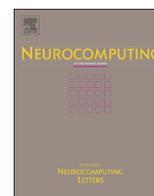




ELSEVIER

Contents lists available at ScienceDirect

Neurocomputing

journal homepage: www.elsevier.com/locate/neucom

Bio-inspired unsupervised learning of visual features leads to robust invariant object recognition



Saeed Reza Kheradpisheh^{a,e}, Mohammad Ganjtabesh^{a,*}, Timothée Masquelier^{b,c,d,e}

^a Department of Computer Science, School of Mathematics, Statistics, and Computer Science, University of Tehran, Tehran, Iran

^b INSERM, U968, Paris F-75012, France

^c Sorbonne Universités, UPMC Univ Paris 06, UMR-S 968, Institut de la Vision, Paris F-75012, France

^d CNRS, UMR-7210, Paris F-75012, France

^e CERCO UMR 5549, CNRS – Université de Toulouse, F-31300, France

ARTICLE INFO

Article history:

Received 17 October 2015

Received in revised form

17 March 2016

Accepted 18 April 2016

Communicated by Mingli Song

Available online 10 May 2016

Keywords:

View-Invariant object recognition

Visual cortex

STDP

Spiking neurons

Temporal coding

ABSTRACT

Retinal image of surrounding objects varies tremendously due to the changes in position, size, pose, illumination condition, background context, occlusion, noise, and non-rigid deformations. But despite these huge variations, our visual system is able to invariantly recognize any object in just a fraction of a second. To date, various computational models have been proposed to mimic the hierarchical processing of the ventral visual pathway, with limited success. Here, we show that the association of both biologically inspired network architecture and learning rule significantly improves the models' performance when facing challenging invariant object recognition problems. Our model is an asynchronous feedforward spiking neural network. When the network is presented with natural images, the neurons in the entry layers detect edges, and the most activated ones fire first, while neurons in higher layers are equipped with spike timing-dependent plasticity. These neurons progressively become selective to intermediate complexity visual features appropriate for object categorization. The model is evaluated on 3D-Object and ETH-80 datasets which are two benchmarks for invariant object recognition, and is shown to outperform state-of-the-art models, including DeepConvNet and HMAX. This demonstrates its ability to accurately recognize different instances of multiple object classes even under various appearance conditions (different views, scales, tilts, and backgrounds). Several statistical analysis techniques are used to show that our model extracts class specific and highly informative features.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Humans can effortlessly and rapidly recognize surrounding objects [1], despite the tremendous variations in the projection of each object on the retina [2] caused by various transformations such as changes in object position, size, pose, illumination condition and background context [3]. This invariant recognition is presumably handled through hierarchical processing in the so-called ventral pathway. Such hierarchical processing starts in V1 layers, which extract simple features such as bars and edges in different orientations [4], continues in intermediate layers such as V2 and V4, which are responsive to more complex features [5], and culminates in the inferior temporal cortex (IT), where the neurons are selective to object parts or whole objects [6]. By moving from the lower layers to the higher layers, the feature complexity,

receptive field size and transformation invariance increase, in such a way that the IT neurons can invariantly represent the objects in a linearly separable manner [7,8].

Another amazing feature of the primates' visual system is its high processing speed. The first wave of image-driven neuronal responses in IT appears around 100 ms after the stimulus onset [1,3]. Recordings from monkey IT cortex have demonstrated that the first spikes (over a short time window of 12.5 ms), about 100 ms after the image presentation, carry accurate information about the nature of the visual stimulus [7]. Hence, ultra-rapid object recognition is presumably performed in a feedforward manner [3]. Moreover, although there exist various intra- and inter-area feedback connections in the visual cortex, some neurophysiological [9,10,3] and theoretical [11] studies have also suggested that the feedforward information is usually sufficient for invariant object categorization.

Appealed by the impressive speed and performance of the primates' visual system, computer vision scientists have long tried to “copy” it. So far, it is mostly the architecture of the visual system that has been mimicked. For instance, using hierarchical

* Corresponding author.

E-mail addresses: kheradpisheh@ut.ac.ir (S.R. Kheradpisheh), mgtabesh@ut.ac.ir (M. Ganjtabesh), timothee.masquelier@alum.mit.edu (T. Masquelier).

feedforward networks with restricted receptive fields, like in the brain, has been proven useful [12–17]. In comparison, the way that biological visual systems learn the appropriate features has attracted much less attention. All the above-mentioned approaches somehow use non biologically plausible learning rules. Yet the ability of the visual cortex to wire itself, mostly in an unsupervised manner, is remarkable [18,19].

Here, we propose that adding bio-inspired learning to bio-inspired architectures could improve the models' behavior. To this end, we focused on a particular form of synaptic plasticity known as spike timing-dependent plasticity (STDP), which has been observed in the mammalian visual cortex [20,21]. Briefly, STDP reinforces the connections with afferents that significantly contributed to make a neuron fire, while it depresses the others [22]. A recent psychophysical study provided some indirect evidence for this form of plasticity in the human visual cortex [23].

In an earlier study [24], it is shown that a combination of a temporal coding scheme – where in the entry layer of a spiking neural network the most strongly activated neurons fire first – with STDP leads to a situation where neurons in higher visual areas will gradually become selective to complex visual features in an unsupervised manner. These features are both salient and consistently present in the inputs. Furthermore, as learning progresses, the neurons' responses rapidly accelerates. These responses can then be fed to a classifier to do a categorization task.

In this study, we show that such an approach strongly outperforms state-of-the-art computer vision algorithms on view-invariant object recognition benchmark tasks including 3D-Object

[25,26] and ETH-80 [27] datasets. These datasets contain natural and unsegmented images, where objects have large variations in scale, viewpoint, and tilt, which makes their recognition hard [28], and probably out of reach for most of the other bio-inspired models [29,30]. Yet our algorithm generalizes surprisingly well, even when “simple classifiers” are used, because STDP naturally extracts features that are class specific. This point was further confirmed using mutual information [31] and representational dissimilarity matrix (RDM) [32]. Moreover, the distribution of objects in the obtained feature space was analyzed using hierarchical clustering [33], and objects of the same category tended to cluster together.

2. Materials and methods

The algorithm we used here is a scaled-up version of the one presented in [24]. Essentially, many more C2 features and iterations were used. Our code is available upon request. We used a five-layer hierarchical network $S_1 \rightarrow C_1 \rightarrow S_2 \rightarrow C_2 \rightarrow \text{classifier}$, largely inspired by the HMAX model [14] (see Fig. 1). Specifically, we alternated simple cells that gain selectivity through a sum operation, and complex cells that gain shift and scale invariance through a max operation. However, our network uses spiking neurons and operates in the temporal domain: when presented with an image, the first layer's S_1 cells, detect oriented edges and the more strongly a cell is stimulated the earlier it fires. These S_1 spikes are then propagated asynchronously through the

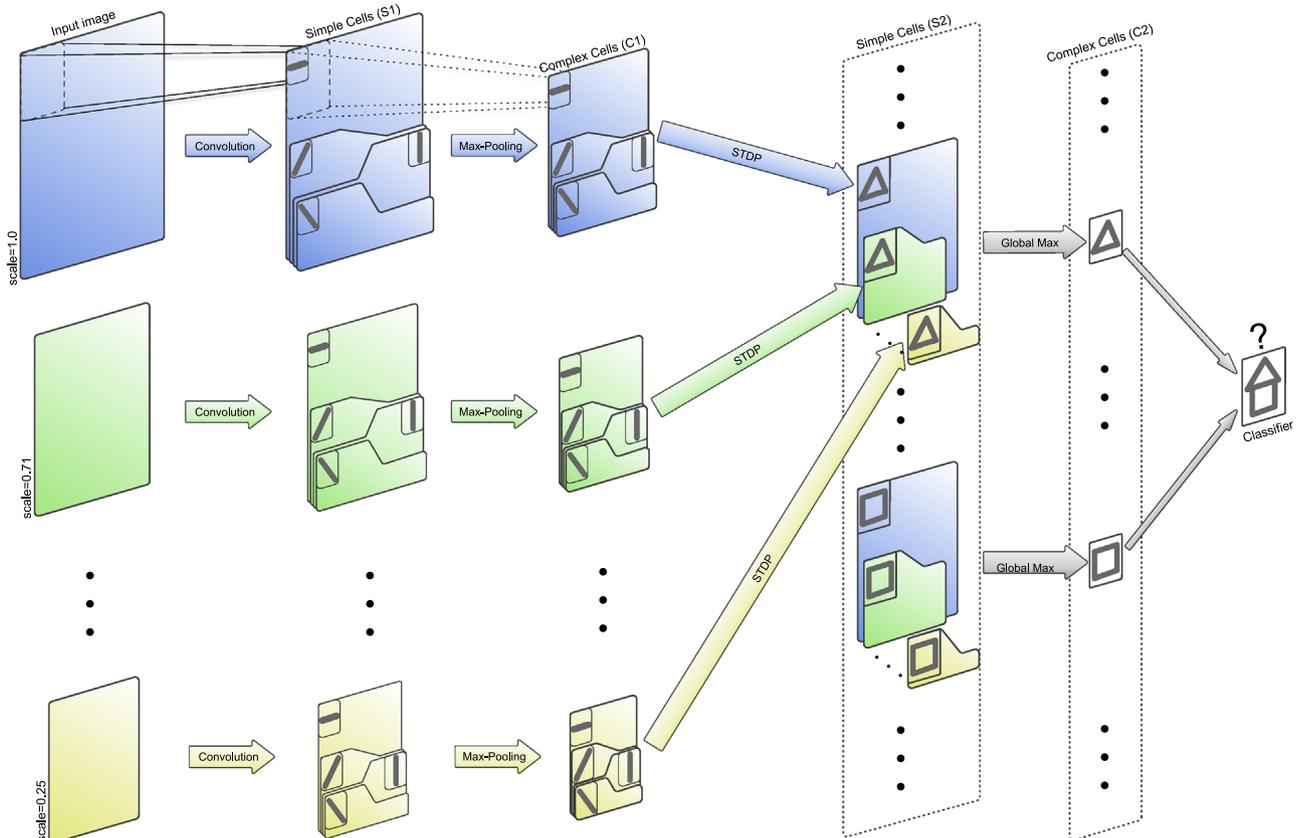


Fig. 1. Overview of our 5 layered feedforward spiking neural network. The network processes the input image in a multi-scale form, each processing scale is shown with a different color. Cells are organized in retinotopic maps until the S_2 layer (included). S_1 cells of each processing scale detect edges from the corresponding scaled image. C_1 maps sub-sample the corresponding S_1 maps by taking the maximum response over a square neighborhood. S_2 cells are selective to intermediate complexity visual features, defined as a combination of oriented edges of a same scale (here we symbolically represented a triangle detector and a square detector). There is one $S_1-C_1-S_2$ pathway for each processing scale. Then C_2 cells take the maximum response of S_2 cells over all positions and scales and are thus shift and scale invariant. Finally, a classification is done based on the C_2 cells' responses (here we symbolically represented a house/non-house classifier). C_1 to S_2 synaptic connections are learned with STDP, in an unsupervised manner. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

feedforward network. We only compute the first spike fired by each neuron (if any), which leads to efficient implementations. The justification for this is that later spikes are probably not used in ultra-rapid visual categorization tasks in primates [34]. We used restricted receptive fields and a weight sharing mechanism (i.e. convolutional network). In our model, images are presented sequentially and the resulting spike waves are propagated through to the S_2 layer, where STDP is used to extract diagnostic features.

More specifically, the first layer's S_1 cells detect bars and edges using Gabor filters. Here we used 5×5 convolutional kernels corresponding to Gabor filters with the wavelength of 5 and four different preferred orientations ($\pi/8, \pi/4 + \pi/8, \pi/2 + \pi/8, 3\pi/4 + \pi/8$). These filters are applied to five scaled versions of the original image: 100%, 71%, 50%, 30%, and 25% (each processing scale declared by a different color in Fig. 1). Hence, for each scaled version of the input image we have four S_1 maps (one for each orientation), and overall, there are $4 \times 5 = 20$ maps of S_1 cells (see the S_1 maps of Fig. 1). Evidently, the S_1 cells of larger scales detect edges with higher spatial frequencies while the smaller scales extract edges with lower spatial frequencies. Indeed, instead of changing the size and spatial frequency of Gabor filters, we are changing the size of input image. This is a way to implement scale invariance at a low computational cost.

Each S_1 cell emits a spike with a latency that is inversely proportional to the absolute value of the convolution. Thus, the more strongly a cell is stimulated the earlier it fires (intensity-to-latency conversion, as observed experimentally [35–37]). To increase the sparsity at a given scale and location (corresponding to one cortical column), only the spike corresponding to the best matching orientation is propagated (i.e. a winner-take-all inhibition is employed). In other word, for each position in the four S_1 orientation maps of a given scale, the S_1 cell with highest convolution value emits a spike and prevents the other three S_1 cells from firing.

For each S_1 map, there is a corresponding C_1 map. Each C_1 cell propagates the first spike emitted by the S_1 cells in a 7×7 square neighborhood of the S_1 map which corresponds to one specific orientation and one scale (see the C_1 maps of Fig. 1). C_1 cells thus execute a maximum operation over the S_1 cells with the same preferred feature across a portion of the visual field, which is a biologically plausible way to gain local shift invariance [38,39]. The overlap between the afferents of two adjacent C_1 cells is just one S_1 row, hence a subsampling over the S_1 maps is done by the C_1 layers as well. Therefore, each C_1 map has $6 \times 6 = 36$ fewer cells than the corresponding S_1 map.

S_2 features correspond to intermediate-complexity visual features which are optimum for object classification [40]. Each S_2 feature has a prototype S_2 cell (specified by a C_1 - S_2 synaptic weight matrix), which is a weighted combination of bars (C_1 cells) with different orientations in a 16×16 square neighborhood. Each prototype S_2 cell is retinotopically duplicated in the five scale maps (i.e. weight-sharing is used). Within those maps, the S_2 cells can integrate spikes only from the four C_1 maps of their corresponding processing scales. This way, a given S_2 feature is simultaneously explored in all positions and scales (see S_2 maps of Fig. 1 with same feature prototype but in different processing scales specified by different colors). Indeed, duplicated cells in all positions of all scale maps integrate the spike train in parallel and compete with each other. The first duplicate reaching its threshold, if any, is the winner. The winner fires and prevents the other duplicated cells in all other positions and scales from firing through a winner-take-all inhibition mechanism. Then, for each prototype, the winner S_2 cell triggers the unsupervised STDP rule and its weight matrix is updated. The changes in its weights are applied over all other duplicate cells in different positions and scales (weight sharing mechanism). This allows the system to learn frequent patterns, independently of their position and size in the training images.

The learning process begins with S_2 features initialized by random numbers drawn from a normal distribution with mean 0.8 and STD 0.05, and the threshold of all S_2 cells is set to 64 ($= 1/4 \times 16 \times 16$). Through the learning process, a local inhibition between different S_2 prototype cells is used to prevent the convergence of different S_2 prototypes to similar features: when a cell fires at a given position and scale, it prevents all the other cells (independently of their preferred prototype) from firing later at the same scale and within a neighborhood around the firing position. Thus, the cell population self-organizes, each cell trying to learn a distinct pattern so as to cover the whole variability of the inputs. Moreover, we applied a k-winner-take-all strategy in S_2 layer to ensure that at most two cells can fire for each processing scale. This mechanism, only used in the learning phase, helps the cells to learn patterns with different real sizes. Without it, there is a natural bias toward “small” patterns (i.e., large scales), simply because corresponding maps are larger, and so likeliness of firing with random weights at the beginning of the STDP process is higher.

A simplified version of STDP is used to learn the C_1 – S_2 weights as follows:

$$\begin{cases} \Delta w_{ij} = a^+ \cdot w_{ij} \cdot (1 - w_{ij}), & \text{if } t_j - t_i \leq 0, \\ \Delta w_{ij} = a^- \cdot w_{ij} \cdot (1 - w_{ij}), & \text{if } t_j - t_i > 0, \end{cases}$$

where i and j respectively refer to the index of post- and pre-synaptic neurons, t_i and t_j are the corresponding spike times, Δw_{ij} is the synaptic weight modification, and a^+ and a^- are two parameters specifying the learning rate. Note that the exact time difference between two spikes ($t_j - t_i$) does not affect the weight change, but only its sign is considered. These simplifications are equivalent to assuming that the intensity-to-latency conversion of S_1 cells compresses the whole spike wave in a relatively short time interval (say, 20–30 ms), so that all presynaptic spikes necessarily fall close to the postsynaptic spike time, and the time lags are negligible. The multiplicative term $w_{ij} \cdot (1 - w_{ij})$ ensures the weights remain in the range [0,1] and maintains all synapses in an excitatory mode. The learning phase starts by $a^+ = 2^{-6}$ which is multiplied by 2 after each 400 postsynaptic spikes up to a maximum value of 2^{-2} . A fixed a^+ / a^- ratio ($-4/3$) is used. This allows us to speed up the convergence of S_2 features as the learning progresses. Initiation of the learning phase with high learning rates would lead to erratic results.

For each S_2 prototype, a C_2 cell propagates the first spike emitted by the corresponding S_2 cells over all positions and processing scales, leading to the global shift- and scale-invariant cells (see the C_2 layer of Fig. 1).

3. Experimental results

3.1. Dataset and experimental setup

To study the robustness of our model with respect to different transformations such as scale and viewpoint, we evaluated it on the 3D-Object and ETH-80 datasets. The 3D-Object is provided by Savarese et al. at CVGLab, Stanford University [25]. This dataset contains 10 different object classes: bicycle, car, cellphone, head, iron, monitor, mouse, shoe, stapler, and toaster. There are about 10 different instances for each object class. The object instances are photographed in about 72 different conditions: eight view angles, three distances (scales), and three different tilts. The images are not segmented and the objects are located in different backgrounds (the background changes even for different conditions of the same object instance). Fig. 2 presents some examples of objects in this dataset.

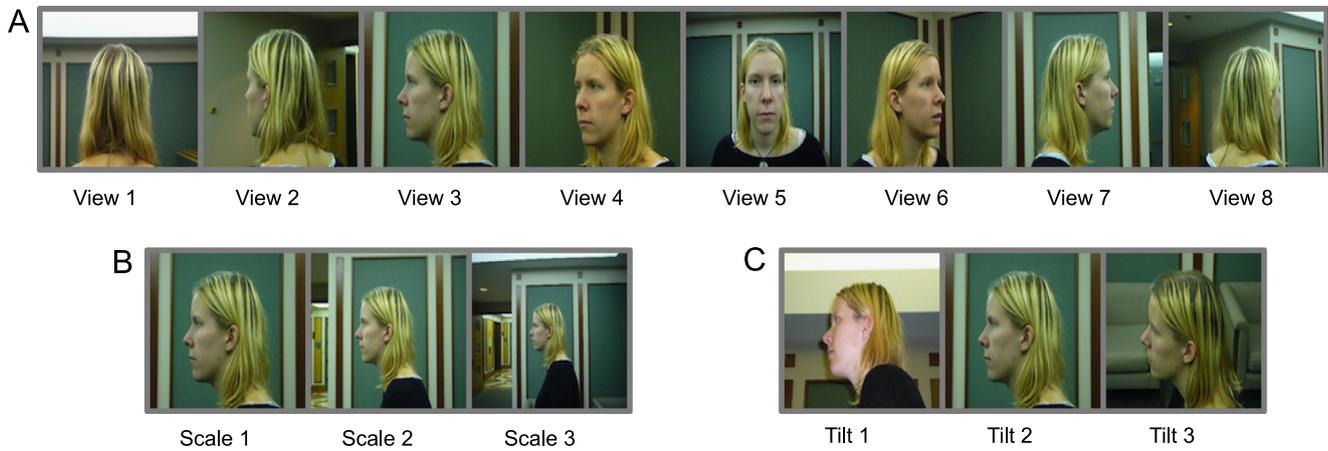


Fig. 2. Some images of the head class of 3D-Object dataset in different A) views, B) scales, and C) tilts.

Table 1

Performance of our model, HMAX, and DeepConvNet with different number of features.

Dataset		Our model				HMAX				DeepConvNet
3D-Object	# Features	200	300	400	500	1000	3000	9000	12000	4096
	Accuracy	76.1%	94.7%	96.0%	96.0%	58.2%	60.1%	61.9%	62.4%	85.8%
ETH-80	# Features	500	750	1000	1250	500	1000	2000	5000	4096
	Accuracy	75.3	79.3%	80.7%	81.1%	66.3%	68.7%	68.9%	69.0%	79.1%

The ETH-80 dataset includes 80 3D objects in eight different category categories including apple, car, toy cow, cup, toy dog, toy horse, pear, and tomato. Each object is photographed in 41 viewpoints with different view angles and different tilts. Fig. S1 in Supplementary Information provides some examples of objects in this dataset from different viewpoints.

For both datasets, five instances of each object category are selected for the training set to be used in the learning phase. The remaining instances constitute the testing set which is not seen during the learning phase, but is used afterward to evaluate the recognition performance. This standard cross-validation procedure allows to measure the generalization ability of the model beyond the specific training examples. Note that for 3D-Object dataset, the original size of all images were preserved, while the images of ETH-80 dataset are resized to 300 pixels in height while preserving the aspect ratio. The images of both datasets were converted to grayscale values.

As already mentioned, the building process of S_2 features is performed in a completely unsupervised manner. Hence, through the execution of the unsupervised STDP-based learning, the training images are randomly fed into the model (without considering their class labels, viewpoints, scales, and tilts). The learning process starts with initial random weights and finishes when 600 spikes have occurred in each S_2 map. Then STDP is turned off, and the ability of the obtained features to invariantly represent different object classes is evaluated. To compute the corresponding C_2 feature vector for each input image, the thresholds of C_2 neurons are set to infinite, and their final potentials are evaluated, after propagating the whole spike train generated by each image. Each final potential can be seen as the number of early spikes in common between the current input and a stored prototype (this is very similar to the tuning operation of S cells in HMAX). Then, a one-versus-one multiclass linear support vector machine (SVM) classifier is trained based on the C_2 features of the training set and it is evaluated on the test set.

We have compared the performance of our model with the HMAX model [14] and deep supervised convolutional network

(DeepConvNet) by Krizhevsky et al. [16]. Comparison with the HMAX model is particularly instructive, since as explained above, we use very similar architecture, tuning and maximum operations. The main difference is that instead of using an unsupervised learning rule like us, the HMAX model uses random crops from the training images to imprint the S_2 features (here of equal size). Then a SVM classifier was trained over the HMAX C_2 features to complete the object recognition process. The employed HMAX model is implemented by Mutch, et al. [41] and the codes are publicly available at <http://cbcl.mit.edu/jmutch/cns/index.html>.

We also compared our model with DeepConvNet which has been shown to be the best algorithm in various object classification tasks including the ImageNet LSVRC-2010 contest [16]. It is comprised of eight consecutive layers (five convolutional layers followed by three fully connected layers) with about 60 millions parameters, learned with stochastic gradient descent. We have used a pre-trained DeepConvNet model implemented by Jia, et al. [42], whose code is also available at <http://caffe.berkeleyvision.org>. The training was done over the ILSVRC2012 dataset (a subset of ImageNet) with about 1.2 million images in 1000 categories. We fed the training and testing images into DeepConvNet and extracted the feature values from the 7th layer. Again, a SVM is used to do the object recognition based on the extracted features.

3.2. Performance analysis

Table 1 provides the accuracy of our model in category classification independently of view, tilt, and scale, when different number of S_2 features are learned by the STDP-based learning algorithm. The results indicate that the model reaches a high classification performance on 3D-Object dataset with about 300 C_2 features only (about 30 features per class). The performance is flattened around 96% for feature vectors of size greater than 400. Also, for the ETH-80 dataset, the model attains to a reasonable recognition accuracy of about 81% with only 1250 extracted features. We have also performed the same experiments on HMAX

and DeepConvNet models which their accuracies are also provided in Table 1.

Performance of the HMAX model was weak on both datasets, which is not too surprising, several previous studies have shown that the performance of the HMAX model extensively decreases when facing significant object transformations [28,29]. Given the structural similarities between our model and HMAX, the superiority of our model is presumably related to the unsupervised feature learning. Indeed, most of the randomly extracted S_2 patches in HMAX tend to be redundant and irrelevant, as we will see in the next section.

DeepConvNet reached a mean performance of about 86% on 3D-Object and about 79% on ETH-80 dataset. Thus, our model outperforms DeepConvNet on both datasets, which itself significantly outperforms HMAX. It should be noted that the images of each object in these two datasets are highly varied (e.g., in 3D-Object dataset, there is a 45° difference between two successive views of an object) and it has previously been shown that the performance of DeepConvNet drops when facing such transformations [30]. Another drawback of DeepConvNet is that, due to the large number of parameters, it needs to be trained over millions of images to avoid overfitting [43] (here we avoided this problem by using a pre-trained version, but doing the training on about 3500 images we used here would presumably lead to massive overfitting). Conversely, our model is able to learn objects using much fewer images.

Consequently, the results indicate that our model has a great ability to learn diagnostic features tolerating transformations and deformations of the presented stimulus.

3.3. Feature analysis

In this section, we demonstrate that unsupervised STDP learning algorithm extracts informative and diagnostic features by comparing them to the randomly picked HMAX features. To this end, we have used several feature analysis techniques: representational dissimilarity matrices, hierarchical clustering, and mutual information. We performed the same analyses on both datasets and obtained similar results. Hence, the results of ETH-80 are presented in Supplementary Information.

Extraction of diagnostic features let our model reach high classification performances with a small number of features (c.f. Table 1). To understand why this is true, we first reconstructed the features' preferred stimuli. Given that each S_2 neuron receives spikes from C_1 neurons responding to bars in different

orientations, the representation of the preferred features of S_2 neurons could be reconstructed by convolving their weight matrices with a set of kernels representing oriented bars. In Fig. 3 the receptive fields of activated S_2 neurons along with the representation of their preferred stimuli are illustrated (Fig. S2 provides the same illustration for the ETH-80 dataset). This demonstrates that only a small number of S_2 neurons are required to represent the input objects. In other words, the obtained features are compatible with the sparse coding theory in visual cortex. In addition, for an input image, the most activated S_2 neurons cover the input objects and they do not respond to the background area. Indeed, the STDP learning algorithm naturally focuses on what are common in the training images, which are the target object features. The backgrounds are generally not learned (at least not in priority), since they are almost always too different from one image to another and the STDP process cannot converge on them.

To characterize the neuronal population coding in the C_2 layer of the model and to study the quality of C_2 features, we used the representational dissimilarity matrix (RDM) [32]. Each element of the RDM reflects the measure of dissimilarity (distance) among the neural activity patterns (i.e., the object representations) associated with two different image stimuli. The distance we used here is $1 - \text{Pearson correlation}$. In an RDM corresponding to a perfect model, the representations of the objects of the same category have low dissimilarities (i.e., highly correlated), whereas objects of different categories are represented highly dissimilarly (i.e., uncorrelated). Hence, if we group the rows and columns of the RDM of a perfect model based on object categories, it is expected to see squares of low dissimilarity values around the main diagonal, each of which corresponds to pairs of same-category images, while other elements have higher values.

Here, to plot the RDM of each view angle, first, the images of all input instances which are taken in that view are picked. Then the corresponding RDM is plotted by computing the pairwise dissimilarity of the values of C_2 features associated with each pair of images. Fig. 4 presents the RDMs of our model for all eight views (see Fig. S3 for ETH-80). In each RDM, rows and columns are sorted based on image categories. Also, a sample image of each category is placed next to the rows and columns which correspond to that category. Here, we used a color-code to represent RDMs which ranges from pure blue to pure yellow demonstrating low to high dissimilarities, respectively. It can be seen that the within-category dissimilarity values (identified by blue squares around the main diagonal) are relatively lower than the between-category dissimilarities (more yellowish areas). As expected, the RDMs indicate

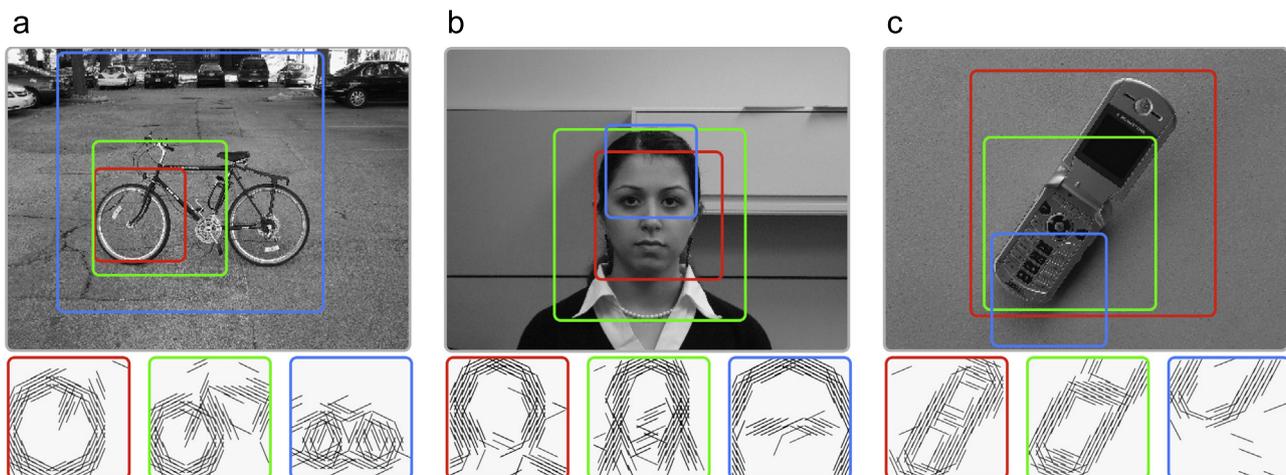


Fig. 3. Three S_2 feature prototypes selective to the a) bicycle, b) face, and c) cellphone classes of 3D-Object dataset along with their reconstructed preferred stimuli. It can be seen that the features converged to specific and salient object parts and neglected the irrelevant backgrounds.

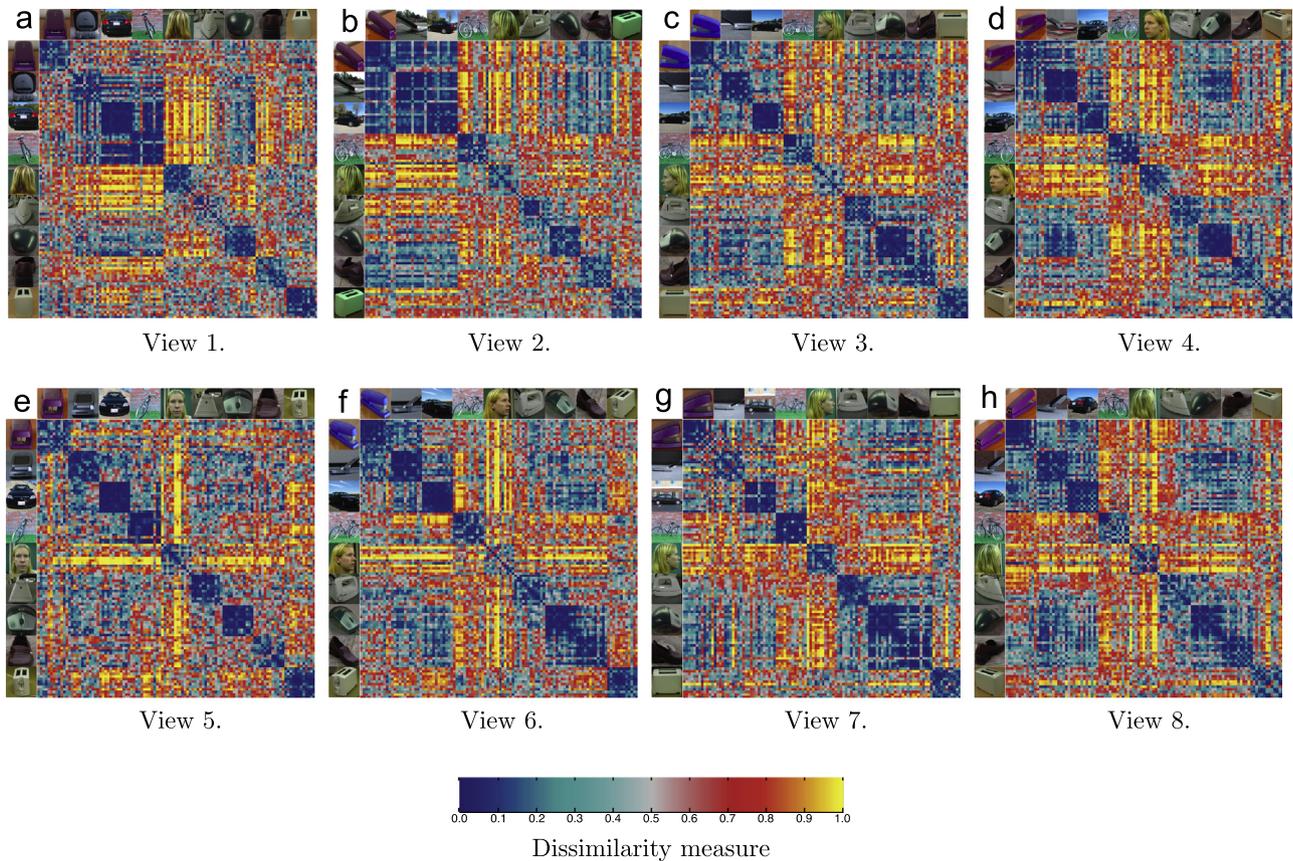


Fig. 4. RDMs of our model on 3D-Object dataset corresponding to different viewpoints. It can be seen that within class dissimilarities are very low (the blue squares around the main diagonal where rows and columns correspond to images of the same category), while between class dissimilarities are higher (more yellowish). Note that due to the absence of image samples for some views of the monitor class, we have eliminated this class from the RDMs, (a) View 1. (b) View 2. (c) View 3. (d) View 4. (e) View 5. (f) View 6. (g) View 7. (h) View 8. Dissimilarity measure. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

that the obtained performance is not due to the capabilities of the classifier, but to the extraction of diagnostic and highly informative C_2 features through STDP.

We have also computed the RDMs of the HMAX model (including 12000 features) for eight views, as provided in Fig. 5 (see Fig. S4 for ETH-80). As it can be seen in this figure, the randomly selected features in the HMAX model are unable to similarly represent within-category objects and dissimilarly represent between-category objects. This is probably due to uninformative features used by the HMAX model. Indeed, in HMAX, the task of selecting the informative features is left to the classifier. We also note the presence of horizontal and vertical yellow lines, indicating “outliers”, whose representation lies far away from all the others. This indicates that the features do not pave well the stimulus space.

To see how well the stimuli are distributed in the high dimensional feature space, we performed hierarchical clustering over the test set. The clustering procedure is started by considering each stimulus as a discrete cluster node, continued by connecting the closest nodes into a new combined cluster node, and completed by connecting all the stimuli to a single node. We performed this analysis on the C_2 feature vectors corresponding to all objects in all views, scales, and tilts. The obtained hierarchy for our model is displayed in Fig. 6 (see Fig. S5 for ETH-80). The distance between a pair of cluster nodes is computed by measuring the dissimilarity among their centers (the average of cluster members). Due to the large number of stimuli, it is not possible to plot the whole hierarchy, hence, only the high level clusters are shown in this figure. For each lowest level cluster, the class with

the highest frequency is illustrated by an image label. The cardinality of this class as well as the cardinality of the cluster are shown below the labels. It can be seen that the instances of each object class are placed in neighboring regions of the C_2 feature space. By considering the obtained hierarchical clustering and the classification accuracies, it can be concluded that the C_2 features are able to invariantly represent the objects in such a way that the classifier can easily separate them.

The same hierarchical clustering is performed for the HMAX feature space (with 12000 features), as depicted in Fig. 7 (see Fig. S6 for ETH-80). As it can be seen, the majority of clusters are small, and contrary to our model, the distances between the clusters are very low. In other words, the objects are densely represented in a small area of such a high dimensional feature space. Furthermore, the mean intra- and inter-class dissimilarities in our model are equal to 0.40 and 0.70, respectively, while these statistics for the HMAX model are equal to 0.27 and 0.29, respectively. In summary, it can be concluded that the distribution of the object classes are dense and highly overlapped in the HMAX feature space, while the object classes are well separated in the feature space of our model.

In an other experiment, we analyzed the class dependency of the C_2 features for our model. To this end, the 50 most informative features, when classifying a specific class against all the other classes, are selected by employing the mutual information technique. In other words, for each class, we selected those 50 features which have the highest activity for samples of that class and have less activity for other classes. Afterwards, the number of common features among the informative features of each pair of classes are computed as provided in Table 2. On average, there are only about 5.4 common features

between pairs of classes. Although there are some common features between any two classes, their co-occurrence with the other features help the classifier to separate them from each other. In this way, our

model can represent various object classes with a relatively small number of features. Indeed, exploiting the intermediate complexity features, which are not common in all classes and are not very rare, can

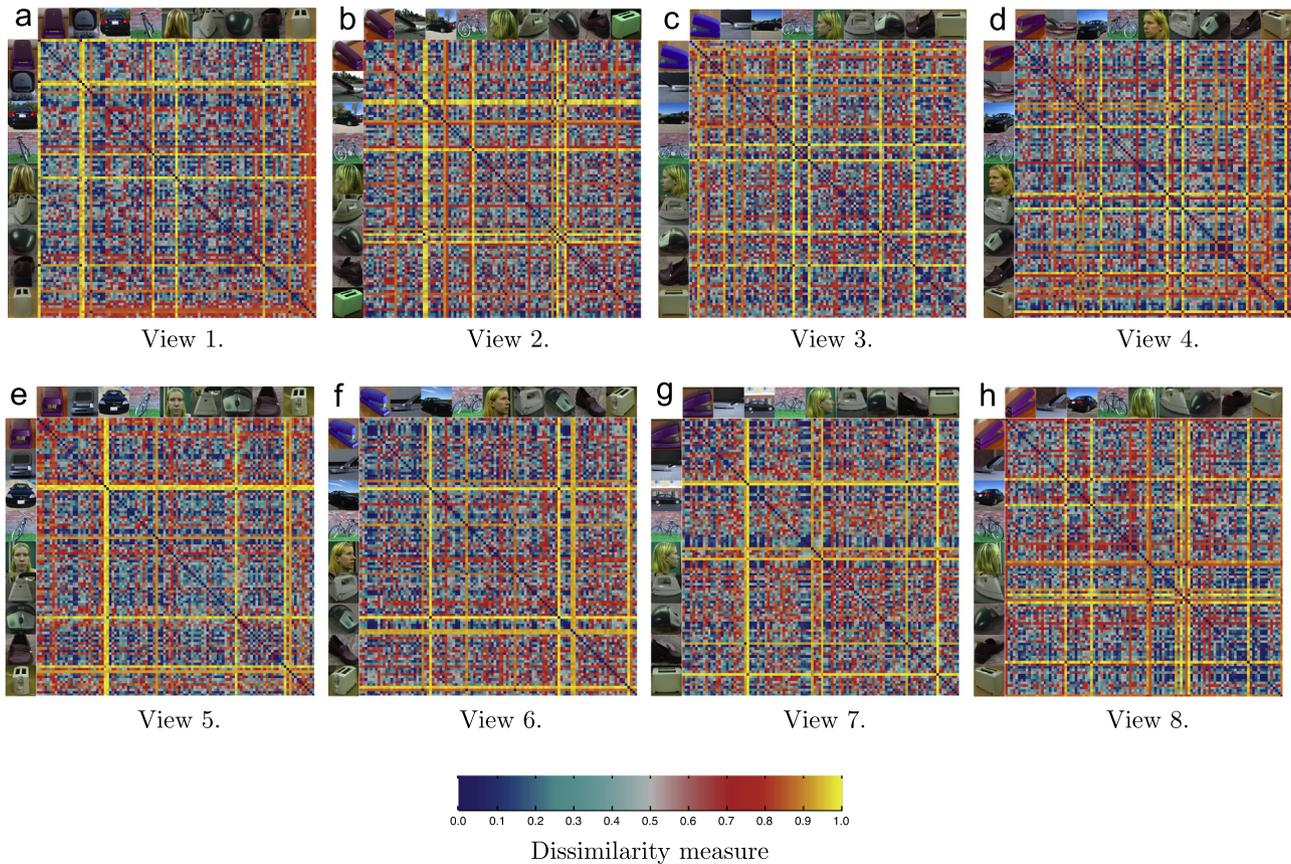


Fig. 5. RDMs of the HMAX model on 3D-Object dataset corresponding to different viewpoints. Randomly selected features in HMAX model are not able to similarly represent within-category objects and dissimilarly represent between-category objects. Note that due to the absence of image samples for some views of the monitor class, we have eliminated this class from the RDMs, (a) View 1. (b) View 2. (c) View 3. (d) View 4. (e) View 5. (f) View 6. (g) View 7. (h) View 8. Dissimilarity measure. (For interpretation of the references to color in this figure, the reader is referred to the web version of this article.)

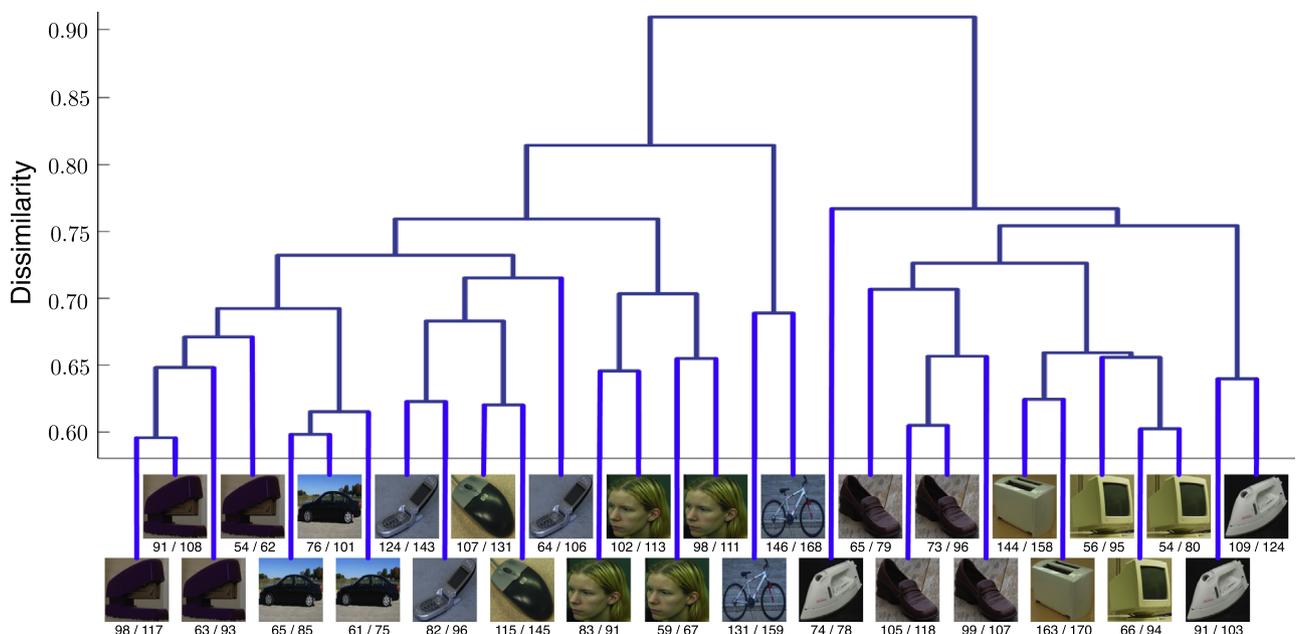


Fig. 6. The hierarchy of clusters and their labels for our model on 3D-object dataset. The label of each cluster indicates the class with the highest frequency in that cluster. It can be seen that the samples of each class are placed in close clusters. The cardinality of each cluster, C , and the cardinality of the class with the highest frequency, H , are placed below the cluster label as H/C .

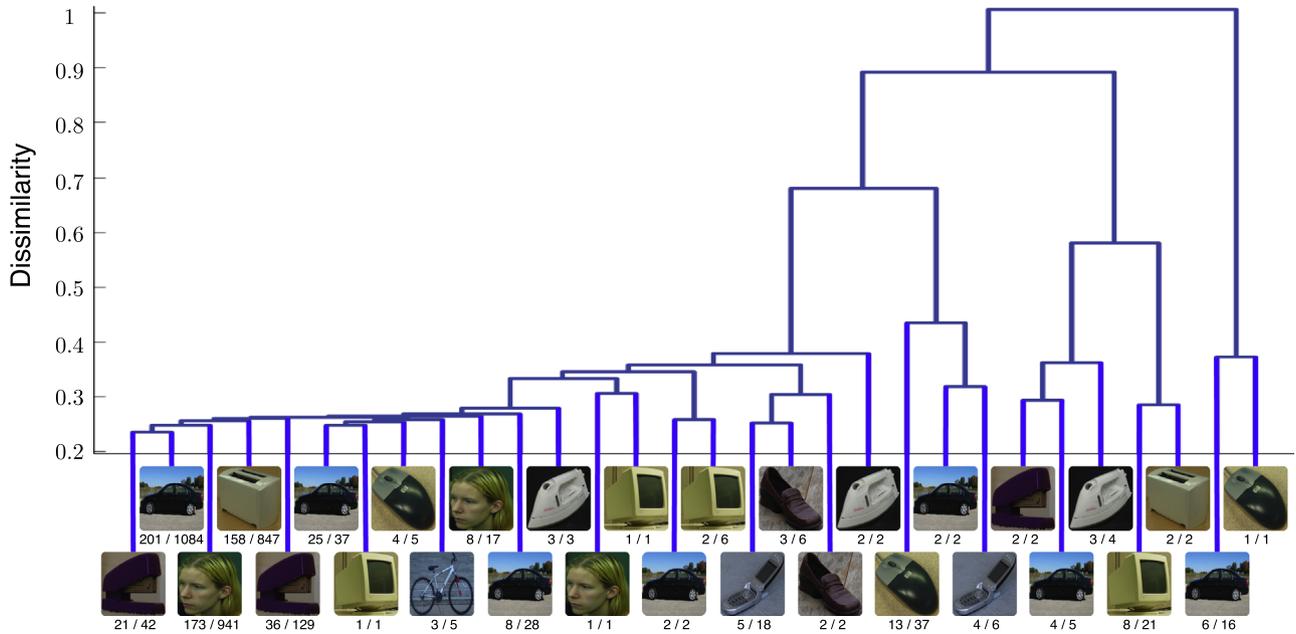


Fig. 7. The hierarchy of clusters and their labels for the HMAX model on 3D-Object dataset. The label of each cluster indicates the class with the highest frequency in that cluster. It can be seen that the majority of the objects are assigned to a small number of clusters and samples of each class are not well placed in close clusters. The cardinality of each cluster, C , and the cardinality of the class with the highest frequency, H , are placed below the cluster label as H/C .

Table 2

The number of common features between each pair of classes of 3D-Object dataset.

Class	Bicycle	Car	Cellphone	Head	Iron	Monitor	Mouse	Shoe	Stapler	Toaster
Bicycle	50	0	1	3	6	4	12	6	5	5
Car	0	50	1	14	10	11	1	2	6	7
Cellphone	1	1	50	0	3	4	2	0	10	9
Head	3	14	0	50	0	0	10	16	2	2
Iron	6	10	3	0	50	21	0	4	12	6
Monitor	4	11	4	0	21	50	0	0	8	14
Mouse	12	1	2	10	0	0	50	2	2	5
Shoe	6	2	0	16	4	0	2	50	0	0
Stapler	5	6	10	2	12	8	2	0	50	20
Toaster	5	7	9	2	6	14	5	0	20	50

help the classifier to discriminate instances of different classes [40].

3.4. Random features and simple classifier

In a previous study [44], it has been shown that using the HMAX model with random dot patterns in the S_2 layer can reach a reasonable performance, comparable to the one obtained with random patches cropped from the training images. It seems that this is due to the dependency of HMAX to the application of a powerful classifier. Indeed, the use of both random dot or randomly selected patches transform the images into a complex and nested feature space and it is the classifier which looks for a complex signature to separate object classes. The deficiencies emerge when the classification problem gets harder (such as invariant or multiclass object recognition problems) and then even a powerful classifier is not able to discriminate the classes [28,29]. Here, we show that the superiority of our model is due to the informative feature extraction through a bio-inspired learning rule. To this end, we have compared the performances on 3D-Object dataset obtained with random features versus STDP features, as well as a very simple classifier versus SVM.

To generate random features, we have set the weight matrix of each S_2 feature of our model with random values. First, we have computed the mean and standard deviation (STD) (253 ± 21) of the number of active (nonzero) weights in the features learned by

STDP. Second, for each random feature, the number of active weights, N , is computed by generating a random number based on the obtained mean and STD. Finally, a random feature is constructed by uniformly distributing the N randomly generated values in the weight matrix.

In addition, we designed a simple classifier comprised of several one-versus-one classifiers. For each binary classifier, two subset of C_2 features with high occurrence probabilities in one of the two classes are selected. In more details, to select suitable features for the first class, the occurrence probabilities of C_2 features in this class are divided by the corresponding occurrence probabilities in the second class. Then, a feature is selected if this ratio is higher than a threshold. The optimum threshold value is computed by a trial and error search in which the performance over the training samples is maximized. To assign a class label to the input test sample, we performed an inner product on the feature value and feature probability vectors. Finally, the class with the highest probability is reported to the combined classifier. The combined classifier selects the winner class based on a simple majority voting.

For 500 random features, using the SVM and the simple classifier, our model reached classification performances of 71% and 21% on average, respectively. Whereas, for the learned S_2 features, both the SVM and simple classifiers attained reasonable performances of 96% and 79%, respectively. Based on these results, it can

be concluded that the features obtained through the bio-inspired unsupervised learning projects the objects into an easily separable space, while the feature extraction by selection of random patches (drawn from the training images) or by generation of random patterns leads to a complex object representation.

4. Discussion

Position and scale invariance in our model are built-in, thanks to weight sharing and scaling process. Conversely, view-invariance must be obtained through the learning process. Here, we used all images of five object instances from each category (varied in all dimensions) to learn the S_2 visual features, while images of all other object instances of each category were used to test the network. Hence, the model was exposed to all possible variations during the learning to gain view-invariance. Moreover, near or opposite views of the same object shares some features which are suitable for invariant object recognition. For instance, consider the overall shape of a head, or close views of a bike wheel which could be a complete circle or an ellipse. Regarding the fact that STDP tends to learn more frequent features in different images, different views of an object could be invariantly represented based on more common features.

Our model appears to be the best choice when dealing with few object classes, but huge variations in view points. As pointed out in previous studies, both HMAX and DeepConvNet models could not handle these variations perfectly [28–30]. Conversely, our model is not appropriate to handle many classes, which requires thousands of features, like in the ImageNet contest, because its time complexity is roughly in N^2 , where N is the number of features (briefly: since the number of firing neurons per image is limited, if the number of features is doubled, reaching convergence will take roughly twice as many images, and the processing time for each of them will be doubled as well). For example, extracting 4096 features in our model, the same number of features in DeepConvNet, would take about 67 times it took us to extract 500. However, parallel implementation of our algorithm could speed-up the computation time by several orders of magnitude [45]. Even in this case, we do not expect to outperform the DeepConvNet model on the ImageNet database, since only the shape similarities are taken into account in our model and the other cues such as color or texture are ignored.

Importantly, our algorithm has a natural tendency to learn salient contrasted regions [24], which is desirable as these are typically the most informative [46]. Most of our C_2 features turned out to be class-specific, and we could guess what they represent by doing the reconstructions (see Fig. 3 and Fig. S2). Since each feature results from averaging multiple input images, the specificity of each instance is averaged out, leading to class archetypes. Consequently, good classification results can be obtained using only a few features, or even using ‘simple’ decision rules like feature counts [24] and majority voting (here), as opposed to a ‘smart classifier’ such as SVM.

There are some similarities between STDP-based feature learning, and non-negative matrix factorization [47], as first intuited in [48], and later demonstrated mathematically in [49]. Within both approaches, objects are represented as (positive) sums of their parts, and the parts are learned by detecting consistently co-active input units.

Our model could be efficiently implemented in hardware, for example using address event representation (AER) [50–53]. With AER, the spikes are carried as addresses of sending or receiving neurons on a digital bus. Time ‘represents itself’ as the asynchronous occurrence of the event [54]. Thus the use of STDP will lead to a system which effectively becomes more and more reactive, in

addition to becoming more and more selective. Furthermore, since biological hardware is known to be incredibly slow, simulations could run several orders of magnitude faster than real time [55]. As mentioned earlier, the primate visual system extracts the rough content of an image in about 100 ms. We thus speculate that some dedicated hardware will be able to do the same in the order of a millisecond or less.

Recent computational [40], psychophysical [56], and fMRI [57] experiments demonstrate that the informative intermediate complexity features are optimal for object categorization tasks. But the possible neural mechanisms to extract such features remain largely unknown. The HMAX model ignores these learning mechanisms and imprints its features with random crops from the training images [14,58], or even uses random filters [44,59]. Most individual features are thus not very informative, yet in some cases, a ‘smart’ classifier such as SVM can efficiently separate the high-dimensional vectors of population responses.

Many other models use supervised learning rules [13,16], sometimes reaching impressive performance on natural image classification tasks [16]. The main drawback of these supervised methods, however, is that learning is slow and requires numerous labeled samples (e.g., about 1 million in [16]), because of the credit assignment problem [60,61]. This contrasts with humans who can generalize efficiently from just a few training examples [43]. We avoid the credit assignment problem by keeping the C_2 features fixed when training the final classifier (that being said, fine-tuning them for a given classification problem would probably increase the performance of our model [17,62]; we will test this in future studies). Even if the efficiency of such hybrid unsupervised-supervised learning schemes has been known for a long time, few alternative unsupervised learning algorithms have been shown to be able to extract complex and high-level visual features (see [15,17]). Finding better representational learning algorithms is thus an important direction for future research and seeking for inspiration in the biological visual systems is likely to be fruitful [43]. We suggest here that the physiological mechanism known as STDP is an appealing start point.

Considering the time relation among the incoming inputs is an important aspect of spiking neural networks. This property is critical to promote the existing models from static vision to continuous vision [63]. A prominent example is the trace learning rule [64], suggesting that the invariant object representation in ventral visual system is instructed by the implicit temporal contiguity of vision. Also, in various motion processing and action recognition problems [65], the important information lies in the appearance timing of input features. Our model has this potential to be extended for continuous and dynamic vision – something that we will further explore.

5. Conclusions

To date, various bio-inspired network architectures for object recognition have been introduced, but the learning mechanism of biological visual systems has been neglected. In this paper, we demonstrate that the association of both bio-inspired network architecture and learning rule results in a robust object recognition system. The STDP-based feature learning, used in our model, extracts frequent diagnostic and class specific features that are robust to deformations in stimulus appearance. It has previously been shown that the trivial models can not tolerate the identity preserving transformations such as changes in view, scale, and position. To study the behavior of our model confronted with these difficulties, we have tested our model over two challenging invariant object recognition databases which includes instances of 10 different object classes photographed in different views, scales,

and tilts. The categorization performances indicate that our model is able to robustly recognize objects in such a severe situation. In addition, several analytical techniques have been employed to prove that the main contribution to this success is provided by the unsupervised STDP feature learning, not by the classifier. Using representational dissimilarity matrix, we have shown that the representation of input images in C_2 layer are more similar for within-category and dissimilar for between-category objects. In this way, as confirmed by the hierarchical clustering, the objects with the same category are represented in neighboring regions of C_2 feature space. Hence, even if using a simple classifier, our model is able to reach an acceptable performance, while the random features fail.

Acknowledgements

We would like to thank Mr. Majid Changi Ashtiani at the Math Computing Center of IPM (<http://math.ipm.ac.ir/mcc>) for letting us to perform some parts of the calculations on their computing cluster. We also thank Dr. Reza Ebrahimpour for his helpful discussions and suggestions.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.neucom.2016.04.029>.

References

- [1] S. Thorpe, D. Fize, C. Marlot, et al., Speed of processing in the human visual system, *Nature* 381 (6582) (1996) 520–522.
- [2] I. Biederman, Recognition-by-components: a theory of human image understanding, *Psychol. Rev.* 94 (2) (1987) 1–15.
- [3] J.J. DiCarlo, D. Zoccolan, N.C. Rust, How does the brain solve visual object recognition? *Neuron* 73 (3) (2012) 415–434.
- [4] P. Lennie, J.A. Movshon, Coding of color and form in the geniculostriate visual pathway (invited review), *J. Opt. Soc. Am. A* 22 (10) (2005) 2013–2033.
- [5] A.S. Nandy, T.O. Sharpee, J.H. Reynolds, J.F. Mitchell, The fine structure of shape tuning in area V4, *Neuron* 78 (6) (2013) 1102–1115.
- [6] K. Tanaka, H.-a. Saito, Y. Fukada, M. Moriya, Coding visual images of objects in the inferotemporal cortex of the macaque monkey, *J. Neurophysiol.* 66 (1) (1991) 170–189.
- [7] C.P. Hung, G. Kreiman, T. Poggio, J.J. DiCarlo, Fast readout of object identity from macaque inferior temporal cortex, *Science* 310 (5749) (2005) 863–866.
- [8] N.C. Rust, J.J. DiCarlo, Selectivity and tolerance (invariance) both increase as visual information propagates from cortical area V4 to IT, *J. Neurosci.* 30 (39) (2010) 12978–12995.
- [9] H. Liu, Y. Agam, J.R. Madsen, G. Kreiman, Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex, *Neuron* 62 (2) (2009) 281–290.
- [10] W.A. Freiwald, D.Y. Tsao, Functional compartmentalization and viewpoint generalization within the macaque face-processing system, *Science* 330 (6005) (2010) 845–851.
- [11] F. Anselmi, J.Z. Leibo, L. Rosasco, J. Mutch, A. Tacchetti, T. Poggio, Unsupervised learning of invariant representations with low sample complexity: the magic of sensory cortex or a new framework for machine learning?, arXiv preprint arXiv:1311.4158, 2014, pp. 1–23.
- [12] K. Fukushima, Neocognitron: a self organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, *Biol. Cybern.* 36 (4) (1980) 193–202.
- [13] Y. LeCun, Y. Bengio, Convolutional networks for images, speech, and time series, in: M.A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks*, Cambridge, MA: MIT Press, 1998, pp. 255–258.
- [14] T. Serre, L. Wolf, S. Bileschi, M. Riesenhuber, T. Poggio, Robust object recognition with cortex-like mechanisms, *IEEE Trans. Pattern Anal. Machine Intel.* 29 (3) (2007) 411–426, <http://dx.doi.org/10.1109/TPAMI.2007.56>.
- [15] H. Lee, R. Grosse, R. Ranganath, A.Y. Ng, Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations, in: Proceedings of the 26th Annual International Conference on Machine Learning, 2009, pp. 1–8, <http://dx.doi.org/10.1145/1553374.1553453>.
- [16] A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep convolutional neural networks, in: *Neural Information Processing Systems (NIPS)*, Lake Tahoe, Nevada, 2012, pp. 1–9.
- [17] Q.V. Le, Building high-level features using large scale unsupervised learning, in: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 2013, pp. 8595–8598, <http://dx.doi.org/10.1109/ICASSP.2013.6639343>.
- [18] G.M. Ghose, Learning in mammalian sensory cortex, *Current Opinion Neurobiol.* 14 (4) (2004) 513–518, <http://dx.doi.org/10.1016/j.conb.2004.07.003>.
- [19] Z. Kourtzi, J.J. DiCarlo, Learning and neural plasticity in visual object recognition, *Current Opinion Neurobiol.* 16 (2) (2006) 152–158, <http://dx.doi.org/10.1016/j.conb.2006.03.012>.
- [20] C.D. Meliza, Y. Dan, Receptive-field modification in rat visual cortex induced by paired visual stimulation and single-cell spiking, *Neuron* 49 (2) (2006) 183–189, <http://dx.doi.org/10.1016/j.neuron.2005.12.009>.
- [21] S. Huang, C. Rozas, M. Trevino, J. Contreras, S. Yang, L. Song, T. Yoshioka, H.-K. Lee, A. Kirkwood, Associative hebbian synaptic plasticity in primate visual cortex, *J. Neurosci.* 34 (22) (2014) 7575–7579, <http://dx.doi.org/10.1523/JNEUROSCI.0983-14.2014>.
- [22] D.E. Feldman, The spike-timing dependence of plasticity, *Neuron* 75 (4) (2012) 556–571, <http://dx.doi.org/10.1016/j.neuron.2012.08.001>.
- [23] D.B.T. McMahon, D.A. Leopold, Stimulus timing-dependent plasticity in high-level vision, *Current Biol.* 22 (4) (2012) 332–337, <http://dx.doi.org/10.1016/j.cub.2012.01.003>.
- [24] T. Masquelier, S.J. Thorpe, Unsupervised learning of visual features through spike timing dependent plasticity, *PLoS Comput. Biol.* 3 (2) (2007) e31.
- [25] S. Savarese, L. Fei-Fei, 3D generic object categorization, localization and pose estimation, in: Proceedings of the 11th IEEE International Conference on Computer Vision (ICCV), 2007, pp. 1–8, <http://dx.doi.org/10.1109/ICCV.2007.4408987>.
- [26] B. Pepik, M. Stark, P. Gehler, B. Schiele, Multi-view priors for learning detectors from sparse viewpoint data, in: Proceedings of the International Conference on Learning Representations (ICLR), Banff, AB, Canada, 2014, pp. 1–13.
- [27] B. Leibe, B. Schiele, Analyzing appearance and contour based methods for object categorization, in: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2, 2003, pp. II–409.
- [28] N. Pinto, D.D. Cox, J.J. DiCarlo, Why is real-world visual object recognition hard? *PLoS Comput. Biol.* 4 (1) (2008) e27, <http://dx.doi.org/10.1371/journal.pcbi.0040027>.
- [29] N. Pinto, Y. Barhomi, D.D. Cox, J.J. DiCarlo, Comparing state-of-the-art visual features on invariant object recognition tasks, in: IEEE workshop on Applications of Computer Vision (WACV), 2011, pp. 463–470.
- [30] M. Ghodrati, A. Farzmaidi, K. Rajaei, R. Ebrahimpour, S.-M. Khaligh-Razavi, Feedforward object-vision models only tolerate small image variations compared to human, *Front. Comput. Neurosci.* 8 (2014) 74.
- [31] J. Pohjalainen, O. Räsänen, S. Kadioglu, Feature selection methods and their combinations in high-dimensional classification of speaker likability, intelligibility and personality traits, *Computer Speech & Language* 29 (1) (2015) 145–171.
- [32] N. Kriegeskorte, M. Mur, P. Bandettini, Representational similarity analysis—connecting the branches of systems neuroscience, *Front. Syst. Neurosci.* 2 (4) (2008) 1–28.
- [33] F. Murtagh, P. Contreras, Algorithms for hierarchical clustering: an overview, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 2 (1) (2012) 86–97.
- [34] S.J. Thorpe, M. Imbert, Biological constraints on connectionist modelling, in: *Connectionism in perspective*, Amsterdam: Elsevier, 1989, pp. 63–92.
- [35] S. Celebri, S. Thorpe, Y. Trotter, M. Imbert, Dynamics of orientation coding in area V1 of the awake primate, *Visual Neurosci.* 10 (5) (1993) 811–825.
- [36] D.G. Albrecht, W.S. Geisler, R.A. Frazor, A.M. Crane, Visual cortex neurons of monkeys and cats: temporal of the contrast response function, *J. Neurophysiol.* 88 (2) (2002) 888–913.
- [37] O. Shriki, A. Kohn, M. Shamir, Fast coding of orientation in primary visual cortex, *PLoS Comput. Biol.* 8 (6) (2012) e1002536, <http://dx.doi.org/10.1371/journal.pcbi.1002536>.
- [38] M. Riesenhuber, T. Poggio, Hierarchical models of object recognition in cortex, *Nat. Neurosci.* 2 (11) (1999) 1019–1025.
- [39] G.A. Rousselet, S.J. Thorpe, M. Fabre-Thorpe, Taking the MAX from neuronal responses, *Trends Cogn. Sci.* 7 (3) (2003) 99–102.
- [40] S. Ullman, M. Vidal-Naquet, E. Sali, Visual features of intermediate complexity and their use in classification, *Nat. Neurosci.* 5 (7) (2002) 682–687.
- [41] J. Mutch, U. Knoblich, T. Poggio, CNS: a GPU-based framework for simulating cortically-organized networks, Tech. Rep. MIT-CSAIL-TR-2010-013 / CBCL-286, Massachusetts Institute of Technology, Cambridge, MA, Feb 2010.
- [42] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, T. Darrell, Caffe: convolutional architecture for fast feature embedding, arXiv preprint arXiv:1408.5093.
- [43] D.D. Cox, T. Dean, Neural networks and neuroscience-inspired computer vision, *Current Biol.* 24 (18) (2014) R921–R929.
- [44] J.Z. Leibo, J. Mutch, S. Ullman, T. Poggio, From primal templates to invariant recognition, MIT-CSAIL-TR-2010-057, CBCL-293, Massachusetts Institute of Technology, Cambridge, MA, 2010, pp. 1–6.
- [45] B. Lemoine, A.S. Maida, Gpu facilitated unsupervised visual feature acquisition in spiking neural networks, in: Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN), 2013, pp. 1–6.
- [46] R. VanRullen, S.J. Thorpe, Rate coding versus temporal order coding: what the retinal ganglion cells tell the visual cortex, *Neural Comput.* 13 (6) (2001) 1255–1283.

- [47] D.D. Lee, H.S. Seung, Learning the parts of objects by non-negative matrix factorization, *Nature* 401 (6755) (1999) 788–791.
- [48] T. Masquelier, S.J. Thorpe, Learning to recognize objects using waves of spikes and Spike Timing-Dependent Plasticity, in: Proceedings of the International Joint Conference on Neural Networks (IJCNN), 2010, pp. 1–8. <http://dx.doi.org/10.1109/IJCNN.2010.5596934>.
- [49] K. Carlson, M. Richert, N. Dutt, J. Krichmar, Biologically plausible models of homeostasis and stdp: Stability and learning in spiking neural networks, in: Proceedings of the International Joint Conference on Neural Networks (IJCNN), 2013, pp. 1–8.
- [50] C. Zamarreño Ramos, L.A. Camuñas Mesa, J.A. Pérez-Carrasco, T. Masquelier, T. Serrano-Gotarredona, B. Linares-Barranco, On spike-timing-dependent-plasticity, memristive devices, and building a self-learning visual cortex, *Front. Neurosci.* 5 (2011) 22.
- [51] O. Bichler, D. Querlioz, S.J. Thorpe, J.-P. Bourgoin, C. Gamrat, Extraction of temporally correlated features from dynamic vision sensors with spike-timing-dependent plasticity, *Neural Network: Official J. Int. Neural Netw. Soc.* 32 (2012) 339–348, <http://dx.doi.org/10.1016/j.neunet.2012.02.022>.
- [52] T. Dorta, M. Zapata, J. Madrenas, G. Sánchez, Aer-srt: scalable spike distribution by means of synchronous serial ring topology address event representation, *Neurocomputing* 171 (2016) 1684–1690.
- [53] C. Diaz, G. Sanchez, G. Duchon, M. Nakano, H. Perez, An efficient hardware implementation of a novel unary spiking neural network multiplier with variable dendritic delays, *Neurocomputing* (2016), <http://dx.doi.org/10.1016/j.neucom.2015.12.086>.
- [54] M. Sivilotti, Wiring considerations in analog VLSI systems with application to field-programmable networks, Ph.D. thesis, Comput. Sci. Div., California Inst. Technol., Pasadena, CA, 1991.
- [55] T. Serrano-Gotarredona, T. Masquelier, T. Prodromakis, G. Indiveri, B. Linares-Barranco, STDP and STDP variations with memristors for spiking neuromorphic learning systems, *Front. Neurosci.* 7 (2013) 2, <http://dx.doi.org/10.3389/fnins.2013.00002>.
- [56] A. Harel, S. Ullman, D. Harari, S. Bentin, Basic-level categorization of intermediate complexity fragments reveals top-down effects of expertise in visual perception, *J. Vis.* 11 (8) (2011) 18.
- [57] Y. Lerner, B. Epshtein, S. Ullman, R. Malach, Class information predicts activation by object fragments in human object areas, *J. Cogn. Neurosci.* 20 (7) (2008) 1189–1206.
- [58] T. Serre, A. Oliva, T. Poggio, A feedforward architecture accounts for rapid categorization, *Proc. Natl Acad Sci U S A* 104 (15) (2007) 6424–6429, <http://dx.doi.org/10.1073/pnas.0700622104>.
- [59] D.L.K. Yamins, H. Hong, C.F. Cadieu, E.a. Solomon, D. Seibert, J.J. Dicarlo, Performance-optimized hierarchical models predict neural responses in higher visual cortex, in: Proceedings of the National Academy of Sciences of the United States of America. <http://dx.doi.org/10.1073/pnas.1403121111>.
- [60] E.T. Rolls, G. Deco, *Computational Neuroscience of Vision*, University Press, Oxford, UK, 2002.
- [61] M.A. Ranzato, F.J. Huang, Y.-L. Boureau, Y. LeCun, Unsupervised learning of invariant feature hierarchies with applications to object recognition, in: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2007, pp. 1–8. <http://dx.doi.org/10.1109/CVPR.2007.383157>.
- [62] H. Goh, N. Thome, M. Cord, J.-H. Lim, Learning deep hierarchical visual feature coding, *IEEE Trans. Neural Netw. Learn. Syst.* (2014), <http://dx.doi.org/10.1109/TNNLS.2014.2307532>, PP 1–1.
- [63] T. Masquelier, Relative spike time coding and STDP-based orientation selectivity in the early visual system in natural continuous and saccadic vision: a computational model, *J. Comput. Neurosci.* 32 (3) (2012) 425–441, <http://dx.doi.org/10.1007/s10827-011-0361-9>.
- [64] P. Földiák, Learning invariance from transformation sequences, *Neural Comput.* 3 (1991) 194–200.
- [65] M.-J. Escobar, G.S. Masson, T. Vieville, P. Kornprobst, Action recognition using a bio-inspired feedforward spiking network, *Int. J. Comput. Vis.* 82 (3) (2009) 284–301.



Saeed Reza Kheradpisheh is currently doing his Ph.D. in Computer Science at University of Tehran. He has received B.Sc. and M.Sc. degrees in Computer Science. His research interests in the area of Computational Visual Neuroscience include invariant object recognition, spiking neural networks, and deep learning.



Mohammad Ganjtabesh received his B.Sc. degree in Pure Mathematics from the University of Tabriz in 2001, and the M.Sc. and Ph.D. degrees in computer science from the University of Tehran in 2003 and 2008, respectively. He has also performed another Ph.D. program in Bioinformatics at Ecole Polytechnique (France). Then he became an assistant professor at the University of Tehran since 2008. His current research interests include Computational Neuroscience (mainly visual cortex modeling based on spiking neural networks) and Computational Biology (all the problems associated to the RNA structures).



the C.N.R.S. in 2012.

Timothée Masquelier is a researcher in Computational Neuroscience. His research is highly interdisciplinary - at the interface between Biology, Computer Science, and Physics. He uses numerical simulations and analytical calculations to gain understanding on how the brain works, and more specifically on how neurons process, encode and transmit information through action potentials (a.k.a spikes), in particular in the visual modality. He is also interested in bio-inspired computer vision and neuromorphic engineering. Timothée Masquelier was trained at Ecole Centrale Paris (Ingénieur 1999), MIT (M. Sc. 2001), and Université Toulouse III (Ph.D. 2008). He was recruited by