

A common variable minimax theorem for graphs

Ronald R. Coifman · Nicholas F. Marshall ·
Stefan Steinerberger

Received: date / Accepted: date

Abstract Let $\mathcal{G} = \{G_1 = (V, E_1), \dots, G_m = (V, E_m)\}$ be a collection of m graphs defined on a common set of vertices V but with different edge sets E_1, \dots, E_m . Informally, a function $f : V \rightarrow \mathbb{R}$ is smooth with respect to $G_k = (V, E_k)$ if $f(u) \sim f(v)$ whenever $(u, v) \in E_k$. We study the problem of understanding whether there exists a nonconstant function that is smooth with respect to all graphs in \mathcal{G} , simultaneously, and how to find it if it exists.

Keywords Spectral graph theory, graph Laplacian, common variable detection

Mathematics Subject Classification (2020) 49J35 · 05C50 · 35P05

1 Introduction

1.1 Introduction

Let $G = (V, E)$ be a graph; loosely speaking, a function $f : V \rightarrow \mathbb{R}$ is smooth with respect to G if it varies little over adjacent vertices meaning that $f(u) \sim f(v)$ whenever $(u, v) \in E$. Let \mathcal{G} be a collection of m graphs on the same set of vertices V

$$\mathcal{G} = \{G_1 = (V, E_1), \dots, G_m = (V, E_m)\}.$$

Communicated by Pencho Petrushev

R.R.C is supported by the NIH (5R01NS100049-04). N.F.M. is supported by the NSF (DMS-1903015). S.S. is supported by the NSF (DMS-2123224) and the Alfred P. Sloan Foundation.

R. R. Coifman
Department of Mathematics, Yale University, New Haven, CT
E-mail: coifman@math.yale.edu

Nicholas F. Marshall
Department of Mathematics, Princeton University, Princeton, NJ
E-mail: nicholas_marshall@princeton.edu

S. Steinerberger
Department of Mathematics, University of Washington, Seattle, WA
E-mail: steinerb@uw.edu

We consider the following problem: among all mean zero unit norm functions $f : V \rightarrow \mathbb{R}$ which is the smoothest with respect to \mathcal{G} (see §1.6 for a formal statement)?

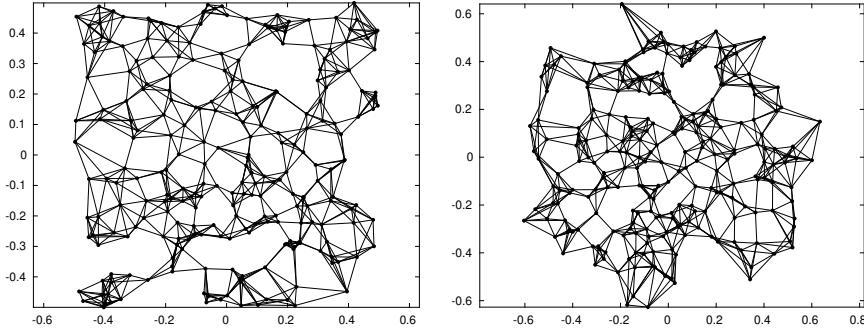


Fig. 1 A 6-nearest neighbor graph of points in the plane (left), and a 6-nearest neighbor graph for the same points after each point has been independently randomly rotated about the origin (right). As the number of points $n \rightarrow \infty$, commonly smooth functions $f : V \rightarrow \mathbb{R}$ are functions of the distance to the origin.

1.2 Motivating example

A geometric example is shown in Figure 1: we are given a set of n uniformly random points in the unit square centered at the origin, and form a graph $G_1 = (V, E_1)$ by connecting each point to its 6-nearest neighbors with respect to Euclidean distance. A second graph $G_2 = (V, E_2)$ is built on the same set of points as follows: each point is randomly rotated about the origin (by independent uniformly random rotations), and the rotated points are connected to their 6-nearest neighbors (see §3.3 for a precise description of this example). Two vertices u and v are close in the graph G_1 if the underlying points are physically close in the plane. Likewise, u and v are close in the graph G_2 if the rotated version of the underlying points are close. It becomes clear that any commonly smooth function $f : V \rightarrow \mathbb{R}$ must be close to a function that only depends on the distance of the underlying points to the origin (in the usual sense as the number of points n becomes large). How can we detect these ‘commonly smooth functions’ or ‘common variables’ if we do not have access to how the graphs were constructed? How can we find them from the graph data alone?

1.3 Problem statement

Suppose that $\mathcal{G} = (V, E_k)_{k=1}^m$ is a collection of m graphs on a common set of V vertices. We address two main problems.

- Is it possible to detect whether there is a nonconstant commonly ‘smooth’ function on the vertices V (that is smooth with respect to all m graphs)?
- Can we determine the ‘smoothest’ nonconstant function on V with respect to the collection of graphs \mathcal{G} ?

The precise nature of these questions will strongly depend on the notion of ‘smoothness’ of a function $f : V \rightarrow \mathbb{R}$. The main purpose of our paper is to define a notion of smoothness inspired by Spectral Graph Theory and to provide an approach which provably solves both problems in the regime where there truly is a common smooth variable shared by all graphs in a certain precise sense. What we observe in practice is that the method is more broadly applicable. We emphasize that the underlying question could be formalized in many different ways (possibly leading to very different mathematics) and many of them might be interesting.

1.4 Related results.

The problem of determining a commonly smooth function for a collection of graphs appears in different contexts, perhaps most frequently in data synthesis. Consider a data synthesis problem where a fixed set of data points is measured in different ways (a multi-view problem). Each measurement of the data points is encoded in a graph $G_k = (V, E_k)$ whose vertices V are the fixed data points and whose edges E_k are determined by the specific measurement. The end goal is to synthesize this data to extract intrinsic information. In particular, is there a common variable (function on V) that is related to how connections between the data points are formed across all of the graphs? This is a well-studied problem, we refer to [1] [2] [3] [4] [5] [8] [9] [10] [12] [14] and references therein. We especially emphasize three papers. Ma and Lee [11] propose working with a sum of Laplacians – this is similar to our approach except for the scaling which is crucial (see below for a longer discussion). Eynard, Kovnatsky, Bronstein, Glashoff, and Bronstein [7] also work within a spectral framework, and discuss the problem of simultaneous diagonalization of Laplacians which is philosophically related to our approach. Yair, Dietrich, Mulayoff, Talmon, and Kevrekidis [13] use the same perspective on smoothness as we will (indeed, their paper directly inspired ours) – they compute smooth functions on each graph and then look for vectors having large correlation with the subspaces of smooth functions.

1.5 Preliminaries

Suppose that G is an undirected connected weighted graph on n vertices described by an $n \times n$ symmetric nonnegative adjacency matrix \mathbf{A} . We use the notion of a graph Laplacian $\mathbf{L} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ throughout the paper; we assume that \mathbf{L} is symmetric, positive semi-definite, and has eigenvalue 0 (of multiplicity 1) corresponding to the eigenvector $\mathbf{1}$. An example of such an operator is the graph Laplacian

$$\mathbf{L} = \mathbf{D} - \mathbf{A}, \quad (1)$$

where \mathbf{D} is the diagonal matrix whose i -th diagonal element d_{ii} is the degree of the vertex $i \in V$. We use the notation

$$0 = \lambda_0(\mathbf{L}) < \lambda_1(\mathbf{L}) \leq \dots \leq \lambda_{n-1}(\mathbf{L}),$$

to denote the eigenvalues of \mathbf{L} , and

$$\frac{1}{\sqrt{n}}\mathbf{1} = \psi_0, \psi_1, \dots, \psi_{n-1},$$

to denote the corresponding eigenvectors which we assume are normalized (so that their ℓ^2 -norm is 1). When \mathbf{L} is given by (1), its associated quadratic form can be expressed by

$$\mathbf{x}^\top \mathbf{L} \mathbf{x} = \sum_{(u,v) \in E} a_{uv} (x_u - x_v)^2, \quad (2)$$

where a_{uv} is the weight associated with the edge (u, v) . In spectral graph theory, this quadratic form is a standard way to measure the smoothness of a function on a graph. In order to use this quadratic form as a smoothness score, we restrict our attention to the set X of vectors with mean zero and unit length

$$X = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{1}^\top \mathbf{x} = 0 \text{ and } \mathbf{x}^\top \mathbf{x} = 1\}.$$

Restricting our attention to X is important since it avoids trivially smooth functions on the vertices of a graph such as functions with a large constant component or functions with a very small norm. We define a smoothness score $s_{\mathbf{L}} : X \rightarrow [1, \infty)$ by

$$s_{\mathbf{L}}(\mathbf{x}) = \frac{1}{\lambda_1(\mathbf{L})} \mathbf{x}^\top \mathbf{L} \mathbf{x}. \quad (3)$$

This normalization ensures that $s_{\mathbf{L}}(\mathbf{x}) \geq 1$ with equality if and only if \mathbf{x} is an eigenvector of \mathbf{L} of eigenvalue $\lambda_1(\mathbf{L})$. The reason that normalizing $s_{\mathbf{L}}$ is important, is that we are going to compare smoothness scores of a given function across different graphs. Dividing by $\lambda_1(\mathbf{L})$ is just one reasonable method of normalization; for some applications it may be advantageous to normalize $s_{\mathbf{L}}$ differently, see Remark 2. The presented results hold for these alternate normalization strategies, as well as more general definitions of $s_{\mathbf{L}}$ whose discussion is delayed until later in the paper to simplify the exposition, see Remark 4

1.6 Main results

Suppose that $\mathcal{G} = \{G_1, \dots, G_m\}$ is a collection of undirected connected weighted graphs on a common set of n vertices V . Informally speaking, a common variable for \mathcal{G} is a function defined on the vertices V which is smooth with respect to the geometry of each graph. More precisely, we can define a score $s_{\mathcal{G}}$ indicating how smooth (in the minimax sense) the smoothest function with respect to \mathcal{G} is by

$$s_{\mathcal{G}} = \min_{\mathbf{x} \in X} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}), \quad (4)$$

where $s_{\mathbf{L}}$ is defined by (3). The score $s_{\mathcal{G}}$ can be used to understand how much ‘common information’ is shared by the collection of graphs \mathcal{G} . If $\psi \in X$ is an argument that minimizes (4), then we call ψ the smoothest function with respect to \mathcal{G} or a common variable of \mathcal{G} .

We can now present our main results. Theorem 1 provides upper and lower bounds on $s_{\mathcal{G}}$ in terms of (explicitly computable) spectral quantities: we show that the smallest eigenvalue of suitably averaged Laplacians serves as a lower bound. Theorem 1 is complemented by Theorem 2 which shows that the lower and upper bounds are equal under an additional assumption, and for a suitable choice of parameters. Numerical examples will show that they indeed coincide in practice.

Theorem 1 (Upper and Lower Bounds) *Let \mathcal{G} be a collection of graphs satisfying the assumptions in §1.5. For any $\mathbf{t} = (t_1, \dots, t_m) \in T$, where*

$$T := \{\mathbf{t} \in [0, 1]^m : t_1 + \dots + t_m = 1\}, \quad (5)$$

define the Laplacian \mathbf{L}_t by the linear combination

$$\mathbf{L}_t = \sum_{k=1}^m t_k \frac{\mathbf{L}_k}{\lambda_1(\mathbf{L}_k)}. \quad (6)$$

Then,

$$\lambda_1(\mathbf{L}_t) \leq s_{\mathcal{G}} \leq \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_t)),$$

where $\lambda_1(\mathbf{L}_t)$ denotes the smallest positive eigenvalue of \mathbf{L}_t , and $\psi_1(\mathbf{L}_t)$ denotes a unit length eigenvector associated with $\lambda_1(\mathbf{L}_t)$.

Theorem 1 provides us with spectral upper and lower bounds on $s_{\mathcal{G}}$. Theorem 2 shows, assuming the first nontrivial eigenvalue is simple, that there is an explicit duality relation which allows us to find the common variable by solving an eigenvalue optimization problem.

Theorem 2 (Common Information Minimax Theorem) *In addition to the hypothesis of Theorem 1, assume that $\lambda_1(\mathbf{L}_t)$ is always simple:*

$$\lambda_2(\mathbf{L}_t) > \lambda_1(\mathbf{L}_t), \quad \text{for all } \mathbf{t} \in T,$$

where T is defined by (5). If

$$\mathbf{t}^* = \underset{\mathbf{t} \in T}{\operatorname{argmax}} \lambda_1(\mathbf{L}_t), \quad (7)$$

then

$$\psi_1(\mathbf{L}_{\mathbf{t}^*}) = \underset{\mathbf{x} \in X}{\operatorname{argmin}} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}).$$

The proofs of Theorems 1 and 2 are given in §2. In practice, these theorems can be used in conjunction, see Remark 1. The optimization problem (7) is straightforward to solve numerically using standard methods, see §3.

1.7 Diffusion geometry interpretation

In the following, we describe how Theorems 1 and 2 can be interpreted in terms of diffusion geometry methods. Given a symmetric positive semi-definite matrix \mathbf{L} which has eigenvalue 0 (of multiplicity 1) associated with the eigenvector $\mathbf{1}$, we can define a diffusion (or averaging) operator \mathbf{H}^τ by

$$\mathbf{H}^\tau = \exp(-\tau \mathbf{L}),$$

where $\exp(\mathbf{A}) = \mathbf{I} + \mathbf{A} + \frac{1}{2!} \mathbf{A}^2 + \dots$ is the matrix exponential and $\tau > 0$ plays the role of time. If \mathbf{L} has eigenvalues

$$0 = \lambda_0(\mathbf{L}) < \lambda_1(\mathbf{L}) \leq \dots \leq \lambda_{n-1}(\mathbf{L}),$$

then by the spectral mapping theorem \mathbf{H}^τ has eigenvalues

$$1 = e^{-\tau \lambda_0(\mathbf{L})} > e^{-\tau \lambda_1(\mathbf{L})} \geq \dots \geq e^{-\tau \lambda_{n-1}(\mathbf{L})}. \quad (8)$$

The following corollary is immediate from Theorem 2 and (8).

Corollary 1 (Diffusion interpretation) *Under the hypothesis of Theorem 2, define the diffusion operator $\mathbf{H}_{\mathbf{t}}^\tau$ by*

$$\mathbf{H}_{\mathbf{t}}^\tau = \exp \left(-\tau \sum_{k=1}^m t_k \frac{\mathbf{L}_k}{\lambda_1(\mathbf{L}_k)} \right).$$

Let

$$\mathbf{t}^* = (t_1^*, \dots, t_m^*) = \operatorname{argmin}_{\mathbf{t} \in T} \mu_1(\mathbf{H}_{\mathbf{t}}^\tau),$$

where $\mu_1(\mathbf{H}_{\mathbf{t}}^\tau)$ is the largest eigenvalue which is less than 1 of $\mathbf{H}_{\mathbf{t}}^\tau$, and T is defined in (5). Then,

$$\varphi_1(\mathbf{H}_{\mathbf{t}^*}^\tau) = \operatorname{argmin}_{\mathbf{x} \in X} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}),$$

where $\varphi_1(\mathbf{H}_{\mathbf{t}^*}^\tau)$ denotes a unit length eigenvector associated with $\mu_1(\mathbf{H}_{\mathbf{t}^*}^\tau)$.

This corollary, which rewrites Theorem 2 in terms of a diffusion operator, has several interesting consequences.

- 1) The parameters $\mathbf{t}^* = (t_1^*, \dots, t_m^*)$ can be interpreted as optimally tuned diffusion times for the graphs G_1, \dots, G_m . The operator $\mathbf{H}_{\mathbf{t}^*}^\tau$ uncovers common information from the graphs by optimally diffusing on these graphs at different rates.
- 2) The operator $\mathbf{H}_{\mathbf{t}^*}^\tau$ can be used to define a diffusion distance on the common set of vertices V on which the graphs \mathcal{G} are defined. Assume that $V = \{1, \dots, n\}$. We can define the diffusion distance $D_{\mathbf{t}^*}^\tau : V \times V \rightarrow \mathbb{R}$ by

$$D_{\mathbf{t}^*}^\tau(i, j) = \|\mathbf{H}_{\mathbf{t}^*}^\tau \delta_i - \mathbf{H}_{\mathbf{t}^*}^\tau \delta_j\|_{\ell^2},$$

where δ_i is the column vector whose i -th entry is 1 and other entries are 0.

- 3) For any chosen dimension $d \geq 1$, the operator $\mathbf{H}_{\mathbf{t}^*}^\tau$ can be used to define a diffusion map $\Psi^\tau : V \rightarrow \mathbb{R}^d$ by

$$\Psi^\tau(j) = \begin{pmatrix} \mu_1(\mathbf{H}_{\mathbf{t}^*}^\tau) \varphi_{1,j}(\mathbf{H}_{\mathbf{t}^*}^\tau) \\ \vdots \\ \mu_d(\mathbf{H}_{\mathbf{t}^*}^\tau) \varphi_{d,j}(\mathbf{H}_{\mathbf{t}^*}^\tau) \end{pmatrix},$$

where $\varphi_{i,j}(\mathbf{H}_{\mathbf{t}^*}^\tau)$ denotes the j -th entry of the eigenvector $\varphi_i(\mathbf{H}_{\mathbf{t}^*}^\tau)$ associated with the eigenvalue $\mu_i(\mathbf{H}_{\mathbf{t}^*}^\tau) = e^{-\tau \lambda_i(\mathbf{L}_{\mathbf{t}^*})}$.

Remark 1 (Using Theorems 1 and 2 in conjunction) In applications, the optimization problem (7) for \mathbf{t}^* in Theorem 2 can be solved numerically using gradient based optimization methods and Lemma 1. Suppose that $\tilde{\mathbf{t}}$ is the numerical solution to (7), and let $\tilde{\psi}_1$ denote a normalized eigenvector corresponding to $\lambda_1(\mathbf{L}_{\tilde{\mathbf{t}}})$. By Theorem 1 we have the following error estimate:

$$\left| \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\tilde{\psi}_1) - s_{\mathcal{G}} \right| \leq \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\tilde{\psi}_1) - \lambda_1(\mathbf{L}_{\tilde{\mathbf{t}}}). \quad (9)$$

This inequality can be used to verify that the numerical optimization is successful. For each of our numerical examples presented in §3 we use (9) to verify that we are able to accurately solve each optimization problem (with error $\lesssim 10^{-6}$); however, in

practice (9) could also be used to stop the optimization process when the error is less than, say, 10^{-1} , if that is sufficient for the application. Furthermore, we note that since Theorem 1 does not require the eigenvalue multiplicity condition of Theorem 2 and Lemma 1, this error estimate can be used to check that a candidate common variable ψ_1 determined using Theorem 2 is close to optimal without having to verify that the spectral gap hypothesis of Theorem 2 holds.

Remark 2 (Alternate normalization methods) Recall that we defined the smoothness score $s_{\mathbf{L}} : X \rightarrow [1, \infty)$ by

$$s_{\mathbf{L}}(\mathbf{x}) = \frac{1}{\lambda_1(\mathbf{L})} \mathbf{x}^\top \mathbf{L} \mathbf{x},$$

such that the ‘smoothest function’ with respect to $s_{\mathbf{L}}$ has smoothness score 1. In practice it may be advantageous to normalize the quadratic form differently. For example, we could define $a_{\mathbf{L}} : X \rightarrow (0, \infty)$ by

$$a_{\mathbf{L}}(\mathbf{x}) = \left(\frac{1}{n-1} \sum_{j=1}^{n-1} \lambda_j(\mathbf{L}) \right)^{-1} \mathbf{x}^\top \mathbf{L} \mathbf{x}, \quad (10)$$

such that the ‘average smoothness’ with respect to $a_{\mathbf{L}}$ is 1. More precisely, with this definition we have $\mathbb{E}[a_{\mathbf{L}}(\mathbf{x})] = 1$, where the expectation is taken over \mathbf{x} chosen uniformly at random from X . Indeed, by writing $\mathbf{x} = \sum_{j=1}^{n-1} c_j \psi_j$, where (c_1, \dots, c_{n-1}) are chosen uniformly at random with respect to the surface measure on \mathbb{S}^{n-2} we have

$$\mathbb{E}[\mathbf{x}^\top \mathbf{L} \mathbf{x}] = \mathbb{E} \left[\sum_{j=1}^{n-1} c_j^2 \lambda_j(\mathbf{L}) \right] = \sum_{j=1}^{n-1} \frac{\lambda_j(\mathbf{L})}{n-1}.$$

Other methods of normalization are conceivable: one could, for example, consider decreasing weights that put more emphasis on lower frequencies. Our theoretical results are independent of the choice of normalization. However, for applications the distinction between choosing to normalize based on the ‘smoothest function’ or ‘average smoothness’ (or some other intermediate normalization method) may be important; we provide such an example in §3.6.

2 Proof of main results

We start by proving Theorem 1. After that, we establish Lemma 1 and use it to establish Theorem 2.

2.1 Proof of Theorem 1

Proof Recall that

$$X = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{1}^\top \mathbf{x} = 0 \text{ and } \mathbf{x}^\top \mathbf{x} = 1\}.$$

We have

$$s_{\mathcal{G}} = \min_{\mathbf{x} \in X} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}) = \min_{\mathbf{x} \in X} \max_{\mathbf{t} \in T} \sum_{k=1}^m t_k s_{\mathbf{L}_k}(\mathbf{x}),$$

where

$$T = \{\mathbf{t} = (t_1, \dots, t_m) \in [0, 1]^m : t_1 + \dots + t_m = 1\}.$$

For any fixed $\mathbf{t} \in T$ (not depending on \mathbf{x}) we have

$$\min_{\mathbf{x} \in X} \max_{\mathbf{t} \in T} \sum_{k=1}^m t_k s_{\mathbf{L}_k}(\mathbf{x}) \geq \min_{\mathbf{x} \in X} \sum_{k=1}^m t_k s_{\mathbf{L}_k}(\mathbf{x}) = \min_{\mathbf{x} \in X} \mathbf{x}^\top \mathbf{L}_t \mathbf{x},$$

where

$$\mathbf{L}_t = \sum_{k=1}^m \frac{t_k}{\lambda_1(\mathbf{L}_k)} \mathbf{L}_k.$$

By the Courant-Fischer Theorem

$$\min_{\mathbf{x} \in X} \mathbf{x}^\top \mathbf{L}_t \mathbf{x} = \lambda_1(\mathbf{L}_t).$$

In combination, the above inequalities give

$$s_{\mathcal{G}} \geq \lambda_1(\mathbf{L}_t).$$

Let $\psi_1(\mathbf{L}_t)$ be a unit length eigenvector corresponding to $\lambda_1(\mathbf{L}_t)$. By using this eigenvector as a test vector we have

$$s_{\mathcal{G}} = \min_{\mathbf{x} \in X} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}) \leq \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_t)).$$

This completes the proof.

Lemma 1 (Gradient of eigenvalue) *Suppose that the assumptions of Theorem 2 hold. Suppose that $\mathbf{t}' = (t_1, \dots, t_{m-1})$ and assume $t_m := 1 - (t_1 + \dots + t_{m-1})$. We have*

$$\nabla_{\mathbf{t}'} \lambda_1(\mathbf{L}_t) = \begin{pmatrix} \psi_1(\mathbf{L}_t)^\top \left(\frac{1}{\lambda_1(\mathbf{L}_1)} \mathbf{L}_1 - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_t) \\ \vdots \\ \psi_1(\mathbf{L}_t)^\top \left(\frac{1}{\lambda_1(\mathbf{L}_{m-1})} \mathbf{L}_{m-1} - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_t) \end{pmatrix}.$$

Moreover, $\lambda_1(\mathbf{L}_t)$ is equal to $s_{\mathcal{G}}$ whenever the gradient vanishes.

Proof Suppose that $\mathbf{t}' = (t_1, \dots, t_{m-1})$ and assume $t_m := 1 - (t_1 + \dots + t_{m-1})$. Under this assumption, we use the notation $\mathbf{L}_{\mathbf{t}'} = \mathbf{L}_t$ interchangeably. We will prove that

$$\frac{\partial}{\partial t_j} \lambda_1(\mathbf{L}_{\mathbf{t}'}) = \psi_1(\mathbf{L}_{\mathbf{t}'})^\top \left(\frac{1}{\lambda_1(\mathbf{L}_j)} \mathbf{L}_j - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_{\mathbf{t}'}),$$

for $j \in \{1, \dots, m-1\}$. Let $\mathbf{e}_j \in \mathbb{R}^{m-1}$ be the j -th standard basis vector (whose j -th entry is 1 and other entries are 0). For $\varepsilon > 0$ we have

$$\lambda_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) = \psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j})^\top (\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) \psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}).$$

By the definition of $\mathbf{L}_{\mathbf{t}'}$ we have

$$\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j} = \mathbf{L}_{\mathbf{t}'} + \varepsilon \left(\frac{1}{\lambda_1(\mathbf{L}_j)} \mathbf{L}_j - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right). \quad (11)$$

We will now argue that it is possible to express the first normalized eigenvector of the perturbed matrix $\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}$ as a small perturbation of the first normalized eigenvector of the unperturbed matrix $\mathbf{L}_{\mathbf{t}'}$. Using that all our eigenvectors are defined to be normalized, we can write

$$\psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) = c\psi_1(\mathbf{L}_{\mathbf{t}'}) + \delta,$$

where c is a coefficient and δ is orthogonal to $\psi_1(\mathbf{L}_{\mathbf{t}'})$. It remains to show that c is close to 1 or, equivalently, that δ is small. For this, we use the Davis-Kahan theorem. If θ denotes the angle between $\psi_1(\mathbf{L}_{\mathbf{t}'})$ and $\psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j})$ then by the Davis-Kahan theorem

$$\sin(\theta) \leq \frac{2\|\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j} - \mathbf{L}_{\mathbf{t}'}\|}{\min_{i \neq 1} |\lambda_1(\mathbf{L}_{\mathbf{t}'}) - \lambda_i(\mathbf{L}_{\mathbf{t}'})|}. \quad (12)$$

The denominator is uniformly bounded away from 0 as part of the assumptions of Theorem 2, which are assumed to hold in the statement of the lemma. Therefore, combining (11) and (12) yields

$$\sin(\theta) \leq \mathcal{O}(\|\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j} - \mathbf{L}_{\mathbf{t}'}\|) = \mathcal{O}(\varepsilon).$$

We can now compute the cosine of θ via an inner product and obtain

$$\cos(\theta) = \langle \psi_1(\mathbf{L}_{\mathbf{t}'}), \psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) \rangle = c.$$

Using

$$1 = \|\psi_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j})\|^2 = c^2 + \|\delta\|^2,$$

we arrive at

$$\|\delta\|^2 = 1 - c^2 = 1 - \cos(\theta)^2 = \sin(\theta)^2 = \mathcal{O}(\varepsilon^2)$$

from which we deduce $\|\delta\| = \mathcal{O}(\varepsilon)$. Using these identities, we can perform an expansion of $\lambda_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j})$ up to order ε . We start by writing

$$\lambda_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) = (c\psi_1(\mathbf{L}_{\mathbf{t}'}) + \delta)^\top \left(\mathbf{L}_{\mathbf{t}'} + \varepsilon \left(\frac{1}{\lambda_1(\mathbf{L}_j)} \mathbf{L}_j - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \right) (c\psi_1(\mathbf{L}_{\mathbf{t}'}) + \delta).$$

Recalling that $c^2 = 1 - \mathcal{O}(\varepsilon^2)$, expanding the right hand side gives

$$\begin{aligned} \lambda_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) &= c^2 \psi_1(\mathbf{L}_{\mathbf{t}'})^\top \mathbf{L}_{\mathbf{t}'} \psi_1(\mathbf{L}_{\mathbf{t}'}) + 2c\delta \mathbf{L}_{\mathbf{t}'} \psi_1(\mathbf{L}_{\mathbf{t}'}) \\ &\quad + \varepsilon \psi_1(\mathbf{L}_{\mathbf{t}'})^\top \left(\frac{1}{\lambda_1(\mathbf{L}_j)} \mathbf{L}_j - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_{\mathbf{t}'}) + \mathcal{O}(\varepsilon^2). \end{aligned}$$

The first term on the right hand side is equal to $\lambda_1(\mathbf{L}'_{\mathbf{t}}) + \mathcal{O}(\varepsilon^2)$, and the second term is equal to zero since δ is orthogonal to the eigenvector $\psi_1(\mathbf{L}'_{\mathbf{t}})$. It follows that

$$\frac{\lambda_1(\mathbf{L}_{\mathbf{t}' + \varepsilon \mathbf{e}_j}) - \lambda_1(\mathbf{L}'_{\mathbf{t}})}{\varepsilon} = \psi_1(\mathbf{L}_{\mathbf{t}'})^\top \left(\frac{1}{\lambda_1(\mathbf{L}_j)} \mathbf{L}_j - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_{\mathbf{t}'}) + \mathcal{O}(\varepsilon).$$

This argument works for all $j \in \{1, \dots, m-1\}$ so the proof is complete.

2.2 Proof of Theorem 2

Proof Suppose that

$$\mathbf{t}^* = \operatorname{argmax}_{\mathbf{t} \in T} \lambda_1(\mathbf{L}_t).$$

We use the notation $\mathbf{t}' = (t_1, \dots, t_{m-1})$ where $t_m := 1 - (t_1 + \dots + t_{m-1})$. First, consider the case where \mathbf{t}^* is contained in the interior of T . In this case,

$$\nabla_{\mathbf{t}'} \lambda_1(\mathbf{L}_t)|_{\mathbf{t}=\mathbf{t}^*} = 0.$$

Thus, by Lemma 1 we have

$$\psi_1(\mathbf{L}_{\mathbf{t}^*})^\top \left(\frac{1}{\lambda_1(\mathbf{L}_k)} \mathbf{L}_k - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_{\mathbf{t}^*}) = 0,$$

for $k \in \{1, \dots, m-1\}$, since this equation can be equivalently written as

$$s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) - s_{\mathbf{L}_m}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = 0,$$

it follows that all of these quadratic forms are equal:

$$s_{\mathbf{L}_1}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = \dots = s_{\mathbf{L}_m}(\psi_1(\mathbf{L}_{\mathbf{t}^*})).$$

Informally speaking, the smoothest function or common variable $\psi_1(\mathbf{L}_{\mathbf{t}^*})$ is indifferent between the different smoothness measures $s_{\mathbf{L}_1}, \dots, s_{\mathbf{L}_m}$. Thus,

$$\lambda_1(\mathbf{L}_{\mathbf{t}^*}) = \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})).$$

We recall Theorem 1 states that

$$\lambda_1(\mathbf{L}_t) \leq s_{\mathcal{G}} \leq \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_t)),$$

from which we can conclude that

$$\max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = s_{\mathcal{G}}.$$

It remains to consider the case where \mathbf{t}^* is not contained in the interior of T . Without loss of generality, suppose that $t_1^* = \dots = t_p^* = 0$ and $t_{p+1}^*, \dots, t_m^* \neq 0$. Suppose first that $p+1 \neq m$. Then there are at least 2 positive entries and, in particular, $0 < t_m < 1$. We can thus apply Lemma 1 and conclude that for $k \in \{p+1, \dots, m\}$

$$\psi_1(\mathbf{L}_{\mathbf{t}^*})^\top \left(\frac{1}{\lambda_1(\mathbf{L}_k)} \mathbf{L}_k - \frac{1}{\lambda_1(\mathbf{L}_m)} \mathbf{L}_m \right) \psi_1(\mathbf{L}_{\mathbf{t}^*}) = 0$$

which, as above, is equivalent to,

$$s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = s_{\mathbf{L}_m}(\psi_1(\mathbf{L}_{\mathbf{t}^*})), \quad \text{for } k \in \{p+1, \dots, m\}. \quad (13)$$

If $p+1 = m$, then (13) holds trivially. It remains to deal with the entries t_1^*, \dots, t_p^* (which are all 0). Fix $k \in \{1, \dots, p\}$. Since \mathbf{t}^* is maximal, the derivative of $\lambda_1(\mathbf{L}_{\mathbf{t}'})$ at $\mathbf{t}' = (t_1^*, \dots, t_{m-1}^*)$ in the direction \mathbf{e}_k must be negative and thus by Lemma 1

$$s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) \leq s_{\mathbf{L}_m}(\psi_1(\mathbf{L}_{\mathbf{t}^*})).$$

It follows that

$$\max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = \max_{k \in \{p+1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) = \lambda_1(\mathbf{L}_{\mathbf{t}^*}).$$

Appealing to Theorem 1 once more gives

$$\lambda_1(\mathbf{L}_{\mathbf{t}^*}) \leq s_{\mathcal{G}} \leq \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})),$$

we can conclude that

$$\lambda_1(\mathbf{L}_{\mathbf{t}^*}) = s_{\mathcal{G}} = \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})).$$

This completes the proof.

3 Numerical examples

3.1 The Laplacian

Our approach is completely general with respect to the underlying notion of Laplacian \mathbf{L} and many different types of Laplacians could be used. We merely require that \mathbf{L} is symmetric positive semi-definite, and that \mathbf{L} has eigenvalue 0 of multiplicity 1 (corresponding to constant functions). For the purpose of consistency, all examples will be computed using the bi-stochastic Laplacian which is defined as follows. Assume that \mathbf{A} is a symmetric non-negative weighted adjacency matrix with a positive main diagonal. By using the Sinkhorn-Kopp algorithm (see Lemma 2) it is possible to determine a symmetric positive definite diagonal matrix \mathbf{D} such that

$$\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2} \mathbf{1} = \mathbf{1},$$

where $\mathbf{1}$ denotes a column vector of ones. Given such a matrix \mathbf{D} we define the bi-stochastic graph Laplacian \mathbf{L} by

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2},$$

where \mathbf{I} is the identity matrix. The bi-stochastic graph Laplacian can be viewed as the graph Laplacian of a graph whose weighted adjacency matrix is $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$, and thus the bi-stochastic graph Laplacian has the same properties as the graph Laplacian discussed in §1.5. We refer to §4.1 for more details on how to compute the bi-stochastic Laplacian.

3.2 Nearest neighbor graph definition

Let $X = \{x_1, \dots, x_n\}$ be a subset of \mathbb{R}^d . We say that $N_k(x_j)$ is a set of k -nearest neighbors of x_j in X if $N_k(x_j)$ is a subset of $X \setminus \{x_j\}$ consisting of k points which has the property

$$\max_{x \in N_k(x_j)} \|x - x_j\| \leq \min_{y \in X \setminus (N_k(x_j) \cup \{x_j\})} \|y - x_j\|.$$

We say that G is a k -nearest neighbor graph for X if its adjacency matrix $\mathbf{A} = (a_{ij})$ satisfies

$$a_{ij} = \begin{cases} 1 & \text{if } i = j, \\ 1 & \text{if } x_i \in N_k(x_j) \text{ or } x_j \in N_k(x_i), \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

for $i, j = 1, \dots, n$, and for some choice of k -nearest neighbors $N_k(x_1), \dots, N_k(x_n)$. Note that our definition of a k -nearest neighbor graph includes self loops for each vertex. This assumption allows us to perform a bi-stochastic normalization of the adjacency matrix. We note that assuming that a graph has self loops is a common assumption when working with stochastic matrices on graphs since it ensures these stochastic matrices are aperiodic.

3.3 Independent rotations in two dimensions

Let $X_1 = \{x_1, \dots, x_n\}$ be a set of $n = 250$ independent uniformly random points from the unit square $[-1/2, 1/2]^2$, and $\theta_1, \dots, \theta_n$ be independent uniformly random points from $[0, 2\pi)$. Set

$$X_2 = \{T_{\theta_1}(x_1), \dots, T_{\theta_n}(x_n)\},$$

where $T_\theta(x)$ denotes the rotation of x by angle θ about the origin. More precisely, if $x = (r \cos \phi, r \sin \phi)$, then $T_\theta(x) = (r \cos(\phi + \theta), r \sin(\phi + \theta))$. To summarize, the set X_2 is created by rotating the points in X_1 about the origin with independent uniformly random rotations. Let G_1 and G_2 be 6-nearest neighbor graphs of X_1 and X_2 , respectively, see Figure 1. For each graph G_1 and G_2 we construct the corresponding bi-stochastic graph Laplacians \mathbf{L}_1 and \mathbf{L}_2 . Next we solve the optimization problem

$$\mathbf{t}^* = \underset{\mathbf{t} \in T}{\operatorname{argmax}} \lambda_1(\mathbf{L}_t),$$

where \mathbf{L}_t is defined in (6). By setting $t_2 := 1 - t_1$ we can optimize over $t_1 \in [0, 1]$. To visualize this optimization problem, we plot $\lambda_1(\mathbf{L}_t)$ against t_1 in Figure 2.

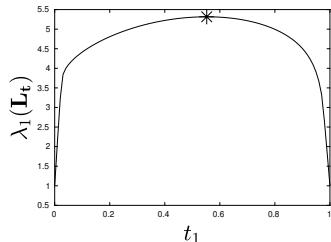


Fig. 2 Parameter t_1 versus $\lambda_1(\mathbf{L}_t)$; the max occurs at the star.

Using numerical optimization, we find that

$$\mathbf{t}^* \approx (0.552330195903778, 0.447669804096222).$$

Next, we use Theorem 1 to validate the results of the optimization, which gives

$$\left| \max_{k \in \{1, 2\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{t^*})) - \lambda_1(\mathbf{L}_{t^*}) \right| \leq 8.802488427051003 \times 10^{-8},$$

indicating that the optimization procedure was successful. Since we know how the graphs G_1 and G_2 were generated, we can further validate the method by checking that $\psi_1(\mathbf{L}_{t^*})$ is a smooth function of the common variable that influences edge creation in both graphs (the distance of a point from the origin). We plot $\psi_1(\mathbf{L}_{t^*})$ versus r (representing the distance of a point from the origin) in Figure 3.3.

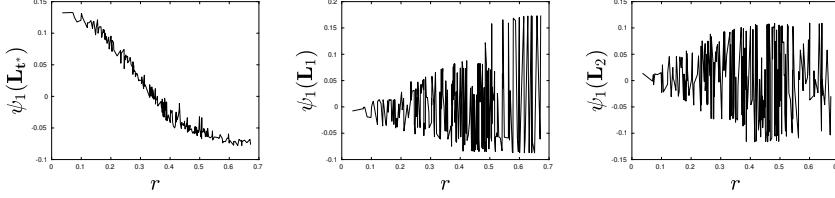


Fig. 3 Plots of $\psi(\mathbf{L}_{t^*})$, $\psi(\mathbf{L}_1)$, and $\psi(\mathbf{L}_2)$ against r .

Observe that in Figure 3.3 the common variable is essentially a re-scaling of the distance to the origin (as would be expected). For comparison, Figure 3.3 also includes plots of $\psi_1(\mathbf{L}_1)$ and $\psi_1(\mathbf{L}_2)$ versus r to demonstrate that neither of them are smooth with respect to the common variable.

In the following section, we will present a similar example of building graphs from randomly rotated points except we start with points in three dimensions, and perform rotations around different axes to demonstrate how the method works when there are three graphs G_1 , G_2 , and G_3 .

3.4 Independent rotations in three dimensions

Let $X_1 = \{x_1, \dots, x_n\}$ be $n = 500$ independent uniformly random points from the unit ball $\{x \in \mathbb{R}^3 : \|x\|_{\ell^2} \leq 1\}$, and let $\theta_1, \dots, \theta_n$ and ϕ_1, \dots, ϕ_n be independent uniformly random angles from $[0, 2\pi)$. Set

$$X_2 = \{T_{\theta_1}(x_1), \dots, T_{\theta_n}(x_n)\},$$

where $T_{\theta_j}(x)$ is a rotation about the z -axis by angle θ : if $x = (r \cos \theta, r \sin \theta, z)$, then

$$T_{\theta_j}(x) = (r \cos(\theta + \theta_j), r \sin(\theta + \theta_j), z),$$

and set

$$X_3 = \{S_{\phi_1}(x_1), \dots, S_{\phi_n}(x_n)\},$$

where $S_{\phi_j}(x)$ is a rotation about the y -axis: if $x = (r \cos \phi, y, r \sin \phi)$, then

$$S_{\phi}(x) = (r \cos(\phi + \phi_j), y, r \sin(\phi + \phi_j)).$$

We construct 6-nearest neighbor graphs G_1 , G_2 and G_3 from the sets X_1 , X_2 , and X_3 , respectively, see Figure 4.

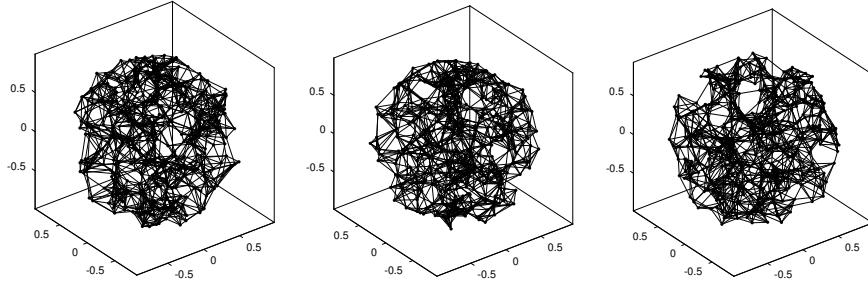


Fig. 4 The graphs G_1 (left), G_2 (middle), and G_3 (right).

For each graph G_1 , G_2 and G_3 we construct the corresponding bi-stochastic graph Laplacians \mathbf{L}_1 , \mathbf{L}_2 , and \mathbf{L}_3 and consider the optimization problem

$$\mathbf{t}^* = \operatorname{argmax}_{\mathbf{t} \in T} \lambda_1(\mathbf{L}_\mathbf{t}).$$

By setting $t_3 = 1 - (t_1 + t_2)$ we can optimize $\lambda_1(\mathbf{L}_\mathbf{t})$ over (t_1, t_2) such that $0 \leq t_1, t_2$ and $t_1 + t_2 \leq 1$, see Figure 5. Using numerical optimization, we find that

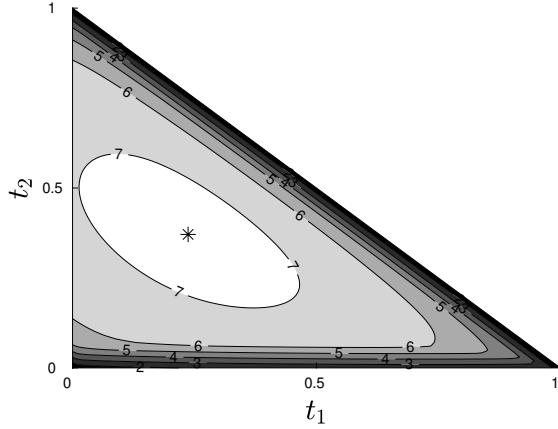


Fig. 5 A contour plot of $\lambda_1(\mathbf{L}_\mathbf{t})$ for $0 \leq t_1, t_2$ and $t_1 + t_2 \leq 1$.

$$\mathbf{t}^* \approx (0.236853469652210, 0.371066650569015, 0.392079879778775).$$

Validating the results of this optimization procedure using Theorem 1 gives

$$\left| \max_{k \in \{1,2,3\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) - \lambda_1(\mathbf{L}_{\mathbf{t}^*}) \right| \leq 9.502285891471729 \times 10^{-8},$$

so the numerical results are very close to optimal. Since we know how the graphs were constructed, we can further interpret the result. As in the previous example, the common variable is the distance of a point to the origin. To demonstrate that $\psi_1(\mathbf{L}_{\mathbf{t}^*})$ is a re-scaling of the common variable, we plot $\psi_1(\mathbf{L}_{\mathbf{t}^*})$ versus the distance

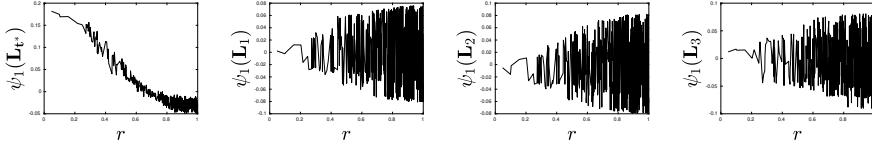


Fig. 6 The first nontrivial eigenvectors of \mathbf{L}_{t^*} , \mathbf{L}_1 , \mathbf{L}_2 , and \mathbf{L}_3 versus the common variable r .

to the origin r ; for comparison, we also plot $\psi_1(\mathbf{L}_1)$, $\psi_1(\mathbf{L}_2)$, and $\psi_1(\mathbf{L}_3)$ against r , see Figure 6.

Finally, we note that this example has some interesting asymmetry. In Figure 5 observe that the level line $\lambda_1(\mathbf{L}_t) = 7$ (the closest level line to the maximum value) almost intersects the line $t_1 = 0$. In contrast, the value of $\lambda_1(\mathbf{L}_t)$ on the lines $t_2 = 0$ and $t_3 = 1 - (t_1 + t_2)$ are close to 1. This indicates that just using the graphs $\{G_2, G_3\}$ could allow us to approximately determine the common variable, while using $\{G_1, G_2\}$ or $\{G_1, G_3\}$ would give bad results. Why is this the case? By definition points in X_1, X_2 have the same z -coordinate, and points in X_1, X_3 have the same y -coordinate, while the only common variable for points in X_2, X_3 is the distance of a point from the origin. For example, if we just consider X_1, X_2 , then the function $f(x, y, z) = z$ is smooth with respect to G_1 and G_2 , but not smooth with respect to G_3 .

3.5 Horizontal and vertical barbell example

Next, we provide a degenerate example, where we are given three graphs G_1 , G_2 , G_3 , and the optimal value of $\mathbf{t} = (t_1, t_2, t_3)$ occurs on the boundary of the region $\{(t_1, t_2) : 0 \leq t_1, t_2 \text{ and } t_1 + t_2 \leq 1\}$. Let $D = \{x \in \mathbb{R}^2 : \|x\|_{\ell^2} \leq 1\}$ be the unit disc, and define the functions $f, g : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ by

$$f(x, y) = (x, y \cdot (1 - \cos \pi x)), \quad \text{and} \quad g(x, y) = (x \cdot (1 - \cos \pi y), y).$$

Informally speaking, the maps f and g squeeze the disc into a horizontal barbell shape and a vertical barbell shape, respectively. Let $X_1 = \{x_1, \dots, x_n\}$ be a set of $n = 250$ independent uniformly random points from the unit disc D . Set $X_2 = f(X_1)$, and $X_3 = g(X_1)$, and let G_1, G_2, G_3 be 6-nearest neighbor graphs of X_1, X_2, X_3 , respectively, see Figure 7.

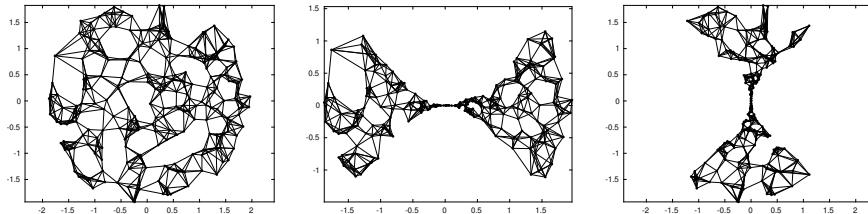


Fig. 7 The graphs G_1 (left), G_2 (middle), and G_3 (right).

Let \mathbf{L}_1 , \mathbf{L}_2 , and \mathbf{L}_3 be the bi-stochastic graph Laplacians of G_1 , G_2 , and G_3 , respectively. We plot $\lambda_1(\mathbf{L}_t)$ versus (t_1, t_2) in Figure 8.

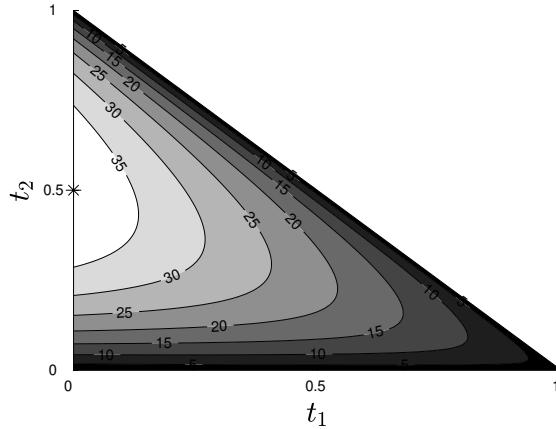


Fig. 8 Contour plot of $\lambda_1(\mathbf{L}_t)$ for (t_1, t_2) such that $0 \leq t_1, t_2$ and $t_1 + t_2 \leq 1$. The maximum is attained at the star.

Using numerical optimization we find that

$$\mathbf{t}^* \approx (2.871019259460022 \times 10^{-15}, .5005153871788890, .4994846128211081);$$

and using Theorem 1 to compute an error estimate gives

$$\left| \max_{k \in \{1,2,3\}} s_{\mathbf{L}_k}(\psi_1(\mathbf{L}_{\mathbf{t}^*})) - \lambda_1(\mathbf{L}_{\mathbf{t}^*}) \right| \leq 2.807697995876879 \times 10^{-6},$$

which verifies that we have solved the optimization problem correctly. Interestingly, the optimal value occurs on the boundary of $\{(t_1, t_2) : 0 \leq t_1, t_2 \text{ and } t_1 + t_2 \leq 1\}$. Since we know how the graphs were created, this behavior makes sense: the sets X_2 and X_3 are modifications of X_1 where points have been squeezed together, which makes the corresponding vertices highly connected in the graphs G_2 and G_3 . This in turn imposes extra conditions for a function to be smooth with respect to G_2 or G_3 . Furthermore, most of the edges appearing in G_1 appear either in G_2 or G_3 .

For this example, the common variable is less straightforward to define. However, one property that is maintained under the deformation is as follows: points in the same quadrant of the plane should remain connected across all graphs. In particular, we can partition the points in X_1 into four groups

$$\begin{aligned} NE &= \{(x, y) \in X_1 : x \geq 0, y \geq 0\}, \\ NW &= \{(x, y) \in X_1 : x < 0, y > 0\}, \\ SW &= \{(x, y) \in X_1 : x < 0, y < 0\}, \text{ and} \\ SE &= \{(x, y) \in X_1 : x \geq 0, y < 0\}. \end{aligned}$$

To understand what the optimal Laplacian $\mathbf{L}_{\mathbf{t}^*}$ is encoding, we plot the first three (nontrivial eigenvectors) of this operator, see Figure 9.

Running k -means clustering on the embedding in Figure 9 would approximately recover the different groups of points NE, NW, SW, and SE.

Remark 3 (Common information spectral clustering) Spectral clustering is a clustering method whose first step is to embed the given data points using the eigenvectors

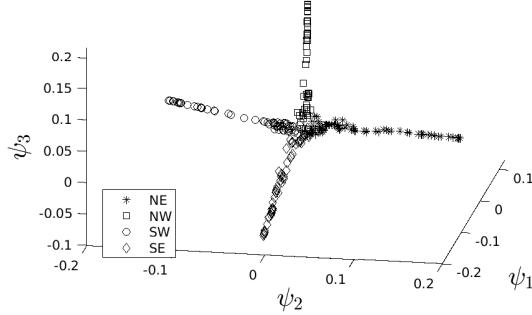


Fig. 9 The first three nontrivial eigenvectors of \mathbf{L}_{t^*} with markers indicating the quadrant of the points in X_1 .

of an operator followed by running the k -means clustering algorithm. The method in this paper can be used to perform a common information spectral clustering algorithm by using the eigenvectors \mathbf{L}_{t^*} to embed the points, and then running k -means clustering. Running k -means on Figure 9 is an example of this common information spectral clustering.

3.6 Spiral and Torus

We conclude with an example illustrating Remark 2: it can be advantageous to change the notion of smoothness. The two graphs in this example are a spiral in the plane and a two-dimensional torus embedded in \mathbb{R}^3 (see Figure 10). Formally, let $\{\theta_1, \dots, \theta_n\}$ and $\{\phi_1, \dots, \phi_n\}$ be $n = 500$ independent uniformly random angles from $[0, 2\pi)$, and let $\{t_1, \dots, t_n\}$ be independent uniformly random points from the interval $[.25, 1.5]$. We define the Spiral set $X_1 = \{x_{1,1}, \dots, x_{n,1}\} \subset \mathbb{R}^2$ by

$$x_{1,j} = (t_j + 0.45\phi_j/(2\pi))(\cos(4\pi(t_j - .25)/1.5), \sin(4\pi(t_j - .25)/1.5)).$$

and define the Torus set $X_2 = \{x_{2,1}, \dots, x_{2,n}\} \subset \mathbb{R}^3$ by

$$x_{2,j} = ((.75 + .25 \cos(\phi_j)) \cos(\theta_j), (.75 + .25 \cos(\phi_j)) \sin(\theta_j), .25 \sin(\phi_j)).$$

For each set X_1 and X_2 we construct 6-nearest neighbor graphs G_1 and G_2 , see Figure 10.

The common variable used to define both graphs is the parameter ϕ_j . The parameter ϕ_j controls the location in the width of the spiral (the width is very small compared to the height), and similarly, controls the location of a point along the smaller circle used to form the torus. The definition

$$s_{\mathbf{L}}(\mathbf{x}) = \frac{1}{\lambda_1(\mathbf{L})} \mathbf{x}^\top \mathbf{L} \mathbf{x} \quad (14)$$

has, in this example, a significant downside: the spiral is only weakly connected and has a very small first Laplacian eigenvalue. The renormalization ensures that $s_{\mathbf{L}}(\mathbf{x})$ is

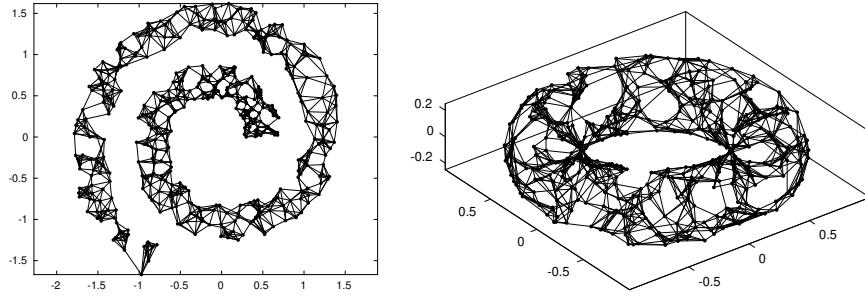


Fig. 10 The graph G_1 (left), and graph G_2 (right).

1, when \mathbf{x} is the first Laplacian eigenvector, however, it will be exceedingly large for all vectors in the orthogonal complement. The degeneracy in the spectrum implies that $s_{\mathbf{L}}(\mathbf{x})$ simply does not accurately capture the spectral geometry of the spiral. The normalization from Remark 2

$$a_{\mathbf{L}}(\mathbf{x}) = \left(\frac{1}{n-1} \sum_{j=1}^{n-1} \lambda_j(\mathbf{L}) \right)^{-1} \mathbf{x}^\top \mathbf{L} \mathbf{x} \quad (15)$$

provides a reasonable alternative: we keep the quadratic form $\mathbf{x}^\top \mathbf{L} \mathbf{x}$ but use a normalization which maintains the global structure of the spectrum better. This is also illustrated in Figure 11.

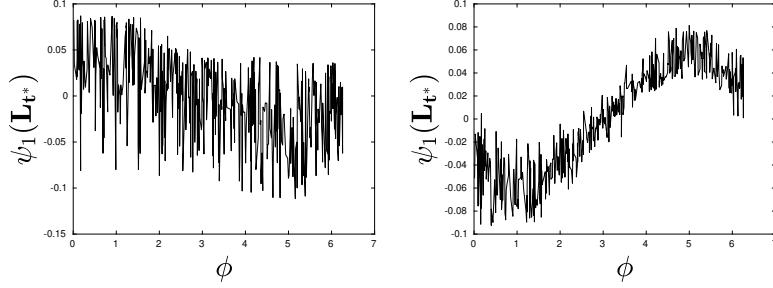


Fig. 11 Using $s_{\mathbf{L}}$ (left) and $a_{\mathbf{L}}$ (right).

Remark 4 (More general definitions of smoothness) In this paper, we considered two notions of smoothness based on the quadratic form $\mathbf{x}^\top \mathbf{L} \mathbf{x}$: the ‘smootherest function’ normalization, and the ‘average smoothness’ normalization, see Remark 2. There are several ways to define intermediate notions of smoothness. For example, given weights w_1, \dots, w_{n-1} one could define a notion of smoothness by normalizing by a weighted sum of the eigenvalues:

$$w_{\mathbf{L}}(\mathbf{x}) = \left(\sum_{j=1}^{n-1} w_j \lambda_j(\mathbf{L}) \right)^{-1} \mathbf{x}^\top \mathbf{L} \mathbf{x}.$$

It is also possible to modify the Laplacians used in the definition; since Theorem 1 and Theorem 2 only require that \mathbf{L} is symmetric positive semi-definite, and has eigenvalue 0 of multiplicity 1 corresponding to constant functions, then it is also possible to define a notion of smoothness by taking a matrix function of the graph Laplacian \mathbf{L} . For example, for $\alpha > 0$ we could define

$$s_{\mathbf{L}}^{\alpha}(\mathbf{x}) = \frac{1}{\lambda_1(\mathbf{L})^{\alpha}} \mathbf{x}^{\top} \mathbf{L}^{\alpha} \mathbf{x}.$$

Such a normalization could be used to adjust for different growth rates of eigenvalues between different graphs. For example, if the given graphs are approximating manifolds of different dimensions, as in Figure 10, then by Weyl's Law the eigenvalues of the Laplace-Beltrami operator on the underlying manifolds will grow at different rates.

4 Technical lemma

4.1 Bi-stochastic normalization

We say that an $n \times n$ nonnegative matrix \mathbf{B} is bi-stochastic if $\mathbf{B}\mathbf{1} = \mathbf{B}^{\top}\mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ denotes the n -dimensional vector whose entries are all 1.

Lemma 2 (Sinkhorn and Kopp [15]) *Let \mathbf{A} be an $n \times n$ nonnegative symmetric matrix with a positive main diagonal. Then, there exists a unique positive definite diagonal matrix \mathbf{D} such that*

$$\mathbf{B} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$$

is a bi-stochastic matrix. Moreover, the matrix \mathbf{B} can be determined by alternating between normalizing the rows and columns (as detailed below).

The iterative procedure of alternating between normalizing the rows and columns of a symmetric matrix can be expressed as follows. We initialize $\mathbf{Q}_0 = \mathbf{I}$, and define

$$\mathbf{Q}_{j+1} = \text{diag}(\mathbf{A}\mathbf{Q}_j^{-1}\mathbf{1}^{\top}), \quad (16)$$

and set $\mathbf{D} = \lim_{k \rightarrow \infty} \mathbf{Q}_{2k+1}\mathbf{Q}_{2k}$, then $\mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ will be bi-stochastic. Given an adjacency matrix \mathbf{A} satisfying the conditions of Lemma 2, and unique positive definite matrix \mathbf{D} from Lemma 2, we define the bi-stochastic graph Laplacian \mathbf{L} by

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}.$$

Remark 5 The bi-stochastic Laplacian is closely related to other operators such as the graph Laplacian $\mathbf{L} = \mathbf{Q} - \mathbf{A}$ (where $\mathbf{Q} = \mathbf{A}\mathbf{1}$), the normalized graph Laplacian $\mathbf{L} = \mathbf{I} - \mathbf{Q}^{-1/2} \mathbf{A} \mathbf{Q}^{-1/2}$, and the random walk graph Laplacian $\mathbf{L} = \mathbf{I} - \mathbf{Q}^{-1} \mathbf{A}$. The bi-stochastic Laplacian \mathbf{L} is symmetric positive semi-definite, has eigenvector $\mathbf{1}$ of eigenvalue 0, and forms a Markov transition matrix when subtracted from the identity matrix.

Remark 6 Numerically, the bi-stochastic graph Laplacian is similar to other graph Laplacians. For example, if $\mathbf{D} = \mathbf{Q}_0 \mathbf{Q}_1$, where \mathbf{Q}_j is defined above in (16), then $\mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ is the normalized graph Laplacian. In practice, the normalized graph Laplacian and random walk graph Laplacian are often used instead of the standard graph Laplacian. The reason we use the bi-stochastic Laplacian for our numerical results is that it is closely related to the normalized graph Laplacian and random walk graph Laplacian, and has all the properties we require: symmetric positive definite with eigenvalue 0 of multiplicity 1 corresponding to constant functions.

5 Comments and Remarks

We conclude with a couple of general comments.

5.1 Extension to multiple functions

Recall we defined the smoothest function or common variable ψ_1 for a collection of graph \mathcal{G} by

$$\psi_1 = \operatorname{argmin}_{\mathbf{x} \in X} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}).$$

By induction we can define

$$X_k = \{\mathbf{x} \in X : \psi_1^\top \mathbf{x} = \dots = \psi_k^\top \mathbf{x} = 0\},$$

and

$$\psi_{k+1} = \operatorname{argmin}_{\mathbf{x} \in X_k} \max_{k \in \{1, \dots, m\}} s_{\mathbf{L}_k}(\mathbf{x}),$$

for $k = 1, 2, 3, \dots$. Informally speaking, ψ_{k+1} is the smoothest function on \mathcal{G} that is orthogonal to ψ_1, \dots, ψ_k . Moreover, by restricting the operators $\mathbf{L}_1, \dots, \mathbf{L}_m$ to X_k the minimax principle of Theorem 2 can be applied to solve this optimization problem. For applications, one might suspect that only the first, or possibly two or three, of these functions would be useful; however, from a theoretical perspective considering the orthogonal basis $\psi_1, \dots, \psi_{n-1}$ may be interesting.

5.2 Sum of diffusions

We quickly mention another approach that is quite similar. Given graphs G_1, \dots, G_m over the same set of vertices V , we can define m different Laplacians $\mathbf{L}_1, \dots, \mathbf{L}_m$ which we restrict to the space orthogonal to constants and normalize via

$$\mathbf{L}_i^* = \frac{1}{\lambda_1(\mathbf{L}_i)} \mathbf{L}_i.$$

This gives rise to diffusion operators

$$\mathbf{H}_i(t) = \exp(-t\mathbf{L}_i^*).$$

The normalization implies that all these m diffusion operators have the same operator norm

$$\|\mathbf{H}_i(t)\| = e^{-t}.$$

If there was a common variable, then it would diffuse slowly among all these different diffusion operators and we would expect that the triangle inequality is almost sharp

$$\left\| \sum_{i=1}^m \mathbf{H}_i(t) \right\| \leq m e^{-t}.$$

This allows us to define a numerical score

$$1 \leq e^t \left\| \sum_{i=1}^m \mathbf{H}_i(t) \right\| \leq m$$

measuring how many of these graphs do indeed have a common variable. Naturally, $\mathbf{H}_i(t)$ will be close to the identity for t small, so the inequality becomes more interesting for t large (and t can play the role of a consistency parameter). This may be interpreted as a simple ‘one-shot’ version of our main idea.

5.3 Random Matrices

We note that the dual version of this idea, the matrix exponential of a linear combination of Laplacians as opposed to a linear combination of matrix exponentials of Laplacians, has a probabilistic interpretation. When we consider applying small multiples of random Laplacians, the main question is the following: if $\mathbf{X}_1, \dots, \mathbf{X}_m \in \mathbb{R}^{n \times n}$ are m matrices what can be said about products

$$\mathbf{X}_s = \left(\mathbf{I} + \frac{\varepsilon}{s} \mathbf{X}_{i_1} \right) \left(\mathbf{I} + \frac{\varepsilon}{s} \mathbf{X}_{i_2} \right) \dots \left(\mathbf{I} + \frac{\varepsilon}{s} \mathbf{X}_{i_s} \right),$$

as $s \rightarrow \infty$. Here we think of i_j as randomly (independently and uniformly) chosen elements from $\{1, 2, \dots, m\}$. An even more general question was studied by Emme and Hubert [6] whose result implies that

$$\lim_{s \rightarrow \infty} \mathbf{X}_s = \exp \left(\frac{\varepsilon}{m} \sum_{k=1}^m \mathbf{X}_k \right).$$

In our setting, we note that

$$\exp \left(-\frac{\varepsilon}{\lambda_1(\mathbf{L}_k)} \mathbf{L}_k \right) = \mathbf{I} - \frac{\varepsilon}{\lambda_1(\mathbf{L}_k)} \mathbf{L}_k + O(\varepsilon^2)$$

implying that the proper limit of random products of matrix exponentials of suitably rescaled Laplacians \mathbf{L}_k chosen with probability proportional to t_k^* would result in

$$\lim_{s \rightarrow \infty} \mathbf{X}_s = \exp(-\varepsilon \mathbf{L}_{\mathbf{t}^*}),$$

which is a natural variant of our approach.

5.4 Summary and discussion

We repeat the main problem: suppose we are given a collection \mathcal{G} of m different graphs over the same set of vertices V

$$\mathcal{G} = \{G_1 = (V, E_1), \dots, G_m = (V, E_m)\}.$$

Among all nonconstant functions $f : V \rightarrow \mathbb{R}$ which is the ‘smoothest’ with respect to \mathcal{G} ? We believe this problem to be of substantial interest. Naturally, there is a certain vagueness in how the problem is posed: 1) what does it mean for a function to be smooth? and 2) what does it mean for a function to be commonly smooth?

In this paper, we propose the classical spectral definition for 1) and a minimax approach for 2). One could, naturally, consider a great many other approaches and we believe it to be a fascinating question for further study. For example, the maximum norm in the definition of smoothness could be replaced by an ℓ^p norm

$$\psi_1 = \underset{\mathbf{x} \in X}{\operatorname{argmin}} \| (s_{\mathbf{L}_1}(\mathbf{x}), \dots, s_{\mathbf{L}_k}(\mathbf{x})) \|_{\ell^p},$$

for some $1 \leq p \leq \infty$, alternatively, the quadratic form in smoothness score $s_{\mathbf{L}_k}(\mathbf{x})$ could be replaced by a different quantity, for example, the Laplacians could be taken to a power as discussed in Remark 4. In summary, there are many potentially interesting ways to formalize our main question; the minimax theorem established in this paper solves the problem for a specific notion of smoothness inspired by spectral graph theory, and provides a basis for further work.

References

1. M. Belkin and P. Niyogi, Laplacian eigenmaps and spectral techniques for embedding and clustering, NIPS 14, pp. 585–591 (2002)
2. X. Cai, F. Nie, H. Huang, and F. Kamangar, Heterogeneous image feature integration via multi-modal spectral clustering, in Proc. Comput. Vis. Pattern Recognit., 2011.
3. R. Coifman, S. Lafon, Diffusion maps. Appl. Comput. Harmon. Anal. 21(1), p. 5–30 (2006)
4. R. Coifman, S. Lafon, Geometric harmonics: a novel tool for multiscale out-of-sample extension of empirical functions. Appl. Comput. Harmon. Anal. 21(1), p. 31–52 (2006)
5. X. Dong, P. Frossard, P. Vandergheynst, N. Nefedov, N. Clustering on multi-layer graphs via subspace analysis on Grassmann manifolds. IEEE Transactions on signal processing, 62, p. 905–918, (2013).
6. J. Emme, P. Hubert, Limit laws for random matrix products, Mathematical Research Letters 25 (2018), p.1205 – 1212
7. D. Eynard, A. Kovnatsky, M. Bronstein, K. Glashoff, A. Bronstein, Multimodal manifold analysis by simultaneous diagonalization of laplacians. IEEE Transactions on Pattern Analysis and Machine Intelligence, 37, p. 2505–2517 (2015).
8. J. Liu, C. Wang, J. Gao, and J. Han, Multi-view clustering via joint nonnegative matrix factorization, in Proc. SDM, 2013.
9. A. Kumar, P. Rai, and H. Daume III, Co-regularized multi-view spectral clustering, in Proc. Neural Inf. Process. Syst., 2011.
10. R. R. Lederman, R. Talmon, Learning the geometry of common latent variables using alternating-diffusion. Appl. Comput. Harmon. Anal. 44(3), p. 509–536 (2018)
11. C. Ma and C.-H. Lee, Unsupervised anchor shot detection using multi-modal spectral clustering, in Proc. ICASSP, 2008.
12. W. Tang, Z. Lu, and I. Dhillon, Clustering with multiple graphs, in Proc. Data Mining, 2009.
13. O. Yair, F. Dietrich, R. Mulayoff, R. Talmon, I. Kevrekidis, Spectral Discovery of Jointly Smooth Features for Multimodal Data, arXiv:2004.04386

14. A. Yeredor, Non-orthogonal joint diagonalization in the leastsquares sense with application in blind source separation, *Trans. Signal Proc.*, vol. 50, no. 7, pp. 1545–1553, 2002.
15. Richard Sinkhorn and Paul Knopp, *Concerning nonnegative matrices and doubly stochastic matrices*, *Pacific J. Math.* **21** (1967), 343–348.