hebbian/anti-Hebbian Neural Network

In the proposed approach, layers of Hebbian/anti-Hebbian network are stacked similarly to the Convolutional DBN [5], and Hierarchical K-means. In the multi-layer Hebbian/anti-Hebbian network, both the weights of the first layer and second layer are continuously updated. Unlike other CNNs, the non-linearity used in each layer is not only due to the positivity constraint, but to the combination of a rectified linear unit activation function and of interneuronal competition. This model combines the powerful architecture of convolutional neural networks using ReLU activation with interneuronal competition, while all synaptic weights are updated using online local learning rules. In between layers, a 2×2 average pooling is used to downsample the feature maps.

3.2 Overcompleteness of the Representation and Multi-resolution

As part of the evaluation of the new model, it is important to assess its performance with different sizes (m) of the hidden layers. If the number of neurons exceeds the size of the input ($m > n$), the representation is called overcomplete. Overcompleteness may be beneficial, but requires increased computation, particularly for deep networks in which the number of neurons has to grow exponentially in order to keep this property. One motivation for overcompleteness is that it may allow more flexibility in matching the output structure with the input. However, not all learning algorithms can learn and take advantage of overcomplete representations. The behaviour of the algorithm is analysed in the transition between undercomplete ($m < n$) and overcomplete ($m > n$) representations.

Although the model might benefit from a large number of neurons, from a practical perspective an increase in the number of neurons is a challenge for such models due to the number of operations required in the coordinate descent. In order to limit the computational cost of training a large network while still benefiting from overcomplete representations, this study proposes to train simultaneously three single-layer neural networks, each of them having different receptive field sizes (4×4 , 6×6 , and 8×8 pixels). Thus, a variation of the model tested here is composed of three different networks. This architecture of parallel networks with different receptive field sizes requires less computational time and memory than a model with only one receptive field size and the same total number of neurons, because the synaptic weights only connect neurons within each neural network. This model will be called multi-resolution in the following.

3.3 Parameters and Preprocessing

The architecture used here has the following tunable parameters: the receptive field size (n) of the neurons and the number of neurons (m). These parameters are standard to CNNs but their influence on this online feed-forward model needs to be investigated.

For computer vision models, understanding the influence of input preprocessing is of critical importance for both biological plausibility and practical applicability. Recent findings [1], confirm partial decorrelation of the input signal in the retinal ganglion cells. The influence of input decorrelation by applying whitening will be investigated.

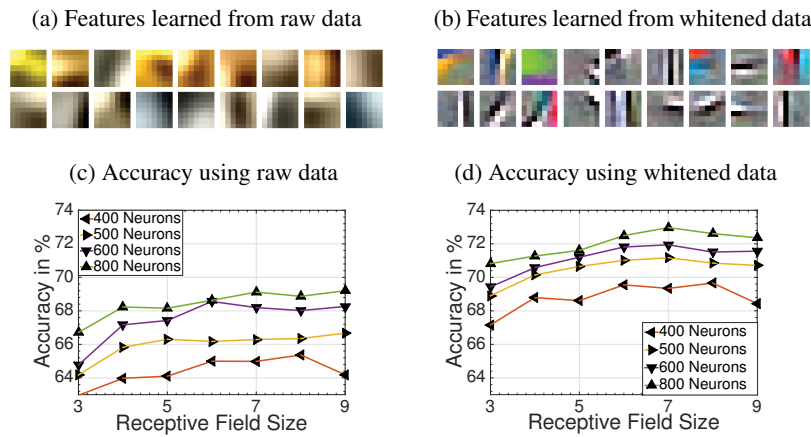
4 Results

The effectiveness of the algorithm is assessed by measuring the performance on an image classification task. We acknowledge that classification accuracy is at best an implicit measure evaluating the performance of representation learning algorithms, but provides a standardised way of comparing them. In the following, single and multi-layer Hebbian/anti-Hebbian neural networks combined with the standard multi-class SVM are trained on the CIFAR-10 dataset [6].

4.1 Evaluation of the Single-layer Model

A first experiment tested the performance of the model with and without whitening of the input data. Although there exist Hebbian networks that can perform online whitening [11], an offline technique based on singular value decomposition [3] is applied in these experiments. Figure 2a and 2b show the features learned by the network from raw input and whitened input respectively. The features learned from raw data (Fig.2a) are neither sharp nor localised filters and just slightly capture edges. With whitened data (Fig.2b), the features are sharp, localised, and resemble Gabor filters, which are observed in the primary visual cortex [10].

Fig. 2: Sample of features learned from raw (2a) and whitened input (2b). Classification accuracy with raw (2c) and whitened input (2d).



In a second set of experiments, the performance of the network was tested for varying receptive field sizes (Fig.2c-2d) and varying network sizes (400, 500, 600, and 800 neurons). The results show that the performance peaks at a receptive field size of 7 pixels and then begins to decline. This property is common to most unsupervised learning algorithms [3], showing the difficulty of learning spatially extended features. Fig-


ures 2c and 2d also show that for every configuration, the performance of the algorithm is largely and uniformly improved when whitening is applied to the input.

4.2 Comparison to State-of-the-art Performances and Online Training

Various unsupervised learning algorithms have been tested on the CIFAR-10 dataset. Spherical K-means, in particular, proved in [3] to outperform autoencoders and restricted Boltzmann machines, providing a very simple and efficient solution for dictionary learning for image classification. Thus, spherical K-means is used here as a benchmark to evaluate the performance of the single-layer network. As with other unsupervised learning algorithms, increasing the number of output neurons to reach over-completeness also improved classification performance (Fig.3a). Although the single-layer neural network has a higher degree of sparsity than the K-means proposed in [3] (results not shown here), they appear to have the same performance in their optimal configurations (Fig.3a).

The classification accuracy of the network during training is shown in Fig.3b. The graph (Fig.3b) suggests that the features learned by the network over time help the system improve the classification accuracy. This is significant because it demonstrates for the first time the effectiveness of features learned with a Hebb-like cost-function minimisation. It is not obvious a priori that the online optimisation of a cost-function for sparse similarity matching (Eq.2) produces features suitable for image classification.

Fig. 3: (a) Proposed model vs K-means, (b) Classification accuracy



As shown in Table 1, the multi-resolution network outperforms the single resolution network and K-means algorithm [3], reaching 80.42% accuracy on the CIFAR-10. The multi-resolution model shows better performance, while requiring less computation and memory than the single resolution model. It also outperforms the single layer NOMP [7], sparse TIRBM [16], CKN-GM and CKN-PM [8], which are more complex models. It was outperformed only by combined models or models with three layers or more. The results reported in Table 1 were averaged over multiple training and testing. The classification accuracy did not deviate by more than $\pm 0.17\%$ for the single resolution and $\pm 0.13\%$ for the multi-resolution networks.

Algorithm	Accuracy
Single-Layer, Single Resolution (4k neurons)	79.58 % ($\pm 0.17\%$)
Single-Layer, Multi-Resolution ($3 \times 1.6k$ neurons)	80.42 % ($\pm 0.13\%$)
Single-layer K-means [3] (4k neurons)	79.60 %
Multi-layer K-means [3] (3 Layers, >4k neurons)	82.60 %
Convolutional DBN [5]	78.90 %
Sparse TIRBM [16] (4k neurons)	80.10%
TIOMP-1/T [16] (combined transformations, 4k neurons)	82.20 %
Single Layer NOMP [7] (5k neurons)	78.00 %
Multi-Layer NOMP [7] (3 Layers, >4k neurons)	82.90 %
Multi-Layer CKN-GM [8]	74.84 %
Multi-Layer CKN-PM [8]	78.30 %
Multi-Layer CKN-CO [8] (combining CKN-GM & CKN-PM)	82.18 %

Table 1: Comparison of the single-layer network with unsupervised learning algorithms on CIFAR-10.

4.3 Evaluation of the Multi-layer Model

A single resolution, double-layer neural network with different numbers of neurons in each layer was trained similarly to the single-layer network in the previous section. In Table 2, ϕ_1 and ϕ_2 correspond respectively to the features learned by the first and second layer. The results show that ϕ_2 alone are less discriminative than ϕ_1 as indicated in Fig. 3a. However, when combined ($\phi_1 + \phi_2$) the model achieves better performance than each layer considered separately. Nevertheless, the preliminary results indicate that the sizes of the two layers unevenly affect the performance of the network. A future test may investigate if a multi-layer architecture can outperform the largest shallow networks. The accuracies in Table 2 are averaged over 7 runs that delivered consistent results with a standard deviation less than 0.2%. Further tests on smaller datasets might help determine if this robustness can be attributed to the model or to the SVM.

		#Neurons Layer 2				
		50	100	200	400	800
100 Neurons Layer 1	ϕ_2	54.9%	59.7%	64.7%	68.7%	71.5%
	$\phi_1 + \phi_2$	67.2%	68.1%	69.9%	72.4%	73.8%
200 Neurons Layer 1	ϕ_2	55.8%	60.6%	65.3%	70.3%	72.7%
	$\phi_1 + \phi_2$	69.9%	70.8%	71.9%	73.7%	75.1%

Table 2: Classification accuracy for a two-layer network.

5 Conclusion

This work proposes single and multi-layer neural networks that exploit Hebbian/anti-Hebbian rules to learn features for image classification. The networks are trained on the CIFAR-10 image dataset prior to feeding a linear classifier. The model successfully learns online more discriminative representations of the data when the number of neurons and the number of layers increase. The overcompleteness of the representation is critical for learning relevant features. The results show that a minimum unsupervised learning time is needed to optimise the network leading to better classification accuracy.

Finally, one key factor in improving image classification is the appropriate choice of the receptive field size used for training the network.

Such findings prove that neural networks can be trained to solve problems as complex as sparse dictionary learning with Hebbian learning rules, delivering competitive accuracy compared to other encoder, including deep neural networks. This makes deep Hebbian networks attractive for building large-scale image classification systems. The competitive performances on the CIFAR-10 suggests that this model can offer an alternative to batch trained neural networks. Ultimately, thanks to its bio-inspired architecture and learning rules, it also stands as a good candidate for memristive devices [13]. Moreover, if a decaying factor is added to the proposed model that might result in an algorithm that can deal with complex datasets with temporal variations of the distributions.

References

1. Abbasi-Asl, R., Pehlevan, C., Yu, B., Chklovskii, D.B.: Do retinal ganglion cells project natural scenes to their principal subspace and whiten them? arXiv preprint arXiv:1612.03483 (2016)
2. Bahroun, Y., Soltoggio, A.: Online representation learning with multi-layer Hebbian networks for image classification tasks. arXiv preprint arXiv:1702.06456 (2017)
3. Coates, A., Lee, H., Ng, A.Y.: An analysis of single-layer networks in unsupervised feature learning. In: AISTATS 2011. vol. 1001 (2011)
4. Cox, T.F., Cox, M.A.: Multidimensional scaling. CRC press (2000)
5. Krizhevsky, A., Hinton, G.: Convolutional deep belief networks on cifar-10. Unpublished manuscript 40 (2010)
6. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images (2009)
7. Lin, T.h., Kung, H.: Stable and efficient representation learning with nonnegativity constraints. In: Proceedings of the 31st International Conference on Machine Learning (ICML-14). pp. 1323–1331 (2014)
8. Mairal, J., Koniusz, P., Harchaoui, Z., Schmid, C.: Convolutional kernel networks. In: Advances in Neural Information Processing Systems. pp. 2627–2635 (2014)
9. Oja, E.: Neural networks, principal components, and subspaces. *International journal of neural systems* 1(01), 61–68 (1989)
10. Olshausen, B.A., et al.: Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381(6583), 607–609 (1996)
11. Pehlevan, C., Chklovskii, D.: A normative theory of adaptive dimensionality reduction in neural networks. In: Advances in Neural Information Processing Systems. pp. 2269–2277 (2015)
12. Pehlevan, C., Chklovskii, D.B.: A Hebbian/anti-Hebbian network derived from online non-negative matrix factorization can cluster and discover sparse features. In: 2014 48th Asilomar Conference on Signals, Systems and Computers. pp. 769–775. IEEE (2014)
13. Poikonen, J.H., Laiho, M.: Online linear subspace learning in an analog array computing architecture. *CNNA 2016* (2016)
14. Rumelhart, D.E., Zipser, D.: Feature discovery by competitive learning. *Cognitive science* 9(1), 75–112 (1985)
15. Sanger, T.D.: Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural networks* 2(6), 459–473 (1989)
16. Sohn, K., Lee, H.: Learning invariant representations with local transformations. In: Proceedings of the 29th International Conference on Machine Learning (ICML-12). pp. 1311–1318 (2012)