



HAL
open science

Unsupervised visual feature learning with spike-timing-dependent plasticity: How far are we from traditional feature learning approaches?

Pierre Falez, Pierre Tirilly, Ioan Marius Bilasco, Philippe Devienne, Pierre Boulet

► To cite this version:

Pierre Falez, Pierre Tirilly, Ioan Marius Bilasco, Philippe Devienne, Pierre Boulet. Unsupervised visual feature learning with spike-timing-dependent plasticity: How far are we from traditional feature learning approaches?. *Pattern Recognition*, 2019, 93, pp.418-429. 10.1016/j.patcog.2019.04.016 . hal-02146284

HAL Id: hal-02146284

<https://hal.science/hal-02146284>

Submitted on 22 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Unsupervised Visual Feature Learning with Spike-timing-dependent Plasticity: How Far are we from Traditional Feature Learning Approaches?

Pierre Falez^a, Pierre Tirilly^{b,*}, Ioan Marius Bilasco^a, Philippe Devienne^a,
Pierre Boulet^a

^aUniv. Lille, CNRS, Centrale Lille, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille, F-59000 Lille, France

^bUniv. Lille, CNRS, Centrale Lille, IMT Lille Douai, UMR 9189 - CRISTAL - Centre de Recherche en Informatique Signal et Automatique de Lille, F-59000 Lille, France

Abstract

Spiking neural networks (SNNs) equipped with latency coding and spike-timing dependent plasticity rules offer an alternative to solve the data and energy bottlenecks of standard computer vision approaches: they can learn visual features without supervision and can be implemented by ultra-low power hardware architectures. However, their performance in image classification has never been evaluated on recent image datasets. In this paper, we compare SNNs to auto-encoders on three visual recognition datasets, and extend the use of SNNs to color images. The analysis of the results helps us identify some bottlenecks of SNNs: the limits of on-center/off-center coding, especially for color images, and the ineffectiveness of current inhibition mechanisms. These issues should be addressed to build effective SNNs for image recognition.

Keywords: feature learning, unsupervised learning, spiking neural networks, spike-timing dependent plasticity, auto-encoders, image recognition.

1. Introduction

Machine learning algorithms require good data representations to be effective [1]. Good data representations can capture underlying correlations of the

*Corresponding author

data, provide invariance properties and help disentangle the data to make it linearly separable. In computer vision, much effort has been put historically into engineering the right visual features for recognizing, organizing and interpreting visual contents [2; 3]. More recently, and especially since the rise of deep learning, those features tend to be learned by algorithms rather than designed by human effort. Learned features have shown their superiority on a number of tasks, such as image classification [4], image segmentation [5], and action recognition [6]. Although effective, feature learning has two major drawbacks:

- it is *data-consuming*, as supervised learning algorithms – especially deep learning ones – require large amounts of annotated data to be trained;
- it is *energy-consuming*, as training large models, e.g., using gradient descent-based algorithms, has a high computational cost, which increases with the amount of training data. These algorithms are usually run on dedicated hardware (typically GPU) that are power-intensive.

The first issue – data consumption – can be mitigated by the use of unsupervised learning models. Unsupervised representation learning is recognized as one of the major challenges in machine learning [1] and is receiving a growing interest in computer vision applications [7; 8; 9]. A number of unsupervised models have been developed through the years, notably auto-encoders [10] and restricted Boltzmann machines (RBMs) [11], and their multi-layer counterparts, stacked auto-encoders [12] and deep belief networks (DBN) [13]. Other lines of work include sparse coding [14] and the use of semi- or weakly supervised learning algorithms [15]. Moreover, in the case of neural networks, initializing a deep neural network with features learned without supervision before training can yield better generalization capabilities than purely supervised training [12].

The second issue – energy consumption – is addressed much less frequently in the literature, but several authors acknowledge its importance [16; 17; 18] which is bound to grow more and more as machine learning becomes overwhelmingly present in a large range of applications: marketing, medicine, finance, education, administration, etc. Most hardware vendors have proposed dedicated machine

learning processor architectures (based on GPU, FPGA, etc.) recently [19]. These hardware improvements help reduce the energy consumption by a small factor (typically one order of magnitude). Reducing further the energy consumption of learning algorithms requires to define new learning models and associated ultra-low power architectures [18; 20; 21]. One promising model is spiking neural networks (SNNs). In this model, artificial neurons communicate information through spikes, as natural neurons do. Initially studied in neuroscience as a model of the brain, SNNs receive constant attention in the fields of machine learning and pattern recognition, from both the theoretical [22] and the applicative [17; 23; 24; 25] perspectives. Dedicated hardware implementing this model can be very energy-efficient [20]. SNNs have already shown their ability to provide near-state-of-the-art results in image classification, but only when they are trained by transferring parameters from pre-trained deep neural networks [21] or by variants of back-propagation [26]. In terms of energy efficiency, the first option is not viable as it still requires to train a standard deep neural network, which is exactly what should be avoided; the second option is not suited either as back-propagation is a global, centralized algorithm – the error must be propagated from the output to all units –, whereas the efficiency of SNNs lies in their ability to perform highly decentralized, parallel processing on sparse spike data. The alternative is to use bio-inspired learning rules, such as Hebbian rules. Among those, rules based on spike-timing dependent plasticity (STDP) [27] have shown promising results for learning visual features; however, they have only been evaluated on datasets with limited challenges (rigid objects, limited number of object instances, uncluttered backgrounds...) such as MNIST, 3D-object, ETH-80 or NORB [28; 29; 30; 31; 32], or on two-class datasets [31; 32]. How they perform on more complex image datasets, what is the performance gap between them and standard approaches, and what needs to be done to bridge this gap is yet to be established.

Aims and scope. In this paper, we evaluate the ability of SNNs equipped with latency coding and STDP to learn features for visual recognition on three stan-

standard datasets (CIFAR-10, CIFAR-100, and STL-10). Our goal is to identify some of the factors that prevent STDP-based SNNs to reach state-of-the-art results on actual computer vision tasks. First, we compare the performance of SNNs on grayscale and color images (Section 5.4), then we compare them to one standard unsupervised feature learning algorithm, sparse auto-encoders (Section 5.5). The resulting models are analyzed with respect to different factors (Section 6): input pre-processing, feature sparsity, feature coherence, and objective functions. It allows us to identify some bottlenecks that should be tackled to bridge the gap from SNNs to state-of-the-art models. In the conclusion (Section 7), we suggest some solutions to help address these bottlenecks.

In this work, we consider only single-layer architectures because multi-layer SNNs with unsupervised STDP are only very recent and difficult to train, due to the loss of spiking activity across layers [33; 31]. To our knowledge, this is the first work that evaluates features learned by unsupervised STDP-based SNNs on recent benchmarks for object recognition and on color images, making one step towards their use for actual vision applications.

2. Unsupervised visual feature learning

A visual feature extractor can be modeled as a function $f : \mathbb{R}^{h \times w} \rightarrow \mathbb{R}^{n_f}$ that maps an image or image region of size $h \times w$ to a real vector of dimension n_f . It defines a dictionary of features of size n_f . In the remaining of the paper, f will denote either the feature extractor function or the resulting dictionary, depending on the context. Early visual feature extractors were handcrafted to capture specific types of visual information (e.g., distributions of edges [3]). Recent approaches rather rely on machine learning to produce features that better fit the data and that can be optimized towards a specific application.

A typical learning-based feature extractor can be seen as a function f_θ whose parameters θ are optimized towards a specific goal by a learning algorithm. The general shape of f_θ can be specified explicitly (e.g., a linear transform [34]), or implicitly, based on the learning algorithm to be used (e.g., in a neural network, the possible shapes for f_θ are defined by the architecture of the network). The

parameters θ are optimized by the learning algorithm by minimizing an objective function L . In a supervised setting, this optimization step can be expressed as:

$$\theta^* = \arg \max_{\theta} L(X, Y; \theta) \quad (1)$$

where θ^* are the parameters returned by the learning algorithm, $X = (x_1, x_2 \dots)$ denotes the training samples and $Y = (y_1, y_2 \dots)$ the ground truth attached to the samples. L is directly set as a performance measure for the specific task to be solved, e.g. the multinomial logistic regression objective for image classification [4], or keypoint matching for image retrieval [34].

In an unsupervised setting, labels are not available. The optimization problem becomes:

$$\theta^* = \arg \max_{\theta} L(X; \theta) \quad (2)$$

In this case, L cannot be formulated towards a specific application. Some surrogate objective must be defined, that is expected to produce features that will fit the problem, e.g. image reconstruction [10], image denoising [35], and maximum likelihood [13]. Learning rules are sometimes defined directly without any explicit objective function, e.g. in k-means clustering [7], but also STDP [27].

Some constraints on the parameters or learning algorithm can be added to regularize the training process and reach better solutions. These constraints reflect assumptions on properties that "good" features should have, such as:

- sparsity: it is often assumed that the extracted features should be sparse, i.e. only a small number of the features can be found in a single image or image region. Sparsity is especially required when the set of features is over-complete¹, to prevent the algorithm from reaching trivial solutions. Sparsity is commonly imposed in sparse coding [14] and auto-encoders [36];
- coherence of features [37]: features should be different to span the space of visual patterns with limited redundancy². Coherence measures the possi-

¹A dictionary of features is over-complete when its dimension (number of features) is larger than the dimension of the input (size in pixels of the images or image regions processed).

²Some authors [38] claim that redundancy should rather be reached to have a good representation, but with limited evidence.

bility to reconstruct a given feature as a weighted sum of a small number of other features, i.e. whether features are locally linearly dependent; coherence should be small for the features to be effective.

In Section 6, these properties will serve as a basis for the analysis of the feature extractors.

3. STDP-based feature learning

While traditional machine learning algorithms represent data as real values, some models inspired by biology use spikes to carry information. These models are called spiking neural networks (SNNs). Working with spikes allows a complete desynchronization of the model since each spike is processed independently of the others. Thus, the system can be more energy-efficient when implemented on dedicated hardware [39] since no master synchronization system is required [40]. Moreover, in these models, memorization and computation are performed locally by neurons, which saves the data exchange cost and, as opposed to traditional Von-Neumann architectures, bypasses the bus bottleneck.

In the following, we detail the preprocessing and learning mechanisms involved in SNN architectures. SNNs have two main components: the spiking neurons, discussed in Section 3.1, and the synapses, which connect the neurons of the network. The synapses are generally responsible for the training of the network. One of the most used learning rule, spike-timing dependent plasticity (STDP) [27], is discussed in Section 3.2. As SNNs are fed with spikes, a preprocessing stage is required to transform pixel values into spike trains. Section 3.3 is focused on the preprocessing required by SNNs to process images. In Section 3.4, we discuss the representation of data as spikes, called the neural coding. The last sections describe other mechanisms required to achieve learning in SNNs: Section 3.5 is focused on lateral inhibition, which sets up competition between neurons to force them to learn different patterns; Section 3.6 presents homeostasis, which prevents certain neurons from taking advantage over the others, enforcing effective spiking activities throughout the network.

3.1. Spiking neuron models

Spiking neurons (in Figure 1a) are defined as processing units receiving spikes from their input connections (or synapses), and emitting spikes towards their output synapses when specific input patterns are received. There are several spiking neuron models, designed to accomplish different objectives [41]. Some aim to faithfully reproduce the behavior of biological neurons (e.g., the Hodgkin–Huxley and Morris–Lecar models), while others are designed to be efficiently implemented on hardware (e.g., the integrate-and-fire and Izhikevich models); the latter are usually simpler. This paper focuses on the integrate-and-fire (IF) model, to facilitate the understanding of the mechanisms of the models and to optimize the simulation speed. Most work on SNN-based image recognition use this kind of model [28; 29; 31]. IF neurons have a single internal variable: their voltage V . When a spike is fed to a neuron, it is integrated to V . V remains constant until the next spike, i.e. no leakage is applied. When V reaches the neuron threshold V_{th} , the neuron fires: it emits a spike towards its output neurons and resets its potential V to V_{rest} (see Figure 1b). Formally,

$$V' = I, V \leftarrow V_{\text{rest}} \text{ if } V \geq V_{\text{th}} \quad (3)$$

with V' the derivative of V and I the input current. V is also reset to V_{rest} between each input sample, to reset the network activity. This model allows the system to perform computations locally, since incoming spikes act on the state of the neuron, but not on the rest of the network (see Figure 1a).

3.2. Spike-timing dependent plasticity

Neurons are connected via synapses, which modulate the intensity of the spikes that they transmit via their synaptic weights W . Adapting these weights allows long-term learning, by reinforcing or weakening connections. To allow energy-efficient hardware implementations, this learning mechanism should use only local information like the input and the output spike trains of the neuron. Finally, each synapse propagates the spikes with a random delay d to introduce some noise and make lateral inhibition more active. In this paper, d is sampled from a uniform distribution in $[d_{\text{min}}; d_{\text{max}}]$.

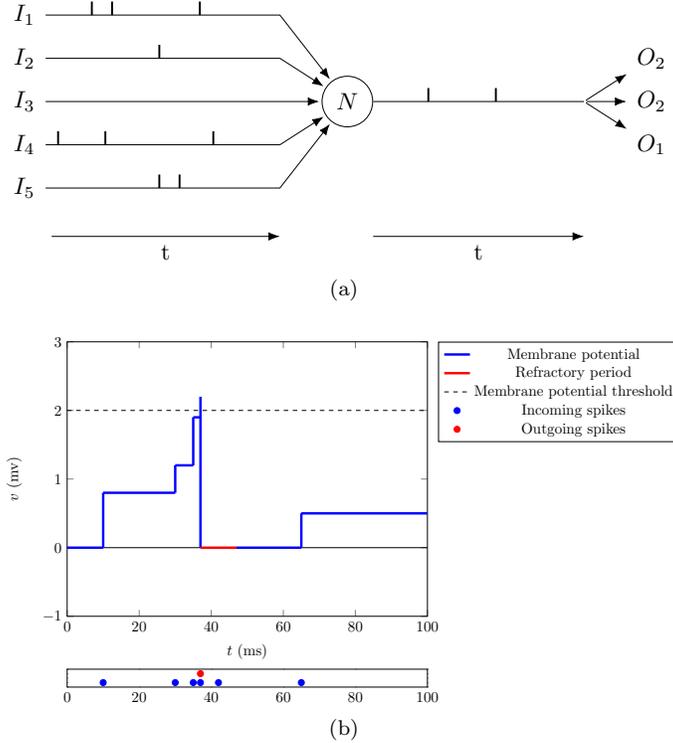


Figure 1: 1a A spiking neuron receives spike trains from incoming synapses, and generates spikes towards outgoing synapses. 1b Evolution of the membrane potential of an integrate-and-fire neuron.

Spike-timing dependent plasticity (STDP, see Figure 2a) [27] is a Hebb-like learning rule observed in biology. It reinforces the connections between neurons that have correlated firing patterns (potentiation), and weakens the others (depression). Multiplicative STDP [28] (see Figure 2b) is used in this paper. When a neuron fires a spike at time t_{post} , all incoming synapses are updated as follows:

$$\Delta w = \begin{cases} \alpha_+ e^{-\beta_+ \frac{w - W_{\min}}{W_{\max} - W_{\min}}} & \text{if } t_{\text{post}} \geq t_{\text{pre}} \\ -\alpha_- e^{-\beta_- \frac{W_{\max} - w}{W_{\max} - W_{\min}}} & \text{otherwise} \end{cases} \quad (4)$$

with Δw the update applied to synapse weight w , t_{pre} the timestamp of the input spike and t_{post} the timestamp of the output spike. α_+ and α_- are respectively the learning rates applied for potentiation and depression. β_+ and β_- control

the slope of the exponential. W_{\min} and W_{\max} are the bounds for synaptic weights. When no spike occurs during the presentation of a sample, t_{pre} is set to $+\infty$, so that weight depression occurs.

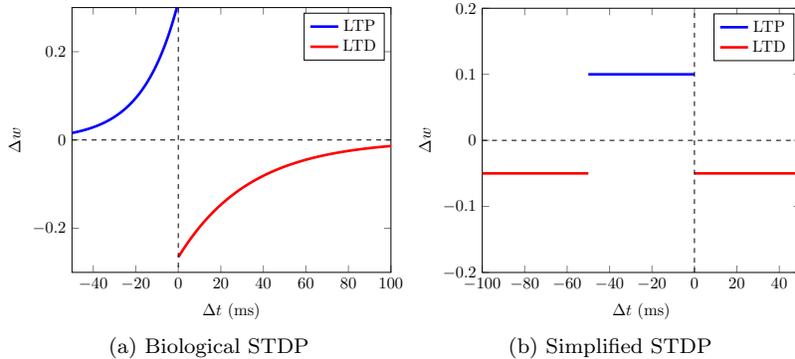


Figure 2: Weight variation w.r.t. the difference in spike timestamps.

3.3. Image pre-processing and color handling

As STDP learns correlations between input spikes, images are usually pre-processed to help STDP find meaningful correlations. Typically, edges are extracted from the grayscale images, e.g. through a difference-of-Gaussian (DoG) filter [31; 42] or Gabor filters [30]. In this paper, we use on-center/off-center coding, which combines DoG filtering with a 2-channel representation necessary to encode the data as spikes. This coding is inspired from bipolar cells of the retina. The output of the DoG filter at position (x, y) of image I is defined as:

$$\text{DoG}(x, y) = I(x, y) * (G_{\text{DoG}_{\text{size}}, \text{DoG}_{\text{center}}} - G_{\text{DoG}_{\text{size}}, \text{DoG}_{\text{surround}}})$$

where $*$ is the convolution operator and $G_{S, \sigma}$ is a normalized Gaussian kernel of size S and scale σ defined as:

$$G_{S, \sigma}(u, v) = \frac{g_{\sigma}(u, v)}{\sum_{i=-\mu}^{\mu} \sum_{j=-\mu}^{\mu} g_{\sigma}(i, j)}, u, v \in [-\mu, \mu], \mu = \frac{S}{2},$$

with g_{σ} the centered 2D Gaussian function of variance σ . The parameters of the filter are its size DoG_{size} and the variances of the Gaussian kernels $\text{DoG}_{\text{center}}$

and $\text{DoG}_{\text{surround}}$. After DoG filtering, the image is split into two channels c_+ and c_- as follows:

$$c_{\star}(x, y) = \max(0, \star \text{DoG}(x, y)), \star \in \{+, -\}$$

where $\text{DoG}(x, y)$ denotes the output of the DoG filter at position (x, y) . c_+ is the positive ("on") channel, and c_- the negative ("off") channel. This coding makes it possible to encode negative DoG values as spikes. No thresholds are applied to c_+ and c_- , i.e. the output values of the DoG filter are not filtered.

Color processing in SNNs is a problem that has not been addressed much in the literature so far. To apply on-center/off-center coding to color images, we define two strategies. In the first strategy, called RGB color opponent channels, the coding is applied to channels computed as differences of pairs of RGB channels: red-green, green-blue, and blue-red. The second strategy is inspired by biological observations: in the lateral geniculate nucleus, which mainly connects the retina to the visual cortex, three types of color channels exist: the black-white opponent channel (which corresponds to the grayscale image), the red-green opponent channel, and the yellow-blue opponent channel [43]. The second strategy applies on-center/off-center coding to the red-green and yellow-blue (computed as $0.5 \times R + 0.5 \times G - B$) channels. This leads to four possible configurations of image coding: grayscale only, RGB opponent channels, biological color opponent channels (referred to as Bio-color), and the combination of the grayscale channel and the Bio-color channels.

3.4. Neural coding

A step called neural coding is required to transform real or integer values from datasets spike trains. Several neural coding techniques have been proposed in the literature. The most common ones are frequency coding, which encodes values directly as frequencies of spike trains, and temporal coding, which encodes values as timestamps of single spikes [44] (see Figure 3). This paper uses latency coding, a form of temporal coding, since it has several advantages over frequency coding in visual tasks [45]. It allows to represent a sample with fewer spikes,

but also to simplify the model since at most one spike per connection is emitted during the presentation of a sample. Each input value x is transformed into a spike timestamp t as follows:

$$t = (1.0 - x) \times T_{\text{duration}} \quad (5)$$

with $x \in [0, 1]$ the input value and T_{duration} the duration of the presentation of a data sample.

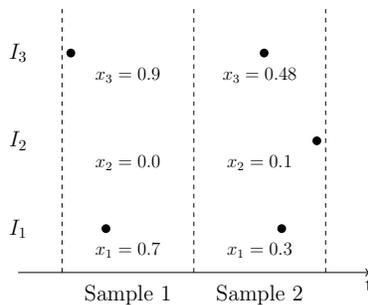


Figure 3: Latency coding.

3.5. Inhibition

STDP, or more generally unsupervised and local learning rules, require some competition mechanisms in order to prevent all neurons from learning the same pattern. Lateral inhibition is a straightforward method to do so in SNNs [28]: when a neuron fires, it sends inhibitory spikes to others neurons of the same layer, which reduce their abilities to fire. Only a small number of neurons is active at each input, which makes it possible to learn different patterns. In this paper, winner-take-all (WTA) competition is used: only one neuron is allowed to spike at each position per sample.

3.6. Homeostasis

Several authors demonstrated experimentally that SNNs need homeostasis to guarantee an effective learning process [28; 29]. It can be done by adapting the thresholds V_{th} to prevent one neuron from dominating the others. The thresholds of neurons that fire often are increased, so they will tend to fire less

often later on; inversely, the thresholds of sub-active neurons are decreased. Since latency coding is used, we train neurons to fire at an objective timestamp t_{obj} . The choice of this parameter can be tricky. An early t_{obj} trains neurons to respond to few input spikes, so that neurons learn only local patterns in few input channels (blobs or edges – see Figure 4a). A late t_{obj} results in larger patterns with a combination of multiple input channels (see Figure 4b).

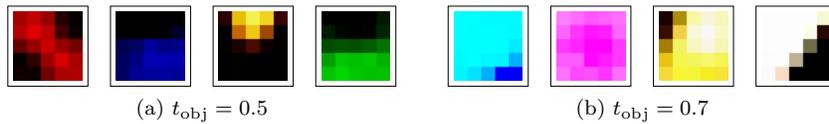


Figure 4: Filters learned with different t_{obj} .

Threshold adaptation is performed as follows:

$$\Delta V_{\text{th}} = -\eta \times (t_{\text{fire}} - t_{\text{obj}}) + \begin{cases} \eta & \text{if the neuron is the first to fire} \\ -\frac{\eta}{N_{\text{neuron}} - 1} & \text{otherwise} \end{cases} \quad (6)$$

with η the learning rate, t_{fire} the timestamp of the spike fired by the neuron, t_{obj} the objective timestamp and N_{neuron} the number of output neurons. Setting large initial values for the thresholds may prevent the neurons from firing. In the absence of neuronal activity, no learning nor threshold adaptation can be performed. It is therefore preferable to initialize the thresholds with small values to promote neuronal activity within the network.

3.7. Spike to feature conversion

Finally, spikes need to be transformed back into feature values that will be fed to the classifier. Feature values are computed as follows:

$$f_i = 1.0 - \frac{t - T_{\text{output}_{\min}}}{T_{\text{output}_{\max}} - T_{\text{output}_{\min}}} \quad (7)$$

with f_i the i -th output feature value, t the spike timestamp (if no spike occurs, then t is set to $T_{\text{output}_{\max}}$), $[T_{\text{output}_{\min}}, T_{\text{output}_{\max}}]$ the range of possible timestamps. $T_{\text{output}_{\min}}$ can be computed as $T_{\text{input}_{\min}} + d_{\min}$, with d_{\min} , and $T_{\text{output}_{\max}}$ as $T_{\text{input}_{\max}} + d_{\max}$.

4. Learning visual features with sparse auto-encoders

Auto-encoders (AEs) [10] are unsupervised neural networks that learn latent representations that allow to best reconstruct the input data. In this work, among all unsupervised feature learning algorithms, we only consider single-layer AEs, for two reasons. First, they belong to the family of neural networks, as SNNs do, and, within this family, they are one of the most representative models for unsupervised learning (its main competitor being RBMs, which have been shown to optimize a similar criterion [35] and yield comparable performance for visual feature learning [7]). Then, we restrict our approach to single-layer networks, as multi-layer SNNs are only starting to emerge [31]; we believe one-layer SNNs should be well mastered before addressing multi-layer architectures.

The typical architecture of an AE is organized in two parts: an encoder enc mapping the input X to its latent representation $Z = \text{enc}(X)$, and a decoder dec computing a reconstruction \tilde{X} of the input from its latent representation: $\tilde{X} = \text{dec}(Z) = \text{dec}(\text{enc}(X))$. The objective function (Eq. 2) thus becomes:

$$\theta^* = \arg \min_{\theta} L(X, \tilde{X}; \theta) \quad (8)$$

where $L(., .; \theta)$ is some measure of the dissimilarity between the input X and its reconstruction \tilde{X} given the model parameterized by θ ; in other words, the auto-encoder aims at reconstructing its input with minimal reconstruction error. In our experiments, the Euclidean distance measures the reconstruction error.

The encoder and the decoder can be defined as single-layer or multilayer (in the case of stacked AEs) neural networks. In the following, we will consider only single-layer models of this form:

$$\begin{aligned} Z &= \text{enc}(X) = \sigma(W_{\text{enc}}X + b_{\text{enc}}) \\ \text{dec}(Z) &= W_{\text{dec}}Z + b_{\text{dec}} \end{aligned} \quad (9)$$

where $W_{\text{enc}} \in \mathbb{R}^{h \times w, n_f}$ (resp. $W_{\text{dec}} \in \mathbb{R}^{n_f, h \times w}$) is the weight matrix of the connections in the encoder (resp. the decoder), $b_{\text{enc}} \in \mathbb{R}^{n_f}$ (resp. $b_{\text{dec}} \in \mathbb{R}^{h \times w}$) is the bias vector of the encoder (resp. the decoder), and $\sigma(\cdot)$ is some activation

function³, in our case the sigmoid activation function $\sigma(x) = \frac{1}{1+e^{-x}}$. The output of the encoder correspond to the visual features learned by the AE: $f = \text{enc}(X)$.

To make the AE learn useful representations, the initial approach was to impose an information bottleneck on the model, by learning representations with dimensionalities lower than the ones of the input data ($n_f < h \times w$). However, such representations cannot capture the richness of the visual information, so current approaches rather use over-complete ($n_f > h \times w$) representations. In this case, some additional constraints must be enforced on the model to prevent it from learning trivial solutions, e.g., the identity function. These constraints generally take the form of an additional term in the objective function, for instance: weight regularization, explicit sparsity constraints (sparse AEs [7; 36], k-sparse AEs [37]) or regularization of the Jacobian of the encoder output Z (contractive AEs [46]). Another approach is to change the objective function from reconstruction to another criterion, for instance data denoising [35].

In this paper, we will consider sparse AEs as a baseline to assess the performance of STDP-based feature learning. More recent models (denoising AEs [35], contractive AEs [46], etc.) can reach better performance, but sparse AEs are closer to current STDP-based SNNs, which also feature explicit sparsity constraints, usually through lateral inhibition. Also, it allows us to set a minimum bound that SNNs should at least reach to be competitive with regular feature learning algorithms, and identify some directions to follow to achieve this goal; it constitutes a first step before taking STDP-based SNNs further.

In the following, we describe the weight regularization and sparsity constraint terms that we used in our experiments :

- L2 weight regularization: $\frac{\lambda}{2}(\|W_{\text{enc}}\|_2^2 + \|W_{\text{dec}}\|_2^2)$, where $\|\cdot\|_2$ denotes the Frobenius norm and λ is the weight decay parameter;
- sparsity term [36]: $\gamma \cdot \text{KL}(\hat{\rho}||\rho)$, where ρ is the desired sparsity level of the system, $\hat{\rho}$ is the vector of average activation values of the hidden neurons

³We only consider models where no activation function is applied to the decoder output as we work on continuous (image) data in [0, 1].

over a batch, $\text{KL}(\cdot|\cdot)$ the Kullback-Liebler divergence, and γ the weight applied to the sparsity term in the objective function.

This yields the final objective function for the AE:

$$L(X, \tilde{X}; \theta) = \frac{1}{2} \|X - \tilde{X}\|_2^2 + \frac{\lambda}{2} (\|W_{\text{enc}}\|_2^2 + \|W_{\text{dec}}\|_2^2) + \gamma \cdot \text{KL}(\hat{\rho}|\rho) \quad (10)$$

5. Experiments

5.1. Experimental protocol

The SNN and AE architectures used in our experiments are single-layer networks with n_f hidden units (see Figure 5). We follow the experimental protocol proposed by Coates *et al.* [7] to compare unsupervised feature extractors. It is organized in two stages, described below: visual feature learning, and the evaluation of the learned features on image classification benchmarks.

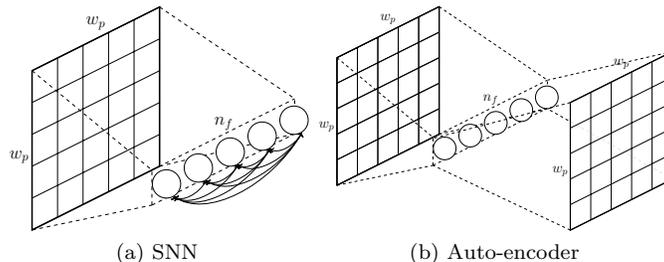


Figure 5: (a) SNN architecture used in the experiments. Solid arrows denote inhibitory connections between hidden units. (b) AE architecture used in the experiments.

Feature learning. From the training image dataset $\mathcal{I} = (I_1, I_2, \dots, I_n)$, we randomly sample n_p patches of size $w_p \times w_p$. The patches are fed to the feature learning algorithm for training, to produce a dictionary of n_f features.

Image recognition. The learned feature dictionary is used to produce image descriptors that are fed to a classifier following this process (Figure 6):

1. Image patches of size $w_p \times w_p$ are densely sampled from the images with stride s , producing $k \times k$ patches per image (Figure 6a).

2. Patches are fed to the feature extractor, producing $k \times k$ feature vectors of dimension n_f per image, organized into feature maps where each position corresponds to one patch of the input image (Figure 6b).
3. We apply sum pooling over a grid of size $r \times r$: the feature vectors of the patches within each grid cell are summed to produce a unique vector of size n_f per cell. These vectors are then concatenated to produce a single feature vector of size $r \times r \times n_f$ for each image (Figure 6c).
4. The feature vectors of the images are fed to a linear support vector machine (SVM) for training (training set) or classification (test set).

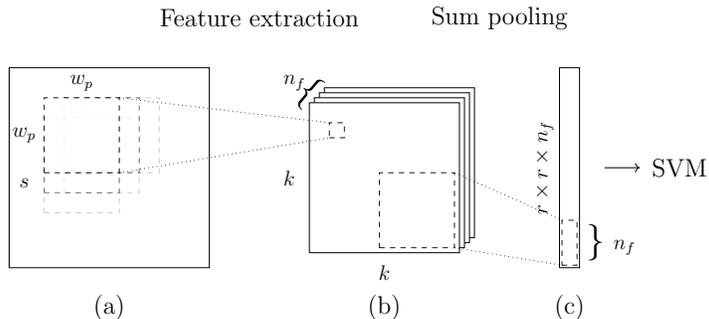


Figure 6: Experimental protocol. (a) Input image, where $k \times k$ patches of size $w_p \times w_p$ are extracted with a stride s . (b) n_f feature maps of size $k \times k$ produced by the feature extractor from its dictionary of n_f features. (c) Output vector constructed by sum pooling over $r \times r$ regions of the feature maps.

5.2. Datasets

We perform experiments on three datasets commonly used to evaluate unsupervised feature learning algorithms: CIFAR-10, CIFAR-100, and STL-10. Table 1 provides the properties of these datasets. Since previous work evaluated SNNs only on grayscale images, we also use grayscale versions of the three datasets, referred to as CIFAR-10-bw, CIFAR-100-bw, and STL-10-bw.

Contrary to MNIST, which is the preferred dataset in the SNN literature [42], these datasets provide color images of actual objects rather than just binary images of digits. It makes it possible to evaluate SNNs in more realistic conditions, in terms of data richness and importance of image pre-processing. Also, unlike

Dataset	Resolution	# classes	Training set size	Test set size
CIFAR-10 [47]	32×32	10	50,000	10,000
CIFAR-100 [47]	32×32	100	50,000	10,000
STL-10 [7]	96×96	10	5,000	8,000

Table 1: Properties of the datasets used in the experiments.

Data	n_f	ρ	γ	λ
Color	64	0.005	0.5	10^{-4}
	1024	0.005	0.1	10^{-5}
Grayscale	64	0.01	0.05	10^{-5}
	1024	0.005	0.1	10^{-5}

Table 2: AE parameters used in the experiments.

MNIST, but also other datasets such as NORB, they are not solved or nearly-solved problems (classification accuracy above 95%), so the results can highlight better the properties of the algorithms.

5.3. Implementation details

We use image patches of size 5×5 pixels ($w_p = 5$) and a stride $s = 1$. We evaluate the algorithms with two sizes of feature dictionaries, $n_f = 64$ and $n_f = 1024$. To produce final image descriptors, features are pooled over 2×2 image regions ($r = 2$), yielding image descriptors of size $4 \times n_f$.

We used a grid search to find the optimal parameters for the AEs and we only report results for the best configuration for each experimental setting. Table 2 provides the values of the parameters that we retained. They were consistently optimal over datasets. The AEs are trained for 1,000 epochs on 200,000 random patches from the training set considered. We use the Adadelta optimizer [48] with an initial learning rate $lr = 1.0$. AEs are implemented using TensorFlow.

Table 3 provides the parameters used to train SNNs. These parameters values were obtained using a greedy search: the optimal value of each parameter was searched while the values of all other parameters were fixed. The large number of parameters in this model did not allow us to perform a full grid search on all parameters. All SNN models are trained on 100,000 random patches from the training sets for 100 epochs.

Neuron				Neural Coding			
$V_{th}(0)$	20 mv	V_{rest}	0 mv	$T_{duration}$	1.0		
STDP				Threshold Adaptation			
W_{min}	0.0	W_{max}	1.0	t_{obj}	0.7	η	0.001
d_{min}	0.0	d_{max}	0.01	Pre-processing			
α_+	0.001	α_-	0.001	DoG_{center}	1.0	$DoG_{surround}$	2.0
β_+	1.0	β_-	1.0	DoG_{size}	7		

Table 3: SNN parameters used in the experiments.

Classification was performed using LibSVM [49] with a linear kernel and default parameters. All reported accuracies are averaged over three runs of the feature learning algorithms.

5.4. Color processing with SNNs

We first evaluate the strategies to encode color discussed in Section 3.3: images are first encoded using one of these strategies, then on-center/off-center coding is applied to each channel. Table 4 shows the classification accuracies yielded by each strategy on every dataset. Both color coding techniques, biological channels and RGB opponent channels, provide similar accuracies. However, using grayscale images yields better results than color images. This is counter-intuitive, since color images contain all the information available from grayscale images. Since the SNN processes all inputs in the same way, on-center/off-center coding must cause this information loss. However, this preprocessing step is currently required to extract edges from the images and feed the SNN inputs with spike trains that represent specific visual information. Training an SNN directly from RGB images could be an alternative, but is very challenging, because the active pixels in the patterns to learn can vary (from zero for a black pattern to the size of a patch when all pixels have the maximum intensity); it cannot be handled by existing homeostasis models. Figure 7 shows examples of filters learned from raw RGB images; since the network has a single layer, the filter image for one neuron is obtained by simply interpreting the normalized weights of its input synapses as RGB values. Many filters converge towards similar or uninformative patterns. It results mostly in dead units and repeated features.

Dataset	Color coding	$n_f = 64$	$n_f = 1024$
CIFAR-10	RGB opponent	37.66 ± 0.73	45.04 ± 0.06
	Bio-color	37.53 ± 0.33	43.54 ± 0.07
	Grayscale	45.37 ± 0.13	52.78 ± 0.41
	Grayscale + color	48.27 ± 0.47	56.93 ± 0.59
CIFAR-100	RGB opponent	17.14 ± 0.22	19.87 ± 0.03
	Bio-color	17.06 ± 0.09	19.19 ± 0.35
	Grayscale	18.43 ± 0.34	22.67 ± 0.36
	Grayscale + color	25.20 ± 0.76	30.44 ± 0.48
STL-10	RGB opponent	44.13 ± 1.30	51.20 ± 0.30
	Bio-color	44.23 ± 0.41	50.95 ± 0.08
	Grayscale	44.66 ± 0.87	51.40 ± 0.69
	Grayscale + color	49.20 ± 1.04	54.34 ± 0.30

Table 4: Classification accuracy (%) w.r.t. to the color coding strategy.

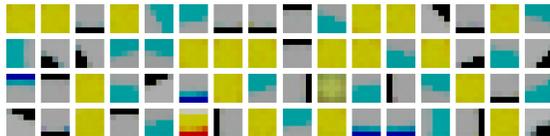


Figure 7: Examples of SNN features learned on raw RGB pixels (trained on CIFAR-10). They are mostly dead units or simple repeated patterns.

Finally, we evaluate the combination of color and grayscale images by training half of the features on each input independently. Results in Table 4 show that it performs best, showing that DoG-filtered color images still contain information that grayscale DoG-filtered images do not contain. In the remaining of the paper, all the runs performed on color images use this strategy.

5.5. SNNs versus AEs

The classification accuracies for each feature learning algorithm and dataset are reported in Table 5. AEs perform consistently better than SNNs⁴. So, how to bridge the gap between STDP learning and standard neural network approaches? Several elements may explain the performance of STDP. The results reported in Table 5 show two trends. First, working with colors always yields

⁴Since we chose to use a simple sparse AE (see Section 4), the actual gap between SNNs and state-of-the-art models should be larger than what these experiments show.

better results than working with grayscale images; a straightforward explanation is that color is significant to recognize objects in the datasets considered, either because natural objects (e.g., animals) represented in the datasets have a limited, meaningful set of colors, either because the contexts of the objects (e.g., the sky behind airplanes) have meaningful colors. The second trend is that the performance gap between SNNs and AEs is larger on color images than on grayscale images, showing that SNNs cannot handle color well, at least not with the straightforward color coding techniques that were used in the experiments. This result highlights the importance of color in object recognition, and therefore the need for a more efficient neural coding of color in SNNs.

Dataset	SNN		AE	
	$n_f = 64$	$n_f = 1024$	$n_f = 64$	$n_f = 1024$
CIFAR-10	48.27±0.47	56.93±0.59	57.56±0.08	66.98±0.33
CIFAR-10-bw	45.37±0.13	52.77±0.41	53.69±0.34	59.50±0.17
CIFAR-100	25.20±0.76	30.45±0.48	37.71±0.19	36.43±0.29
CIFAR-100-bw	18.43±0.34	22.67±0.36	23.62±0.18	26.56±0.05
STL-10	49.20±1.04	54.34±0.30	52.28±0.47	55.74±0.25
STL-10-bw	44.66±0.87	51.40±0.69	50.63±0.23	52.88±0.29

Table 5: Average classification accuracy (%) and its standard deviation w.r.t. to the datasets and feature learning algorithms.

Looking at the filters learned by SNNs and AEs provides additional information about the properties of features learned by STDP and potential reasons for the performance gap. Figures 9 and 10 show samples of filters learned by SNNs and AEs, respectively. The filters are different in nature. Filters learned by STDP are mostly edges, and some blobs, that are well-defined, with one or two dominant colors. By contrast, AEs learn more complex features; edges and blobs can still be observed, but they include a larger range of color or gray levels and are not as elementary as the ones learned by SNNs. Simple, well-defined features like the ones learned by STDP are conceptually pleasing because they represent elementary object shapes that can easily be understood. They suggest better generalization abilities from the feature extractor, and correspond to biological observations [31]. However, they are not as effective in practice. AEs

can also produce features closer to the ones obtained with SNNs (although with larger ranges of tones and intensities), but we could observe such features only by increasing the weight of L2 regularization, usually at some cost in accuracy.

The specific looks of SNN features can be explained in two ways. First, the use of on-center/off-center coding as a preprocessing step biases the models towards edge-like filters, as it highlights the edges in the images. Moreover, the learned features contain exclusively black or saturated colors because STDP rules tend towards a saturating regime for weights: once a given unit has learned a pattern, repeated expositions to this pattern will reinforce the sensitivity to this pattern until the weights reach either 1 or 0. This is illustrated in Figure 8a, which shows the distribution of weights in an SNN after training: most weight values are close to 0 or 1. Since AEs perform better and have more staggered weights, one may believe that saturated weights are detrimental to the performance of SNNs. To check this, we performed experiments with different values for the parameters β_+ and β_- : increasing their values allow the weights to "escape" more easily from their limit values W_{\min} and W_{\max} . Figure 8b shows that the weights are indeed more staggered, but the classification accuracy decreases as β_+ and β_- get larger (see Table 6). So, the fact that STDP leads to saturated weights is not the reason for the performance gap with AEs.

Dataset	$\beta = 1$	$\beta = 2$	$\beta = 3$	$\beta = 4$
CIFAR-10	48.27±0.47	46.56±0.68	43.18±1.60	41.03±0.21
CIFAR-10-bw	45.37±0.13	44.55±0.57	41.74±1.50	38.90±1.57

Table 6: SNN recognition rate according to STDP β parameter ($n_f = 64$).

Finally, the filters shown in Figure 9 also show a good property of SNNs: they do not raise any dead units, i.e. features that get stuck in a state with average weights that do not correspond to any significant pattern. By contrast, AEs tend to learn a fair amount of such features, especially when the number of features increases (see Figure 10). This behavior of SNNs can be due to two factors: lateral inhibition, which prevents neurons from learning similar patterns (here, becoming dead units), and the saturated regime of STDP.

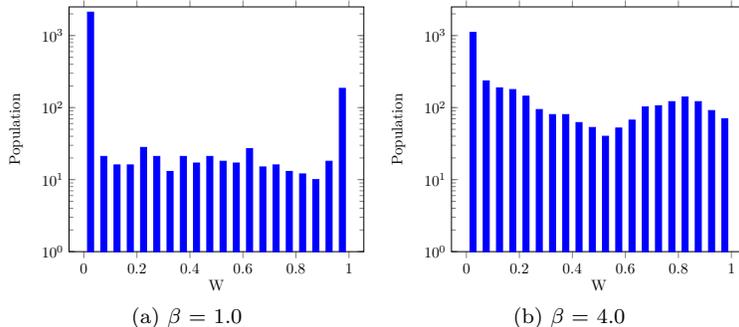


Figure 8: Distribution of weights (log. scale) in an SNN ($n_f = 64$) after training w.r.t. β . Most weights have values close to 0 or 1 when β decreases.

6. Result Analysis and Properties of the Networks

6.1. On-center/off-center coding

In this section, we investigate the impact of on-center/off-center coding on classification accuracy. As mentioned in Section 5.5, this image coding is responsible for the type of visual features learned by STDP, but does it impact the final accuracy of the system? We compared the accuracy of two systems, each with and without preprocessing images: an AE, under the same protocol as before, and an SVM performing classification directly from image pixels. The AE parameters for the on-center/off-center coding runs are: $\rho = 0.005$, $\gamma = 1.0$, and $\lambda = 10^{-4}$. Results on CIFAR-10 and CIFAR-10-bw are reported in Table 7. Using on-center/off-center coding decreases the accuracy of the classification in both configurations, which confirms that this coding is one of the causes of the limited performance of SNNs in image classification. This is due to the fact that extracting edges with DoG has the effect of selecting only a subrange of frequencies. In addition, the accuracies obtained on filtered color images are only on par with (in the case of AEs) or worse than (with SVM) the results obtained using grayscale images; it highlights the fact that on-center/off-center coding cannot handle color effectively. One reason is that edge information is effectively represented by grayscale pixels, and the additional information brought by color is essentially located in uniform image regions. Interestingly, the un-

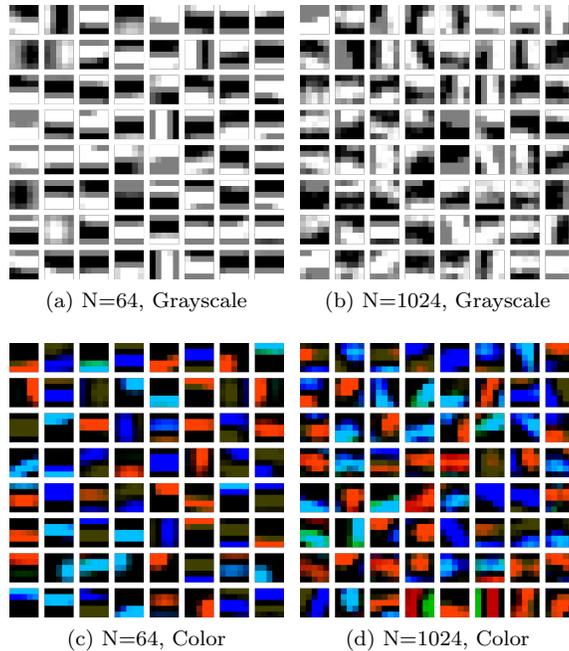


Figure 9: Grayscale and color filters learned by SNNs on CIFAR-10-bw and CIFAR-10. For $n_f = 1024$, random samples are shown.

supervised SNN models of the literature that are competitive with traditional approaches are only evaluated on the MNIST dataset [42], which does not require on-center/off-center coding as the images are only made of edges (white handwritten digits on black backgrounds).

Therefore, to be effective, SNNs require the design of a suited image coding that preserves as much visual information as possible. Using alternative methods to extract edges (such as the image gradient or the image Laplacian) could capture slightly different types of edge information, which could be processed within a single SNN for improved performance, in a feature fusion approach. However, this would only process edge information, which is insufficient to reach optimal classification performances. Ideally, SNNs should be able to handle raw RGB pixels; however, this is not straightforward, as we showed in Section 5.4.

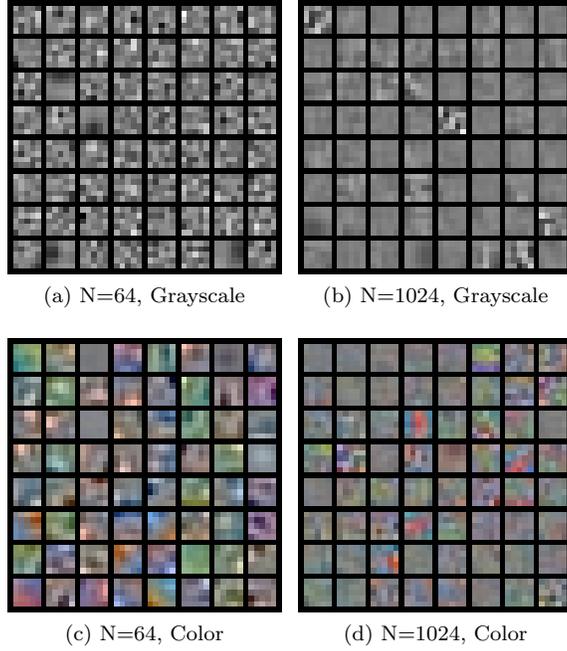


Figure 10: Grayscale and color filters learned by AEs on CIFAR-10-bw and CIFAR-10. For $n_f = 1024$, random samples are shown.

6.2. Sparsity

We investigate here the sparsity properties of SNNs and AEs. To do so, we use the following sparseness measure [50]:

$$\text{sp}(f) = \frac{\sqrt{n_f} - \frac{\sum_i^{n_f} |f_i|}{\sqrt{\sum_i^{n_f} f_i^2}}}{\sqrt{n_f} - 1} \quad (11)$$

where f is the vector of activations of hidden units (i.e. the visual feature vector) and n_f is the number of hidden units. $\text{sp}(f) \in [0, 1]$; larger values indicate sparser activations.

Table 8 shows the mean sparseness of features computed on the test set of CIFAR-10. The sparseness is much higher in SNNs than in AEs. Indeed, the specialization of features in SNNs relies mostly on lateral inhibition, which prevents units from integrating spikes, leading to very sparse activations of the

Dataset	Raw pixels	AE features $n_f = 64$
CIFAR-10	37.79	57.56±0.08
CIFAR-10-dog	21.07	52.65±0.30
CIFAR-10-bw	28.38	53.69±0.34
CIFAR-10-dog-bw	25.29	52.76±0.08

Table 7: Classification accuracy (%) obtained with raw pixels and AE features w.r.t. pre-processing methods. Only one run is performed on raw pixels as SVM training is deterministic.

features. Sparsity is often cited as a necessary condition for good representations [1], and has been shown to be correlated to classification accuracy on image datasets [36]. However, some results in [36] show that maximizing sparsity does not always lead to improvements in classification accuracy in AEs. Similarly, we observed experimentally that enforcing too much sparsity on the AEs (e.g., by lowering ρ) is detrimental to the classification accuracy. To push it further, we performed five runs on CIFAR-10 with $n_f = 64$ and different values for parameters λ , γ , and ρ . The AE parameters were set so that the sparseness would be close to the sparseness that we measured in SNNs (i.e., in the range [0.8;0.9]). In these runs, the classification accuracy varies from 35.53% to 41.03%, much lower than our 57.56% baseline. To check whether high levels of sparseness are an issue for SNNs too, we ran an experiment where lateral inhibition is deactivated during the feature extraction phase⁵. As expected, deactivating inhibition decreased the sparseness of the model (from 0.869 to 0.638 on CIFAR-10). However, the classification rate decreased too (from 48.27% to 47.35%). It shows that, although sparsity is a desirable feature for good representations, an excessive level of sparseness can be detrimental, and that the right amount of sparsity should be enforced during training. This calls for the use of other, less restrictive, inhibition strategies than WTA.

6.3. Coherence

One measure of the quality of the learned feature is their incoherence, i.e. the fact that one feature cannot be obtained by a sparse linear combination of other

⁵Inhibition is still maintained during feature training, because SNNs cannot converge if there is no competition between neurons, as explained in Section 3.5.

Model	$n_f = 64$	$n_f = 1024$
SNN	$0.869 \pm 1.96e-5$	$0.967 \pm 3.04e-5$
AE	0.352 ± 0.116	0.112 ± 0.077

Table 8: Mean and standard deviation of feature sparseness (test set of CIFAR-10).

features in the vocabulary. If the incoherence is low, features are redundant, which is harmful for classification as redundant features will overweight other features. Inspired by the measure introduced in [37], we measure the coherence μ_{ij} of two features f_i and f_j as their cosine similarity:

$$\mu_{ij} = \frac{|\langle f_i, f_j \rangle|}{\|f_i\|_2 \cdot \|f_j\|_2} \quad (12)$$

where f_i is the i -th feature, $\langle \cdot, \cdot \rangle$ is the dot product operator, and $\|\cdot\|_2$ is the L2 vector norm. $\mu \in [0, 1]$; 0 corresponds to orthogonal (incoherent) features and 1 to similar (coherent) features. The weights span different ranges of values depending on the feature extractor considered; feature normalization makes coherence measures comparable from one extractor to another.

Table 9 displays the mean and the standard deviation of coherence measure μ under all experimental settings. Overall, STDP-based SNNs produce more coherent features, which is one of the factors that can explain their lower performance. Moreover, the maximum pairwise coherence between two SNN feature is higher ($\max(\mu_{ij}) = 0.999$ in most experimental configurations) than the maximum coherence between AE-produced features ($\max(\mu_{ij}) \in [0.898, 0.998]$), i.e. SNNs can learn almost identical features; in AEs, such features mostly correspond to dead units, whereas in SNNs they are significant features that are repeated. This result shows the limits of WTA inhibition, which should prevent features from reacting to the same patterns but fails to do so in practice. This calls for more work on understanding inhibition mechanisms and designing inhibition models that better prevent the co-adaptation of features.

6.4. Objective Function

One issue with STDP learning is that the objective function optimized by the system is not explicitly expressed, unlike AEs, which minimize reconstruc-

Dataset	SNN		Autoencoder	
	$n_f = 64$	$n_f = 1024$	$n_f = 64$	$n_f = 1024$
CIFAR-10	0.252±0.252	0.285±0.249	0.154±0.144	0.145±0.109
CIFAR-10-bw	0.313±0.271	0.340±0.234	0.119±0.138	0.225±0.161
CIFAR-100	0.256±0.230	0.289±0.238	0.154±0.149	0.143±0.199
CIFAR-100-bw	0.320±0.238	0.343±0.223	0.121±0.137	0.234±0.166
STL-10	0.263±0.293	0.293±0.246	0.177±0.164	0.151±0.114
STL-10-bw	0.263±0.293	0.293±0.246	0.119±0.132	0.236±0.169

Table 9: Mean and standard deviation of feature coherence μ under all experimental settings.

tion error. Identifying the criteria that are optimized by STDP rules would help better understand the related learning process and design learning rules for specific tasks. In this section, we check whether STDP rules embed reconstruction as a training criteria, by checking whether features learned through STDP are suited for image reconstruction, as those learned by AEs do. To do so, we reconstructed the test images from the visual features. First, we reconstruct individual patches: in AEs, the reconstructed patches are directly provided by the decoder; in SNNs, patches are reconstructed as a linear combination of the filters weighted by their activations for the current sample, like in an AE with tied weights. Images are reconstructed from patches by averaging the values of overlapping patches at each location. Table 10 shows the reconstruction error of each feature extractor on the test set of CIFAR-10, computed as the sum of squared errors between input images and reconstructed images, averaged over the samples. The reconstruction error is much higher for SNNs than AEs, which suggests that STDP does not learn features that allow reconstruction. However, qualitatively, the results look different (see samples in Figure 11): the edges of the objects are reconstructed, although with less details than in the original images, but the global illumination is degraded. The degradation of pixel intensities explains for a large part the increased reconstruction error. This is best illustrated by the best and worst reconstructions (in the sense of the MSE) that we obtained using SNNs (see Figure 12): edges are reconstructed correctly in both, but not pixel intensities. The reason for this is that SNNs process DoG-filtered images, in which color intensities are discarded and only edge in-

formation is retained. One could expect the reconstruction error of SNNs to be much lower if they were able to process raw images directly. Also, the lack of details around the edges could be blamed on the learned features being too elementary and sparse, which prevents the reconstruction of complex patterns.

These results show that, although this is not explicit in the learning rules, STDP learns to reconstruct images, among other potential criteria. However, it is known that minimizing reconstruction error is not sufficient to provide meaningful representations [35]. This is why recent AE models include additional criteria such as sparsity penalties [36] or Jacobian regularization [46]. How such criteria could be implemented within STDP rules, as well as which other criteria are already embedded in the STDP rules, are still open questions.

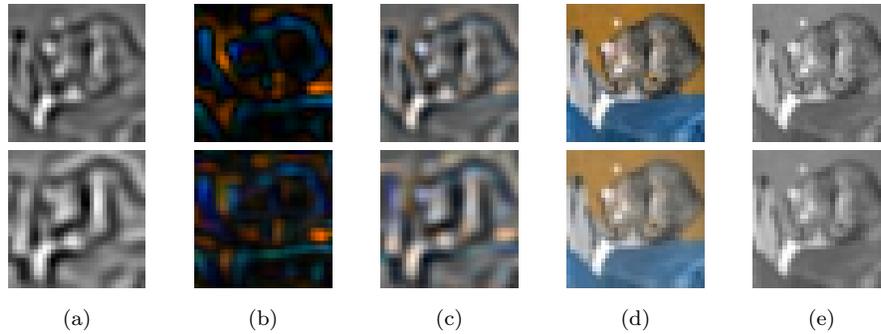


Figure 11: Image reconstruction samples from the test sets of CIFAR-10 and CIFAR-10-bw (top: pre-processed input images, bottom: reconstructed images). (a) SNN features, DoG-filtered grayscale image (b) SNN features, DoG-filtered color image (c) SNN features, grayscale and color DoG-filtered image (d) AE filters, color image (e) AE filters, grayscale image.

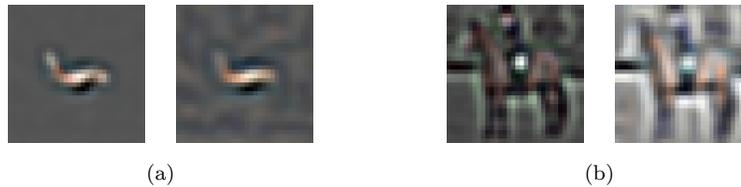


Figure 12: (a) Best (error: 1.60) and (b) worst (error: 15.38) reconstructions from SNN features from the test set of CIFAR-10 (left: input images, right: reconstructions).

	SNN		AE	
	$n_f = 64$	$n_f = 1024$	$n_f = 64$	$n_f = 1024$
CIFAR-10	4.9429	4.4179	0.0802	0.0742
CIFAR-10-bw	4.9797	4.4628	0.00407	0.00472

Table 10: Average reconstruction errors on the test set of CIFAR-10.

7. Conclusion

In this paper, we compared spiking neural networks (SNNs) equipped with STDP to auto-encoders (AEs) for unsupervised visual feature learning. Experiments on three image classification datasets showed that STDP cannot currently compete with classical neural networks trained with gradient descent, but also highlighted a number of properties of SNNs and provided specific directions towards effective feature learning with SNNs. Specifically, we showed that:

- STDP-based SNNs are unable to deal naturally with RGB images;
- the common on-center/off-center image coding used in SNNs results in an information loss, thus harming the classification accuracy; this information loss is even more pronounced on color images;
- winner-take-all inhibition results in overly sparse features and does not prevent the co-adaptation of features in practice;
- STDP-based learning rules produce features that enable to reconstruct images from the learned features, as AEs do, even though the features are not explicitly optimized for this task.

These conclusions suggest two research directions to bridge the gap with standard neural networks. The first is to address the information loss caused by on-center/off-center coding and the conversion to grayscale images. One straightforward solution is to use standard pre-processing methods that have shown to be effective in computer vision, such as applying DoG filters in a scale-space fashion [2], or replacing them with whitening [7], which highlights edges while retaining some color information. However, the resulting filters

would still be limited to edge regions; dealing with non-edge information would require the design of alternatives to DoG filtering or new neural coding schemes that are not based on contrast only. The second direction is to ensure that the features have adequate amounts of sparsity and redundancy. To do so, one could design more suited inhibition mechanisms. Such mechanisms should be "soft", i.e. allow several neurons to fire at the same time, as demonstrated in [37] in the case of AEs. The methods in [37] could be good candidates but should be adapted to preserve the locality of computations. Another option is to design adaptive inhibition rules that maintain the homeostasis of the system w.r.t to some sparsity and co-adaptation objectives. Instead of inhibition rules, or in addition to them, learning rules with homeostatic plasticity [51] can allow to reach the targeted sparsity and co-adaptation levels by acting on the synapses rather than the states of the neurons.

Acknowledgments

This work was supported in part by IRCICA (Univ. Lille, CNRS, USR 3380 – IRCICA, F-59000 Lille, France) under the Bioinspired Project.

References

- [1] Y. Bengio, A. Courville, P. Vincent, Representation learning: A review and new perspectives, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35 (8) (2013) 1798–1828.
- [2] D. G. Lowe, Distinctive Image Features from Scale-Invariant Keypoints, *International Journal of Computer Vision* 60 (2) (2004) 91–110.
- [3] K. Mikolajczyk, C. Schmid, A performance evaluation of local descriptors, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 27 (10) (2005) 1615–1630.
- [4] G.-S. Xie, X.-Y. Zhang, W. Yang, M. Xu, S. Yan, C.-L. Liu, LG-CNN: From local parts to global discrimination for fine-grained recognition, *Pattern Recognition* 71 (2017) 118–131.

- [5] F. Liu, G. Lin, C. Shen, CRF learning with CNN features for image segmentation, *Pattern Recognition* 48 (10) (2015) 2983–2992.
- [6] Z. Tu, W. Xie, Q. Qin, R. Poppe, R. C. Veltkamp, B. Li, J. Yuan, Multi-stream CNN: Learning representations based on human-related regions for action recognition, *Pattern Recognition* 79 (2018) 32–43.
- [7] A. Coates, A. Ng, H. Lee, An analysis of single-layer networks in unsupervised feature learning, in: *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 215–223, 2011.
- [8] D. Wang, X. Tan, Unsupervised feature learning with C-SVDDNet, *Pattern Recognition* 60 (2016) 473–485.
- [9] Y. Yuan, J. Wan, Q. Wang, Congested scene classification via efficient unsupervised feature learning and density estimation, *Pattern Recognition* 56 (2016) 159–169.
- [10] H. Bourlard, Y. Kamp, Auto-association by multilayer perceptrons and singular value decomposition, *Biological Cybernetics* 59 (4-5) (1988) 291–294.
- [11] P. Smolensky, *Information Processing in Dynamical Systems: Foundations of Harmony Theory*, vol. 1, chap. 6, MIT Press, Cambridge, 194–281, 1986.
- [12] Y. Bengio, P. Lamblin, D. Popovici, H. Larochelle, Greedy layer-wise training of deep networks, in: *Advances in Neural Information Processing Systems*, 153–160, 2007.
- [13] G. E. Hinton, S. Osindero, Y.-W. Teh, A fast learning algorithm for deep belief nets, *Neural Computation* 18 (7) (2006) 1527–1554.
- [14] S. Zhang, J. Wang, X. Tao, Y. Gong, N. Zheng, Constructing deep sparse coding network for image classification, *Pattern Recognition* 64 (2017) 130–140.

- [15] P. Tang, X. Wang, Z. Huang, X. Bai, W. Liu, Deep patch learning for weakly supervised object classification and discovery, *Pattern Recognition* 71 (2017) 446–459.
- [16] R. Bekkerman, M. Bilenko, J. Langford, Scaling Up Machine Learning: Parallel and Distributed Approaches, in: *Proceedings of the 17th ACM SIGKDD International Conference Tutorials*, 4, 2011.
- [17] Y. Cao, Y. Chen, D. Khosla, Spiking deep convolutional neural networks for energy-efficient object recognition, *International Journal of Computer Vision* 113 (1) (2015) 54–66.
- [18] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, Y. Bengio, Quantized neural networks: Training neural networks with low precision weights and activations, *Journal of Machine Learning Research* 18 (187) (2018) 1–30.
- [19] S. Tang, Deep Learning Processor List, URL <https://github.com/chethiya/Deep-Learning-Processor-List>, Accessed on 11/7/2018.
- [20] C. D. James, J. B. Aimone, N. E. Miner, C. M. Vineyard, F. H. Rothganger, K. D. Carlson, S. A. Mulder, T. J. Draelos, A. Faust, M. J. Marinella, et al., A historical survey of algorithms and hardware architectures for neural-inspired and neuromorphic computing applications, *Biologically Inspired Cognitive Architectures* 19 (2017) 49–64.
- [21] S. K. Esser, P. A. Merolla, J. V. Arthur, A. S. Cassidy, R. Appuswamy, A. Andreopoulos, D. J. Berg, J. L. McKinstry, T. Melano, D. R. Barch, C. di Nolfo, P. Datta, A. Amir, B. Taba, M. D. Flickner, D. S. Modha, Convolutional networks for fast, energy-efficient neuromorphic computing, *Proceedings of the National Academy of Sciences* 113 (41) (2016) 11441–11446.
- [22] Y. Bengio, A. Fischer, T. Mesnard, S. Zhang, Y. Wu, From STDP towards biologically plausible deep learning, in: *Deep Learning Workshop, International Conference on Machine Learning*, 2015.

- [23] F. Kamaruzaman, A. A. Shafie, Recognizing faces with normalized local Gabor features and spiking neuron patterns, *Pattern Recognition* 53 (2016) 102–115.
- [24] J. Dennis, H. D. Tran, H. Li, Combining robust spike coding with spiking neural networks for sound event classification, in: *International Conference on Acoustics, Speech and Signal Processing*, 176–180, 2015.
- [25] M.-J. Escobar, G. S. Masson, T. Vieville, P. Kornprobst, Action recognition using a bio-inspired feedforward spiking network, *International Journal of Computer Vision* 82 (3) (2009) 284.
- [26] J. H. Lee, T. Delbruck, M. Pfeiffer, Training deep spiking neural networks using backpropagation, *Frontiers in Neuroscience* 10 (2016) 508.
- [27] G.-G. Bi, M.-M. Poo, Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type, *Journal of Neuroscience* 18 (24) (1998) 10464–10472.
- [28] D. Querlioz, O. Bichler, C. Gamrat, Simulation of a memristor-based spiking neural network immune to device variations, in: *International Joint Conference on Neural Networks*, 1775–1781, 2011.
- [29] P. U. Diehl, M. Cook, Unsupervised learning of digit recognition using spike-timing-dependent plasticity, *Frontiers in Computational Neuroscience* 9 (2015) 99.
- [30] S. R. Kheradpisheh, M. Ganjtabesh, T. Masquelier, Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition, *Neurocomputing* 205 (2016) 382–392.
- [31] S. R. Kheradpisheh, M. Ganjtabesh, S. J. Thorpe, T. Masquelier, STDP-based spiking deep convolutional neural networks for object recognition, *Neural Networks* 99 (2018) 56–67.

- [32] M. Mozafari, S. R. Kheradpisheh, T. Masquelier, A. Nowzari-Dalini, M. Ganjtabesh, First-spike-based visual categorization using reward-modulated STDP, *IEEE Transactions on Neural Networks and Learning Systems* 29 (12) (2018) 6178–6190.
- [33] P. Falez, P. Tirilly, I. M. Bilasco, P. Devienne, P. Boulet, Mastering the Output Frequency in Spiking Neural Networks, in: *International Joint Conference on Neural Networks*, 1–8, 2018.
- [34] J. Philbin, M. Isard, J. Sivic, A. Zisserman, Descriptor learning for efficient retrieval, in: *European Conference on Computer Vision*, 677–691, 2010.
- [35] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, *Journal of Machine Learning Research* 11 (Dec) (2010) 3371–3408.
- [36] N. Jiang, W. Rong, B. Peng, Y. Nie, Z. Xiong, An empirical analysis of different sparse penalties for autoencoder in unsupervised feature learning, in: *International Joint Conference on Neural Networks*, 1–8, 2015.
- [37] A. Makhzani, B. Frey, k-Sparse autoencoders, in: *International Conference on Learning Representations*, 2014.
- [38] K. Gupta, A. Majumdar, Learning autoencoders with low-rank weights, in: *International Conference on Image Processing*, 3899–3903, 2017.
- [39] J. L. Krichmar, P. Coussy, N. Dutt, Large-scale spiking neural networks using neuromorphic hardware compatible models, *ACM Journal on Emerging Technologies in Computing Systems* 11 (4) (2015) 36.
- [40] B. Han, A. Sengupta, K. Roy, On the energy benefits of spiking deep neural networks: A case study, in: *International Joint Conference on Neural Networks*, 971–976, 2016.

- [41] E. M. Izhikevich, Which model to use for cortical spiking neurons?, *IEEE Transactions on Neural Networks* 15 (5) (2004) 1063–1070.
- [42] A. Tavanaei, M. Ghodrati, S. R. Kheradpisheh, T. Masquelier, A. S. Maida, Deep learning in spiking neural networks, *CoRR* abs/1804.08150.
- [43] M. S. Livingstone, D. H. Hubel, Anatomy and physiology of a color system in the primate visual cortex, *Journal of Neuroscience* 4 (1) (1984) 309–356.
- [44] R. Brette, Philosophy of the spike: rate-based vs. spike-based theories of the brain, *Frontiers in Systems Neuroscience* 9 (2015) 151.
- [45] R. VanRullen, S. J. Thorpe, Rate coding versus temporal order coding: What the retinal ganglion cells tell the visual cortex, *Neural Computation* 13 (6) (2001) 1255–1283.
- [46] S. Rifai, P. Vincent, X. Muller, X. Glorot, Y. Bengio, Contracting auto-encoders: Explicit invariance during feature extraction, in: *International Conference on Machine Learning*, 833–840, 2011.
- [47] A. Krizhevsky, Learning Multiple Layers of Features from Tiny Images, *Tech. Rep.*, University of Toronto, 2009.
- [48] M. D. Zeiler, ADADELTA: An adaptive learning rate method, *CoRR* abs/1212.5701.
- [49] C.-C. Chang, C.-J. Lin, LIBSVM: a library for support vector machines, *Transactions on Intelligent Systems and Technology* 2 (3) (2011) 27.
- [50] P. O. Hoyer, Non-negative matrix factorization with sparseness constraints, *Journal of Machine Learning Research* 5 (2004) 1457–1469.
- [51] A. Watt, N. Desai, Homeostatic Plasticity and STDP: Keeping a Neuron’s Cool in a Fluctuating World, *Frontiers in Synaptic Neuroscience* 2 (2010) 5.