

Spectral Embedding Norm: Looking Deep into the Spectrum of the Graph Laplacian*

Xiuyuan Cheng[†] and Gal Mishne[‡]

Abstract. The extraction of clusters from a dataset which includes multiple clusters and a significant background component is a nontrivial task of practical importance. In image analysis this manifests for example in anomaly detection and target detection. The traditional spectral clustering algorithm, which relies on the leading K eigenvectors to detect K clusters, fails in such cases. In this paper we propose the *spectral embedding norm* which sums the squared values of the first I normalized eigenvectors, where I can be significantly larger than K . We prove that this quantity can be used to separate clusters from the background in unbalanced settings, including extreme cases such as outlier detection. The performance of the algorithm is not sensitive to the choice of I , and we demonstrate its application on synthetic and real-world remote sensing and neuroimaging datasets.

Key words. spectral clustering, graph Laplacian, spectral theory, outlier detection, calcium imaging

AMS subject classifications. 62H30, 62H35, 35P05, 58J50, 35J05

DOI. 10.1137/18M1283160

1. Introduction. In unsupervised learning and data analysis, one of the most common goals is to group the data points into clusters. A variant task is to extract interesting clusters from the data when, in practice, data points do not perfectly fall into K clusters. We consider a nontrivial setting in which data consist of not only interesting subgroups, namely, “clusters”, but also a large component containing points which are less structured or of less interest, which we call “background.” Important examples in imaging data analysis include image segmentation and saliency detection where the clusters are regions of interest in the image, and the background consists of the rest of the image [4, 28]. Another example is the task of anomaly (or outlier) detection where anomalous samples (small clusters) in the dataset differ from the normal ones (background) and indicate that something important has happened or a problem has occurred. By the very nature of the problem, most data points belong to the background and only a small fraction of data points are anomalies. Anomaly detection in images is an important task in a variety of applications such as target detection in remote sensing imagery, detecting abnormalities such as tumors in biomedical imagery, and for quality

*Received by the editors August 23, 2019; accepted for publication (in revised form) February 4, 2020; published electronically June 30, 2020.

<https://doi.org/10.1137/18M1283160>

Funding: The work of the first author was partially supported by the National Science Foundation grants DMS-1818945, DMS-1820827, by the National Institute of Health grant R01GM131642, and by the Alfred P. Sloan Foundation. The work of the second author was partially supported by the National Institute of Neurological Disorders and Stroke grant R01EB026936.

[†]Department of Mathematics, Duke University, Durham, NC 27708 (xiuyuan.cheng@duke.edu).

[‡]Halcioğlu Data Science Institute, University of California, San Diego, La Jolla, CA 92093 (gmishne@ucsd.edu).

inspection in production lines. An automated solution highlighting only suspicious regions to be reviewed by an expert would save greatly on time.

In theory, one may view the background component as an extra cluster; however the unbalanced size of the clusters versus background components poses challenges for traditional clustering methods. The popular spectral clustering algorithm [25, 33, 41] reduces the dimensionality of the data using a spectral embedding, and then performs clustering in the low-dimensional space. The method originally proposed to cluster data into K clusters by applying k -means to the leading K eigenvectors (the low-lying eigenvectors of the graph Laplacian) computed from an affinity matrix built from the data [24, 25, 41, 45]. A question is then how to set the parameter K , and this is especially important for exploratory data analysis when the number of clusters underlying the data is not known a priori. The traditional solution is to use the spectral gap of the eigenvalues to determine K [39], yet in practical settings, such a gap may not exist. In particular, when the cluster sizes are unbalanced, or a large background component is present, there is no spectral gap after the K th eigenvalue and the leading eigenvectors do not localize on K given clusters, but rather tend to be supported mostly on the large component due to the slow mixing time of the diffusion process restricted to it. The unbalanced case of outlier detection is a classical scenario where traditional spectral clustering fails to identify the existing clusters [23, 43, 46]. As has been shown in [7, 18, 26, 42] and will be demonstrated, the eigenvectors which indicate clusters or outliers may lie deep within the spectrum of the affinity matrix. This gives rise to the notion of abandoning the guideline of focusing on K eigenvectors and rather choosing to look deeper into the spectrum in such settings.

In this paper we consider a cluster-background splitting model of the graph, including anomaly detection as a special case. The model is motivated by applications and will be tested on real-world datasets. We propose a quantity called the *spectral embedding norm*, which maps each node in the graph to a positive number, and separates clusters from background with a theoretical guarantee. The idea is closely related to the “localization” pattern of the eigenvectors, namely, what they are supported on—either mainly on the cluster block or on the background block—and this pattern maintains even when the spectral gap vanishes. Viewing the affinity matrix as a perturbed one from a baseline affinity where the background and clusters are completely disconnected, one can analyze the consequent deformation of the spectrum of the graph Laplacian matrix. However, the instability of eigenvectors under the deformation poses difficulty to the use of individual eigenvectors in this environment. The spectral embedding norm, on the other hand, improves the stability by using a summation over multiple eigenmodes and provides guaranteed detection of the clusters by simple thresholding. The algorithm involves a parameter which is the number of eigenvectors summed over, and the performance is not sensitive to the parameter choice.

Our result thus provides a way to go beyond dominating eigenvectors of the graph Laplacian to unbalanced data clustering tasks with theoretical verification. It suggests that, in the presence of cluttered background samples, it is beneficial to look deep into the spectrum to identify important and subtle structures. By providing a simple measure by which to separate clusters from background, the method allows the application of signal processing and machine learning methods to the cluster samples only without contamination from the background. Our analysis spans a setting of multiple clusters against background to the extremely unbalanced case of outlier detection.

In the rest of the paper, we review more related literature before ending the current section. The spectral embedding norm is introduced in section 2, with an illustrative example to show the main idea. The theoretical result is presented in section 3, experiments on synthetic and real-world image datasets are presented in section 4, proofs given in section 5, and further remarks made in the final section.

Notations. $|\cdot|$ stands for the cardinal number of a set. A^c means the complement of a set A .

1.1. Related works. As spectral clustering and variants have been intensively studied in literature, we list the most relevant works to our problem.

The spectral embedding for clustered data has been previously analyzed in many places. While Schiebinger, Wainwright, and Yu [32] analyzed a nonparametric mixture model to show that under certain conditions the embedded points lie in an orthogonal cone structure and k -means succeeds in clustering the data, Nadler and Galun [23] showed that even for well-separated Gaussians the top K eigenvectors do not necessarily localize on K clusters, based on the analysis of a diffusion process in a multiwell potential [24, 34]. Different approaches, e.g., [10, 45], attempted to align the eigenvectors axes with the different clusters and improve the robustness of cluster identification. Zelnik-Manor and Perona [45] proposed to estimate the number of clusters from the eigenvectors instead of from the spectral gap, and empirically demonstrated improved performance when a background cluster is present. Damle, Minden, and Ying [10] considered the case of a balanced blocklike affinity matrix. The *embedding norm* studied in the current work differs from the above approaches, and it involves a simple algorithm with a theoretical guarantee under the specified settings.

Spectral embeddings have been used for anomaly detection in several modified ways: based on the first nontrivial eigenvector of an affinity matrix [16, 26], eigenvector selection [18, 42], out-of-sample extension [1, 19, 20, 22], the algebraic structure of the weighted magnitude sum of Laplacian eigenfunctions [6], and multiscale constructions of spectral embeddings [19, 29], and usually requiring tuning of multiple parameters. For a general review of anomaly detection methods, the interested reader is referred to [2, 5, 13, 14]. In particular, “eigenvector selection” has been proposed to determine eigenvectors that localize on clusters or specifically on anomalies, using unsupervised spectral ranking [26], kurtosis [42], relevance learning [43], entropy [11, 46], the L1 norm [18], tensor product [9], and local linear regression [12]. Both [43, 46] proposed calculating Km eigenvectors, where $m > 1$, and then select informative eigenvectors. Wu et al. [42] analyzed the adjacency matrix of a graph, while Miller, Bliss, and Wolfe [18] considered the modularity matrix of a graph. In the current paper, we analyze the spectrum of the (normalized) random-walk Laplacian matrix which has a more stable spectrum with finite samples [40].

The proposed notion of embedding norm formally resembles the probability amplitude in quantum mechanics, which sums the squared modulus of the low-energy wave functions. It is also similar to the *leverage score* in statistics, defined to be the squared sum of principal components, which has been used as an indicator of outlier samples in linear regression and other statistical applications [15]. The generalized form of the embedding norm with exponentially decaying weights has been previously suggested as a tool to identify salient features in shape analysis, called the *heat kernel signature* [37]. However, the settings in the previous statistical

and computer graphical studies are different from our consideration of the cluster-background separation and, particularly, the equal weights on the truncated sum over eigenvectors which more closely resembles probability amplitude has its own motivation; see more in the last section.

2. Spectral embedding norm. Given n data points in the feature space, an undirected weighted graph can be constructed which has n nodes, denoted by \mathcal{V} , and the weight on edge (x, y) is the affinity between nodes x and y denoted by $W(x, y)$. W is an n -by- n real-symmetric matrix of nonnegative entries $W(x, y) = W(y, x) \geq 0$, called the graph affinity matrix. In applications, W is built as a pairwise affinity between data points in a feature space, e.g., $W(x, y) = k(x, y)$, where k is a symmetric kernel function applied to the feature vectors of data points x and y . In our analysis we assume that W has been constructed.

2.1. Cluster-background splitting in the graph. Suppose that \mathcal{V} can be divided into two disjoint subsets, background and clusters, denoted by \mathcal{B} and \mathcal{C} , respectively. The typical scenario which we consider is when data points in \mathcal{C} are concentrated in the feature space and well-clustered into K subclusters, whereas those in \mathcal{B} can be “manifold-like” and spread over the space. The precise assumptions will be formulated in terms of the graph-Laplacian spectra constrained to the subgraphs of \mathcal{C} and \mathcal{B} (see Assumption 1). We also assume the connections between \mathcal{C} and \mathcal{B} are weak. As a result, the submatrix of W constrained to \mathcal{C} is close to having K blocks, and is almost separated from the submatrix of \mathcal{B} . Define matrix W_0 by removing all the connections between \mathcal{B} and \mathcal{C} from W , i.e., W_0 is a block-diagonal matrix consisting of two blocks of \mathcal{C} and \mathcal{B} , respectively, and E is the matrix consisting of the \mathcal{B} - \mathcal{C} connections. We introduce a pseudodynamic parametrized by time t as

$$(2.1) \quad W(t) = W_0 + tE, \quad t \in [0, 1],$$

so that $W(0) = W_0$ and $W(1) = W$. To simplify the analysis, we assume that the K subclusters are of equal size, that is,

$$|\mathcal{C}| = \delta|\mathcal{V}|, \quad |\mathcal{B}| = (1 - \delta)|\mathcal{V}|,$$

and each of the K clusters in \mathcal{C} has $\frac{\delta|\mathcal{V}|}{K}$ nodes. The result extends to the unequal-size case.

2.2. Graph Laplacian and embedding norm. Consider the normalized random-walk graph Laplacian of W ,

$$L = I - D^{-1}W := I - P,$$

where D is a diagonal matrix defined by $D_{ii} = \sum_j W_{ij}$ and P is a Markov matrix. We shall see that D is always invertible. P is similar to $D^{-1/2}WD^{-1/2}$ which is real-symmetric, thus P is diagonalizable and has n real eigenvalues. Let

$$(2.2) \quad P\psi_k = \lambda_k\psi_k, \quad k = 1, \dots, n, \quad \psi_k^T D\psi_j = \delta_{kj},$$

where $\{\lambda_k\}_k$ are the eigenvalues of P , $\{\psi_k\}_k$ are the corresponding right eigenvectors, and $\delta_{kj} = 1$ when $k = j$ and 0 otherwise. Given that $|\lambda_k| \leq 1$ (Perron–Frobenius theorem), and when W is a positive-definite kernel matrix then all the eigenvalues are between 0 and 1. The

largest eigenvalue of P is 1 and the associated eigenvectors are the constant vector. Note that time dependence is omitted in the above notations: as we introduced the deformation of W in (2.1), D , P , and, consequently, ψ_k and λ_k also depend on t . We assume that at $t = 1$ the eigenvalues are sorted to be decreasing, and for other t , the indexing k is arranged so that ψ_k and λ_k are differentiable with respect to t [17].

The spectral embedding norm of every node $x \in \mathcal{V}$ is defined to be

$$(2.3) \quad S(x) = \sum_{k \in I} \psi_k(x)^2,$$

where I is a subset of the eigenvalue indices $\{1, \dots, n\}$. $S(x)$ is the (squared) Euclidean norm of the embedded vector of x in the spectral embedding space using eigenvectors of indices in I . When needed, we include the time dependence in the notation written as $S(x, t)$, $t \in [0, 1]$. A typical choice of I is $I = \{1, \dots, |I|\}$, when the eigenvalues are sorted to be descending. The cardinal number $|I|$ is a parameter of the method and in the scenario of outlier detection it is typically larger than K . In practice, estimates of K can be used if K is not known. We will show that the result is not sensitive to the choice of $|I|$, in analysis and experiments.

The embedding norm S is able to separate \mathcal{C} from \mathcal{B} by a provable margin under certain assumptions (Theorem 3.5). A general weighted form of S can be introduced and the result extends directly. This is naturally related to the *diffusion distance* [8], and we will explain more on this in the last section.

2.3. A prototypical toy example. The prototypical scenario which motivates the proposed method is illustrated in the toy example in Figure 1. (a) shows data points in \mathbb{R}^2 consisting of two groups: a large group, denoted by \mathcal{B} , which lie close to the unit circle (blue) and a small one, denoted by \mathcal{C} , which form a small cluster lying close to the circle (red), and $\frac{|\mathcal{C}|}{|\mathcal{V}|} = \delta = 0.01$. Traditional spectral clustering and k -means will fail to separate the cluster (\mathcal{C} component) from the circle (\mathcal{B} component), as shown in Figure A.4(a)–(c). Variations of this model involve multiple subclusters in \mathcal{C} (see Figure 3 and the applications on real-world data), and the qualitative picture is the same.

The affinity matrix W built from the data is shown in (b). The first few eigenvalues, evolving over time t , are shown in (c), and the associated eigenvectors evaluated on two nodes, one in \mathcal{C} and one in \mathcal{B} , are plotted in (d) and (e). We sort the eigenvalues of P from large to small, and the first eigenvalue is always 1. The embedding norm $S(x)$ takes the squared sum of the first 40 eigenvectors on each node (cf., (2.3)) and is plotted in (f) over time. The figure demonstrates that

1. though two blocks \mathcal{C} and \mathcal{B} exist in the graph, there is no clear eigengap between the second and third eigenvalues. Actually, the leading eigenvalues are all very close to 1 throughout time 5 (the eighth eigenvalue is greater than 0.998);
2. while the second eigenvector ψ_2 distinguishes \mathcal{C} at short time t , once t is greater than 0.01 ψ_2 fails to indicate the cluster \mathcal{C} . (d) shows the value on one node and it is typical for the value of ψ_k 's on \mathcal{C} . The transition actually happens when the initial gap between λ_2 and λ_3 almost vanishes. However, the high-index eigenvectors may take a large value on \mathcal{C} (the eighth eigenvector starts to take a large magnitude on \mathcal{C} around $t = 0.1$, and the trend of high-indexed eigenvectors localizing on \mathcal{C} continues,

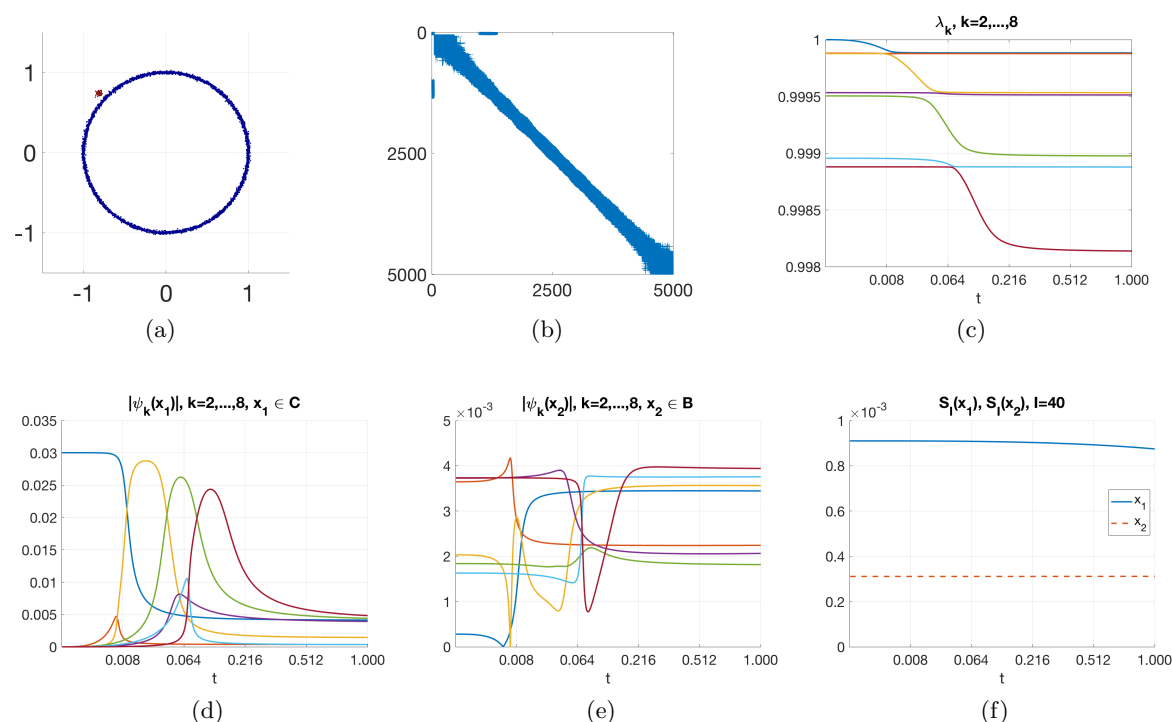


Figure 1. Plots of eigenvalue and eigenvectors of $P = D^{-1}W$ over time. (a) $n = 5000$ data points in \mathbb{R}^2 sampled on $\mathcal{B} \cup \mathcal{C}$, where points in \mathcal{B} lie close to a circle (blue) and points in \mathcal{C} form the small cluster lying close to the circle (red). In this case, $K = 1$, $\delta = 0.01$. (b) The affinity matrix $W(t)$ at $t = 1$, cf., (2.1). (c) Plot of the first 8 eigenvalues as t increases from 0 to 1 (excluding $\lambda_1 = 1$). (d) The absolute values of the associated first 8 eigenvectors at $x_1 \in \mathcal{C}$ over time. (e) Same plot at $x_2 \in \mathcal{B}$. (f) The values of the embedding norm $S(x)$ defined in (2.3) at x_1 and x_2 over time, where $|I| = 40$.

which is not shown). This is evident by $S(x)$ consistently distinguishing \mathcal{C} from \mathcal{B} over time, as shown in (f).

This suggests that when the leading eigenvectors fail to identify the cluster \mathcal{C} , the information of the location of \mathcal{C} may be contained in higher-indexed eigenvectors, and looking deep into the spectrum may be helpful. However, the selection of informative eigenvectors is generally a challenging problem. In particular, as shown in (d), (e), the deformation of eigenvectors is not stable when eigenvalues get close, which makes it difficult to study them individually. Instead, the embedding norm we proposed varies smoothly over time and preserves a gap between \mathcal{C} and \mathcal{B} . As a result, one can detect \mathcal{C} from \mathcal{B} by thresholding the value of S at $t = 1$. Note that 40 eigenvectors are used in the summation, which is much larger than 2. We will justify this improved stability in the analysis.

3. Theoretical analysis of cluster detection. At $t = 0$ in (2.1), the matrix W_0 has a two-block structure, and the spectrum of the graph Laplacian of W_0 also splits into two groups, one residing on \mathcal{C} and the other on \mathcal{B} , respectively. However, as t increases, interactions among the eigenvectors develop and the perfect splitting pattern is no longer preserved. The embedding




Figure 2. Diagram showing the evolution of eigenvalues of the Markov matrix as $W(t)$ changes over time as in (2.1). The trivial eigenvalue $\lambda = 1$ is omitted. At $t = 0$, circles indicate \mathcal{B} eigenvalues, and crosses indicate \mathcal{C} ones. In this example, $K = 3$ and $|I| = 10$. Eigenvalues of the \mathcal{B} -submatrix are shown in circles, and those of the \mathcal{C} -submatrix in crosses. Note that \mathcal{B} can have eigenvalues close to 1 even at $t = 0$. As t increases, at most times the eigenvalues are all of multiplicity one. Eigencrossings may happen within I and I^c but not in-between, and the I -spectral gap denoted by $\Delta(t)$ is preserved (Proposition 3.3). The differential equation (3.15) is obtained by the contour integral Γ , which exists for all t due to positive $\Delta(t)$.

norm varies more stably than individual eigenvectors over time, and serves as a measure by which to separate \mathcal{C} from \mathcal{B} up to time $t = 1$.

3.1. Initial separation by S and assumptions. Since we will use S to separate \mathcal{B} and \mathcal{C} , we need it to do so at least at $t = 0$ when the two blocks \mathcal{B} and \mathcal{C} are perfectly separated by removing all the edges connecting them. Note that this does not necessarily happen unless certain assumptions are made: because the eigenvalues of the \mathcal{B} block can be close to 1 and the clustering in the \mathcal{C} block may not be perfect, the first $|I|$ eigenvectors may be supported either on \mathcal{C} or on \mathcal{B} , and there is generally no guarantee that the squared sum (2.3) will distinguish the two blocks. We make the following two assumptions:

(1) At $t = 0$, the eigenvectors in I which are supported on \mathcal{B} are sufficiently delocalized (“flat”) and those on \mathcal{C} are close to the well-clustered case.

(2) The fraction δ of $|\mathcal{C}|$ is sufficiently small so that the eigenvectors on \mathcal{C} are of sufficiently larger magnitude than those on \mathcal{B} , due to the eigenvector normalization (2.2). The precise condition depends on the choice of $|I|$, the node degrees, and so on.

We denote the volume of set A at time t by $\nu(A, t)$ defined as the sum of the degrees at time t :

$$(3.1) \quad \nu(A, t) = \sum_{x \in A} d(x, t), \quad d(x, t) = \sum_{y \in \mathcal{V}} W(x, y; t).$$

We also define lower and upper bounds

$$(3.2) \quad \underline{d}_0 := \min_{x \in \mathcal{V}} d(x, 0), \quad \overline{d}_0 := \max_{x \in \mathcal{V}} d(x, 0),$$

and assume that $\underline{d}_0 > 0$. By construction (2.1), the degree $d(x, t)$ of any node monotonically increases over time. Thus \underline{d}_0 is the universal degree lower bound.

Lemma 3.1. For all $x \in \mathcal{V}$ and all $0 \leq t \leq 1$, $d(x, t) \geq \underline{d}_0 > 0$.

At $t = 0$, since the affinity matrix decomposes into two separated blocks \mathcal{B} and \mathcal{C} , so do the eigenvectors. We call them initial eigenvectors, and the set of eigenvectors which are only

supported on \mathcal{B} are called the \mathcal{B} -eigenvectors, denoted by $\Psi^{\mathcal{B}}$, and similarly for \mathcal{C} -eigenvectors and $\Psi^{\mathcal{C}}$. The assumption on these eigenvectors and the index set I is the following.

Assumption 1 (\mathcal{B} and \mathcal{C} -eigenvectors). At $t = 0$,

- (a) the index I includes K \mathcal{C} -eigenvectors and $|I| - K$ \mathcal{B} -eigenvectors;
- (b) each of the K eigenvectors in $I \cap \Psi^{\mathcal{C}}$ (up to a K -by- K rotation of these K vectors) is associated with one of the K clusters in the following sense: there exists $0 \leq \varepsilon_1 < 1$, and for each $\psi \in I \cap \Psi^{\mathcal{C}}$, there is a unique j , $1 \leq j \leq K$, s.t.

$$(3.3) \quad \frac{1 - \varepsilon_1}{\nu(C_j, 0)} \leq \psi(x)^2 \leq \frac{1 + \varepsilon_1}{\nu(C_j, 0)} \quad \forall x \in C_j,$$

$$(3.4) \quad \psi(x)^2 \leq \frac{\varepsilon_1}{\nu(C, 0)} \quad \forall x \in C \setminus C_j;$$

- (c) there exists $\varepsilon_2 \geq 0$, s.t. for any $\psi \in I \cap \Psi^{\mathcal{B}}$,

$$\psi(x)^2 \leq \frac{1 + \varepsilon_2}{\nu(B, 0)} \quad \forall x \in B.$$

The above assumption, while appearing to be complicated, poses only generic conditions on the subgraphs \mathcal{B} and \mathcal{C} :

In the perfectly separated case the largest K \mathcal{C} -eigenvalues are 1, and the $(K + 1)$ th one is strictly less than 1 and depends on the mixing time of the Markov chain within each cluster. This spectral gap is usually significant since we primarily work with a well-clustered \mathcal{C} which takes a small fraction of nodes and is localized in the graph; e.g., \mathcal{C} is an outlier cluster; or several localized regions of interest. As a result, even when the clustering is not perfect, the $(K + 1)$ th \mathcal{C} -eigenvalue is still sufficiently far away from the first K ones, and they can be excluded from the index set I , since I selects the largest $|I|$ eigenvalues. This fulfills (a).

If the K clusters in \mathcal{C} are perfectly separated, one can verify that $\varepsilon_1 = 0$ in (b). Thus (b) holds when \mathcal{C} (without the \mathcal{B} component) is not far from being well-clustered.

Assumption 1(c) requires that the eigenvector ψ is sufficiently delocalized, or “flattened” on \mathcal{B} . Recall that (2.2),

$$\sum_{x \in \mathcal{B}} \psi(x)^2 d(x, 0) = 1,$$

and the first eigenvector (associated with eigenvalue 1) takes the constant value $\psi(x)^2 = \frac{1}{\nu(B, 0)}$. If all the other eigenvectors are flattened, then (c) holds with some small ε_2 . The delocalization widely applies when the \mathcal{B} are built from data vectors lying on certain regular manifolds: assuming that the discrete eigenvectors well-approximate the continuous limits which are eigenfunctions of the manifold Laplacian, the delocalization of the former inherit from that of the latter (quantum ergodicity theorem [44, 47]). When the spectral convergence is poor, the finite-sample effects may create some localized pattern in the “noisy” eigenvectors; however, since $|I|$ is typically a small number compared to n , we assume that the selected \mathcal{B} -eigenvectors are sufficiently close to the population ones.

The second assumption is on the proportion of cluster nodes: Recall that $\delta = \frac{|\mathcal{C}|}{|\mathcal{V}|}$.

Assumption 2. The constants δ , $|I|$, and K satisfy that

$$(3.5) \quad \frac{\delta}{1-\delta} \frac{|I| - K}{K} < \frac{d_0(1-\varepsilon_2)}{\underline{d}_0(1+\varepsilon_1)},$$

where $\varepsilon_1, \varepsilon_2$ are as in Assumption 1.

Prototypical cases where Assumptions 1 and 2 are satisfied include the example in section 2.3 (Figures 1) and in section 4.1 (Figures 3, A.1). Theoretically, the above two assumptions guarantee that the embedding norm $S(x, 0)$ separates the blocks \mathcal{C} and \mathcal{B} at time $t = 0$ together with an upper bound of $S(x, 0)$ over \mathcal{V} .

Proposition 3.2 (initial separation by $S(x)$). Under Assumption 1, at time $t = 0$,

$$(3.6) \quad \frac{1}{n\underline{d}_0} \frac{K}{\delta} (1 - \varepsilon_1) \leq S(x) \leq \frac{1}{n\underline{d}_0} \frac{K}{\delta} (1 + 2\varepsilon_1) \quad \forall x \in \mathcal{C},$$

$$(3.7) \quad S(x) \leq \frac{1}{n\underline{d}_0} \frac{(1 + \varepsilon_2)(|I| - K)}{1 - \delta} \quad \forall x \in \mathcal{B}.$$

If furthermore, Assumption 2 holds, then

(1) the initial gap between \mathcal{B} and \mathcal{C} is at least

$$(3.8) \quad g_0 := \frac{1}{n\underline{d}_0} \frac{K}{\delta} (\#), \quad (\#) := \frac{d_0(1 - \varepsilon_1)}{\underline{d}_0} - \frac{\delta}{1 - \delta} \frac{|I| - K}{K} (1 + \varepsilon_2),$$

that is, $\forall x \in \mathcal{C}$ and $y \in \mathcal{B}$, $S(x) - S(y) \geq g_0 > 0$;

(2) at $t = 0$,

$$(3.9) \quad \sup_{x \in \mathcal{V}} S(x) \leq \frac{1}{n\underline{d}_0} \frac{K}{\delta} (1 + 2\varepsilon_1).$$

Proof is given in section 5.

3.2. Stable deformation of S and separation. We will prove the stability of $S(x, t)$ over time making use of the Hadamard variation formula for the eigenvalues and eigenvectors, after properly indexing them. Specifically, since we assume that $\underline{d}_0 > 0$, the diagonal matrix D is invertible throughout time, and the Markov matrix $P = \bar{D}^{-1}W$ is diagonalizable and similar to $D^{-1/2}WD^{-1/2}$. Under the matrix perturbation model (2.1) which is linear in t , the n eigenvalues of the Markov matrix P can be indexed as $\lambda_1(t), \dots, \lambda_n(t)$, so that they are descending at $t = 0$, i.e., $\lambda_{k+1}(0) \leq \lambda_k(0)$, and differentiable with respect to t for $0 \leq t \leq 1$ (Chapter 2 of [17]). Similarly to the classical Hadamard variation formula, the evolution equation of λ_k can be shown to be

$$(3.10) \quad \dot{\lambda}_k = \psi_k^T (\dot{W} - \lambda_k \dot{D}) \psi_k,$$

and the equation of the associated eigenvector ψ_k is, when valid,

$$(3.11) \quad \dot{\psi}_k = -\frac{1}{2}(\psi_k^T \dot{D} \psi_k) \psi_k + \sum_{j \neq k} \frac{\psi_j^T (\dot{W} - \lambda_k \dot{D}) \psi_k}{\lambda_k - \lambda_j} \psi_j,$$

that is, (3.11) holds on time intervals when no eigencrossing of any pair of λ_k and λ_j happens. The derivation of (3.10), (3.11) is left to Appendix B.

Though the n eigenvalues are ordered from large to small at $t = 0$, an eigencrossing (or neighboring eigenvalues becoming very close) may happen as t increases, as illustrated in the diagram in Figure 2, and numerically in the toy example in Figure 1. This voids a direct adoption of (3.11) unless one shows that the singularity does not affect the differentiability of the eigenvector branches before and after the crossing, which is still possible in our setting [17]. However, even if (3.11) can be made valid with such an effort, when an eigencrossing or a near crossing happens there is generally no control on the speed of change of the associated pair of eigenvectors. Some steep changes of eigenvectors are shown in the toy example in Figure 1, at times of (near) eigencrossings. This instability of eigenvectors under matrix perturbation underlies the main difficulty of justifying the use of leading eigenvectors in this environment, for both theoretical analysis and for algorithms.

The main observation of this work is to overcome such instability by considering the spectral embedding norm instead of individual eigenvectors. A key quantity needed in the stability bounds (of both the eigenvalues and the embedding norm) is the \mathcal{C} - \mathcal{B} “connection strength”, measured by

$$(3.12) \quad C := \sum_{\substack{x \in B \\ y \in C}} W(x, y).$$

The analysis needs C to be a small compared to the magnitude of the node degrees, specifically, $\frac{C}{d_0}$ needs to be a small constant. We note that the condition may be much stronger than encountered in applications due to the reliance on a spectral gap between I and I^c eigenvalues. To be specific, we define the I -eigengap (depending on time t) to be

$$(3.13) \quad \Delta(t) := \min_{\substack{i \in I \\ j \notin I}} |\lambda_i(t) - \lambda_j(t)|, \quad t \in [0, 1].$$

Such an “ I -eigengap” prevents eigenvalues from I and I^c from getting too close, but allows arbitrary eigencrossings within I and within I^c . While needed in the perturbation analysis, we note that $\Delta(t)$ should be viewed as an artifact due to the limitation of our theory (see remark after Theorem 3.5). However, this is essentially different from the traditional spectral gap assumed after the K th eigenvalue. All proofs for this section are given in section 5.

The following proposition proves the preserved I -eigengap assuming an initial one, based upon the stable evolution of eigenvalues; cf., (3.10).

Proposition 3.3 (preservation of I -eigengap). *Under Assumption 1, C as in (3.12), if for some constant $\Delta > 0$, $\Delta(0) \geq 2\Delta$, and*

$$(3.14) \quad \frac{C}{d_0} \leq \frac{1}{8}\Delta,$$

then

$$\Delta(t) \geq \Delta \quad \forall 0 \leq t \leq 1.$$

As shown in the proof, the constant C in (3.14) can be improved to be a smaller one which only involves the maximum of the row sum of E rather than the summation of all the entries in E . We however keep the stronger condition with C here as control by C is needed later in Theorem 3.5 to control the deformation of $S(x)$ for each x .

The significance of the preserved I -eigengap is that we can derive the evolution equation of the embedding norm $S(x, t)$ without being concerned with the eigencrossings within I (and within I^c). This is possible by relying on $S(x)$ being the (x, x) th diagonal entry of the spectral projection matrix $P_I := \sum_{k \in I} \psi_k \psi_k^T$, which can be written in the form of a contour integral of the resolvent in the complex plane where the contour circles the eigenvalues in I throughout $t \in [0, 1]$, as illustrated in Figure 2. The evolution equation below only requires the eigenvalue difference $\lambda_k - \lambda_j$ to be nonvanishing when one is from I and the other is from I^c . Actually, this difference is bounded from below by the constant Δ by Proposition 3.3.

Proposition 3.4 (evolution of S). *When Proposition 3.3 applies, for $0 < t < 1$,*

$$(3.15) \quad \frac{\partial}{\partial t} S(x, t) = - \sum_{\substack{k \in I \\ j \in I}} (\psi_j^T \dot{D} \psi_k) \psi_k(x) \psi_j(x) + 2 \sum_{\substack{k \in I \\ j \notin I}} \frac{\psi_j^T (\dot{W} - \lambda_k \dot{D}) \psi_k}{\lambda_k - \lambda_j} \psi_k(x) \psi_j(x).$$

We then derive the main result.

Theorem 3.5 (separation at $t = 1$). *Under Assumptions 1 and 2, if for some constant $\Delta > 0$, the following conditions are satisfied:*

- (i) $\Delta(0) \geq 2\Delta$, $\Delta(t)$ as in (3.13);
- (ii) C as in (3.12),

then

$$(3.16) \quad \frac{C}{\underline{d}_0} \leq \frac{\Delta}{8} \frac{1}{1 + \frac{\Delta}{4}} \log \left(1 + \frac{1}{2} \cdot \frac{(\#)}{1 + 2\varepsilon_1} \right),$$

where $(\#)$ is defined in (3.8), and $(\#) > 0$ under Assumption 2.

Then the two parts \mathcal{B} and \mathcal{C} can be separated by thresholding the embedding norm, i.e., there exists a constant τ s.t. at $t = 1$

$$\begin{aligned} S(x) &> \tau \quad \forall x \in \mathcal{C}, \\ S(x) &< \tau \quad \forall x \in \mathcal{B}. \end{aligned}$$

In practice, τ can be set to a certain quantile of the empirical values of $S(x)$ on all the nodes. The proof controls the pointwise change of $S(x)$ using the condition (3.16), which can be viewed as bounding the change after adding one weighted edge in the \mathcal{C} - \mathcal{B} connection (one entry in E) and then summing over accumulatively for all such edges.

We now make a few comments on the assumptions needed in Theorem 3.5.

First, we show that in the typical setting the right-hand side (r.h.s.) of (3.16) can be simplified to be a multiple of Δ similar to the form of (3.14) ((3.16) implies the latter, as will be shown in the proof). Note that the r.h.s. is greater than (using that $\log(1 + \frac{x}{2}) > \frac{2}{5}x$ for $0 < x \leq 1$)

$$(3.17) \quad \frac{\Delta}{8} \cdot \frac{1}{1 + \frac{\Delta}{4}} \cdot \frac{2}{5} \cdot \frac{(\#)}{1 + 2\varepsilon_1},$$

and thus unless $(\#)$ is too small, this term would be comparable to $\frac{\Delta}{8}$. To be specific, suppose

that δ is so small that the first term in the formula of $(\#)$ (3.8) dominates, which makes $(\#)$ approximately $\frac{d_0}{\delta}$, assuming that ε_1 and ε_2 are small constants. This is reasonable since we typically apply the proposed method when the initial separation is large, where the initial gap $g_0 = \frac{K}{nd_0}(\#)$. Furthermore, in such cases, if the graph has balanced degree, i.e., $\underline{d}_0 \approx \overline{d}_0$, then $(\#)$ would be close to 1. Combined with Δ being small, e.g., $\Delta < 0.1$, (3.17) is then approximately $0.39 \cdot \frac{\Delta}{8}$.

Second, our theory relies on the I -eigengap condition (3.16). Empirically, we have observed that the matrix $P = I - L$ may have a bulk of eigenvalues near 1 and, when moving from 1 (towards 0), namely, going “deeper” into the spectrum, the eigengaps may be larger compared to immediately near 1. The eigenvalue distribution also depends on the construction of the graph affinity W . For example, when the bandwidth parameter k_{ST} used in the self-tuning kernel is large the eigengaps are usually also larger. The empirical eigenvalues and eigengaps for the data example in Figure 1 are shown in Figure A.4(d) and (e). Our theory thus improves over the traditional eigengap constraint on the K th eigengap, where K is the number of clusters.

Nevertheless, we conjecture that the requirement on C in (3.16) and the need for an I -eigengap are more restrictive than what occurs in practical applications. We have observed that in practice the embedding norm S can successfully separate \mathcal{C} from \mathcal{B} even when these conditions are not satisfied. This suggests that the analysis here is likely to be not tight: for one thing, the relaxation of the term $\frac{1}{(\lambda_k - \lambda_j)}$ for $k \in I, j \in I^c$ by $\frac{1}{\Delta}$ is crude, and can be improved, e.g., under proper assumptions of the eigenvalue distribution. We think that further analysis should be able to relax the constraints in (3.16) and the I -eigengap.

3.3. Extensions of the analysis. The main result Theorem 3.5 extends to the following cases, with proof sketches given.

1. *Weighted embedding norm.* The definition of the embedding norm S can be generalized as

$$(3.18) \quad S(x) = \sum_{k \in I} f(\lambda_k) \psi_k(x)^2,$$

where $f(\lambda)$ is a (complex) analytic function which is real valued on real λ . With f being a power of λ and a certain exponential function, S is related to diffusion distance [8] and heat kernel signature [37], respectively, to be discussed more in the last section.

We have been addressing the special case where $f = 1$. To extend the analysis to any analytic f , consider the contour integral of $f(z)R(z)$, R being the resolvent (defined in (5.10)), and then the time-evolution equation of $S(x, t)$ can be shown to be

$$\begin{aligned} \dot{S}(x) = & 2 \sum_{\substack{k \in I \\ j \notin I}} \frac{f(\lambda_k)}{\lambda_k - \lambda_j} (\psi_k^T (\dot{W} - \lambda_k \dot{D}) \psi_j) \psi_k(x) \psi_j(x) \\ & + \sum_{\substack{k, j \in I \\ k \neq j}} \left(\frac{f(\lambda_k) - f(\lambda_j)}{\lambda_k - \lambda_j} \psi_k^T \dot{W} \psi_j - \frac{\lambda_k f(\lambda_k) - \lambda_j f(\lambda_j)}{\lambda_k - \lambda_j} \psi_k^T \dot{D} \psi_j \right) \psi_k(x) \psi_j(x) \\ & + \sum_{k \in I} (f'(\lambda_k) \psi_k^T \dot{W} \psi_k - (zf(z))'(\lambda_k) \psi_k^T \dot{D} \psi_k) \psi_k(x)^2, \end{aligned}$$

where in the case that the eigenvalues λ_k and λ_j coincide, the term $\frac{f(\lambda_k)-f(\lambda_j)}{\lambda_k-\lambda_j}$ is replaced by $f'(\lambda_k)$ and $\frac{\lambda_k f(\lambda_k)-\lambda_j f(\lambda_j)}{\lambda_k-\lambda_j}$ by $(zf(z))'(\lambda_k)$. So the r.h.s. is well-defined when an eigencrossing within I happens, and the terms $\frac{f(\lambda_k)-f(\lambda_j)}{\lambda_k-\lambda_j}$ and $\frac{\lambda_k f(\lambda_k)-\lambda_j f(\lambda_j)}{\lambda_k-\lambda_j}$ are uniformly bounded due to the analyticity of f . When $f = 1$, the equation reduces to (3.15). Proceeding with the same technique as in the proof of the main result, the deformation bound of $S(x, t)$ will then involve constant factors which depend on the boundedness of f and f' on $[0, 1]$. Specifically, the constant \tilde{C} will need to be redefined to be $(c_1 + \frac{c_2}{\Delta})\frac{2C}{d_0}$, where c_1 and c_2 are absolute constants. E.g., when $f(\lambda) = \lambda^p$, $p > 0$, c_2 remains 4 (which is the dominating term with small Δ) and $c_1 = (p + 1)$.

2. *Unequal cluster size in \mathcal{C} .* The requirement of equal cluster size of the K clusters in \mathcal{C} can be relaxed. Specifically, suppose that the K clusters have varying sizes $|C_j| = \delta_j n$ and $\sum_{j=1}^K \delta_j = \delta$. Let $\delta_{\min} = \min_{1 \leq j \leq K} \delta_j$ and similarly define δ_{\max} . Then under Assumption 1, (3.6) and (3.7) become

$$(3.19) \quad \frac{1 - \varepsilon_1}{nd_0 \delta_j} \leq S(x, 0) \leq \frac{1}{nd_0} \left(\frac{\varepsilon_1}{\delta/K} + \frac{1 + \varepsilon_1}{\delta_j} \right) \quad \forall x \in C_j \text{ for some } j,$$

$$(3.20) \quad S(x, 0) \leq \frac{1}{nd_0} \frac{(1 + \varepsilon_2)(|I| - K)}{1 - \delta} \quad \forall x \in \mathcal{B}.$$

Define for $j = 1, \dots, K$,

$$(3.21) \quad g_{j,0} := \frac{1}{nd_0} \left(\frac{d_0(1 - \varepsilon_1)}{d_0 \delta_j} - \frac{(1 + \varepsilon_2)(|I| - K)}{1 - \delta} \right),$$

and the minimum of $g_{j,0}$ is

$$g_{\min,0} = \frac{1}{nd_0} \left(\frac{d_0(1 - \varepsilon_1)}{d_0 \delta_{\max}} - \frac{(1 + \varepsilon_2)(|I| - K)}{1 - \delta} \right).$$

Modify Assumption 2 so that $g_{\min,0} > 0$, then the initial separation of $S(x, 0)$ on \mathcal{C} and \mathcal{B} is at least $g_{\min,0}$ (and more precisely $g_{j,0}$ between C_j and \mathcal{B}), and (3.9) becomes

$$(3.22) \quad \bar{S}(0) = \sup_{x \in \mathcal{V}} S(x) \leq \frac{1}{nd_0} \left(\frac{\varepsilon_1}{\delta/K} + \frac{1 + \varepsilon_1}{\delta_{\min}} \right).$$

Note that Propositions 3.3 and 3.4 and claims (i) and (ii) in the proof of Theorem 3.5 do not rely on Assumptions 1 or 2 and are valid. As a result, it can be shown that the $t = 1$ separation between \mathcal{C} and \mathcal{B} by S holds as long as

$$(3.23) \quad g_{\min,0} \geq 2(e^{\tilde{C}} - 1) \frac{1}{nd_0} \left(\frac{\varepsilon_1}{\delta/K} + \frac{1 + \varepsilon_1}{\delta_{\min}} \right), \quad \tilde{C} = \left(1 + \frac{4}{\Delta} \right) \frac{2C}{d_0}.$$

This condition is more restrictive when the cluster sizes in \mathcal{C} are less balanced, namely, when the difference $\delta_{\max} - \delta_{\min}$ becomes larger. In our numerical experiments, all the subclusters are

of comparable sizes (in the outlier detection in images, $K = 1$ or 2 , and in image segmentation the clusters are of similar sizes), while we note that an extremely unbalanced cluster size, e.g., very small δ_{\min} , could affect the performance of the method.

3. *Detection of parts of \mathcal{C} .* The above argument leads to a “personalized” detection condition for each cluster C_j in \mathcal{C} , that is, even when $S(x, 1)$ fails to separate some clusters in \mathcal{C} from \mathcal{B} it may still successfully detect the rest. To see this, note that the proof of the theorem actually gives the following: for any subsets E_1 and E_2 of \mathcal{V} , if

$$\min_{x \in E_1} S(x, 0) - \max_{x \in E_2} S(x, 0) \geq g_{E_1, E_2} > 0,$$

then E_1 and E_2 can be separated by $S(x, 1)$ by some threshold as long as

$$g_{E_1, E_2} \geq 2(e^{\tilde{C}} - 1)\bar{S}(0).$$

The previous results correspond to $E_1 = \mathcal{C}$ and $E_2 = \mathcal{B}$. Let E_1 be any individual cluster C_j , then since $\bar{S}(0)$ is upper bounded by (3.22), we have that each cluster C_j can be separated from \mathcal{B} by $S(x, 1)$ if $g_{j,0}$ as defined in (3.21) is larger than the r.h.s. of (3.23).

4. *Initial inclusion of I .* Assumption 1(a) can be relaxed by only requiring K' \mathcal{C} -eigenvectors in I , $0 < K' \leq K$, as long as they contribute to a sufficiently large $S(x, 0)$ on \mathcal{C} , or any subset of \mathcal{C} such as an individual cluster C_j . The separation guarantee at time 1 follows the same argument as in item 2 above, where the quantities (3.19), (3.20) and consequently (3.21), (3.22) need to be modified. The precise condition is not pursued here. In practice, this means that even if less than K “nearly” \mathcal{C} -eigenvectors are included in I , the method may still be able to detect part of \mathcal{C} from \mathcal{B} .

4. Experiments. In this section we will apply the spectral embedding norm to both synthetic and real-world datasets, in scenarios of both single outliers and multiple clusters in a cluttered background. Codes are available at <https://github.com/xycheng/EmbeddingNorm>.

4.1. Manifold data toy example. We begin with a simulated dataset in \mathbb{R}^2 composed of a manifold-like background \mathcal{B} and clusters in \mathcal{C} according to the following model. The background \mathcal{B} consists of independent and identically distributed (i.i.d.) samples x_i distributed as $x_i = y_i + n_i$, where y_i are uniformly distributed on the unit circle, which is a one-dimensional manifold, and $n_i \sim \mathcal{N}(0, \epsilon_{\mathcal{B}}^2 I)$ with $\epsilon_{\mathcal{B}} = 0.01$. \mathcal{C} contain K equal-sized subclusters, each has i.i.d. samples drawn from $\mathcal{N}(\mu_j, \epsilon_{\mathcal{C}}^2 I)$, where the μ_j are centered close to the circle, and $\epsilon_{\mathcal{C}} = 0.02$. We generate $n = 5000$ points, and the number of points in \mathcal{C} is set to be δn for positive δ , rounded to the closest integer. To measure the accuracy of the detection of \mathcal{C} we compute the F1 score, $F1 = \frac{2pr}{p+r}$, where $p := \frac{TP}{TP+FP}$, $r := \frac{TP}{TP+FN}$, and TP, FP, and FN stand for true positive, false positive, and false negative, respectively. The δ -quantile of the empirical values of S is used as the threshold τ when computing the classification.

Figure 3 shows results for a typical realization of the dataset with $K = 10$, $\delta = 0.1$. From (c) it can be seen that the eigenvalues do not reveal any clear eigengap at $K = 10$. The first K eigenvectors do not give a clear indication of the cluster \mathcal{C} , but are mainly supported on the \mathcal{B} -eigenvectors, as shown in (e) for $k = 2$ and 8 . Examining up to the first $k \sim 40$ ones, certain eigenvectors are more localized on \mathcal{C} when $k > K$, e.g., $k = 24$ and 26 . The

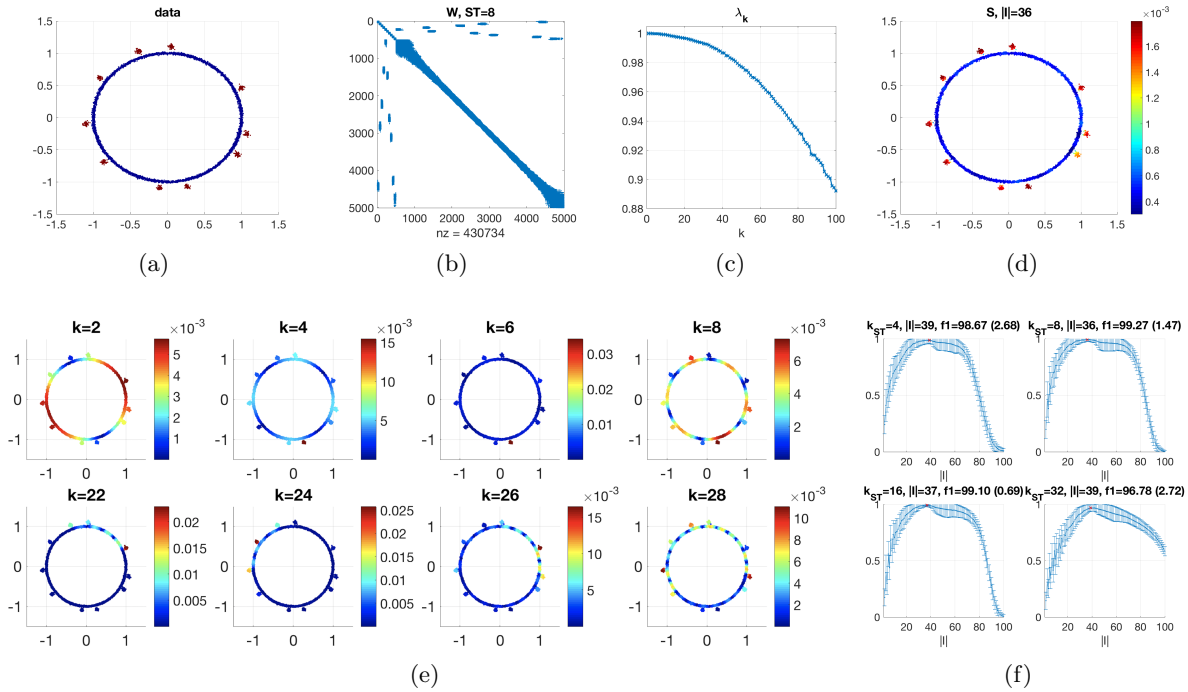


Figure 3. Detection of \mathcal{C} from a manifold-like \mathcal{B} : (a) $n = 5000$ data points in \mathbb{R}^2 sampled on $\mathcal{B} \cup \mathcal{C}$, where \mathcal{B} points lie close to a circle (blue) and \mathcal{C} points form $K = 10$ clusters lying nearby (red), $\delta = 0.1$. (b) The affinity matrix $W(t)$ at $t = 1$; cf., (2.1). (c) The first 100 eigenvalues of the Markov matrix. (d) The plot of S . (e) k th eigenvectors of multiple k 's of the Markov matrix. (f) F1 score of the detection of \mathcal{C} by thresholding the values of S_I , where $|I|$ varies from 2 to 100, and for multiple choices of self-tuning parameter (k -nearest neighbor in self-tuning, denoted by k_{ST}). Mean and standard deviation of the F1 score are shown, and optimal values of $|I|$ are indicated by a red cross.

embedding norm S_I clearly separates \mathcal{C} from \mathcal{B} , as shown in (d). The results are not sensitive to algorithmic parameter choices. Let k_{ST} be the k -nearest neighbor used to set the local self-tuning scale [45] in constructing the affinity matrix W . Then, throughout varying values of the parameter k_{ST} , the F1 score of the detection by thresholding S reveals a “plateau” of valid values of $|I|$, e.g., when $k_{ST} = 8$, the range of $|I|$ is about $30 \sim 45$, with the optimal F1 score obtained at $|I| = 36$. The best F1 scores for $k_{ST} = 4, 8$, or 16 are all greater than 0.98.

Similar results are obtained for smaller $K = 2$, where $\delta = 0.02$, as shown in Figure A.1. The condition (3.5) in Assumption 2 suggests that $|I|$ is chosen to be proportional to $\frac{K}{\delta}$, and this is revealed in Figures 3 and A.1 (in these two examples $\frac{K}{\delta}$ is kept to be the same) as the plateaus of valid $|I|$ are at about the same range, across values of k_{ST} .

4.2. Image anomaly detection. Anomaly detection can be seen as a special case of clustering in which there is a vast imbalance in the size of clusters, i.e., background versus anomaly, and the density of each cluster. In image anomaly detection, the goal is to detect a small compact region (subset of connected pixels) that differs from the normal image background. In general, we can assume the number of anomalies K to be small or even 1.

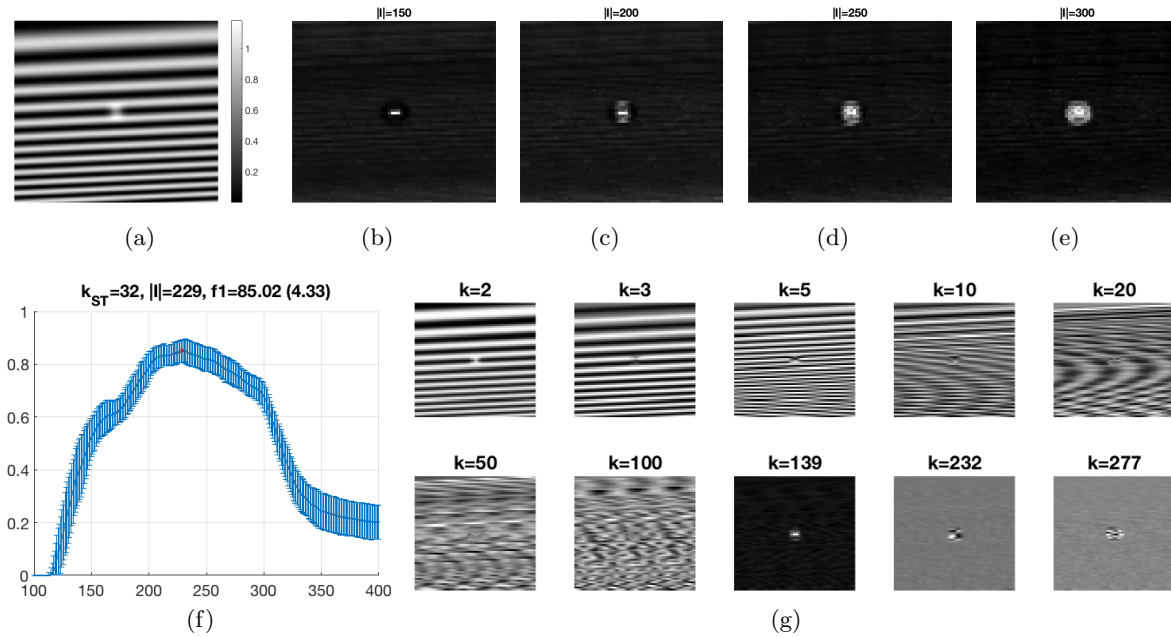


Figure 4. Detection of outlier patches in a synthetic image. (a) Image consisting of background stripes and an anomaly region in the center, $\delta \approx 0.01$. (b)–(e) Embedding norm S_I computed with $|I| = 150, 200, 250, 300$, respectively, plotted as images, where the pixel value indicates the value of S_I and the pixel location is the center of every image patch. (f) F1 score of the detection of outlier patches by thresholding the values of S_I , where $|I|$ varies from 100 to 400. Randomness is attained by subsampling image patches and running independent replicas of the experiment. Mean and standard deviation of the F1 score are shown, and optimal values of $|I|$ are indicated by a red cross. (g) Eigenvectors for various index k , plotted as images. The last three show high-index exemplar eigenvectors which are localized on the outlier patches, while the first 100 eigenvectors are mainly supported on the background.

Figure 4 shows the numerical result on a synthetic image which consists of anomaly patches against a slowly varying background, as a model of patterned images. The image is a function on $[-1, 1]^2$ expressed as

$$I(x, y) = \left(1 + \frac{1}{2} \cos((0.05x + y + 1.5)^2 \cdot 2\pi) \right) + 0.6 \exp \left\{ -\frac{x^2 + y^2}{2 \cdot 0.05^2} \right\},$$

where the first term models the background stripes, and the second term models the outlier region in the center, as shown in (a). The image size is 200×200 , and $n = 4096$ image patches of size 9×9 are extracted with a stride of 3. True outlier patches are identified by thresholding the value of the Gaussian bump with a fraction of $\delta \approx 0.01$. The graph affinity is computed with self-tuning bandwidth [45] and $k_{ST} = 32$. The spectral embedding norm computed for various values of $|I|$ is shown in (b)–(e), all of which identify the outlier region with improved performance for $|I| = 200$ and 250. The leading eigenvectors up to $k = 100$ fail to be indicative of the outlier region, while those with higher k may be, as shown in (g). To quantitatively evaluate the performance of the proposed method, we threshold S_I at




Figure 5. Detecting anomalies in side-scan sonar images. (A), (C) Side-scan sonar images with sea mines indicated by a red circle. (B) Spectral embedding norm S_I for image (a) with increasing $|I|$. $|I| = 50$ reveals the sea mine and the separation is stable up to $|I| = 100$. (D) Spectral embedding norm S_I for image in (C) with increasing $|I|$. The sea mine first “appears” for $|I| = 6$. For $|I| = 13$ the sea mine separates cleanly from the background, and as $|I|$ increases to 100 the embedding norm starts to slightly reveal background components as well. (E) Exemplar eigenvectors computed from the image in (C) plotted as an image demonstrating localization on the background in the three left plots and localization on the sea mine in the two right plots.

the 0.99 quantile to detect the anomaly patches, and compute the F1 score. We repeat 100 experiments by randomly subsampling 3000 patches and the average F1 score is shown in (f) with standard deviation. The method achieves an average best $F1 = 85.02$ when $|I| = 229$. The results are similar with $k_{ST} = 16$ and 64.

In Figure 5 we demonstrate on real-world images that eigenvectors localizing on the anomaly can be buried deep within the spectrum of the image, and that by calculating the spectral embedding norm we can separate the anomalies from a cluttered background. We examine two side-scan sonar images containing a single sea mine, displayed in (A) and (C), where we consider the sea mine to be an anomaly (indicated by a red circle). The sea mine can appear as either a bright highlight (A) or as only a dark shadow (C), which is due to the object blocking the sonar waves from reaching the seabed. The background is composed of sea-bed reverberations and exhibits great variability in appearance. For the side-scan sonar images, in practice δ is in the range of $5 \times 10^{-3} \sim 5 \times 10^{-4}$. To construct the affinity matrix W , we extract all overlapping patches of size 8, and build a nearest neighbor graph with 64 neighbors, setting $k_{ST} = 32$. (B) and (D) display $S(x)$ for all pixels x in the image for increasing values of $|I|$. For both images the sea mine is revealed consistently for a wide range of values, while the background is suppressed. Note that in both cases this requires looking deep enough in the spectrum, and summing over the first few eigenvectors brings out background structures.

Finally, (E) displays eigenvectors ψ_k for the side-scan sonar image in (C), where the three left eigenvectors localize on the background \mathcal{B} , revealing its periodic nature at different scales and orientations, while the eigenvectors in the right two plots localize on the sea mine \mathcal{C} .

4.3. Calcium imaging. Calcium imaging is an experimental method in neuroscience that enables imaging the individual activity of hundreds of neurons in an awake behaving animal, at cellular resolution [38]. The acquired data are composed of a spatiotemporal volume, where, after motion correction, the neuron locations are fixed and the temporal activity consists of hundreds to tens of thousands of time frames. There is also varying temporal activity in the background (neuropil). Thus, these data can also be viewed as an image whose pixels lie in a high-dimensional space (time frames), consisting of hundreds of clusters (neurons) in an image plane with a nontrivial background, which matches our problem setting.

The analysis pipeline of calcium imaging typically includes calculating a two-dimensional (2D) image that depicts the structure that exists in this volume and highlights the existing neurons. Such images serve for manual segmentation, to align volumes across days (where the field of view may shift), to display neurons detected by automatic and manual means, and even for initialization of automatic region of interest extraction algorithms methods [27, 30]. A common choice is the temporal correlation image [35], or the temporal mean image. Here we show that the spectral embedding norm provides a meaningful visualization of the data, with sharp morphology and suppression of noise from the background clutter. For these datasets, depending on the brain region and neuron type being imaged, K is in the range of dozens to a few hundreds, and $\delta \lesssim 0.1$.

In Figure 6, we analyze a publicly available dataset from Neurofinder [3]. The images are 512×512 pixels and 8000 time frames have been recorded at 8 Hz. Ground truth labels provided with the dataset include 197 identified neurons; however, note that recent papers point out that the ground truth on Neurofinder datasets is probably lacking, i.e., not all neurons are labeled [30, 36]. The affinity matrix W is calculated using a nearest neighbor graph for all pixels, represented as high-dimensional vectors in time, with 50 nearest neighbors. To accelerate the nearest neighbor search, dimensionality is reduced from 8000 to 300 using principal component analysis. (A) displays examples of eigenvectors from both the background (top) and localizing on single neurons (bottom).

In (B) we compare the spectral embedding norm (right) to the temporal mean (left) and temporal correlation image (middle). In each image, the values (mean/correlation/norm) appear in the green channel, while we overlay in the red channel a mask of the ground truth labels that were manually detected (where the two overlap it appears as yellow). The mean image exhibits a strong background, while neurons appear as typical “donuts” [27]. In the correlation image, the background mostly appears as noise. In comparison, the background has been suppressed in the spectral embedding norm image, while neurons which are barely or not at all visible in the correlation image appear as bright clusters.

To quantify, the separation of background and clusters, we segment the spectral embedding norm image for increasing $|I|$, and compare the overlap between the segmented clusters and the given ground-truth mask. We set the threshold τ to be the value of the 93rd percentile of S values for each value of $|I|$. In (C) we plot the F1 score for this segmentation and demonstrate a plateau of stable F1-score values for $|I|$ in the range 200–250. To demonstrate

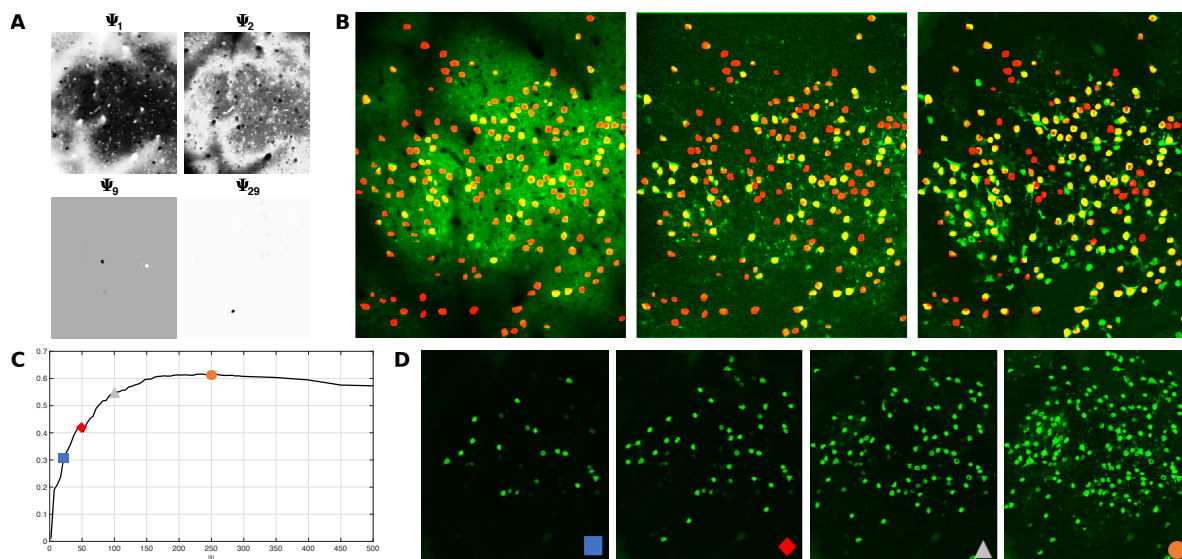


Figure 6. (A) Example of Laplacian eigenvectors on a calcium imaging dataset from Neurofinder, either mainly supported on the background clutter (top) or localizing on neurons which are clusters (bottom). (B) Images of the temporal mean (left), temporal correlation (middle), and spectral embedding norm (right) for a Neurofinder dataset. The spectral embedding norm has both removed the background (which is present in the mean image) and enhanced the appearance of the structure in the image: neuronal soma and dendrites with sharp morphology. The correlation image is much noisier with fewer visual neurons. (C) F1 score of segmenting neurons from background based on the spectral embedding norm image, for increasing $|I|$ values. For a range of values (200–250) the F1 score plateaus, and then decreases as the number of included eigenvectors increases. (D) $S(x)$ for $|I| = 20, 49, 100, 250$ demonstrates how more and more clusters are revealed.

the property of the spectral embedding norm to perform partial detection of \mathcal{C} , we display $S(x)$ for multiple values of $|I|$ in (D). Note that we are not performing clustering here, but rather demonstrating how the embedding norm can be used to separate meaningful structure from background clutter. Thus, beyond visualization, this approach can then serve to remove the background and focus only on the remaining clusters in \mathcal{C} , thus simplifying subsequent clustering and data analysis tasks.

5. Proofs.

Proof of Proposition 3.2. It suffices to prove (3.6) and (3.7) because Assumption 2 implies that the r.h.s of (3.7) is strictly less than the left-hand side of (3.6) by g_0 , and then claims (1) and (2) directly follow.

To prove (3.6), note that for any $x \in \mathcal{C}$, at $t = 0$,

$$S(x) = \sum_{k \in I} \psi_k(x)^2 = \sum_{k \in I \cap \Psi^{\mathcal{C}}} \psi_k(x)^2.$$

By Assumption 1(a), up to a possible K -by- K rotation among the K eigenvectors in $I \cap \Psi^{\mathcal{C}}$, we assume that ψ_j is the eigenvector associated with the subcluster C_j , $j = 1, \dots, K$, and

then

$$(5.1) \quad S(x) = \sum_{j=1}^K \psi_j(x)^2,$$

as the rotation preserves the squared sum. Furthermore, suppose that $x \in C_{j_x}$, Assumption 1(b) gives

$$\begin{aligned} \psi_{j_x}(x)^2 &\in \frac{1}{\nu(C_{j_x}, 0)} [1 - \varepsilon_1, 1 + \varepsilon_1], \quad j = j_x, \\ \psi_j(x)^2 &\leq \frac{\varepsilon_1}{\nu(C, 0)}, \quad j \neq j_x. \end{aligned}$$

Plugging this into (5.1), it shows that

$$S(x) \leq \frac{1 + \varepsilon_1}{\nu(C_{j_x}, 0)} + (K - 1) \frac{\varepsilon_1}{\nu(C, 0)},$$

and together with $\nu(C_j) = \sum_{x \in C_j} d(x, 0) \geq \underline{d}_0 |C_j| = \underline{d}_0 \frac{\delta n}{K}$ for any j (the K subclusters are of equal size) and similarly $\nu(C) \geq \underline{d}_0 |C| = \underline{d}_0 \delta n$, it gives the upper bound in (3.6). Considering the lower bound, (5.1) continues as

$$S(x) \geq \psi_{j_x}(x)^2 \geq \frac{1 - \varepsilon_1}{\nu(C_{j_x}, 0)}.$$

Combined with $\nu(C_j) = \sum_{x \in C_j} d(x, 0) \leq \overline{d}_0 |C_j| = \overline{d}_0 \frac{\delta n}{K}$ for any j , this gives the lower bound in (3.6).

Proof of (3.7): For any $x \in \mathcal{B}$, at $t = 0$,

$$(5.2) \quad S(x) = \sum_{k \in I} \psi_k(x)^2 = \sum_{k \in I \cap \Psi^{\mathcal{B}}} \psi_k(x)^2 \leq \sum_{k \in I \cap \Psi^{\mathcal{B}}} \frac{1 + \varepsilon_2}{\nu(B, 0)},$$

where $\nu(B, 0) = \sum_{x \in \mathcal{B}} d(x, 0) \geq \underline{d}_0 |\mathcal{B}| = \underline{d}_0 n(1 - \delta)$, and the last inequality is by Assumption 1(c). Meanwhile, $|I \cap \Psi^{\mathcal{B}}| = |I| - K$ by Assumption 1(a). Plugging this into (5.2), this proves (3.7). ■

Proof of Proposition 3.3. We will establish that for any $k = 1, \dots, n$,

$$(5.3) \quad |\lambda_k(t) - \lambda_k(0)| \leq \frac{4C}{\underline{d}_0} t, \quad 0 \leq t \leq 1.$$

Given that this inequality holds, then by (3.14),

$$(5.4) \quad |\lambda_k(t) - \lambda_k(0)| \leq \frac{4C}{\underline{d}_0} \leq \frac{1}{2} \Delta \quad \forall 1 \leq k \leq n.$$

This means that it is impossible for $\Delta(t) < \Delta$; otherwise, there exists $|\lambda_{k_1}(t) - \lambda_{k_2}(t)| < \Delta$, where $k_1 \in I$ and $k_2 \notin I$, and then (5.4) implies that $|\lambda_{k_1}(0) - \lambda_{k_2}(0)| < 2\Delta$ which contradicts the assumption that $\Delta(0) \geq 2\Delta$.

It suffices to show (5.3) to finish the proof. To do so, we prove the following bound:

$$(5.5) \quad |\dot{\lambda}_k| \leq \frac{4C}{d_0} \quad \forall k \forall t.$$

From (3.10),

$$|\dot{\lambda}_k| = |\psi_k^T (\dot{W} - \lambda_k \dot{D}) \psi_k| \leq |\psi_k^T \dot{W} \psi_k| + |\lambda_k| |\psi_k^T \dot{D} \psi_k|.$$

As $|\lambda_k| \leq 1$ (Perron–Frobenius), then

$$(5.6) \quad |\dot{\lambda}_k| \leq |\psi_k^T \dot{W} \psi_k| + |\psi_k^T \dot{D} \psi_k|.$$

If the following claim is true, then (5.5) follows directly from (5.6):

$$(5.7) \quad |\psi_l^T \dot{W} \psi_k| |\psi_l^T \dot{D} \psi_k| \leq \frac{2C}{d_0} \quad \forall k, l = 1, \dots, n, \forall 0 \leq t \leq 1.$$

Proof of (5.7): To bound $|\psi_l^T \dot{W} \psi_k|$, note that $\dot{W} = E$, and then

$$(5.8) \quad |\psi_l^T E \psi_k| \leq \sum_{x \in \mathcal{C}, y \in \mathcal{B}} |W(x, y)| |\psi_l(x)| |\psi_k(y)| + \sum_{x \in \mathcal{B}, y \in \mathcal{C}} |W(x, y)| |\psi_l(x)| |\psi_k(y)|.$$

We prove a stronger claim which replaces C with

$$C_1 := \max \left\{ \max_{x \in \mathcal{C}} \sum_{y \in \mathcal{B}} W(x, y), \max_{y \in \mathcal{B}} \sum_{x \in \mathcal{C}} W(x, y) \right\}$$

in (5.7), and $C_1 \leq C$. To proceed, by Cauchy–Schwarz,

$$\begin{aligned} \sum_{x \in \mathcal{C}, y \in \mathcal{B}} W(x, y) |\psi_l(x)| |\psi_k(y)| &\leq \left(\sum_{x \in \mathcal{C}, y \in \mathcal{B}} W(x, y) \psi_l(x)^2 \right)^{1/2} \left(\sum_{x \in \mathcal{C}, y \in \mathcal{B}} W(x, y) \psi_k(y)^2 \right)^{1/2} \\ &\leq \left(C_1 \sum_{x \in \mathcal{C}} \psi_l(x)^2 \right)^{1/2} \left(C_1 \sum_{y \in \mathcal{B}} \psi_k(y)^2 \right)^{1/2}. \end{aligned}$$

We then use that $\forall k$,

$$(5.9) \quad 1 = \sum_{x \in \mathcal{V}} \psi_k(x)^2 d(x) \geq d_0 \sum_{x \in \mathcal{V}} \psi_k(x)^2,$$

which gives that

$$\sum_{x \in \mathcal{C}} \psi_l(x)^2, \sum_{y \in \mathcal{B}} \psi_k(y)^2 \leq \frac{1}{d_0}.$$

This proves that

$$\sum_{x \in \mathcal{C}, y \in \mathcal{B}} W(x, y) |\psi_l(x)| |\psi_k(y)| \leq \frac{C_1}{d_0},$$

and using the same argument in the second term in (5.8), we have that

$$|\psi_l^T E \psi_k| \leq \frac{2C_1}{d_0}.$$

For $|\psi_l^T \dot{D} \psi_k|$, similarly,

$$\begin{aligned} |\psi_l^T \dot{D} \psi_k| &\leq \sum_{x \in \mathcal{V}} |\psi_l(x)| |\psi_k(x)| |\dot{d}(x)| \\ &= \sum_{x \in \mathcal{V}} |\psi_l(x)| |\psi_k(x)| \sum_y E(x, y) \\ &\leq C_1 \sum_{x \in \mathcal{V}} |\psi_l(x)| |\psi_k(x)| \\ &\leq \frac{C_1}{d_0}, \end{aligned}$$

where Cauchy–Schwarz and (5.9) are used to obtain the last inequality.

Note that while time dependence has been omitted in all the notations, the above arguments hold throughout time $t \in [0, 1]$. ■

Proof of Proposition 3.4. As explained in the text, one may first establish the validity of (3.11) and then verify the formula (3.15) based on the former by observing the cancelation of terms. As an alternative approach, we use the contour integral of the resolvent.

For $z \in C$ and not an eigenvalue of $P = D^{-1}W$, define

$$R(z) = (W - zD)^{-1},$$

where the time dependence is omitted. Since $P = \Psi \Lambda \Phi^T$, $\Lambda = \text{diag}\{\lambda_1, \dots, \lambda_n\}$, $\Phi = D\Psi$, and $\Psi^T \Phi = I$, one can verify the equivalent form of R as

$$(5.10) \quad R(z) = \Psi(\Lambda - zI)^{-1}\Psi^T = \sum_{k=1}^n \frac{\psi_k \psi_k^T}{\lambda_k - z}.$$

This means that

$$P_I = \sum_{k \in I} \psi_k \psi_k^T = -\frac{1}{2\pi i} \oint_{\Gamma} R(z) dz,$$

where the contour Γ is such that the eigenvalues in I (I^c) stay inside (outside) Γ throughout time t (Figure 2), and such Γ exists due to Proposition 3.3. Thus the above expression of P_I holds for all time t and, as a result,

$$\dot{P}_I = -\frac{1}{2\pi i} \oint_{\Gamma} \dot{R}(z) dz.$$

By differencing both sides of

$$(W - zD)R = I,$$

one obtains that

$$\dot{R} = -R(\dot{W} - z\dot{D})R.$$

This means that

$$\begin{aligned} \dot{P}_I &= \frac{1}{2\pi i} \oint_{\Gamma} R(z)(\dot{W} - z\dot{D})R(z)dz \\ &= \sum_{k=1}^n \sum_{l=1}^n \frac{1}{2\pi i} \oint_{\Gamma} \frac{\psi_k^T(\dot{W} - z\dot{D})\psi_l}{(\lambda_k - z)(\lambda_l - z)} \psi_k \psi_l^T dz \\ (5.11) \quad &= \sum_{k=1}^n \sum_{l=1}^n (\psi_k^T(\alpha_{kl}\dot{W} - \beta_{kl}\dot{D})\psi_l) \psi_k \psi_l^T, \end{aligned}$$

where

$$\alpha_{kl} = \frac{1}{2\pi i} \oint_{\Gamma} \frac{1}{(\lambda_k - z)(\lambda_l - z)} dz, \quad \beta_{kl} = \frac{1}{2\pi i} \oint_{\Gamma} \frac{z}{(\lambda_k - z)(\lambda_l - z)} dz.$$

By Cauchy's integral formula, one can verify the following:

- (a) When $k \in I, l \in I$, $\alpha_{kl} = 0, \beta_{kl} = 1$.
- (b) When $k \in I, l \notin I$, $\alpha_{kl} = \frac{1}{\lambda_k - \lambda_l}, \beta_{kl} = \frac{\lambda_k}{\lambda_k - \lambda_l}$.
- (c) When $k \notin I, l \in I$, $\alpha_{kl} = \frac{-1}{\lambda_k - \lambda_l}, \beta_{kl} = \frac{-\lambda_l}{\lambda_k - \lambda_l}$.
- (d) When $k \notin I, l \notin I$, $\alpha_{kl} = 0, \beta_{kl} = 0$.

Then (5.11) continues as

$$\begin{aligned} \dot{P}_I &= \sum_{\substack{k \in I \\ l \in I}} -(\psi_k^T \dot{D} \psi_l) \psi_k \psi_l^T + \sum_{\substack{k \in I \\ l \notin I}} \frac{\psi_k^T (\dot{W} - \lambda_k \dot{D}) \psi_l}{\lambda_k - \lambda_l} \psi_k \psi_l^T + \sum_{\substack{k \notin I \\ l \in I}} \frac{\psi_k^T (-\dot{W} + \lambda_l \dot{D}) \psi_l}{\lambda_k - \lambda_l} \psi_k \psi_l^T \\ &= \sum_{\substack{k \in I \\ l \in I}} -(\psi_k^T \dot{D} \psi_l) \psi_k \psi_l^T + \sum_{\substack{k \in I \\ l \notin I}} \frac{\psi_k^T (\dot{W} - \lambda_k \dot{D}) \psi_l}{\lambda_k - \lambda_l} (\psi_k \psi_l^T + \psi_l \psi_k^T). \end{aligned}$$

Since $S(x) = P_I(x, x)$, the claim follows by evaluating at the entry (x, x) on both sides. ■

Proof of Theorem 3.5. We first show that condition (ii) implies (3.14). Note that

$$(\#) \leq \frac{d_0(1 - \varepsilon_1)}{d_0} \leq \frac{d_0}{d_0} \leq 1,$$

thus

$$\log \left(1 + \frac{1}{2} \cdot \frac{(\#)}{1 + 2\varepsilon_1} \right) \leq \log \left(1 + \frac{1}{2}(\#) \right) \leq \log \left(1 + \frac{1}{2} \right) < 0.5.$$

Together with $\frac{1}{1 + \frac{\Delta}{4}} < 1$, this means that the r.h.s. of (3.16) is less than $0.5 \frac{\Delta}{8}$.

As a result, under these assumptions, Propositions 3.3 and 3.4 apply. By (3.15),

$$(5.12) \quad S(x, t) - S(x, 0) = \int_0^t \left\{ - \sum_{\substack{k \in I \\ j \in I}} (\psi_j^T \dot{D} \psi_k) \psi_k(x) \psi_j(x) + 2 \sum_{\substack{k \in I \\ j \notin I}} \frac{\psi_j^T (\dot{W} - \lambda_k \dot{D}) \psi_k}{\lambda_k - \lambda_j} \psi_k(x) \psi_j(x) \right\} d\tau,$$

where in the integrand all the variables involving time take value at time τ .

Introducing the notation

$$(5.13) \quad \bar{S}(t) := \sup_{x \in \mathcal{V} = \mathcal{B} \cup \mathcal{C}} S(x, t),$$

we are going to prove the following two claims: For any $t \in [0, 1]$,

(1) $\forall x \in \mathcal{V} = \mathcal{B} \cup \mathcal{C}$,

$$|S(t, x) - S(0, x)| \leq \tilde{C} \int_0^t \bar{S}(\tau) d\tau, \quad \tilde{C} := \left(1 + \frac{4}{\Delta}\right) \frac{2C}{\underline{d}_0}.$$

(2) $\bar{S}(t) \leq e^{\tilde{C}t} \bar{S}(0)$.

If true, then

$$|S(t, x) - S(0, x)| \leq \tilde{C} \int_0^t \bar{S}(\tau) d\tau \leq \tilde{C} \int_0^t \bar{S}(0) e^{\tilde{C}\tau} d\tau = \bar{S}(0) (e^{\tilde{C}t} - 1).$$

Meanwhile, by Proposition 3.2(2), $\bar{S}(0) \leq \frac{1}{n\underline{d}_0} \frac{K}{\delta} (1 + 2\varepsilon_1)$, thus

$$|S(t, x) - S(0, x)| \leq (e^{\tilde{C}t} - 1) \frac{K(1 + 2\varepsilon_1)}{n\delta\underline{d}_0} \quad \forall x \in \mathcal{V}.$$

By Proposition 3.2(1), the initial separation on \mathcal{B} and \mathcal{C} by $S(x)$ is at least g_0 , so the threshold claimed in the theorem exists as long as

$$g_0 \geq 2(e^{\tilde{C}t} - 1) \frac{K(1 + 2\varepsilon_1)}{n\delta\underline{d}_0}.$$

This is reduced to

$$2(e^{\tilde{C}t} - 1) \leq \frac{(\#)}{1 + 2\varepsilon_1},$$

which is guaranteed by condition (ii).

To prove claim (1): By (5.12), $\forall x \in \mathcal{V}$,

$$|S(x, t) - S(x, 0)| \leq \int_0^t \text{I}(x, \tau) + \text{II}(x, \tau) + \text{III}(x, \tau) d\tau,$$

where

$$(5.14) \quad \mathrm{I}(x, \tau) = \left| \sum_{\substack{k \in I \\ j \in I}} (\psi_j^T \dot{D} \psi_k) \psi_k(x) \psi_j(x) \right|,$$

$$(5.15) \quad \mathrm{II}(x, \tau) = 2 \left| \sum_{\substack{k \in I \\ j \notin I}} \frac{\psi_j^T \dot{W} \psi_k}{\lambda_k - \lambda_j} \psi_k(x) \psi_j(x) \right|,$$

$$(5.16) \quad \mathrm{III}(x, \tau) = 2 \left| \sum_{\substack{k \in I \\ j \notin I}} \frac{\lambda_k \psi_j^T \dot{D} \psi_k}{\lambda_k - \lambda_j} \psi_k(x) \psi_j(x) \right|.$$

For $\mathrm{I}(x, \tau)$, note that

$$\begin{aligned} \mathrm{I}(x, \tau) &\leq \sum_{\substack{k \in I \\ j \in I}} \sum_{y \in \mathcal{V}} |\dot{D}(y)| |\psi_j(y)| |\psi_k(y)| |\psi_k(x)| |\psi_j(x)| \\ &= \sum_{y \in \mathcal{V}} |\dot{D}(y)| \left(\sum_{k \in I} |\psi_k(y)| |\psi_k(x)| \right)^2 \\ &\leq \sum_{y \in \mathcal{V}} |\dot{D}(y)| \left(\sum_{k \in I} |\psi_k(y)|^2 \right) \left(\sum_{l \in I} |\psi_l(x)|^2 \right) \\ &= \sum_{y \in \mathcal{V}} |\dot{D}(y)| S(y) S(x). \end{aligned}$$

Utilizing the relation that $D\Psi\Psi^T = I$, and thus $\Psi\Psi^T = D^{-1}$, we have that

$$S(x) = \sum_{j \in I} \psi_j(x)^2 \leq \sum_{j=1}^n \psi_j(x)^2 = \frac{1}{d(x)} \leq \frac{1}{\underline{d}_0} \quad \forall x \in \mathcal{V},$$

and then

$$\mathrm{I}(x, \tau) \leq S(x) \sum_{y \in \mathcal{V}} \left| \dot{D}(y) \right| \frac{1}{\underline{d}_0} = S(x) \frac{1}{\underline{d}_0} \sum_{y \in \mathcal{V}} \left| \sum_{z \in \mathcal{V}} E(y, z) \right| = S(x) \frac{2C}{\underline{d}_0}.$$

By the definition of \bar{S} in (5.13), this gives

$$(5.17) \quad \mathrm{I}(x, \tau) \leq \bar{S}(\tau) \frac{2C}{\underline{d}_0}.$$

For $\Pi(x, \tau)$, by Proposition 3.3, in the denominator $|\lambda_k - \lambda_j| \geq \Delta$, and then

$$\begin{aligned}
 \Pi(x, \tau) &\leq 2 \sum_{\substack{k \in I \\ j \notin I}} \frac{|\psi_j^T \dot{W} \psi_k|}{|\lambda_k - \lambda_j|} |\psi_k(x)| |\psi_j(x)| \\
 &\leq \frac{2}{\Delta} \sum_{\substack{k \in I \\ j \notin I}} |\psi_j^T E \psi_k| |\psi_k(x)| |\psi_j(x)| \quad (\dot{W} = E) \\
 &\leq \frac{2}{\Delta} \sum_{\substack{k \in I \\ j \notin I}} \sum_{\substack{y \in \mathcal{V} \\ z \in \mathcal{V}}} E(y, z) |\psi_j(y)| |\psi_k(z)| |\psi_k(x)| |\psi_j(x)| \\
 &= \frac{2}{\Delta} \sum_{\substack{y \in \mathcal{V} \\ z \in \mathcal{V}}} E(y, z) \sum_{k \in I} |\psi_k(z)| |\psi_k(x)| \sum_{j \notin I} |\psi_j(y)| |\psi_j(x)| \\
 &\leq \frac{2}{\Delta} \sum_{\substack{y \in \mathcal{V} \\ z \in \mathcal{V}}} E(y, z) \left(\sum_{k \in I} \psi_k(z)^2 \right)^{1/2} \left(\sum_{l \in I} \psi_l(x)^2 \right)^{1/2} \\
 &\quad \times \left(\sum_{j=1}^n \psi_j(y)^2 \right)^{1/2} \left(\sum_{m=1}^n \psi_m(x)^2 \right)^{1/2} \\
 &\leq \frac{2}{\Delta} \sum_{\substack{y \in \mathcal{V} \\ z \in \mathcal{V}}} E(y, z) \sqrt{S(z)} \sqrt{S(x)} \frac{1}{\sqrt{d(y)}} \frac{1}{\sqrt{d(x)}} \\
 &\quad \left(\text{definition of } S, \sum_{j=1}^n \psi_j(x)^2 = \frac{1}{d(x)} \right) \\
 &\leq \frac{2}{\Delta} \bar{S}(\tau) \frac{1}{\underline{d}_0} \sum_{\substack{y \in \mathcal{V} \\ z \in \mathcal{V}}} E(y, z) \\
 &\quad (\text{by that } S(x), S(z) \leq \bar{S}(\tau), \text{ and } d(y), d(x) \geq \underline{d}_0)
 \end{aligned}$$

which means that

$$(5.18) \quad \Pi(x, \tau) \leq \frac{2}{\Delta} \bar{S}(\tau) \frac{2C}{\underline{d}_0}.$$

For $\text{III}(x, \tau)$, since $|\lambda_k| \leq 1$, and $\dot{D}(y) = \sum_{z \in \mathcal{V}} E(y, z)$, one can show that

$$(5.19) \quad \text{III}(x, \tau) \leq \frac{2}{\Delta} \bar{S}(\tau) \frac{2C}{\underline{d}_0}$$

using a similar argument as in bounding $\Pi(x, \tau)$. Putting (5.17), (5.18), (5.19) together, one has that

$$|S(x, t) - S(x, 0)| \leq \int_0^t \bar{S}(\tau) \frac{2C}{\underline{d}_0} \left(1 + \frac{4}{\Delta} \right) d\tau = \int_0^t \bar{S}(\tau) \tilde{C} d\tau,$$

namely, claim (1).

To prove claim (2): Note that

$$\bar{S}(t) - \bar{S}(0) \leq \sup_{x \in \mathcal{V}} (S(x, t) - S(x, 0))$$

(suppose that $\bar{S}(t) = S(x_0, t)$ for some x_0 , then $S(x_0, t) - \bar{S}(0) \leq S(x_0, t) - S(x_0, 0)$). Since claim (1) holds uniformly for x , this implies that

$$\bar{S}(t) - \bar{S}(0) \leq \tilde{C} \int_0^t \bar{S}(\tau) d\tau,$$

and the claim then follows by Gronwall's inequality. ■

6. Further comments. *Eigenvector selection.* Related works have devised different methods to perform eigenvector selection to identify anomalies [18, 42]. The proposed spectral embedding norm can also be used for eigenvector selection. [21] demonstrated that it can be used to identify pixels which define clusters and find the embedding coordinates that best separate them from the background. An example in the anomaly detection case is given in Appendix A.

Viewed as diffusion distance. With $f(\lambda) = \lambda^p$, $S(x)$ can be interpreted as the (squared) diffusion distance between node x and the origin at diffusion time $\frac{p}{2}$ [8]. The diffusion distance can be interpreted as a geometric distance between two nodes when the affinity graph is built from data points lying on a manifold embedded in the ambient space, and the distance is intrinsic to the manifold geometry and invariant to the specific embedding. Since $\lambda_k^p \rightarrow 0$ when p is large (except for $\lambda_k = 1$), the origin point is the limiting point of the diffusion map embedding. Thus for x in a subcluster in \mathcal{C} , the weighted norm $S(x)$ with positive p can be viewed as a measurement of the extent of metastability (the depth of the well) of the potential well associated with the subcluster. In view of the diffusion distance, under the setting of this paper, nodes in \mathcal{B} are very similar to one another and, in comparison, nodes in \mathcal{C} are distinct from those in \mathcal{B} (as well as from other subclusters in \mathcal{C} , which is not reflected in S). A similar weighted form has also been studied in [6] for graph-based outlier detection. In the primary application considered in this paper, the leading eigenvalues are all close to 1, which means that the weighted form (3.18) is not very different unless p is large. On the other hand setting p to be large may suppress the high-index eigenvectors by small weights while they are actually the informative ones to indicate \mathcal{C} . Due to these reasons, we mainly consider $f = 1$ in the current paper, though the analysis directly extends.

Relation to heat kernel signature. With $f(\lambda) = e^{-(1-\lambda)t}$ from some positive t , $S(x)$ takes the form of the heat kernel signature (HKS) proposed by Sun, Ovsjanikov, and Guibas [37]. HKS is used in the shape analysis community as a concise multiscale feature descriptor on manifolds or three-dimensional (3D) meshes in tasks of shape matching, correspondence, and retrieval. Note that the HKS has been used locally to identify salient features on a given shape, similarly to the spectral embedding norm used for outlier detection in this paper. The application setting for HKS focuses on Riemannian manifolds and 3D mesh, which differs from here as we analyze the affinity matrix of a high-dimensional dataset composed of clusters and background. In particular, for reasons explained above, we mainly consider $f = 1$ in the definition of $S(x)$. We experimentally compare the embedding norm with HKS in Appendix C.

Indexing eigenvectors by support regions. The phenomenon studied here also suggests that sorting by the magnitude of eigenvalues may not be the most informative way to index the eigenvectors, a problem recently addressed in [31]. Here we study the special case where eigenvectors can be grouped by what they are mainly supported on. In the pseudodynamic (2.1), the eigenvectors begin with being exactly supported on either \mathcal{C} or \mathcal{B} at $t = 0$, and as time develops this pattern is nearly preserved as long as the \mathcal{C} - \mathcal{B} interblock connections are not too strong. The distinct support regions of eigenvectors appears to be irrelevant to the magnitude of the eigenvalues nor the existence of spectral gaps. This suggests that grouping eigenvectors by their localization regions may be a better way to arrange them in such cases. However, one still needs to be careful with the instability of eigenvectors: As shown in the numerical example, when two eigenvalues get close in the pseudodynamic, the associated pair of eigenvectors “swap” their values. (The swapping may be analyzed by the differential equation (3.11): assuming that among all the pairs of neighboring eigenvalues only one pair $(\lambda_k - \lambda_j)$ is approaching zero, then the dynamic of ψ_k evolution is dominated by that pair, which approximates a rotation among the indices j and k .) Our analysis in the current paper handles this by the summation in S over the index group I , which makes S invariant to such swaps as long as j and k both belong to I . Generally, since eigencrossings only happen at isolated times in the deformation dynamics [17], these special times can be excluded. Then one can say that the eigenvectors continue to almost localize on one of the two blocks most of the time.

Appendix A. Selection of eigenvectors in anomaly detection. Eigenvector selection can be used to better visualize and characterize an anomaly in the data, by finding a subspace

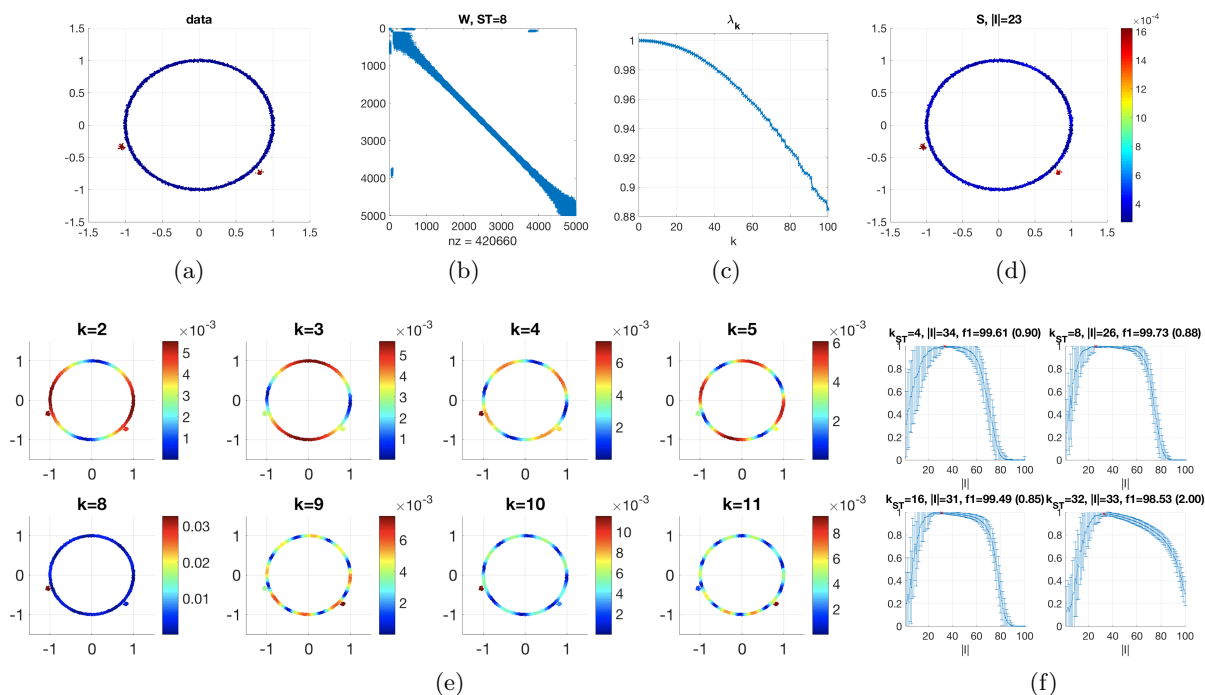


Figure A.1. Same plot as Figure 3. $K = 2$ clusters, $\delta = 0.02$.




Figure A.2. Selected eigenvectors of the graph Laplacian computed from the side-scan sonar image in Figure 5(a). (Left) Plot of $|\Psi_k(x_{\max})|$, where $x_{\max} = \arg \max_x S(x)$, and the x -axis is the index k . (Middle) Each pixel colored according to RGB colors assigned to the three embedding coordinates, (right) with maximum absolute value on x_{\max} .




Figure A.3. Same plot as Figure A.2 for the side-scan sonar image in Figure 5(c).

in which it is separated from the normal data. Let $x_{\max} = \arg \max_x S(x)$ be the pixel with the maximal embedding norm. In Figures A.2–A.3 we select the three eigenvectors with maximum absolute value on x_{\max} , plotting $|\Psi_k(x_{\max})|$ on the left. In the middle plot we color the pixels in the image according to pseudo-RGB values assigned to the three selected coordinates (right plot). Note that for Figure A.3, the selected coordinates are quite deep in the spectrum: $k = 59, 47, 48$.

Appendix B. Derivation of (3.10) and (3.11). Recall that $\{\lambda_k\}_{k=1}^n$ are the n real-valued eigenvalues of P , and ψ_k is the eigenvector associated with λ_k . By definition,

$$W\psi_k = \lambda_k D\psi_k, \quad \psi_k^T D\psi_l = \delta_{kl},$$

where W , D , λ_k , and ψ_k all depend on t (the dependence is omitted in the notation) and are differentiable over time. Taking the derivative w.r.t t of both sides,

$$(B.1) \quad \dot{W}\psi_k + W\dot{\psi}_k = \dot{\lambda}_k D\psi_k + \lambda_k \dot{D}\psi_k + \lambda_k D\dot{\psi}_k,$$

$$(B.2) \quad 0 = \dot{\psi}_k^T D\psi_l + \psi_k^T \dot{D}\psi_l + \psi_k^T D\dot{\psi}_l.$$

Applying ψ_k^T to both sides of (B.1) gives

$$\psi_k^T \dot{W}\psi_k = \dot{\lambda}_k \psi_k^T D\psi_k + \lambda_k \psi_k^T \dot{D}\psi_k = \dot{\lambda}_k + \lambda_k \psi_k^T \dot{D}\psi_k,$$

where in the first equality we use $W\psi_k = \lambda_k D\psi_k$. This proves (3.10).

To prove (3.11), since the $\{\psi_k\}_k$ form a D -orthonormal basis of \mathbb{R}^n , let

$$\dot{\psi}_k = \sum_{j=1}^n b_j \psi_j, \quad b_j = \dot{\psi}_k^T D \psi_j.$$

Applying ψ_j^T , $j \neq k$, to both sides of (B.1) gives

$$\psi_j^T \dot{W} \psi_k = (\lambda_k - \lambda_j) \psi_j^T D \dot{\psi}_k + \lambda_k \psi_j^T \dot{D} \psi_k,$$

thus

$$\beta_j = \frac{\psi_j^T (\dot{W} - \lambda_k \dot{D}) \psi_k}{\lambda_k - \lambda_j}, \quad j \neq k,$$

whenever $\lambda_k - \lambda_j \neq 0$. Letting $l = k$ in (B.2) gives

$$2\dot{\psi}_k^T D \psi_k + \psi_k^T \dot{D} \psi_k = 0,$$

thus

$$\beta_k = -\frac{1}{2} \psi_k^T \dot{D} \psi_k.$$

This proves (3.11).

Appendix C. Comparison to classical methods and HKS.

k -means and spectral clustering. Traditional k -means or spectral clustering fails to separate \mathcal{C} from \mathcal{B} in the synthetic 2D examples in the paper; cf., Figures 1, A.1, 3, and similarly on the real-world data examples.

The results on the data in Figure 1 are shown in Figure A.4. In this case, the difficulty of k -means to cluster the data is due to the unbalanced size of the clusters, both in terms of cluster diameter (the circle in \mathcal{B} can be large) and the number of nodes in each cluster (the parameter δ , defined to be ratio of number of nodes in \mathcal{C} and the total number of nodes, is small). Spectral clustering is problematic due to the unbalanced cluster and the manifold-like component \mathcal{B} , as has been pointed out in [23]. Specifically, for this example in section 2.3, none of the leading 10 eigenvectors will indicate the small cluster \mathcal{B} , thus using the leading eigenvectors will not detect the cluster \mathcal{B} .

HKS. For the scenario considered in the current paper, the embedding norm can outperform HKS on detecting small clusters. The results of HKS on the side-scan sonar image data in Figure 5 with multiple values of the time t are shown in Figure A.5: the best results are with $t = 8$ or 16, while the embedding norm removes the background in a much “cleaner” fashion in comparison; cf., $|I| = 13$ in Figure 5(D). The weights $f(\lambda_k)$ given to the k th eigenvector by HKS is plotted in Figure A.5(f).

In this example, the first eigenvector which gives a significant indication of the small cluster \mathcal{C} (the sea mine) starts at $k = 6$; cf., Figure 5(E), while the first 5 leading eigenvectors are mainly supported on the background. Due to the exponential decay of the weights in HKS, when t is small, the weights given to higher-indexed eigenvectors are large, which gives a poor result as shown in Figure A.5(a), apart from the computational issue of computing many more eigenvectors. Increasing t will suppress the tail, as shown in (f), while

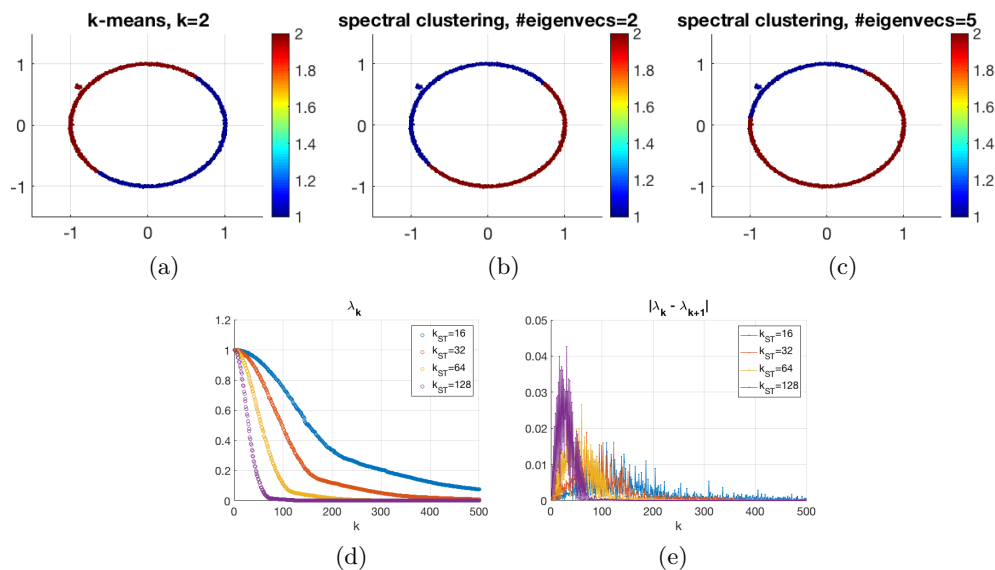


Figure A.4. (a) Results of k -means, $k = 2$, applied to the example in Figure 1, and (b), (c) spectral clustering using the first 2 and 5 eigenvectors. (d), (e) Eigenvalues and differences between consecutive eigenvalues of the normalized graph Laplacian for different settings of the graph construction parameter k_{ST} .




Figure A.5. (a)–(e): Plots of HKS on the side-scan sonar image data in Figure 5(C) with multiple values of the time t . (f): The weights $f(\lambda_k)$ given to k th eigenvector by HKS with different t and by the Embedding norm.

giving much smaller weights to $k = 6$ than the first 5 eigenvectors which are not indicating \mathcal{C} . This results in relatively large values on the background in the HKS, as shown in (b)–(c) and, when $t \geq 32$, mainly the first 3 eigenvectors remain to contribute to HKS and the \mathcal{C} disappears from the plot as in (d), (e). In comparison, the embedding norm gives the same weights to the first $|I|$ eigenvectors which help to enhance \mathcal{C} when $|I| \sim 10$, as shown in Figure 5(D). The negative effects of decaying weights is similar to the diffusion distance scheme $f(\lambda_k) = \lambda_k^t$.

Acknowledgments. We thank Ronald R. Coifman for useful discussions, and the anonymous reviewers for their constructive comments and useful suggestions.

REFERENCES

- [1] Y. AIZENBUD, A. BERMANIS, AND A. AVERBUCH, *PCA-Based Out-of-Sample Extension for Dimensionality Reduction*, preprint, <https://arxiv.org/abs/1511.00831> (2015).
- [2] L. AKOGLU, H. TONG, AND D. KOUTRA, *Graph based anomaly detection and description: A survey*, Data Min. Knowl. Discov., 29 (2015), pp. 626–688.
- [3] P. BERENS ET AL., *Standardizing and benchmarking data analysis for calcium imaging*, in COSYNE, 2017, pp. 66–67.
- [4] A. BORJI, M.-M. CHENG, H. JIANG, AND J. LI, *Salient object detection: A benchmark*, IEEE Trans. Image Process., 24 (2015), pp. 5706–5722.
- [5] V. CHANDOLA, A. BANERJEE, AND V. KUMAR, *Anomaly detection: A survey*, ACM Comput. Surv., 41 (2009), 15.
- [6] X. CHENG, G. MISHNE, AND S. STEINERBERGER, *The geometry of nodal sets and outlier detection*, J. Number Theory, 185 (2018), pp. 48–64, <https://doi.org/10.1016/j.jnt.2017.09.021>.
- [7] A. CLONINGER AND W. CZAJA, *Eigenvector localization on data-dependent graphs*, in 2015 International Conference on Sampling Theory and Applications (SampTA), IEEE, Piscataway, NJ, 2015, pp. 608–612.
- [8] R. R. COIFMAN AND S. LAFON, *Diffusion maps*, Appl. Comput. Harmon. Anal., 21 (2006), pp. 5–30.
- [9] M. CUCURINGU AND M. W. MAHONEY, *Localization on Low-Order Eigenvectors of Data Matrices*, preprint, <https://arxiv.org/abs/1109.1355> (2011).
- [10] A. DAMLE, V. MINDEN, AND L. YING, *Simple, direct, and efficient multi-way spectral clustering*, Inf. Inference, 8 (2019), pp. 181–203.
- [11] C. R. DE OLIVEIRA, *Eigenfunction entropy as a criterion of localization*, Phys. Lett. A, 296 (2002), pp. 165–170.
- [12] C. J. DSILVA, R. TALMON, R. R. COIFMAN, AND I. G. KEVREKIDIS, *Parsimonious representation of nonlinear dynamical systems through manifold learning: A chemotaxis case study*, Appl. Comput. Harmon. Anal., 44 (2018), pp. 759–773.
- [13] T. EHRET, A. DAVY, J.-M. MOREL, AND M. DELBRACIO, *Image anomalies: A review and synthesis of detection methods*, J. Math. Imaging Vision, 61 (2019), pp. 710–743.
- [14] M. GOLDSTEIN AND S. UCHIDA, *A comparative evaluation of unsupervised anomaly detection algorithms for multivariate data*, PloS One, 11 (2016), e0152173.
- [15] D. C. HOAGLIN AND R. E. WELSCH, *The hat matrix in regression and ANOVA*, Amer. Statist., 32 (1978), pp. 17–22.
- [16] T. IDÉ AND H. KASHIMA, *Eigenspace-based anomaly detection in computer systems*, in Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, 2004, pp. 440–449.
- [17] T. KATO, *Perturbation Theory for Linear Operators*, 2nd ed., Grundlehren Math. Wiss. 132, Springer, Berlin, 1976.
- [18] B. MILLER, N. BLISS, AND P. J. WOLFE, *Subgraph detection using eigenvector L1 norms*, in Advances in Neural Information Processing Systems, Curran Associates, Red Hook, NY, 2010, pp. 1633–1641.
- [19] G. MISHNE AND I. COHEN, *Multiscale anomaly detection using diffusion maps*, IEEE J. Sel. Topics Signal Process., 7 (2013), pp. 111–123.

- [20] G. MISHNE AND I. COHEN, *Iterative diffusion-based anomaly detection*, in International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, Piscataway, NJ, 2017, pp. 1682–1686.
- [21] G. MISHNE, R. R. COIFMAN, M. LAVZIN, AND J. SCHILLER, *Automated Cellular Structure Extraction in Biological Images with Applications to Calcium Imaging Data*, preprint, <https://doi.org/10.1101/313981> (2018).
- [22] G. MISHNE, U. SHAHAM, A. CLONINGER, AND I. COHEN, *Diffusion Nets*, Appl. Comput. Harmon. Anal., 47 (2019), pp. 259–285, <https://doi.org/10.1016/j.acha.2017.08.007>.
- [23] B. NADLER AND M. GALUN, *Fundamental limitations of spectral clustering*, in Advances in Neural Information Processing Systems (NIPS) 19, MIT Press, Cambridge, MA, 2007, pp. 1017–1024.
- [24] B. NADLER, S. LAFON, AND R. COIFMAN, *Diffusion Maps - A Probabilistic Interpretation for Spectral Embedding and Clustering Algorithms*, in Principal Manifolds for Data Visualization and Dimension Reduction, Springer, Berlin, 2007, pp. 242–264.
- [25] A. Y. NG, M. I. JORDAN, AND Y. WEISS, *On spectral clustering: Analysis and an algorithm*, in Advances in Neural Information Processing Systems (NIPS) 14, MIT Press, Cambridge, MA, pp. 849–896, 2002.
- [26] K. NIAN, H. ZHANG, A. TAYAL, T. COLEMAN, AND Y. LI, *Auto insurance fraud detection using unsupervised spectral ranking for anomaly*, J. Finance Data Sci., 2 (2016), pp. 58–75.
- [27] M. PACHITARIU, A. M. PACKER, N. PETTIT, H. DALGLEISH, M. HAUSER, AND M. SAHANI, *Extracting regions of interest from biological images with convolutional sparse block coding*, in Advances in Neural Information Processing Systems (NIPS), Curran Associates, Red Hook, NY, 2013, pp. 1745–1753.
- [28] B. PENG, L. ZHANG, AND D. ZHANG, *A survey of graph theoretical approaches to image segmentation*, Pattern Recogn., 46 (2013), pp. 1020–1038, <https://doi.org/10.1016/j.patcog.2012.09.015>.
- [29] N. RABIN AND A. AVERBUCH, *Detection of anomaly trends in dynamically evolving systems*, in AAAI Fall Symposium Series, AAAI, Palo Alto, CA, 2010, pp. 44–49.
- [30] S. REYNOLDS, T. ABRAHAMSSON, R. SCHUCK, P. JESPER SJÖSTRÖM, S. R. SCHULTZ, AND P. L. DRAGOTTI, *ABLE: An activity-based level set segmentation algorithm for two-photon calcium imaging data*, eNeuro, 4 (2017), 12, <https://doi.org/10.1523/ENEURO.0012-17.2017>.
- [31] N. SAITO, *How can we naturally order and organize graph Laplacian eigenvectors?*, <https://arxiv.org/abs/1801.06782>, 2018.
- [32] G. SCHIEBINGER, M. J. WAINWRIGHT, AND B. YU, *The geometry of kernelized spectral clustering*, Ann. Statist., 43 (2015), pp. 819–846.
- [33] J. SHI AND J. MALIK, *Normalized cuts and image segmentation*, IEEE Trans. Pattern Anal. Mach. Intell., 22 (2000), pp. 888–905.
- [34] A. SINGER, Y. SHKOLNISKY, AND B. NADLER, *Diffusion interpretation of nonlocal neighborhood filters for signal denoising*, SIAM J. Imaging Sci., 2 (2009), pp. 118–139.
- [35] S. L. SMITH AND M. HÄUSSER, *Parallel processing of visual space by neighboring neurons in mouse visual cortex*, Nature Neurosci., 13 (2010), pp. 1144–1149.
- [36] Q. SPAEN, D. S. HOCHBAUM, AND R. ASÍN-ACHÁ, *HNCcorr: A novel combinatorial approach for cell identification in calcium-imaging movies*, eNeuro, 6 (2019), 304.
- [37] J. SUN, M. OVSJANIKOV, AND L. GUIBAS, *A concise and provably informative multi-scale signature based on heat diffusion*, Computer Graphics Forum, 28 (2009), pp. 1383–1392.
- [38] L. TIAN, J. AKERBOOM, E. R. SCHREITER, AND L. L. LOOGER, *Neural activity imaging with genetically encoded calcium indicators*, Progr. Brain Res., 196 (2012), pp. 79–94.
- [39] U. VON LUXBURG, *A tutorial on spectral clustering*, Stat. Comput., 17 (2007), pp. 395–416.
- [40] U. VON LUXBURG, M. BELKIN, AND O. BOUSQUET, *Consistency of spectral clustering*, Ann. Statist., 36 (2008), pp. 555–586.
- [41] Y. WEISS, *Segmentation using eigenvectors: A unifying view*, in Proceedings of the ICCV 1999, Vol. 2, IEEE Computer Society, Los Alamitos, CA, 1999, pp. 975–982.
- [42] L. WU, X. WU, A. LU, AND Z.-H. ZHOU, *A spectral approach to detecting subtle anomalies in graphs*, J. Intell. Inform. Systems, 41 (2013), pp. 313–337.
- [43] T. XIANG AND S. GONG, *Spectral clustering with eigenvector selection*, Pattern Recogn., 41 (2008), pp. 1012–1029.
- [44] S. ZELDITCH, *Local and global analysis of eigenfunctions on Riemannian manifolds*, in Handbook of Geometric Analysis, No. 1, Adv. Lect. Math. (ALM) 7, International Press, Somerville, MA, 2008, pp. 545–658.

- [45] L. ZELNIK-MANOR AND P. PERONA, *Self-tuning spectral clustering*, in NIPS 17, MIT Press, Cambridge, MA, 2005, pp. 1601–1608.
- [46] F. ZHAO, L. JIAO, H. LIU, X. GAO, AND M. GONG, *Spectral clustering with eigenvector selection based on entropy ranking*, Neurocomputing, 73 (2010), pp. 1704–1717.
- [47] M. ZWORSKI, *Semiclassical Analysis*, Grad. Stud. Math. 138, American Mathematical Society, Providence, RI, 2012.