

BASiS: Batch Aligned Spectral Embedding Space

Or Streicher Ido Cohen Guy Gilboa

Viterbi Faculty of Electrical and Computer Engineering
 Technion - Israel Institute of Technology, Haifa, Israel

orr.shtr@gmail.com ido.coh@gmail.com guy.gilboa@ee.technion.ac.il

Abstract

Graph is a highly generic and diverse representation, suitable for almost any data processing problem. Spectral graph theory has been shown to provide powerful algorithms, backed by solid linear algebra theory. It thus can be extremely instrumental to design deep network building blocks with spectral graph characteristics. For instance, such a network allows the design of optimal graphs for certain tasks or obtaining a canonical orthogonal low-dimensional embedding of the data. Recent attempts to solve this problem were based on minimizing Rayleigh-quotient type losses. We propose a different approach of directly learning the graph's eigenspace. A severe problem of the direct approach, applied in batch-learning, is the inconsistent mapping of features to eigenspace coordinates in different batches. We analyze the degrees of freedom of learning this task using batches and propose a stable alignment mechanism that can work both with batch changes and with graph-metric changes. We show that our learnt spectral embedding is better in terms of NMI, ACC, Grassman distance, orthogonality and classification accuracy, compared to SOTA. In addition, the learning is more stable.

1. Introduction

Representing information by using graphs and analyzing their spectral properties has been shown to be an effective classical solution in a wide range of problems including clustering [8, 21, 32], classification [13], segmentation [26], dimensionality reduction [5, 10, 23] and more. In this setting, data is represented by nodes of a graph, which are embedded into the eigenspace of the graph-Laplacian, a canonical linear operator measuring local smoothness.

Incorporating analytic data structures and methods within a deep learning framework has many advantages. It yields better transparency and understanding of the network, allows the use of classical ideas, which were thoroughly investigated and can lead to the design of new architectures, grounded in solid theory. Spectral graph algo-

gorithms, however, are hard to incorporate directly in neural networks since they require eigenvalue computations which cannot be integrated in back-propagation training algorithms. Another major drawback of spectral graph tools is their low scalability. It is not feasible to hold a large graph containing millions of nodes and to compute its graph-Laplacian eigenvectors. Moreover, updating the graph with additional nodes is cumbersome and one usually resorts to graph-interpolation techniques, referred to as Out Of Sample Extension (OOSE) methods.

An approach to solve the above problems using deep neural networks (DNNs), firstly suggested in [24] and recently also in [9], is to train a network that approximates the eigenspace by minimizing Rayleigh quotient type losses. The core idea is that the Rayleigh quotient of a sum of n vectors is minimized by the n eigenvectors with the corresponding n smallest eigenvalues. As a result, given the features of a data instance (node) as input, these networks generate the respective coordinate in the spectral embedding space. This space should be equivalent in some sense to the analytically calculated graph-Laplacian eigenvector space. A common way to measure the equivalence of these spaces is using the Grassman distance. Unfortunately, applying this indirect approach does not guarantee convergence to the desired eigenspace and therefore the captured might not be faithful.

An alternative approach, suggested in [18] for computing the diffusion map embedding, is a direct supervised method. The idea is to compute the embedding analytically, use it as ground-truth and train the network to map features to eigenspace coordinates in a supervised manner. In order to compute the ground truth embedding, the authors used the entire training set. This operation is very demanding computationally in terms of both memory and time and is not scalable when the training set is very large.

Our proposed method is to learn directly the eigenspace in batches. We treat each batch as sampling of the full graph and learn the eigenvector values in a supervised manner. A major problem of this kind of scheme is the inconsistency in the embedding coordinates. Thus, two instances

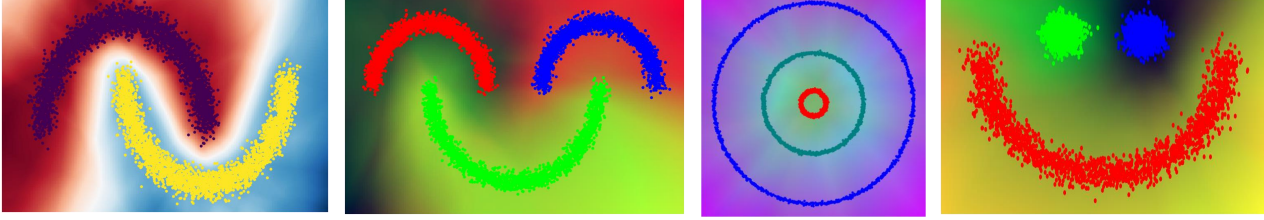


Figure 1. **Toy examples.** An Illustration of trained BASiS models over common spectral-clustering toy examples. Each figure describes the embedding values given by the network to each point in the space and the clustering results over selected points. BASiS performs successful clustering and is able to interpolate and extrapolate the training data smoothly.

in different batches with the same features can be mapped to very different eigenspace coordinates. Our solution is to use affine registration techniques to align the batches. Further, we use this alignment strategy to also allow changes in the graph affinity metric. Our proposed method retains the following main qualities: 1) **Scalability.** Data is learnt in batches, allowing a training based on large and complex input sets; 2) **OOSE.** Out of sample extension is immediate. 3) **High quality approximation of the eigenspace.** Since our learning method is direct and fully supervised, an excellent approximation of the graph eigenspace is obtained. 4) **Robustness to features change.** We can train the model also when features and affinities between nodes change; All the above properties yield a spectral building block which can be highly instrumental in various deep learning algorithms, containing an inherent orthogonal low dimensional embedding of the data, based on linear algebraic theory.

2. Settings and Notations

Let $\{x_i\}_{i=1}^n$ be a set of data instances denoted as X which is a finite set in \mathbb{R}^d . These samples are assumed to lie on a lower dimensional manifold \mathcal{M} .

These instances are represented as nodes on an undirected weighted graph $G = (V, E, W)$, where V and E are sets of the vertices and edges, respectively, and W is the adjacency matrix. This matrix is symmetric and defined by a distance measure between the nodes. For example, a common choice is a Gaussian kernel and Euclidean distance,

$$W_{ij} = \exp\left(-\frac{\|x_i - x_j\|_2^2}{2\sigma^2}\right), \quad (1)$$

where σ is a soft-threshold parameter.

The degree matrix D is a diagonal matrix where D_{ii} is the degree of the i -th vertex, i.e., $D_{ii} = \sum_j W_{ij}$. The graph-Laplacian operator is defined by,

$$L := D - W. \quad (2)$$

The graph-Laplacian is a symmetric, positive semi-definite matrix, its eigenvalues are real, and its eigenvectors form

an orthogonal basis. The eigenvalues of L are sorted in ascending order $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$, where the corresponding eigenvectors are denoted by u_1, u_2, \dots, u_n . The sample x_i is represented in the spectral embedding space as the i th row of the matrix $U = \begin{bmatrix} u_1 & \dots & u_K \end{bmatrix} \in \mathbb{R}^{n \times K}$, denoted as ϕ_i . Thus, more formally, the dimensionality reduction process can be formulated as

$$x_i \mapsto \phi_i = [u_1(i), u_2(i), \dots, u_K(i)] \in \mathbb{R}^K, \quad (3)$$

where $K \ll d$. This representation preserves well essential data information [10, 15, 17, 22]

Alternatively, one can replace the Laplacian definition (2) with

$$L_N := D^{-\frac{1}{2}} L D^{-\frac{1}{2}} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}. \quad (4)$$

This matrix may yield better performances for certain tasks and datasets [25–27].

3. Related Work

OOSE and scalability of graph-based learning methods are ongoing research topics. Mathematical analyses and analytical solutions to these problems can be found, for example, in [1, 4, 7, 12, 30]. However, neural networks learning the latent space of the data usually yield an efficient, robust and reliable solution for these problems. Moreover, neural network modules can be easily integrated in larger networks, employing this embedding. For a recent use of learnable graphs in semi-supervised learning and data visualization see [2]. The effectiveness of modeling PDE's and certain eigenproblems in grid-free, mesh-free manner was shown in [3, 6, 29]. We review below the main advances in eigenspace embedding.

Diffusion Nets [18]. Diffusion Maps (DM) is a spectral embedding, resulting from the eigendecomposition of

$$P := W D^{-1}, \quad (5)$$

known as the random-walk matrix [10]. More formally, similarly to Eq. (3), diffusion maps is defined by

$$x_i \mapsto \phi_i = [\gamma'_1 \Phi_1(i), \gamma'_2 \Phi_2(i), \dots, \gamma'_K \Phi_K(i)] \in \mathbb{R}^K, \quad (6)$$

where $\{\Phi_j\}_{j=1}^K$ are the first non-trivial eigenvectors of P , $\{\gamma\}_{i=j}^K$ are the corresponding eigenvalues and $t > 0$ is the diffusion time. Note, that P and L_N have the same eigenvectors, in reverse order with respect to their eigenvalues.

Diffusion Net (DN) is an autoencoder trained to map between the data and the DM. The loss function of the encoder is defined by,

$$\mathcal{L}_{DN}^e(\theta^e) = \frac{1}{2n} \sum_{i=1}^n \|f_{\theta^e}^e(x_i) - \phi_i\|^2 + F(\theta^e, X), \quad (7)$$

where θ^e denotes the encoder's parameters, $f_{\theta^e}^e(x_i)$ is the encoder output and $F(\theta^e, X) = \frac{\mu}{2} \sum_{l=1}^{L-1} \|\theta_{e_l}^e\|_F^2 + \frac{\eta}{2m} \sum_{j=1}^d \|(P - \gamma_j I)(o_j^e)^T\|^2$ is a regularization term such that $\theta_{e_l}^e$ are the weights of the l -th layer, o_j^e is the j -th column of the output matrix, μ and η are regularization parameters. Note, Diffusion Net requires to compute the embedding of the training set in advance, meaning *it cannot be trained with mini-batches* and therefore has difficulty dealing with large datasets.

SpectralNet1 [24] (SpecNet1). This DNN learns the embedding corresponds to L by minimizing the *ratio-cut* loss of Ng *et al.* [21], without adding an orthogonality constraint on the solution, with the loss

$$\mathcal{L}_{SN1}(\theta) = \frac{1}{m^2} \sum_{i,j=1}^m W_{i,j} |y_i - y_j|^2 = \frac{2}{m^2} \text{tr}(Y^T L Y), \quad (8)$$

where $y_i = f_{\theta}(x_i)$ is the network output, m is the batch size, and tr is the trace operator. In order to calculate the eigenvectors of L_N , one should normalize y_i, y_j with the corresponding node degree. In SpectralNet1 orthogonality of the training is gained by defining the last layer of the network as a linear layer set to orthogonalize the output. The last layer's weights are calculated during training with QR decomposition over the DNN's outputs. The authors point out that in order to get good generalization and approximate orthogonal output at inference, large batches are required.

SpectralNet2 [9] (SpecNet2). In this recent work the authors suggested to solve the eigenpair problem of the matrix pencil (W, D) . The loss function is defined by,

$$\mathcal{L}_{SN2}(\theta) = \frac{1}{m^2} \text{tr} \left(-2Y^T W Y + \frac{1}{m^2} Y^T D Y Y^T D Y \right), \quad (9)$$

where Y is the network's output. Given the output Y , an approximation to the eigenvectors of P , Eq. (5), can be calculated as $\hat{U} = Y O$ where $O \in \mathbb{R}^{K \times K}$ satisfies

$$Y^T W Y O = Y^T D Y O \Lambda, \quad (10)$$

where Λ is a refined approximation of the eigenvalue matrix of (W, D) . Note that Eq. (10) requires a batch for its computation, which may be problematic at inference. The authors show qualitatively a successful approximation to the analytical embedding.

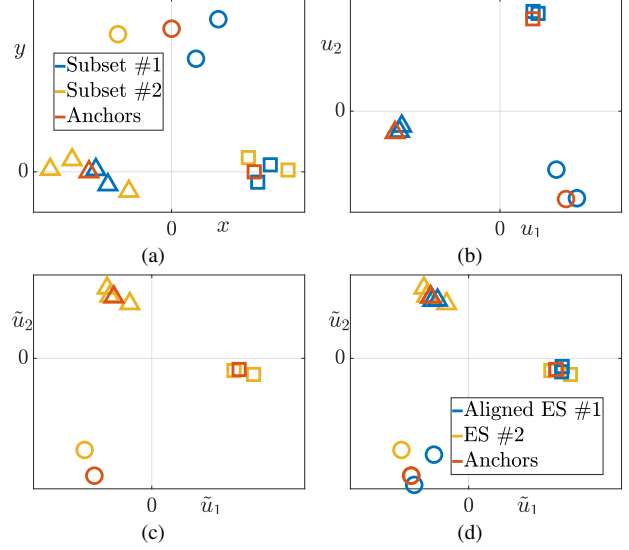


Figure 2. **Illustration.** The full dataset of three separated clusters, divided into two subsets and anchors is shown in Fig. 2a. Figs. 2b-2c show the embedding spaces of subset #1 and subset #2, respectively. Fig. 2d shows the embedding space of the entire data after aligning the embedding of subset #1 to that of subset #2.

4. Our Method

4.1. Motivation

Our goal is to learn a model $f: \mathcal{M} \rightarrow \mathbb{R}^K$, where given a sample $x \in \mathcal{M}$ approximates well the corresponding ϕ of Eq. (3). As is common with DNN learning, given a large training set, we would like to train the model by splitting the dataset into batches. A batch can be viewed as sampling the graph of the training set. A straightforward approach would be to compute the eigenspace of each batch and to learn a mapping from x to ϕ , using a data loss similar to Diffusion Nets. The problem is that different samples of the training set most often lead to different embeddings. Specifically, the same instance x_i can be mapped very differently in each batch.

This can be demonstrated in a very simple toy example, shown in Fig. 2, which illustrates the core problem and our proposed solution. Three distinct clusters in Euclidean space are sampled in two trials (batches) and the eigenspace embedding is computed analytically. Three samples appear in both subsets, one for each cluster (red color). We refer to the common samples as *anchors*. The plots of the instances in the embedding space for the two subsets are shown in Figs. 2b-2c. One can observe the embeddings are different. Specifically, all anchors, which appear in both samplings, are mapped differently in a substantial way. It is well known that eigenvector embedding has a degree of freedom of rotation (as shown for example in [32]). However, in the case of uneven sampling of clusters there may be also some scale

changes and slight translation (caused by the orthonormality constraints). We thus approximate these degrees of freedom in the general case as an affine rigid transformation according to the anchors. Aligning one embedding space to the other one, using this transformation, yields a consistent embedding, as can be seen in Fig. 2d. Following the alignment process, the embedding can be learnt well using batches.

In the toy example of the Three Moons, see Fig. 3, we show the mapping of 9 anchor-points, as shown in Fig. 3b. In Figs. 3c-3d the analytic computation of the first two non-trivial eigenvectors are plotted for 20 different batch samples of size 256 (out of 9000 nodes in the entire dataset), all of which contain the 9 points. In this simple example anchors located on the same moon receive approximately the same value. However, in different batches the embedding of the anchors is clearly inconsistent. Surely, a network cannot be expected to generalize such a mapping. After learning the transformation and performing alignment, the embedding values are consistent. In Figs. 3e-3f the values are shown after our correction procedure. This consistency allows to train DNN to learn the desired embedding, by dividing the data into batches. The result of the trained DNN model for the Three Moons dataset appears in Fig. 1 (second from left). These toy examples lead us to the detailed description of our algorithm.

4.2. BASiS Algorithm

We propose to calculate the embedding space with batches. To obtain consistency in this representation, we calculate the first-order approximation of the distortion obtained in the eigenvector values between different samples of the data. The main steps of our algorithm are as follows: First we perform two preliminary initialization steps.

Defining an anchor set. Draw l samples from the data. This subset is denoted as V^a and will be added to any batch in the learning process.

Defining the reference embedding space. We would like to define the embedding space of the anchor set as a reference space. However, to get more information about the manifold \mathcal{M} , we add $m-l$ samples (randomly) and term it as the reference set V^{ref} . After calculating the embedding $V^{ref} \rightarrow \varphi^{ref}$ (as in Eq. (3)), one can extract the coordinates of the anchor samples,

$$V^a \rightarrow \{\varphi_i^{a,ref}\}_{i=1}^l. \quad (11)$$

Following this initialization, the main steps of the training are as follows:

Calculate the embedding space over a new batch. Draw $m-l$ new samples and add them to the anchor set. Let us denote the union set as V^b . We calculate $\{\varphi_i\}_{i=1}^m$, the embedding of V^b and extract the embedding $\{\varphi_i^a\}_{i=1}^l$ corresponding to the anchors.

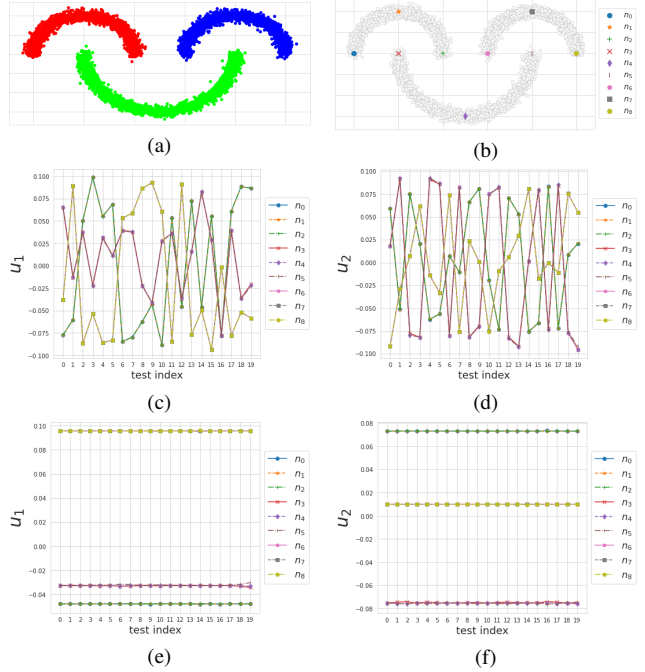


Figure 3. **Three-Moons toy example.** The full dataset is shown in Fig. 3a and the chosen anchor-nodes in Fig. 3b. Figs. 3c- 3d show the values of the two leading eigenvectors for the anchors, for 20 different graph-samples. Figs. 3e- 3f show those values after our proposed alignment.

Calculate the alignment transformation. Now, we calculate the alignment transformation between $\{\varphi_i^a\}_{i=1}^l$ to $\{\varphi_i^{a,ref}\}_{i=1}^l$. More formally, for $\varphi^a, \varphi^{a,ref} \in \mathbb{R}^K$ we find $A \in \mathbb{R}^{K \times K}$ and $b \in \mathbb{R}^K$ which minimize

$$\min_{A,b} \sum_{i=1}^l \left\| \varphi_i^{a,ref} - (A\varphi_i^a + b) \right\|^2. \quad (12)$$

Alternatively, one can define $\hat{\varphi}^a = [\varphi^a, 1]$ and find the transformation $T \in \mathbb{R}^{K \times (K+1)}$ such that

$$\min_T \sum_{i=1}^l \left\| \varphi_i^{a,ref} - T\hat{\varphi}_i^a \right\|^2. \quad (13)$$

In this case there are $K \times (K+1)$ degrees of freedom. Each anchor provides K constraints, that means at least $K+1$ anchors are needed in order to solve this least squares problem. Since in many real-world problem there is noise in the measurements, it is customary to solve such problems in an overdetermined setting, using a higher number of anchors. In addition, given a large number of anchors, the transformation can be calculated using best matches - for example by using the RANDOM SAMPLE CONSENSUS (RANSAC) algorithm [11].

Batch Alignment. Given the transformation T , we can align the embedding $\{\varphi_i\}_{i=1}^m$ of all the instances of V^b . We

define $\hat{\phi} = [\phi, 1]$ and update

$$\phi \leftarrow T\hat{\phi}. \quad (14)$$

Gradient Step. Now that we have a mechanism that allows to get consistent embedding, we can train the DNN by dividing the data into batches and use a simple MSE lose function

$$\mathcal{L}_{BASIS}(\theta) = \frac{1}{m} \sum_{i=1}^m \|y_i - \phi_i\|^2, \quad (15)$$

where $y_i = f_{\theta}(x_i)$ is the DNN's output and ϕ_i is the embedding of x_i after alignment.

The full training scheme is detailed in Algorithm 1.

Algorithm 1 BASiS Training Scheme

- 1: **Inputs:**
data features $\{x_i\}_{i=1}^n$, number of eigenvectors K , batch size m , number of anchors l .
 - 2: **Outputs:**
Spectral embedding model $f: \mathcal{M} \rightarrow \mathbb{R}^K$
 - 3: **Initialize:**
Define anchor set V^a .
Extract $\{\varphi_i^{a,ref}\}_{i=1}^l$, the reference embedding of V^a using Eq. (11)
 - 4: **while** $\mathcal{L}_{BASIS}(\theta)$ not converged **do**
 - 5: Draw $m - l$ new samples.
 - 6: Define node set V^b as union of the anchors with the new sampled nodes.
 - 7: Calculate the embedding $\{\phi_i\}_{i=1}^m$ of V^b .
 - 8: Calculate the optimal transformation T , Eq. (13).
 - 9: Align $\{\phi_i\}_{i=1}^m$ with T , Eq. (14).
 - 10: Do a gradient step of \mathcal{L}_{BASIS} , Eq. (15).
 - 11: **end while**
-

4.3. BASiS for feature perturbation

In the process of network training, the features are often not fixed and slightly change each iteration during training. In this case the adjacency values change and hence naturally also the embedding space. We suggest an algorithm to allow iterative changes in the feature space (inducing different graph metric). This algorithm is also based on an alignment technique. Similar to Algorithm 1, we define anchor nodes. Given the current features, we extract $\{\varphi_i^{a,prev}\}_{i=1}^l$, the current spectral embedding of the anchors. When the features are updated, we find the new embedding of the anchors $\{\varphi_i^{a,update}\}_{i=1}^l$. In order to maintain consistency in the learning process, we find a transformation T_G , as in Eq. (13), that aligns the updated anchor embedding to the previous one. Then we align the entire embedding space according to the calculated transformation. Algorithm 2 summarizes the proposed method.

Algorithm 2 BASiS under iterative feature change

- 1: **Inputs:**
 $\{\varphi_i^{a,prev}\}_{i=1}^l$ the anchors embedding over previous features $\{x_i\}_{i=1}^n$, updated features $\{\hat{x}_i\}_{i=1}^n$.
 - 2: **Outputs:**
 $\{\varphi_i^{update}\}_{i=1}^n$ the spectral embedding over the updated features, aligned to $\{\varphi_i^{a,prev}\}_{i=1}^l$.
 - 3: Calculate the embedding $\{\varphi_i^{update}\}_{i=1}^n$ over the updated features.
 - 4: Extract $\{\varphi_i^{a,update}\}_{i=1}^l$ the embedding correspond to the anchors.
 - 5: Calculate the transformation T_G between $\{\varphi_i^{a,prev}\}_{i=1}^l$ and $\{\varphi_i^{a,update}\}_{i=1}^l$.
 - 6: Align $\{\varphi_i^{update}\}_{i=1}^n$ with T_G .
-

5. Experimental Results

In this section we examine the ability of BASiS to learn the graph-spectral embedding over different datasets quantifying its success using several performance measures. Our method is able to learn any desired eigen embedding (since it is supervised by analytic calculations). To fairly compare our method to the ones mentioned in Sec. 3 we calculate the eigenspace of L_N (Eq. (4)). For all methods the DNN's architecture includes 5 fully connected (FC) layers with ReLU activation in between (see full details in the supplementary).

5.1. Evaluation Metrics

We evaluate our results using several measures. We calculate the Grassmann distance (projection metric) [14] between the network output and the analytically calculated eigenvectors. The squared Grassmann distance between two orthonormal matrices $Y_1, Y_2 \in \mathbb{R}^{n \times K}$ is defined as:

$$d_G(Y_1, Y_2) = K - \sum_{i=1}^K \cos^2 \theta_i, \quad (16)$$

where $0 \leq \theta_1 \leq \dots \leq \theta_K \leq \frac{\pi}{2}$ are the principal angles between the two subspaces $span(Y_1)$ and $span(Y_2)$. The distance is in $[0, K]$ where lower values indicate greater proximity between the subspaces.

A second measure is the degree of orthogonality of the DNN's output Y . We use the following orthogonality measure:

$$d_{\perp}(Y) = \|Y^T Y - I\|_F^2, \quad (17)$$

where I is an identity matrix and $\|\cdot\|_F$ is the Frobenius norm. For Y containing columns of orthonormal vectors we get $d_{\perp}(Y) = 0$. In general, we expect this measure to be close to zero in proper embeddings.

Measures	Networks	MNIST	Fashion-MNIST	SVHN	CIFAR-10
$d_G \downarrow$	Diffusion-Net	0.204 ± 0.058	0.488 ± 0.238	1.909 ± 0.238	1.022 ± 0.250
	SpecNet1	0.386 ± 0.074	0.375 ± 0.132	3.526 ± 0.529	2.256 ± 0.471
	SpecNet2	1.388 ± 0.262	1.976 ± 0.210	1.903 ± 0.242	2.970 ± 0.682
	BASiS (Ours)	0.107 ± 0.038	0.284 ± 0.073	1.656 ± 0.170	0.803 ± 0.085
$d_{\perp} \downarrow$	Diffusion-Net	0.535 ± 0.365	0.823 ± 0.664	1.532 ± 0.354	2.957 ± 1.837
	SpecNet1	6.296 ± 0.922	6.384 ± 0.899	4.507 ± 0.821	5.169 ± 0.775
	SpecNet2	9.486 ± 0.001	8.561 ± 1.397	4.104 ± 0.269	4.922 ± 0.102
	BASiS (Ours)	0.247 ± 0.076	0.590 ± 0.144	0.488 ± 0.098	0.407 ± 0.095
NMI \uparrow	Diffusion-Net	0.944 ± 0.041	0.759 ± 0.085	0.645 ± 0.016	0.466 ± 0.034
	SpecNet1	0.911 ± 0.008	0.761 ± 0.011	0.665 ± 0.018	0.443 ± 0.012
	SpecNet2	0.925 ± 0.012	0.759 ± 0.010	0.701 ± 0.009	0.466 ± 0.013
	BASiS (Ours)	0.961 ± 0.001	0.798 ± 0.001	0.736 ± 0.001	0.501 ± 0.001
ACC \uparrow	Diffusion-Net	0.944 ± 0.030	0.781 ± 0.179	0.687 ± 0.303	0.620 ± 0.062
	SpecNet1	0.963 ± 0.005	0.815 ± 0.029	0.811 ± 0.039	0.637 ± 0.029
	SpecNet2	0.966 ± 0.007	0.801 ± 0.023	0.813 ± 0.015	0.606 ± 0.039
	BASiS (Ours)	0.986 ± 0.001	0.865 ± 0.003	0.880 ± 0.001	0.688 ± 0.001
Accuracy($\%$) \uparrow	Diffusion-Net	95.508 ± 1.449	86.207 ± 0.196	86.850 ± 1.386	67.316 ± 2.112
	SpecNet1	92.278 ± 4.776	84.123 ± 1.229	85.154 ± 0.377	65.336 ± 0.626
	SpecNet2	97.026 ± 0.546	85.953 ± 0.240	87.469 ± 0.130	67.093 ± 0.644
	BASiS (Ours)	98.522 ± 0.065	87.202 ± 0.187	88.021 ± 0.064	68.887 ± 0.128

Table 1. **Spectral embedding performance comparison.** Average performance obtained over 10 different installations of each of the four methods. In each experiment we learn an embedding space in dimension of 10 for 1000 training iterations using batches of size 512.

To evaluate the potential clustering performance we examined two common metrics: Normalized mutual information (NMI) and unsupervised clustering accuracy (ACC). The clustering result is achieved by performing the K-Means algorithm over the spectral embedding. Both indicators are in the range $[0, 1]$, where high values indicate a better correspondence between the clustering result and the true labels. NMI is defined as,

$$NMI(c, \hat{c}) = \frac{I(c, \hat{c})}{\max\{H(c), H(\hat{c})\}}, \quad (18)$$

where $I(c, \hat{c})$ is the mutual information between the true labels c and the clustering result \hat{c} and $H(\cdot)$ denotes entropy. ACC is defined as,

$$ACC(c, \hat{c}) = \frac{1}{n} \max_{\pi \in \Pi} \sum_{i=1}^n \mathbb{1}\{c_i = \pi(\hat{c}_i)\}, \quad (19)$$

where Π is the set of possible permutations of the clustering results. To choose the optimal permutation π we followed [24] and used the Kuhn-Munkres algorithm [19].

Finally, we examine how suitable the embedding is for classification. We trained (with Cross-Entropy loss) a supervised linear regression model containing a single fully connected layer without activation, and examined its accu-

racy:

$$Accuracy(c, \hat{c}) = \frac{1}{n} \sum_{i=1}^n \mathbb{1}\{c_i = \hat{c}_i\}. \quad (20)$$

5.2. Spectral Clustering

In this section we examine the ability of our method to learn the spectral embedding for clustering of different datasets. First, we examine the performance for well-known spectral clustering toy examples, appearing in Fig. 1. In these examples the dataset includes 9000 nodes and the model is trained by calculating the first non-trivial eigenvectors (sampling 256 nodes, using 1000 iterations). In all the experiments NMI and ACC of 1.0 were obtained over the test set (i.e., perfect clustering results). In addition, as demonstrate in Fig. 1, the learnt model is able to generalize the clustering result and performs smooth interpolation and extrapolation of the space mapping. We note that no explicit regularization loss is used in our method, generalization and smoothness are obtained naturally through the neural training process.

Next we compare the performance of BASiS to those of the models mentioned in Sec. 3. We examine the results over 4 well-known image datasets: MNIST, Fashion-MNIST [31], SVHN [20] and CIFAR-10 [16]. For each dataset we first learn an initial low-dimensional represen-

tation, found to be successful for graph-based learning, by a Siamese network, a convolutional neural network (CNN) trained in a supervised manner using Contrastive loss

$$\mathcal{L}_{Com}(x_i, x_j, \theta) = \mathbb{1}\{y_i = y_j\} \|f_{\theta}^{rep}(x_i) - f_{\theta}^{rep}(x_j)\|_2^2 + \mathbb{1}\{y_i \neq y_j\} \max(0, \varepsilon - \|f_{\theta}^{rep}(x_i) - f_{\theta}^{rep}(x_j)\|_2)^2, \quad (21)$$

where $f_{\theta}^{rep}(x_i)$ is the Siamese network’s output for input image x_i labeled as y_i , $\varepsilon \in \mathbb{R}^+$. More details on the architecture of the Siamese network are in the supplementary material. We use this representations as inputs to the spectral embedding models. In all the experiments the graph affinity matrix W is defined by Eq. (1), using the 50 nearest neighbors of each node. We use 50 neighbors in order that all methods could perform well (see sensitivity analysis hereafter). The model is trained to learn the first K eigenvectors, where K is equal to the number of clusters. The batch size is set to $m = 512$. For our method, we randomly select 25 anchor-nodes from each cluster and use RANSAC to find the best transformation.

Comparison between the methods is summarized in Table 1. The numbers are obtained by showing the average and empirical standard deviation of each measure, as computed based on 10 experiments.

Since the training process in Diffusion Net is not scalable, in each initialization we randomly sampled a single batch from the training set and trained the model with the analytically calculated spectral embedding. In relation to SpecNet2, as indicated in Sec. 3, to obtain an approximation of the spectral space, SpecNet2 requires a post-processing stage over the network’s output. In order to maintain consistency and obtain reasonable performance for all measures, the post-processing is performed over the entire test set (this naturally limits the method and increases the level of complexity at inference). More implementation details are in the supplementary. Table 1 shows that our method is superior and more stable in approximating the analytical embedding and in clustering.

We further examined the robustness of the methods to changes in the number of neighbors for each node. This parameter defines the connectivity of the graph in the affinity matrix. Fig. 4 shows the average and the empirical standard deviation of the performance measures, for 50 training procedures over the MNIST dataset. It is shown that our method is highly robust and consistent. We note the high sensitivity of Diffusion Net to this meta-parameter.

5.3. Diffusion Maps Encoding

We examine the ability of our model to learn the DM embedding (6). For this purpose we use the dataset from [17] which includes 2000 snapshots of toy bunny located on a rotating surface. Six examples of such frames are shown in Fig. 5a. We define a graph using the 20 nearest neighbors

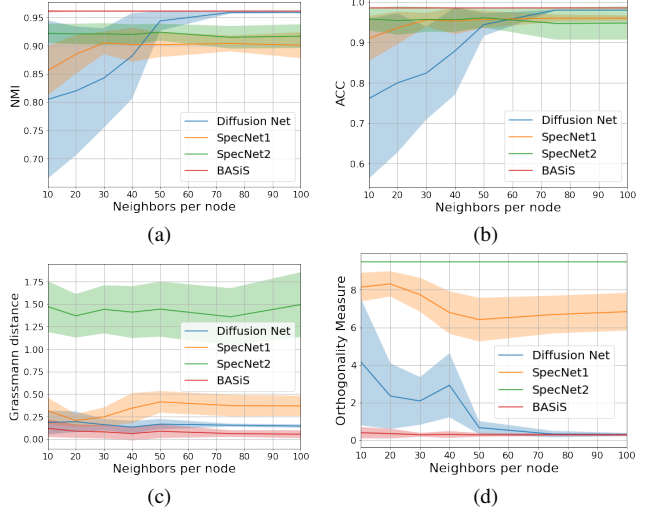


Figure 4. **Robustness to the node neighborhood.** The average and standard deviation of 50 different training experiments over the MNIST dataset, for different number of neighbors per node.

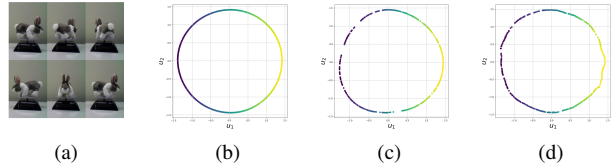


Figure 5. **Diffusion Maps encoding.** Data set of 2000 snapshots of bunny on rotating display [17]. Fig. 5a shows snapshot examples. Fig. 5b present the analytically calculated DM embedding for the full dataset. The test set analytical embedding is shown in Fig. 5c and the network output for the test set in Fig. 5d.

of each node, and calculate the random-walk matrix P (5). Raw pixels are used as features (dimension 288,000). Dimension reduction is performed with DM to \mathbb{R}^2 . In Fig. 5b the analytically calculated embedding obtained based on the entire dataset is shown. Our model was trained to approximate this embedding using 1500 images. Test is performed over 500 images. Fig. 5c shows the analytically calculated embedding over the test snapshots. Fig. 5d shows the embedding obtained by our trained model. Our method approximate well the analytic calculation.

5.4. Iterative Change of Features

In this section we illustrate the performance of Algorithm 2 for aligning the spectral embedding space under changing features. We define two DNN models. The first one is for feature extraction, trained to minimize the Contrastive loss (21). The second is trained for calculating the spectral embedding, using Algorithm 1, based on the output of the features model. Both models are trained simultaneously. The feature model is trained for 1500 iterations,

where every tenth iteration we perform an update step for the spectral embedding model. To maintain consistency under the feature change, we align the spectral space using Algorithm 2 before performing an update step for the spectral model. Fig. 6 shows the results of the training process over the MNIST dataset where the learnt features are of dimension 16 and the eigenvectors are of dimension 9. Fig. 7 shows a visualization (TSNE [28]) of the test set embedding at the beginning and the end of the training process. The two modules were able to converge and to reach good clustering results. In addition, when the loss of the spectral module is sufficiently low (around iteration 800, marked with a red line in Fig. 6) the clustering performance of the spectral module is comparable to the analytic embedding, calculated with the current features (the orange and the green plots tightly follow the blue plot).

To illustrate the role of the transformation T_G , we show in Fig. 8 the results of another experiment using a similar setting. For better visualization, the training is only for three digits of MNIST: 4, 7 and 9. The embedding (and visualization) is two dimensional. It can be seen that T_G imposes consistency of the resulting embedding under feature change, allowing convergence and good approximation of the eigenvector space.

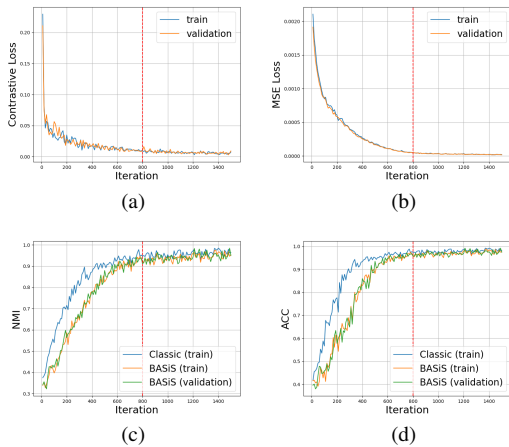


Figure 6. **Training under feature change.** Evolution of measures during training (MNIST). 6a-6b losses of the features module and the spectral embedding module, respectively. 6c-6d – clustering measures. Blue, analytic calculation of the eigenvectors based on current features. Orange and green, output of spectral embedding module (training and validation sets, respectively).

6. Conclusion

In this paper we introduced BASiS, a new method for learning the eigenspace of a graph, in a supervised manner, allowing the use of batch training. Our proposed method has shown to be highly robust and accurate in approximat-

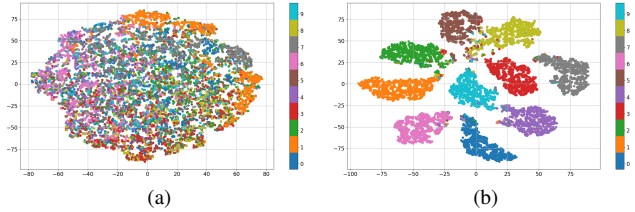


Figure 7. **Test set embedding.** Visualization (TSNE) of the spectral embedding, MNIST test set. Fig. 7a shows the spectral embedding at the beginning of training, Fig. 7b at the end.

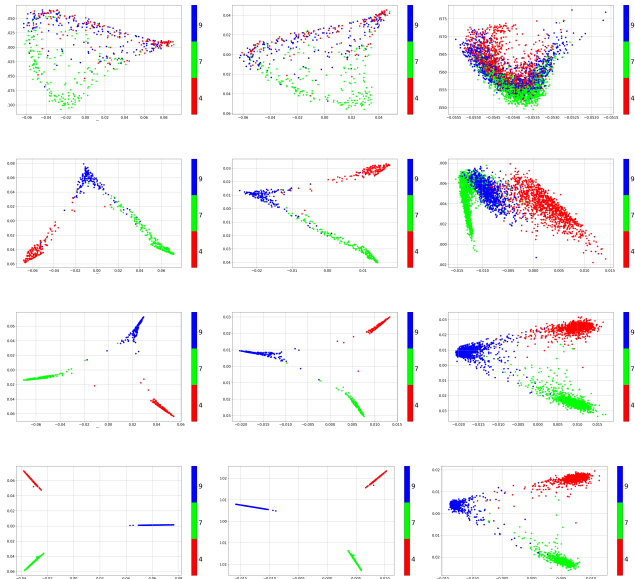


Figure 8. **Features change demonstration** Left column: the analytically calculated spectral embedding of V^{ref} obtained over the features module output. Distortion is a consequence of feature change. Middle column: Spectral embedding of V^{ref} after alignment with T_G . Right column: Network spectral embedding of the test set. Each row represents the embeddings for the 10th, 100th, 500th and 1500th iteration, respectively.

ing the analytic spectral space, surpassing all other methods with respect to Grassman distance, orthogonality, NMI, ACC and accuracy, over various benchmarks. In addition, we proposed an adaptation of our procedure for learning the eigenspace during iterative changes in the graph metric (as common in neural training). Our method can be viewed as a useful building block for integrating analytical spectral methods in deep learning algorithms. This enables to effectively use extensive theory and practices available, related to classical spectral embedding.

Acknowledgements. We acknowledge support by the Israel Science Foundation (Grant No. 534/19), by the Ministry of Science and Technology (Grant No. 5074/22) and by the Ollendorff Minerva Center.

References

- [1] Carlos Alzate and Johan AK Suykens. Multiway spectral clustering with out-of-sample extensions through weighted kernel pca. *IEEE transactions on pattern analysis and machine intelligence*, 32(2):335–347, 2008. 2
- [2] Angelica I Aviles-Rivero, Philip Sellars, Carola-Bibiane Schönlieb, and Nicolas Papadakis. Graphxcovid: explainable deep graph diffusion pseudo-labelling for identifying covid-19 on chest x-rays. *Pattern Recognition*, 122:108274, 2022. 2
- [3] Leah Bar and Nir Sochen. Strong solutions for pde-based tomography by unsupervised learning. *SIAM Journal on Imaging Sciences*, 14(1):128–155, 2021. 2
- [4] Mohamed-Ali Belabbas and Patrick J Wolfe. On landmark selection and sampling in high-dimensional data analysis. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906):4295–4312, 2009. 2
- [5] Mikhail Belkin and Partha Niyogi. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396, 2003. 1
- [6] Ido Ben-Shaul, Leah Bar, and Nir Sochen. Solving the functional eigen-problem using neural networks. *arXiv preprint arXiv:2007.10205*, 2020. 2
- [7] Yoshua Bengio, Jean-françois Paiement, Pascal Vincent, Olivier Delalleau, Nicolas Roux, and Marie Ouimet. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. *Advances in neural information processing systems*, 16, 2003. 2
- [8] Xavier Bresson, Thomas Laurent, David Uminsky, and James H Von Brecht. Multiclass total variation clustering. *arXiv preprint arXiv:1306.1185*, 2013. 1
- [9] Ziyu Chen, Yingzhou Li, and Xiuyuan Cheng. Specnet2: Orthogonalization-free spectral embedding by neural networks. *arXiv preprint arXiv:2206.06644*, 2022. 1, 3
- [10] Ronald R Coifman and Stéphane Lafon. Diffusion maps. *Applied and computational harmonic analysis*, 21(1):5–30, 2006. 1, 2
- [11] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 4
- [12] Charless Fowlkes, Serge Belongie, Fan Chung, and Jitendra Malik. Spectral grouping using the nystrom method. *IEEE transactions on pattern analysis and machine intelligence*, 26(2):214–225, 2004. 2
- [13] Cristina Garcia-Cardona, Ekaterina Merkurjev, Andrea L Bertozzi, Arjuna Flenner, and Allon G Percus. Multiclass data segmentation using diffuse interface methods on graphs. *IEEE transactions on pattern analysis and machine intelligence*, 36(8):1600–1613, 2014. 1
- [14] Jihun Hamm and Daniel D Lee. Grassmann discriminant analysis: a unifying view on subspace-based learning. In *Proceedings of the 25th international conference on Machine learning*, pages 376–383, 2008. 5
- [15] Ori Katz, Ronen Talmon, Yu-Lun Lo, and Hau-Tieng Wu. Alternating diffusion maps for multimodal data fusion. *Information Fusion*, 45:346–360, 2019. 2
- [16] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 6
- [17] Roy R Lederman and Ronen Talmon. Learning the geometry of common latent variables using alternating-diffusion. *Applied and Computational Harmonic Analysis*, 44(3):509–536, 2018. 2, 7
- [18] Gal Mishne, Uri Shaham, Alexander Cloninger, and Israel Cohen. Diffusion nets. *Applied and Computational Harmonic Analysis*, 47(2):259–285, 2019. 1, 2
- [19] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38, 1957. 6
- [20] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011. 6
- [21] Andrew Ng, Michael Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. *Advances in neural information processing systems*, 14, 2001. 1, 3
- [22] Antonio Ortega, Pascal Frossard, Jelena Kovačević, José MF Moura, and Pierre Vandergheynst. Graph signal processing: Overview, challenges, and applications. *Proceedings of the IEEE*, 106(5):808–828, 2018. 2
- [23] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500):2323–2326, 2000. 1
- [24] Uri Shaham, Kelly Stanton, Henry Li, Boaz Nadler, Ronen Basri, and Yuval Kluger. Spectralnet: Spectral clustering using deep neural networks. In *Proceedings of the 6th International Conference on Learning Representations*, 2018. 1, 3, 6
- [25] Jianbo Shi and Jitendra Malik. Motion segmentation and tracking using normalized cuts. In *Sixth international conference on computer vision (IEEE Cat. No. 98CH36271)*, pages 1154–1160. IEEE, 1998. 2
- [26] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000. 1, 2
- [27] Suman Tatiraju and Avi Mehta. Image segmentation using k-means clustering, em and normalized cuts. *Department of EECS*, 1:1–7, 2008. 2
- [28] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 8
- [29] E Weinan, Jiequn Han, and Arnulf Jentzen. Algorithms for solving high dimensional pdes: from nonlinear monte carlo to machine learning. *Nonlinearity*, 35(1):278, 2021. 2
- [30] Christopher Williams and Matthias Seeger. Using the nystrom method to speed up kernel machines. *Advances in neural information processing systems*, 13, 2000. 2
- [31] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. 6

- [32] Lihi Zelnik-Manor and Pietro Perona. Self-tuning spectral clustering. *Advances in neural information processing systems*, 17, 2004. 1, 3