



A Neural Network for Semi-supervised Learning on Manifolds

Alexander Genkin¹(✉), Anirvan M. Sengupta^{2,3}, and Dmitri Chklovskii^{1,2}

¹ Neuroscience Institute, NYU Langone Medical Center,
435 E 30th St, New York, NY 10016, USA
alexander.genkin@nyumc.org

² Flatiron Institute, Simons Foundation, 162 5th Ave, New York, NY 10010, USA
dchklovskii@flatironinstitute.org

³ Department of Physics and Astronomy, Rutgers University,
136 Frelinghuysen Rd, Piscataway, NJ 08854, USA
anirvans@physics.rutgers.edu

Abstract. Semi-supervised learning algorithms typically construct a weighted graph of data points to represent a manifold. However, an explicit graph representation is problematic for neural networks operating in the online setting. Here, we propose a feed-forward neural network capable of semi-supervised learning on manifolds without using an explicit graph representation. Our algorithm uses channels that represent localities on the manifold such that correlations between channels represent manifold structure. The proposed neural network has two layers. The first layer learns to build a representation of low-dimensional manifolds in the input data as proposed recently in [8]. The second learns to classify data using both occasional supervision and similarity of the manifold representation of the data. The channel carrying label information for the second layer is assumed to be “silent” most of the time. Learning in both layers is Hebbian, making our network design biologically plausible. We experimentally demonstrate the effect of semi-supervised learning on non-trivial manifolds.

Keywords: Semi-supervised learning · Online learning · Manifold learning

1 Introduction

When labeled data are scarce or expensive to obtain, we often resort to semi-supervised learning which exploits the abundance of unlabeled data. For data concentrating on a lower dimensional manifold, it is often reasonable to assume smoothness, i.e., that data points adjacent on the manifold are likely to have similar values of the target variable (the label). Then, learning the manifold structure from both labeled and unlabeled data can assist in label prediction. [1–3, 11].

In machine learning, an online method updates the model incrementally as it receives training data in a sequential manner. This approach is to be contrasted with offline machine learning, which generates the best model by learning on the entire training data set at once. Online learning is used either because it is computationally infeasible to train over the entire dataset, or it is used where the algorithm has to dynamically adapt to new patterns in the data, e.g. when the data itself is generated in real time. The last situation is particularly relevant in the context of neuronal networks.

Our brains likely rely on online semi-supervised learning to generate behavior. As our sensory organs stream data about the world they are analyzed in real time to produce behaviorally relevant output. While most of the sensory data lack labels, some supervision is available from other sources such as inter-personal communication.

To represent a data manifold, semi-supervised learning algorithms typically construct an adjacency graph whose vertices are labeled and unlabeled data points and edge weights represent their adjacency on the manifold. However, such representation is impractical in the online setting where the data are streamed sequentially and the labels are predicted on the fly. Furthermore, the online setting does not have the memory capacity to store the past data.

Thus, there is a need for online semi-supervised algorithms both for modeling neural computation and solving general machine learning tasks. Whereas existing online algorithms [4] can rely on a sparse representation, they still require memory quadratic in the dimensionality of data. In addition, these algorithms rely on the availability of an adjacency measure between new and stored data points.

In this paper, we propose a biologically plausible neural network for online semi-supervised learning (Fig. 1, left). By avoiding explicit representation of the adjacency graph our network can process unlimited-size datasets in online setting. Moreover, as required by biology, the network relies only on local learning rules meaning that synaptic weight update depends on the activity of only the two neurons this synapse connects.

The network has two layers. The first layer learns the manifold structure of the data by representing each datum as a sparse vector whose components represent overlapping localities on the manifold. The manifold structure is captured by the correlations between the components carried by corresponding channels. Because most existing algorithms for sparse representations, such as [10], do not have natural neural implementations we base our work on the recently developed manifold tiling algorithm [8]. Inspiration for such design comes from biological neural networks such as place cells in the rodent hippocampus.

The second layer learns a classifier using both occasional supervision and the similarity of the manifold representation of the data provided by the first layer. In our neural network, the supervision signal is not fed back from downstream layers of the network like in perceptron or back-propagation networks, but comes along and synchronously with the data from the previous layer. To make it semi- (rather than fully) supervised, the label signal is assumed to be “silent” most of

the time. The output attempts to predict the correct label when that signal is not available, otherwise it just reproduces the label.

We derive both the activity dynamics and the learning rules in each layer from the principle of similarity preservation [6] which was previously used in the unsupervised setting. Starting with a similarity preserving objective function allows us to analyze the output of the algorithm and obtain biologically plausible local learning rules.

We demonstrate experimentally the effectiveness of this semi-supervised network compared to fully supervised online learning. Moreover, we observe that online semi-supervised learning may be competitive with offline methods, especially on smaller samples. This is an important advantage allowing our network to adapt quickly when the manifold shape or the labels are changing with time.

2 Review of the Manifold-Tiling Network Derived from Non-negative Similarity Matching

To introduce our notation, let the input to the network be a set of vectors, $\mathbf{x}_t \in R^n, t = 1, \dots, T$, coming from n channels at time t . In response, the manifold learning network layer outputs an activity vector, $\mathbf{h}_t \in R^m, t = 1, \dots, T$, m being the number of output channels, or hidden units in our two-layer network, Fig. 1, left.

Manifold-tiling networks have been derived [8] from similarity-preserving objectives [5] with a non-negativity constraint. Similarity preservation postulates that similar input pairs, \mathbf{x}_t and $\mathbf{x}_{t'}$, evoke similar output pairs, \mathbf{h}_t and $\mathbf{h}_{t'}$. Similarity of a pair of vectors can be quantified by their scalar product. Nonlinear manifolds can be learned by constraining the sign of the output and introducing a similarity threshold α . [8] propose an optimization problem:

$$\begin{aligned} & \min_{\substack{\mathbf{H} \geq 0 \\ \text{diag} \mathbf{H}^\top \mathbf{H} \leq \mathbf{I}}} -\text{Tr}((\mathbf{X}^\top \mathbf{X} - \alpha \mathbf{E}) \mathbf{H}^\top \mathbf{H}) \tag{1} \\ & = \min_{\mathbf{h}_t \geq 0, \|\mathbf{h}_t\|_2 \leq 1} - \sum_{t,t'} (\mathbf{x}_t^\top \mathbf{x}_{t'} - \alpha) \mathbf{h}_t^\top \mathbf{h}_{t'}. \end{aligned}$$

Here matrix notation was introduced: $\mathbf{X} \equiv [\mathbf{x}_1, \dots, \mathbf{x}_T] \in R^{n \times T}$ and $\mathbf{H} \equiv [\mathbf{h}_1, \dots, \mathbf{h}_T] \in R^{m \times T}$, and \mathbf{E} is a matrix of all ones.

Intuitively, (1) attempts to preserve similarity for similar pairs of input samples but orthogonalizes the outputs corresponding to dissimilar input pairs. Indeed, if the input similarity of a pair of samples t, t' is above a specified threshold, $\mathbf{x}_t^\top \mathbf{x}_{t'} > \alpha$, then the output vectors \mathbf{h}_t and $\mathbf{h}_{t'}$ would prefer to have $\mathbf{h}_t^\top \mathbf{h}_{t'} \approx \mathbf{x}_t^\top \mathbf{x}_{t'} - \alpha$, i.e., they would be similar. If, however, $\mathbf{x}_t^\top \mathbf{x}_{t'} < \alpha$, then they would tend to be orthogonal, $\mathbf{h}_t^\top \mathbf{h}_{t'} = 0$, since the lowest value of $\mathbf{h}_t^\top \mathbf{h}_{t'}$ for $\mathbf{h}_t, \mathbf{h}_{t'} \geq 0$ is zero. As \mathbf{h}_t and $\mathbf{h}_{t'}$ are nonnegative, to achieve orthogonality, the output activity patterns for dissimilar patterns would have non-overlapping sets of active output channels. In the context of manifold representation, (1) strives to preserve in the \mathbf{h} -representation the local geometry of the input data cloud




Fig. 1. *Left:* Two-layer network for semi-supervised learning. The first layer learns manifolds (the upper layer with interneurons in red), the second layer is one neuron that learns a classifier in the semi-supervised manner. The output of the network predicts the label value when the labels channel is silent, otherwise it reproduces the label. *Right:* Receptive fields of manifold tiling. Data are 2000 points sampled from two arcs in 2D (black dots); in the third dimension we show responses of individual channels to the corresponding data points, each channel is assigned a color (Color figure online).

in \mathbf{x} -space and let the global geometry emerge out of the nonlinear optimization process.

Figure 1 illustrates manifold tiling on a two spiral arcs in two dimensions, showing the receptive fields of output channels in the third dimension. Receptive fields tile the arcs with overlaps, but there is no overlap between separate arcs.

To derive a neural network that optimizes (1), we express the norm constraint in the Lagrangian form:

$$\min_{\forall t: \mathbf{h}_t \geq 0} \max_{\forall t: \mathbf{u}_t \geq 0} - \frac{1}{T} \sum_{t,t'} (\mathbf{x}_t^\top \mathbf{x}_{t'} - \alpha) \mathbf{h}_t^\top \mathbf{h}_{t'} + \sum_t \mathbf{u}_t^\top \mathbf{u}_t (1 - \mathbf{h}_t^\top \mathbf{h}_t). \tag{2}$$

Here, unconventionally, the non-negative Lagrange multipliers that impose the inequality constraints are factorized into inner products of two non-negative vectors $(\mathbf{u}_t^\top \mathbf{u}_t)$. In the second step, we introduce auxiliary variables, \mathbf{W} , \mathbf{b} , \mathbf{V}_t [7]:

$$\min_{\forall t: \mathbf{h}_t \geq 0} \max_{\forall t: \mathbf{u}_t \geq 0} \min_{\mathbf{W}} \max_{\mathbf{b}} \min_{\forall t: \mathbf{V}_t \geq 0} T \text{Tr}(\mathbf{W}^\top \mathbf{W}) - T \|\mathbf{b}\|_2^2 + \sum_t \left(-2\mathbf{x}_t \mathbf{W}^\top \mathbf{h}_t + 2\sqrt{\alpha} \mathbf{h}_t^\top \mathbf{b} + \|\mathbf{u}_t\|_2^2 - 2\mathbf{u}_t \mathbf{V}_t \mathbf{h}_t + \text{Tr}(\mathbf{V}_t^\top \mathbf{V}_t) \right) \tag{3}$$

The equivalence of (3) to (2) can be seen by performing the \mathbf{W} , \mathbf{b} , and \mathbf{V}_t optimizations explicitly and plugging back in the optimal values. Equation (3) suggests a two-step online algorithm (see [8] for full derivation). For each input \mathbf{x}_t , in the first step, one solves for \mathbf{h}_t , \mathbf{u}_t and \mathbf{V}_t , by projected gradient descent-descent,

$$\begin{bmatrix} \mathbf{h}_t \\ \mathbf{u}_t \\ \mathbf{V}_t \end{bmatrix} \leftarrow \begin{bmatrix} \mathbf{h}_t + \gamma_h (\mathbf{W}\mathbf{x}_t - \mathbf{V}_t^\top \mathbf{u}_t - \sqrt{\alpha}\mathbf{b}) \\ \mathbf{u}_t + \gamma_u (-\mathbf{u}_t + \mathbf{V}_t \mathbf{h}_t) \\ \mathbf{V}_t + \gamma_V (\mathbf{u}_t \mathbf{h}_t^\top - \mathbf{V}_t) \end{bmatrix}_+, \quad (4)$$

where $\gamma_{h,u,V}$ are step sizes. This iteration can be interpreted as the dynamics of a biologically plausible neural circuit (Fig. 1, right, the upper layer), where components of \mathbf{h}_t are activities of excitatory neurons, \mathbf{b} is a bias term, components of \mathbf{u}_t are activities of inhibitory neurons (shown in red), and \mathbf{W} is the feedforward connectivity matrix. \mathbf{V}_t is the synaptic weight matrix from excitatory to inhibitory neurons, which undergoes a fast time-scale anti-Hebbian plasticity, which in computer simulation means repeated updates within one t step. In the second step, \mathbf{W} and \mathbf{b} are updated by gradient descent-ascent:

$$\mathbf{W} \leftarrow \mathbf{W} + \eta (\mathbf{h}_t \mathbf{x}_t^\top - \mathbf{W}), \quad \mathbf{b} \leftarrow \mathbf{b} + \eta (\sqrt{\alpha} \mathbf{h}_t - \mathbf{b}),$$

where \mathbf{W} is going through a slow time-scale Hebbian plasticity and \mathbf{b} through homeostatic plasticity. The parameter η is a learning rate.

3 A Neural Network for Semi-supervised Learning

In this section, we propose a neural network architecture for semi-supervised learning. In our approach, contrary to the widely accepted schemes, the label signal is not fed back from downstream layers of the network but comes along and synchronously with the rest of the data. To make it semi- (rather than fully) supervised, the signal is assumed to be “silent” most of the time.

Consider a classification problem with the input stream of data, $\{\mathbf{h}_1, \dots, \mathbf{h}_t, \dots\}$, where $\mathbf{h}_t \in R^m$, and the corresponding class labels $\{\tilde{z}_1, \dots, \tilde{z}_t, \dots\}$, where in a binary case $\tilde{z}_t \in \{-1, +1\}$. The labels are occasionally signalled by a channel carrying values $z_t = \theta_t \tilde{z}_t$, where $\theta_t \in \{0, 1\}$ either masks or reveals the true label. The data from the previous layer and the label channel are combined in the semi-supervised learning neuron, Fig. 1, right, bottom layer.

Consider a time period of $1, \dots, T$, where the inputs are organized into a matrix $\mathbf{H} = [\mathbf{h}_1, \dots, \mathbf{h}_T]$ and a vector of (partly hidden) labels $\mathbf{z}^\top = (z_1, \dots, z_T)$. The output $\mathbf{y}^\top = (y_1, \dots, y_T)$ needs to reproduce the label signal, so we employ a quadratic loss function $\|\mathbf{y} - \mathbf{z}\|^2$. We express the assumption of smoothness of predicted label \mathbf{y} on the manifold using the similarity alignment [7] between the input and output Gramians: $\text{Tr}(\mathbf{H}^\top \mathbf{H} \mathbf{y} \mathbf{y}^\top)$. Also, as the label only takes values 1 and -1 , we restrict the output to stay within those limits. This gives rise to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{y}} \|\mathbf{y} - \mathbf{z}\|^2 - \frac{\mu}{T} \text{Tr}(\mathbf{H}^\top \mathbf{H} \mathbf{y} \mathbf{y}^\top), \\ \text{s.t.} \quad -1 \leq y_t \leq 1, \quad t = 1, \dots, T, \end{aligned} \quad (5)$$

where we also introduced a regularization coefficient μ controlling the relative importance of the two parts of the objective function.

To derive an online algorithm, following [7], we introduce an auxiliary variable \mathbf{w} and expand in time:

$$\min_{\mathbf{y}} \min_{\mathbf{w}} \frac{1}{T} \sum_t \left[\frac{1}{2} (y_t - z_t)^2 - \mu y_t \mathbf{w}^\top \mathbf{h}_t \right] + \frac{\mu}{2} \mathbf{w}^\top \mathbf{w} \tag{6}$$

s.t. $-1 \leq y_t \leq 1, \quad t = 1 \dots T.$

Optimizing over \mathbf{w} , we obtain: $\mathbf{w} = \frac{1}{T} \sum_t y_t \mathbf{h}_t$, which makes it clear the new formulation is equivalent to Eq. (5). The advantage of this formulation is that it suggests a two-step online algorithm. For each input \mathbf{h}_t , on the first step, one solves for the instantaneous output y_t under fixed \mathbf{w} :

$$y_t = \max(-1, \min(1, \mu \mathbf{w}_t^\top \mathbf{h}_t + z_t)) \tag{7}$$

On the second step, \mathbf{w} is updated as:

$$\mathbf{w} \leftarrow \frac{t}{t+1} \mathbf{w} + \frac{1}{t+1} y_t \mathbf{h}_t \tag{8}$$

This also maps well onto a biologically plausible neural network where components of \mathbf{w} are interpreted as synapse weights, updated by local Hebbian rule. We assume that the synapse weight of the z channel is not changing, thus differentiating it from the other input channels. We set this weight to be equal to 1 without loss of generality. The algorithm is initialized with $\mathbf{w} = 0$, assuming no prior information.

An alternative objective function can be obtained by expressing the loss as $-\mathbf{y}^\top \mathbf{z}$ and adding an entropy-like regularizer treating $(y_t + 1)/2$ as a probability estimate for $\tilde{z}_t = 1$:

$$\min_{\mathbf{y}} -\mathbf{y}^\top \mathbf{z} - \frac{\mu}{2T} \text{Tr}(\mathbf{H}^\top \mathbf{H} \mathbf{y} \mathbf{y}^\top) - \sum_t \left[\frac{1+y_t}{2} \log\left(\frac{1+y_t}{2}\right) + \frac{1-y_t}{2} \log\left(\frac{1-y_t}{2}\right) \right] \tag{9}$$

The solution of this optimization problem is the familiar sigmoidal neuron rule:

$$y_t = \tanh(\mu \mathbf{w}_t^\top \mathbf{h}_t + z_t) \tag{10}$$

with the same update for \mathbf{w} as in Eq. (8). The behavior of both algorithms is almost indistinguishable, so we only report the results from Eq. (7).

4 Numerical Experiments

We apply our algorithm to the synthetic dataset designed as “two moons”: two classes are sets of points in 2D, each concentrated around a spiral arc, Fig. 2, top. Such a synthetic dataset is widely used as a test for semi-supervised learning algorithms (see, e.g., [2, 4]). Note that the classes are not linearly separable,

and can be separated only when their manifold structure is discovered. Upon discovering the manifold structure, intuitively, the data can be classified using only one labeled example for each class, see red asterisks in Fig. 2, top.

Our network solves this highly non-linear classification problem. The first layer learns units that tile each “moon” with overlaps while no unit is shared between the two moons. The second layer propagates labeling information along links formed by correlations in the tile responses. We generated data points randomly and uniformly, only placing two labeled points early in the data stream. We used tiling layer with 40 neurons and semi-supervised neuron with $\mu = 1000$. The Fig. 2, bottom row, illustrates the working of the semi-supervised neuron. As seen on Fig. 2, bottom left, the output is zero until labeled points arrive. Then there is a transition period, during which the label signals propagate along correlated tiles. Finally, the responses stabilize to correct values: 1 for green, -1 for blue. Figure 2, bottom right, illustrates propagation of the labeling information. Initially, all weights are zero. When a labeled point arrives, weights corresponding to the tiles overlapping this point increase in absolute value. That signal gradually propagates until all synapses corresponding to “green moon” get positive weights and, all “blue” ones - negative weights.

Next, we apply our network to a larger dataset, a 3D Chessboard on a Swiss roll, Fig. 3, left. All the data live on a 2D Swiss roll manifold and the two classes are defined by the squares of the chessboards. We consider chessboards with varying square sizes with the most fine-grained chessboard being most difficult for classification. Whereas linear classifiers per se can not solve this problem, after learning the manifold classification is linear.

We compare our semi-supervised algorithm with an online fully supervised classifier – logistic regression. Both algorithms get the same input stream of 2000 data points, of which 50, 100, or 200 randomly selected are labeled and the rest are unlabeled. The input for both classification algorithms is the output of tiling with 200 neurons. Parameters μ for our neuron and learning rate for logistic regression are selected for best results of each algorithm. All runs repeated 10 times to obtain error bars.

Both algorithms classify each input using their current weights. However, the fully supervised algorithm cannot update its weights when an unlabeled example arrives, unlike the semi-supervised algorithm. Indeed, experiments show that the semi-supervised neural network performs better than the supervised classifier (Fig. 3, center), demonstrating its ability to take advantage of unlabeled data.

Next, we compare our online algorithm with an offline semi-supervised learning algorithm. For the latter, we use a state of the art linear SVM with Laplacian penalty following [2], but with a twist ¹: for the linear case we assume smoothness of weights rather than labels. This means that components of the separating vector \mathbf{w} should have similar values when corresponding tiling components have highly overlapping receptive fields. The degree of overlap between receptive fields can be measured by dot products between tiling components, and can be calculated on both labeled and unlabeled data. Then the Gramian $S = \frac{1}{T} \mathbf{H} \mathbf{H}^\top$

¹ In a separate experiment we made sure this twist only improves the results.

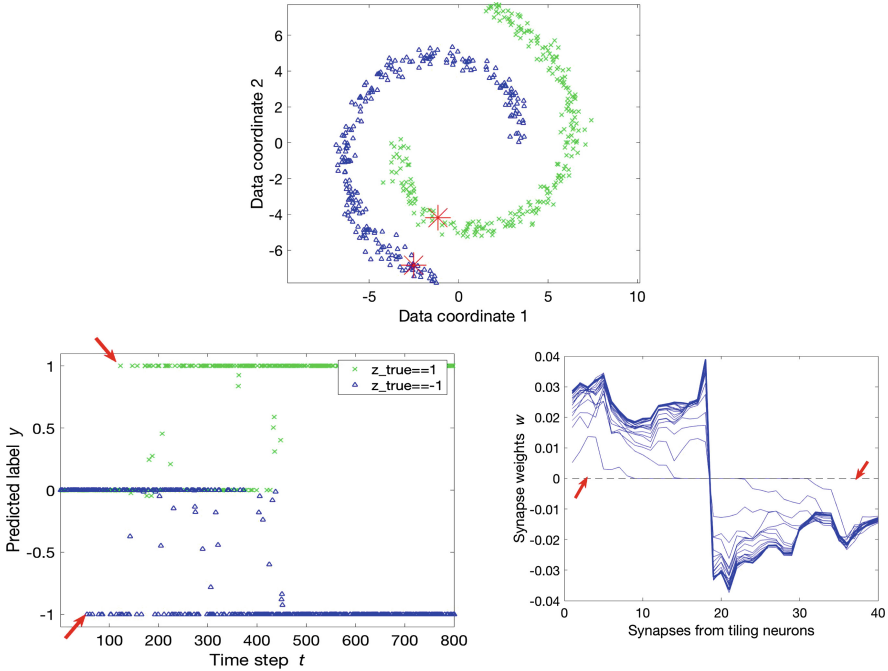


Fig. 2. Semi-supervised learning on the “two moons” dataset with two labeled points. *Top:* “Two moons” in 2D, classes - green crosses and blue triangles; red asterisks indicate two labeled points. *Bottom left:* Label predicted by the semi-supervised neuron (y_t values) in time, green crosses and blue triangles as the true class of the input. Labeled points indicated by arrows. *Bottom right:* Propagation of labels is shown by the weights of the tiles in time. Tiles are ordered along the x axis according to their locations on the “moons”: “crosses” on the left and “triangles” - on the right. Lines show temporal dynamics of the weights, and only every 100th time point is shown. Arrows indicate tiles where labeled points fall.

can be thought of as the adjacency matrix of a graph where vertices are tiling components. The graph will be fairly sparse due to the nature of tiling. The graph Laplacian penalty is then:

$$\sum_{i,j} (w_i - w_j)^2 S_{i,j} = 2\mathbf{w}^\top L\mathbf{w}, \quad \text{where } L = \text{diag}(S\mathbf{1}) - S, \quad (11)$$

and the objective function of linear SVM with Laplacian penalty takes the form:

$$\min_{\mathbf{w}, b} \sum_t [1 - z_t(\mathbf{w}^\top \mathbf{h}_t + b)]_+ + \lambda \|\mathbf{w}\|^2 + \mu \mathbf{w}^\top L\mathbf{w} \quad (12)$$

where the index t runs through the labeled samples only, with $z_t \in \{-1, 1\}$ being labels.

In this experiment, the online algorithm is fed a data stream with 0.05% of samples randomly labeled. Then at every 500th step the classification rule

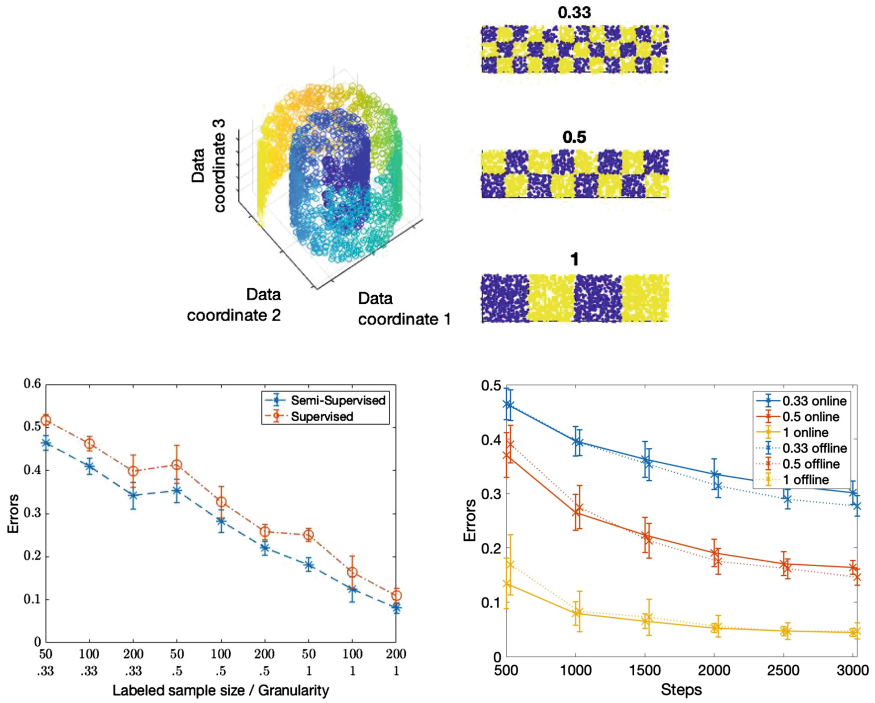


Fig. 3. *Left:* 3D Swiss roll manifold and three binary classification problems on the unrolled manifold. Granularity of the chessboard decreases from top to bottom. *Center:* Semi- vs. fully supervised: comparison of our semi-supervised algorithm with a fully supervised algorithm in the online setting. *Right:* Online vs. offline: comparison of our online semi-supervised learning algorithm with the offline semi-supervised algorithm (see text) for different granularity levels of the chessboard.

obtained up to this point is applied to a separate test set of 2000 samples. At the same step, the SVM with Laplacian regularization, Eq. (12), is trained on all data seen online so far and tested against the same test set. As before, the input for both algorithms is the output of tiling with 200 neurons. Parameters μ for our neuron and learning rate for logistic regression are selected for best results of each algorithm. All runs repeated 10 times to obtain error bars.

Offline algorithm has an advantage of considering all data samples before taking decision on labeling, while online algorithm has to assign a label estimate to each data sample as it appears. Results on Fig. 3, right, show, however, that with enough smoothness (i.e., coarser granularity in the “Chessboard” example), the online algorithm perform closely to the offline one. Moreover, online algorithm can perform better than the offline one while the number of presented data points is small (e.g., less than approximately 1200 with granularity 0.5). But small sample sizes is exactly the situation where semi-supervised learning is supposed to be helpful. The ability of the online algorithm to adapt quickly is also important when there is a drift in the manifold shape or the labels.

5 Relation to Graph Laplacian

Existing algorithms for semi-supervised learning on manifolds typically utilize the graph Laplacian for smoothness regularization [2, 3, 11], see the last term in Eq. (12). This follows from the analysis of [9], which showed that graph Laplacian regularization results in classifier corresponding to normalized graph cut, which helps avoid heavily imbalanced classes. In contrast, our smoothness term, last term, in Eq. (6), lacks diagonal normalization of Laplacian. When optimized exactly, it should lead to the minimum cut of the graph, which is prone to generate classes of very different size [9]. Consider a simple example of a square, where two labeled points for two classes are close to diagonally opposite corners. Laplacian regularization would cut the square in half approximately along the other diagonal, Fig. 4, left, while the minimum graph cut would lead to highly asymmetric solution: one predicted label concentrates closely around one of the labeled points, all the rest occupied by the other label.

However in our experiments we very rarely observe this trend towards asymmetrical solutions. To develop an intuition for why this happens, consider a period in the learning process during which labels channel is silent ($z_t = 0$), and y_t is not reaching the limits yet. This is the decisive period, where the label information propagates between the synapse weights, see Fig. 2, bottom left. Then (7) becomes simply $y_t = \mu \mathbf{w}_t^\top \mathbf{h}_t$. Assume the input points arrive i.i.d., then so are \mathbf{h}_t vectors. Then substituting expression for y_t into (8) we can write an expectation for one component of \mathbf{w} :

$$\mathbb{E}(\Delta w_i) = \eta \left(\sum_j \mu \mathbb{E}(h_i h_j) w_j - w_i \right) = \eta \left(\mu \sum_j s_{i,j} w_j - w_i \right), \quad (13)$$

where we defined $s_{i,j} \equiv \mathbb{E}(h_i h_j)$, and $\eta = 1/(t + 1)$. Now $(s_{i,j})$ can be seen as the adjacency matrix of a weighted graph, where vertices are tiling channels, in a manner analogous to the matrix $(S_{i,j})$, appearing in Eq. (11) in the previous section. The term $\mu \sum_j s_{i,j} w_j$, appearing in the right hand side of Eq. (13), has the effect of “smoothing” out \mathbf{w} in over the tiling channels, in a manner analogous to the effect of the Laplacian penalty presented in (12). Essentially, the Laplacian penalty causes the components of \mathbf{w} diffuse over the graph [11]. So, in expectation, the evolution of \mathbf{w} in our algorithm would share some features with the gradient descent of \mathbf{w} to optimize the expression in (12). “Smoothness” of the resulting \mathbf{w} over channels translate to smoothness of prediction over input space, thereby reducing the likelihood of extremely imbalanced solutions.

We illustrate this with a simulation experiment on the square in Fig. 4. Ideally, there should be equal number of predicted labels for both classes. We, therefore, measure the imbalance, by looking at fraction associated with the majority class among predictions. This measure of imbalance ranges from 0.5 to 1.0. For each run, we generate 2000 unlabeled sample points uniformly on the square, plus 2 labeled points near the corners. These data were fed to our network with 50 tiling channels and $\mu = 10$ in (6). For comparison, the output of tiling layer is also used as input to linear classifier with Laplacian regularization. The histogram of results, after 100 such runs, is presented in Fig. 4, right.

While indeed the results for our network fluctuate more, compared to those of the Laplacian regularization approach, the extreme imbalances are rare in both approaches.

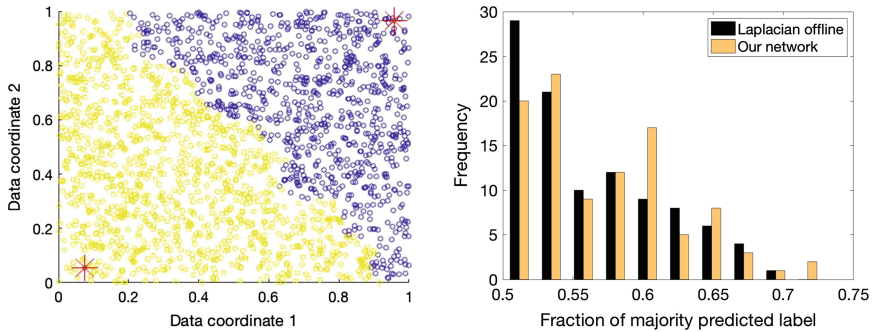


Fig. 4. *Left:* Linear classification with Laplacian regularization, points colored by predicted label; labeled points are red arrows. *Right:* The histogram of 100 repetitions of simulation of our measure of imbalance, namely, the fraction of the majority predicted label for Laplacian solution offline and for our network. Note that the frequencies fall off in both methods as the fraction moves up from 0.5.

6 Conclusion

We presented a neural network that learns low-dimensional manifolds in the data stream, then learns a classifier in a semi-supervised setting, where only small part of inputs are labeled. The network operates in an online fashion, producing an output immediately after seeing every input. Weights are updated by a biologically plausible local Hebbian-type rule. We demonstrated the effectiveness of the network in simulations, comparing it with fully supervised online algorithm and with a semi-supervised offline algorithm.

Acknowledgements. The authors are grateful to Victor Minden and Mariano Tepper for their insightful comments. We thank Johannes Friedrich, Tiberiu Tesileanu and Charles Windolf for helpful discussions.

References

1. Ando, R.K., Zhang, T.: Learning on graph with laplacian regularization. In: Advances in Neural Information Processing Systems, pp. 25–32 (2007). <https://doi.org/10.7551/mitpress/7503.003.0009>
2. Belkin, M., Niyogi, P., Sindhwani, V.: Manifold regularization: a geometric framework for learning from labeled and unlabeled examples. *J. Mach. Learn. Res.* **7**, 2399–2434 (2006)
3. Bengio, Y., Delalleau, O., Le Roux, N.: Label propagation and quadratic criterion. In: Semi-Supervised Learning. MIT Press (2006). <https://doi.org/10.7551/mitpress/9780262033589.001.0001>

4. Goldberg, A.B., Li, M., Zhu, X.: Online manifold regularization: a new learning setting and empirical study. In: Daelemans, W., Goethals, B., Morik, K. (eds.) ECML PKDD 2008. LNCS (LNAI), vol. 5211, pp. 393–407. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-87479-9_44
5. Pehlevan, C., Chklovskii, D.: A normative theory of adaptive dimensionality reduction in neural networks. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 28, pp. 2269–2277. Curran Associates, Inc. (2015)
6. Pehlevan, C., Hu, T., Chklovskii, D.: A hebbian/anti-hebbian neural network for linear subspace learning: a derivation from multidimensional scaling of streaming data. *Neural Comput.* **27**, 1461–1495 (2015). https://doi.org/10.1162/neco_a.00745
7. Pehlevan, C., Sengupta, A.M., Chklovskii, D.B.: Why do similarity matching objectives lead to hebbian/anti-hebbian networks? *Neural Comput.* **30**(1), 84–124 (2018). https://doi.org/10.1162/neco_a.01018
8. Sengupta, A., Pehlevan, C., Tepper, M., Genkin, A., Chklovskii, D.: Manifold-tiling localized receptive fields are optimal in similarity-preserving neural networks. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) Advances in Neural Information Processing Systems, vol. 31, pp. 7080–7090. Curran Associates, Inc. (2018)
9. Shi, J., Malik, J.: Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(8) (2000). <https://doi.org/10.1109/cvpr.1997.609407>
10. Yu, K., Zhang, T., Gong, Y.: Nonlinear learning using local coordinate coding. In: Advances in Neural Information Processing Systems, pp. 2223–2231 (2009)
11. Zhu, X., Ghahramani, Z., Lafferty, J.D.: Semi-supervised learning using Gaussian fields and harmonic functions. In: Proceedings of the 20th International Conference on Machine Learning (ICML-03), pp. 912–919 (2003)