

Geometric structure of graph Laplacian embeddings

Nicolás García Trillos

*Department of Statistics
University of Wisconsin-Madison
Madison, WI 53706, USA*

NICOLASGARCIA@STAT.WISC.EDU

Franca Hoffmann

Bamdad Hosseini

*Computing and Mathematical Sciences
California Institute of Technology
Pasadena, CA 91125, USA*

FKOH@CALTECH.EDU

BAMDADH@CALTECH.EDU

Editor: Ingo Steinwart

Abstract

We analyze the spectral clustering procedure for identifying coarse structure in a data set $\mathbf{x}_1, \dots, \mathbf{x}_n$, and in particular study the geometry of graph Laplacian embeddings which form the basis for spectral clustering algorithms. More precisely, we assume that the data are sampled from a mixture model supported on a manifold $\mathcal{M}$ embedded in $\mathbb{R}^d$, and pick a connectivity length-scale $\varepsilon > 0$ to construct a kernelized graph Laplacian. We introduce a notion of a well-separated mixture model which only depends on the model itself, and prove that when the model is well separated, with high probability the embedded data set concentrates on cones that are centered around orthogonal vectors. Our results are meaningful in the regime where $\varepsilon = \varepsilon(n)$ is allowed to decay to zero at a slow enough rate as the number of data points grows. This rate depends on the intrinsic dimension of the manifold on which the data is supported.

Keywords: Unsupervised learning, spectral clustering, graph Laplacian, mixture models, continuum limit

1. Introduction

The goal of this article is to analyze the *spectral clustering* procedure for identifying coarse structure in a data set by means of appropriate spectral embeddings. Specifically, we combine ideas from spectral geometry, metastability, optimal transport, and spectral analysis of elliptic operators to describe the geometry of the embeddings that form the basis of spectral clustering in the case where the data generating model is sufficiently “well-separated”.

Throughout this article we assume the data set $\mathcal{M}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ consists of points that are distributed according to an underlying probability measure ν supported on a m -dimensional manifold $\mathcal{M}$ embedded in some Euclidean space $\mathbb{R}^d$. The first step in spectral clustering is to construct a similarity graph $G := \{\mathcal{M}_n, \mathbf{W}_n\}$ with vertices $\mathcal{M}_n$ and edge weight matrix $\mathbf{W}_n$ whose entries indicate how similar two points in $\mathcal{M}_n$ are; typical constructions are proximity graphs, where a length scale $\varepsilon > 0$ is chosen and high weights are given to pairs of points that are within distance ε from each other. Once the graph has been constructed a *graph Laplacian* operator is introduced, and its first few eigenvectors are used

to define an embedding that maps the data points $\mathcal{M}_n$ into a low-dimensional Euclidean space. More precisely, the low-lying spectrum of the graph Laplacian induces an embedding $F_n : \{\mathbf{x}_1, \dots, \mathbf{x}_n\} \rightarrow \mathbb{R}^N$ mapping a data point $\mathbf{x}_i$ into a vector $F_n(\mathbf{x}_i)$ in $\mathbb{R}^N$ whose entries are the first N eigenfunctions of the graph Laplacian evaluated at $\mathbf{x}_i$. Here we refer to the map F_n as the *Laplacian embedding*. The embedding step is the crucial step in spectral clustering, which afterwards will typically proceed by running an algorithm like k -means so as to recover the clustering structure of the data set (see von Luxburg (2007)). While k -means is a standard choice in the literature, our analysis will be agnostic to the actual algorithm used to cluster the embedded data set, and we will only focus on the geometry of the embedded data set. We refer the reader to Section 1.2 and in particular to Ling and Strohmer (2019) for alternatives to k -means.

How can one motivate the use of the previous spectral clustering algorithm? Consider the simple scenario where the manifold $\mathcal{M}$ has N connected components. For a data set consisting of uniform samples from such $\mathcal{M}$ one can build the proximity graph G (at least for n large enough) in such a way that its associated graph Laplacian has N connected components and whose corresponding first N eigenvectors coincide with rescaled versions of indicator functions of the N connected components of $\mathcal{M}$ restricted to $\mathcal{M}_n$. The resulting Laplacian embedding F_n is seen to map the original data set into a set of N orthogonal vectors in $\mathbb{R}^N$. The simple geometric structure of the embedded data set allows one to readily identify the “true” coarse structure of the original data set. The key question that we attempt to answer in this work is the following: what is the geometric structure of graph Laplacian embeddings in more general settings where $\mathcal{M}$ is connected, but only weakly? We attempt to answer this question by:

1. Relaxing the notion of N connected components of $\mathcal{M}$, and replacing it with a notion of well-separatedness for the mixture model ν as defined in terms of three quantities described in Section 2.2. We highlight that our notion of well-separatedness only depends on the mixture model ν and not on extrinsic objects like a user chosen kernel in the construction of the graph G .
2. Describing the geometry of the Laplacian embedding in terms of what we will call an *orthogonal cone structure*, a generalization of the orthogonality condition of the embedded data set observed when $\mathcal{M}$ has N connected components. We consider this notion for arbitrary probability measures and not just for point clouds, thus extending the notion introduced in Schiebinger et al. (2015).

In order to study the geometry of the embedded data set, we split our analysis into two parts. First, we study the geometry of a continuum analogue Laplacian embedding by using techniques from spectral theory of elliptic PDEs. The continuum Laplacian embedding can be thought as an ideal object that is recovered (spectrally) from the graph Laplacian in the limit $n \rightarrow \infty$, $\varepsilon = \varepsilon(n) \rightarrow 0$ (at a slow enough rate). In the second part, we quantify the difference between the two embeddings, the graph based and its continuum analogue. To achieve this we use recent results quantifying the approximation error of the spectra of continuum Laplacian operators with the spectra of appropriate graph Laplacians. It is only in the second step that we study statistical properties of random objects. The first part of the paper is, to some extent, of independent interest in connection to theory of diffusion processes and PDE analysis. The structure of our analysis illustrates a more general

approach for the analysis of learning methods that rely on the construction of proximity graphs. Namely, first study the ideal “continuum” problem and then use approximating results to show that the graph problem inherits properties from its continuum counterpart.

Overall, our ideas, definitions, and results provide new theoretical insights on spectral clustering. They further justify the use of spectral clustering beyond the trivial setting where a manifold is disconnected. In particular, we show that when the ground-truth measure ν has well-separated components, the geometry of the embedded data set $F_n(\mathcal{M}_n)$ is simple enough so that an algorithm like k -means can successfully identify the coarse structure of the data set.

1.1 Overview of main results

Let us make our setup more precise. We assume that $\mathcal{M}$ is a connected manifold and that ν is a mixture model with N components supported on $\mathcal{M}$ (see Section 2 for detailed assumptions on this model). We work in a non-parametric setting and assume that the components of ν have sufficiently regular densities with respect to the volume form of $\mathcal{M}$. We will primarily work with graph Laplacian operators corresponding to *kernelized* Laplacians (see (4) for the definition and Section 4.2.2 for some reasons behind this choice where in particular we observe that the kernelized Laplacian already covers many other cases of interest). Before describing the geometry of the embedded data set $F_n(\mathcal{M}_n)$ we study a continuum limit analogue of spectral clustering which we refer to as *measure clustering*. This can be thought as an ideal clustering problem where an infinite amount of data is available. Studying this continuum analogue is motivated by recent results concerning the convergence of graph Laplacian operators to differential operators in the large data limit as n , the number of vertices, goes to infinity and ε , the *connectivity length scale* specifying G , goes to zero sufficiently slowly. The low-lying spectrum of the limiting differential operator induces an embedding $F : \mathcal{M} \rightarrow \mathbb{R}^N$ in a manner similar to F_n , but now at the continuum level. We characterize the geometry of the measure $F_\# \nu$ (i.e. the push-forward of ν through the continuum limit Laplacian embedding F) which forms the basis of the measure clustering procedure. In particular, we identify conditions on the ground-truth model that ensure that the measure $F_\# \nu$ has an orthogonal cone structure. This is the content of our first key result Theorem 10 which is informally stated below.

Theorem 1 (Informal) *If the mixture ν has well-separated components (see Definition 4), then $F_\# \nu$, the push forward of ν through the Laplacian embedding F , has an orthogonal cone structure, i.e., most of its mass is concentrated within disjoint cones centered at orthogonal vectors (see Definition 8).*

Let us now turn our attention to the geometry of the embedded point cloud $F_n(\mathcal{M}_n)$. As mentioned above, when constructing the graph G from the data $\mathcal{M}_n$, one often scales ε with n in such a way that when $n \rightarrow \infty$, $\varepsilon \rightarrow 0$ (i.e. the weight matrix $\mathbf{W}_n$ localizes) at a slow enough rate. Analyzing this regime is motivated by the desire to balance between graph connectivity and computational efficiency. Recent results in the literature provide explicit rates of convergence for the spectra of graph Laplacians towards the spectrum of an elliptic differential operator (see (6) and Section A.3 below) in such regimes. Moreover, the convergence of the eigenvectors is with respect to a strong enough topology, implying that the empirical measure of $F_n(\mathcal{M}_n)$ and the measure $F_\# \nu$ get closer in the Wasserstein

sense as n grows. In turn, this convergence guarantees that if $F_{\sharp}\nu$ has an orthogonal cone structure, then $F_n(\mathcal{M}_n)$ will have one as well (i.e., orthogonal cone structures are stable under Wasserstein perturbations). In Theorem 15 we make the previous discussion mathematically precise. For the moment we provide an informal statement of that result.

Theorem 2 (Informal) *Under suitable conditions, if $F_{\sharp}\nu$ has an orthogonal cone structure then, with very high-probability, the empirical measure of the embedded cloud $F_n(\mathcal{M}_n)$ also has an orthogonal cone structure.*

1.2 Related work

In this paper we focus on clustering. More specifically, clustering by means of appropriate embeddings using the spectrum of a graph Laplacian operator. In the geometric graph setting (our setting) von Luxburg et al. (2008) were the first to rigorously establish a statistical consistency result for graph-based spectral embeddings when the number of data points n goes to infinity and the connectivity $\varepsilon > 0$ remains constant. Up until then most papers focused on proving pointwise convergence estimates (see Belkin and Niyogi (2008); Hein et al. (2005) and references within). The work of von Luxburg et al. (2008) however, did not address the question of determining scalings for $\varepsilon := \varepsilon(n) \rightarrow 0$ that implied spectral convergence towards differential operators in the continuum limit. Moreover, it was still a pending task to analyze the entire spectral clustering algorithm, and in particular study the geometry of the resulting spectral embedding.

García Trillos and Slepčev (2015) introduced a framework using notions from optimal transport and the calculus of variations to study large data limits of optimization problems defined on random geometric graphs. Many large data limits of relevant functionals for machine learning have been studied using similar approaches: Slepčev and Thorpe (2019); Dunlop et al. (2019); García Trillos and Murray (2017); García Trillos and Sanz-Alonso (2018). In particular, the optimal transport framework was used by García Trillos and Slepčev (2018) to analyze the entire spectral clustering procedure (embedding step and k -means step) and the authors were able to show statistical consistency for scalings $\log(n)^{p_m}/n^{1/m} \ll \varepsilon \ll 1$ where m is the intrinsic dimension of the manifold $\mathcal{M}$ and $p_m > 0$ is an appropriate constant depending on m . Nevertheless, the spectral convergence was asymptotic and no rate was provided until Tao and Shi (2020). Burago et al. (2014) took a somewhat similar approach to García Trillos and Slepčev (2018), and in a non-probabilistic setting, conducted a careful analysis which ultimately allowed to provide very precise rates of convergence for the spectrum of graph Laplacians. García Trillos et al. (2020) extended this analysis to cover the relevant probabilistic setting where point clouds are samples from a ground truth measure supported on a manifold. The rates obtained by García Trillos et al. (2020) were recently improved in Calder and García Trillos (2019): these are to the best of our knowledge the best rates of spectral convergence available. The work of Calder and García Trillos (2019) relies on a priori convergence rates (such as the ones in García Trillos et al. (2020)) which are then cleverly improved by relying on pointwise consistency of graph Laplacians for smooth enough functions. In Appendix A.3 we revisit some of the results and ideas introduced by García Trillos et al. (2020); as a matter of fact, the probabilistic estimates that we use in this paper are proved in a similar fashion to the work of García Trillos et al. (2020) after making some slight modifications that are discussed in the appendix.

In principle, we could pursue the improvement of these rates following the framework presented by Calder and García Trillos (2019), but we decided not to do so in order to avoid making the already lengthy exposition much longer and tedious: this is technical work that in our opinion is out of the scope of this paper.

All of the publications listed above focus on spectral convergence of graph Laplacians, but do not study the geometry of Laplacian embeddings themselves, i.e., what is the geometry of the embedded data set, and how such geometry may affect the subsequent clustering step. Schiebinger et al. (2015) formulates this problem in a precise mathematical setting, by introducing a notion of a well-separated data generating mixture model. Roughly speaking, the authors obtain a result that, in spirit, reads similarly to our Theorem 2. However, the notion of a well-separated model introduced in that work depends on extrinsic quantities such as the kernel used to construct the graph at the finite data level, and in particular considers the regime where n is thought to be large, but the length-scale ε remains constant. In other words, the average number of neighbors per point is comparable to the total number of points. The work of Schiebinger et al. (2015) has inspired and motivated our work.

Different scenarios where one can define a notion of well-separated components were studied by Little et al. (2020) and Ling and Strohmer (2019). Both of these papers are closely related to our work, and we believe are complementary of each other. In both works, the respective notion of “well-separated components” is used to help validate proposed methodologies for clustering. Little et al. (2020) focus on understanding the stability of spectral embeddings in situations where the original data set has a geometric structure that is highly anisotropic. The authors propose the use of a general family of distances (other than the Euclidean one) to construct graphs which ultimately induce corresponding spectral embeddings. The model considered in that paper is more geometric in nature than the one we consider here. Ling and Strohmer (2019) introduce a notion of well-separated components at the finite data level and use it to validate a new algorithm for clustering which is based on a convex relaxation for a multi-way cut problem and does not require a k -means step, making it quite appealing.

We believe further exploration of the connections between the various mathematical ideas and models presented by Little et al. (2020) and Ling and Strohmer (2019) with the ones we present here is a worthwhile avenue of future research.

Finally, it is worth reiterating that this work assumes a random *geometric* graph model for the data, i.e. randomness comes from the fact that nodes in the graph are assumed to be random draws from some underlying distribution supported on a manifold, and edges are directly determined from the distance relationship between these nodes. This type of model should be contrasted to other random graph models such as Erdős-Rényi type graphs or planted partition models. Several authors have investigated spectral clustering in those settings (e.g., see Rohe et al. (2011); Chaudhuri et al. (2012) and references within). In a way, an analysis of spectral clustering for *any* random graph model relies on an application, whether direct or indirect, of a Davis-Kahan type argument (see Davis and Kahan (1970)); the only distinction being the operator norms that are feasible in each setting and the use of some specific tools to build a bridge between the sample and population settings. In our case for example, the theory developed by García Trillos et al. (2020) allows us to provide a discrete-to-continuum link, and the Davis-Kahan like argument comes from analyzing a suitable elliptic PDE operator.

1.3 Outline

The rest of the paper is divided as follows. Section 2 covers the preliminaries required for our theory and introduces the precise set-up of the paper. In particular, in Section 2.2 we introduce our notion of well-separated mixture models. We present our main results in Section 3. Theorem 10 on the geometry of continuum Laplacian embeddings (i.e. the measure clustering embedding) and Theorem 15 describing the geometry of graph Laplacian embeddings, are presented in Subsections 3.1 and 3.2 respectively. The proofs of these results are postponed to the appendix.

Section 4 is devoted to discussing our notion of well-separated mixture models. We present some illustrative examples, discuss the implications of our results, and finally comment on possible extensions and generalizations.

Detailed proofs of our main results are presented in Appendix A. In particular, Appendix A.1 focuses on stability properties of orthogonal cone structures, allowing us to prove Theorem 10 in Appendix A.2 after showing several preliminary results. In Appendix A.3 we make the connection between the eigenspaces generated by the first N eigenvectors of the graph Laplacian and those of the continuum Laplacian. We conclude with Appendix A.4, where we present the proof of Theorem 15. Appendix B contains further technical lemmata and theorems that are used in the proofs of Appendix A and the examples in Subsection 4.1.

2. Preliminaries

Here we collect some definitions and notation that are used throughout the article. Discrete and continuum limit graph Laplacian operators are introduced in Subsection 2.1; the notion of a well-separated mixture model is defined in Subsection 2.2.

Let $\mathcal{M}$ be a smooth, connected, orientable, m -dimensional manifold embedded in $\mathbb{R}^d$. For the moment we think of $\mathcal{M}$ as either a compact manifold without boundary or a subset of $\mathbb{R}^d$. Let $\rho_1, \dots, \rho_N$ be a family of $C^1(\mathcal{M})$ probability density functions on $\mathcal{M}$ (densities with respect to the volume form of $\mathcal{M}$) and let $w_1, \dots, w_N$ be positive weights adding to one. The densities $\rho_1, \dots, \rho_N$ are weighted to produce the *mixture model*

$$\rho(x) := \sum_{k=1}^N w_k \rho_k(x), \quad x \in \mathcal{M}. \quad (1)$$

We assume that $\rho(x) > 0$ for all $x \in \mathcal{M}$ and let ν be the probability measure on $\mathcal{M}$ induced by ρ , that is,

$$d\nu(x) = \rho(x)dx,$$

where in the above and in the remainder dx denotes integration with respect to $\mathcal{M}$'s volume form and $\mathcal{P}(\mathcal{M})$ denotes the space of complete Borel probability measures on $\mathcal{M}$.

Given a density ϱ on $\mathcal{M}$ we define the weighted function spaces

$$L^2(\mathcal{M}, \varrho) := \left\{ u : \mathcal{M} \mapsto \mathbb{R} \mid \langle u, u \rangle_{\varrho} < +\infty \right\},$$

equipped with the weighted inner product

$$\langle u, v \rangle_{\varrho} := \int_{\mathcal{M}} u(x)v(x)\varrho(x)dx.$$

Throughout the article we routinely write $L^2(\varrho)$ instead of $L^2(\mathcal{M}, \varrho)$ when the domain of the function space is clear from the context. In Appendix A.3 we modify our notation slightly and use $L^2(\mu)$ to denote $L^2(\mathcal{M}, \varrho)$ where $\mu = \varrho dx$ and use $\langle \cdot, \cdot \rangle_{L^2(\mu)}$ to denote $\langle \cdot, \cdot \rangle_\rho$.

We consider a point cloud $\mathcal{M}_n = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ that are i.i.d. samples from ν , denote by ν_n the empirical measure associated to the samples $\mathbf{x}_1, \dots, \mathbf{x}_n$,

$$\nu_n := \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}. \quad (2)$$

Let $L^2(\nu_n)$ be the space of functions $u_n : \mathcal{M}_n \rightarrow \mathbb{R}$, and we identify $u_n \in L^2(\nu_n)$ with a column vector $(u_n(\mathbf{x}_1), \dots, u_n(\mathbf{x}_n))^T$ in $\mathbb{R}^n$.

2.1 Graph Laplacian operators

We associate a weighted graph structure to $\mathcal{M}_n$ as follows. Let $\eta : [0, \infty) \rightarrow [0, \infty)$ be a non-increasing Lipschitz function with compact support and let

$$\alpha_\eta := \int_0^\infty \eta(r) r^{m+1} dr, \quad (3)$$

where we recall $m \leq d$ is the intrinsic dimension of the manifold $\mathcal{M}$ (typically we think of $m \ll d$ for applications, but this is unimportant for the purposes of the results discussed here). We assume that the kernel has been normalized so that

$$\int_{\mathbb{R}^m} \eta(|x|) dx = 1.$$

For a given $\varepsilon > 0$ we define the weight matrix $\mathbf{W}_n$ with entries $(\mathbf{W}_n)_{ij}$,

$$(\mathbf{W}_n)_{ij} := \frac{2\eta_\varepsilon(|\mathbf{x}_i - \mathbf{x}_j|)}{\alpha_\eta \varepsilon^2 (d_\varepsilon(\mathbf{x}_i))^{1/2} \cdot (d_\varepsilon(\mathbf{x}_j))^{1/2}}, \quad (4)$$

measuring how similar points $\mathbf{x}_i$ and $\mathbf{x}_j$ are. Here, $|\cdot|$ denotes the Euclidean distance in the ambient space $\mathbb{R}^d$, η_ε is defined as

$$\eta_\varepsilon(r) := \frac{1}{\varepsilon^m} \eta\left(\frac{r}{\varepsilon}\right),$$

and

$$d_\varepsilon(y) := \sum_{j=1}^n \eta_\varepsilon(|y - \mathbf{x}_j|) \quad y \in \mathcal{M}.$$

Finally, we introduce the graph Laplacian operator $\Delta_n : L^2(\nu_n) \rightarrow L^2(\nu_n)$,

$$\Delta_n := \mathbf{D}_n - \mathbf{W}_n, \quad (5)$$

where $\mathbf{D}_n$ is the degree matrix associated to the weight matrix $\mathbf{W}_n$, that is, $\mathbf{D}_n$ is a diagonal matrix whose diagonal entries $(\mathbf{D}_n)_{ii}$ are given by $(\mathbf{D}_n)_{ii} = \sum_{j=1}^n (\mathbf{W}_n)_{ij}$. The dependence of Δ_n on ε has been suppressed for notational convenience.

The core idea in our analysis is the convergence of the spectrum of Δ_n to that of the differential operator

$$\Delta_\rho(u) := -\frac{1}{\rho} \operatorname{div}(\rho \nabla u) = -\Delta u - \frac{1}{\rho} \nabla \rho \cdot \nabla u, \quad (6)$$

which is defined for smooth functions $u : \mathcal{M} \mapsto \mathbb{R}$. In the above, ∇ stands for the gradient operator on $\mathcal{M}$, div is the divergence on $\mathcal{M}$, and Δ is the Laplace-Beltrami operator on $\mathcal{M}$. The term $\nabla \rho \cdot \nabla u$ evaluated at a point $x \in \mathcal{M}$ is the inner product of the vectors $\nabla \rho(x), \nabla u(x)$ in the tangent space $\mathcal{T}_x \mathcal{M}$; seen as vectors in $\mathbb{R}^d$, $\nabla \rho(x) \cdot \nabla u(x)$ is just their inner product in $\mathbb{R}^d$. We establish the convergence of Δ_n to Δ_ρ in Appendix A.3 when the number of samples $n \rightarrow \infty$ while the local connectivity parameter $\varepsilon \rightarrow 0$ at an appropriate rate.

We notice that the weights in (4) have been appropriately rescaled so that the eigenvalues of the graph Laplacian Δ_n defined below approximate those of Δ_ρ . We also point out that $(\mathbf{W}_n)_{ij}$ can be written as

$$(\mathbf{W}_n)_{ij} = \frac{2\eta_\varepsilon(|\mathbf{x}_i - \mathbf{x}_j|)}{\alpha_\eta n \varepsilon^2 (d_\varepsilon(\mathbf{x}_i)/n)^{1/2} \cdot (d_\varepsilon(\mathbf{x}_j)/n)^{1/2}},$$

and that $d_\varepsilon(\cdot)/n$ is nothing but a kernel density estimator for the density ρ due to the fact that η is normalized.

Remark 3 *The normalization terms appearing in (4) induce the kernelized graph Laplacian Δ_n which in turn converges to Δ_ρ in the continuum limit. There are other possible normalizations that one can consider such as the ones in the parametric family introduced in Coifman and Lafon (2006). The effect of the different normalizations can be observed asymptotically (as $n \rightarrow \infty$) in the way the density ρ affects the differential operator Δ_ρ (which in general will differ from (6)). We expand our discussion on the choice of different normalizations in Section 4.2.2; for an overview of a more general family of normalizations see Section 5 of Hoffmann et al. (2019). In the context of this paper, it is enough to say that Δ_n and Δ_ρ (as defined in (6)) are closely related to each other following Theorem 30.*

Similarly to (6) we also define the differential operators Δ_{ρ_k} in association with the densities ρ_k , which for smooth functions u , are defined by

$$\Delta_{\rho_k}(u) := -\frac{1}{\rho_k} \operatorname{div}(\rho_k \nabla u) = -\Delta u - \frac{1}{\rho_k} \nabla \rho_k \cdot \nabla u. \quad (7)$$

2.2 Well-separated mixture models

Let us now introduce the notion of a well-separated mixture model. This is a notion defined in terms of three quantities associated to the model ν that we refer to as *overlapping*, *coupling*, and *indivisibility* parameters.

For $k = 1, \dots, N$, define the functions

$$q_k := \sqrt{\frac{w_k \rho_k}{\rho}}. \quad (8)$$

The *overlapping* parameter of the mixture model ν is the quantity

$$\mathcal{S} := \max_{i \neq j} \left\langle \left(\frac{q_i}{\sqrt{w_i}} \right)^2, \left(\frac{q_j}{\sqrt{w_j}} \right)^2 \right\rangle_\rho = \max_{i \neq j} \int_{\mathcal{M}} \frac{\rho_i \rho_j}{\rho} dx, \quad (9)$$

Note that $\mathcal{S} \in \left(0, \frac{1}{\min_j w_j}\right)$ since for any $k \in \{1, \dots, N\}$ we have the bound

$$(\min_j w_j) \rho_k \leq \sum_{j=1}^N w_j \rho_j = \rho,$$

and so $\frac{\rho_k}{\rho} \leq \frac{1}{\min_j w_j}$.

The *coupling* parameter of the mixture model is defined as

$$\mathcal{C} := \max_{k=1, \dots, N} \mathcal{C}_k, \quad (10)$$

where

$$\mathcal{C}_k := \frac{1}{4} \int_{\mathcal{M}} \left| \frac{\nabla \rho_k}{\rho_k} - \frac{\nabla \rho}{\rho} \right|^2 \rho_k dx, \quad k = 1, \dots, N. \quad (11)$$

This quantity is the Fisher information of the measure $\rho_k dx$ with respect to the measure ρdx . We use the convention that if the density $\rho_k(x)$ vanishes at some point $x \in \mathcal{M}$, then we interpret the integrand in (11) as zero at that point; in other words we only integrate over the set $\{x \in \mathcal{M} : \rho_k(x) > 0\}$.

Finally, the *indivisibility* parameter for the model is defined as

$$\Theta := \min_{k=1, \dots, N} \Theta_k, \quad (12)$$

where

$$\Theta_k := \min_{u \perp \mathbb{1}} \frac{\int_{\mathcal{M}} |\nabla u|^2 \rho_k dx}{\langle u, u \rangle_{\rho_k}}. \quad (13)$$

In (13), $\mathbb{1}$ stands for the function that is identically equal to one, and $\perp$ indicates orthogonality with respect to the inner product $\langle \cdot, \cdot \rangle_{\rho_k}$. This may as well be defined as the first non-trivial eigenvalue of the operator Δ_{ρ_k} .

In Subsection 4.1 we discuss different interpretations of $\mathcal{S}, \mathcal{C}, \Theta$ and present some examples that illustrate the significance of these parameters. Intuitively, the overlapping and coupling parameters $\mathcal{S}$ and $\mathcal{C}$ are small whenever the mixture components ρ_k are individual bumps that are mostly concentrated on disjoint sets (see Figure 1) while the indivisibility parameter Θ is large when the components ρ_k cannot be further divided into smaller components. We are now ready to introduce the notion of a well-separated mixture model.

Definition 4 (Informal) *The mixture model (1) is well-separated if*

$$\mathcal{S} \ll 1, \quad \frac{\mathcal{C}}{\Theta} \ll 1.$$

These conditions are to be read as: the overlapping $\mathcal{S}$ is small enough, and the coupling parameter $\mathcal{C}$ is small in comparison to the indivisibility parameter Θ .

We highlight that the above definition is only qualitative in accordance with the informal version of our first main result (Theorem 1). However, this definition can be made rigorous in the relative sense. For example, given two mixture models ν and ν' , the model for which $\mathcal{S}$ and $\mathcal{C}/\Theta$ are smaller is the better separated model and in turn the Laplacian embedding F can cluster that model more effectively. This is a direct consequence of Theorem 10 that gives precise quantitative bounds on how tightly the measure $F_{\sharp}\nu$ concentrates around orthogonal vectors depending on the quantities $\mathcal{S}$ and $\mathcal{C}/\Theta$.

Remark 5 *We emphasize that both of the conditions in Definition 4 are crucial for the mixture model to be well-separated and to exclude pathologies. Having $\mathcal{S} \ll 1$ ensures that the components ρ_k are not repeated while $\mathcal{C} \ll \Theta$ ensures that components that should have otherwise been merged or components that should have otherwise been split into smaller components are excluded. See also Examples 1 and 2 and Remark 16.*

Remark 6 *We will use the functions $q_1, \dots, q_N$ frequently in the proofs of Appendix A so it is worth mentioning a simple interpretation for them. Suppose that we wanted to obtain one sample from the mixture model introduced in (1): $\mathbf{x} \sim \rho$. One way to obtain $\mathbf{x}$ is to first choose a random number $\mathbf{k} \in \{1, \dots, N\}$ where $\mathbb{P}(\mathbf{k} = k) = w_k$ and then sample $\mathbf{x}$ from $\rho_{\mathbf{k}}$. Conversely given the value $\mathbf{x} = x$, we can check that:*

$$\mathbb{P}(\mathbf{k} = k | \mathbf{x} = x) = \frac{w_k \rho_k(x)}{\rho(x)}.$$

Thus, $q_k(x)$ is simply the square root of the likelihood of $\mathbf{x}$ being sampled from ρ_k .

3. Geometric structure of Laplacian embeddings

In this section we discuss detailed versions of our main results. In Subsection 3.1 we present Theorem 10 as the rigorous version of Theorem 1. This theorem gives quantitative bounds on the concentration properties of the pushforward measure $F_{\sharp}\nu$ around orthogonal cones in low dimensions as a function of $\mathcal{S}, \mathcal{C}$, and Θ . The map F (defined in (14)) is the continuum analogue of the spectral embedding commonly used in spectral clustering.

Subsection 3.2 contains Theorem 15 as the quantitative analogue of Theorem 2 which extends Theorem 10 from the continuum to the discrete setting. More precisely, this theorem shows that the embedded point cloud $F_n(\mathcal{M}_n)$ inherits the concentration properties of $F_{\sharp}\nu$ with very high probability provided that (n, ε) are within an appropriate regime. Here F_n is the usual discrete Laplacian embedding defined in (18).

In summary, Theorems 10 and 15 together state that whenever the mixture ν has N well-separated components and (n, ε) are in the appropriate regime, the point cloud $F_n(\mathcal{M}_n)$ concentrates within N disjoint orthogonal cones with very high probability. Further, the better ν is separated, the smaller the angle of these cones, and the more mass lies within the cones. This means that in classification tasks, the embedded cloud $F_n(\mathcal{M}_n)$ can be classified easily using $N - 1$ dimensional hypersurfaces in $\mathbb{R}^N$ that separate the orthogonal cones.

3.1 The continuum setting

In what follows we describe the geometry of the ground-truth spectral embedding (i.e. the measure clustering embedding in the continuum). First, let us collect assumptions on the mixture model ν .

Assumption 7 *The components $\rho_1, \dots, \rho_N$ of the mixture ν belong to $C^1(\mathcal{M})$ and are such that the operators Δ_{ρ_k} and Δ_ρ have discrete point spectrum with associated orthonormal eigenfunctions spanning $L^2(\rho_k)$ and $L^2(\rho)$ respectively.*

Now let $u_1, \dots, u_N$ be the first N orthonormal (with respect to $\langle \cdot, \cdot \rangle_\rho$) eigenfunctions of Δ_ρ corresponding to its N smallest eigenvalues. Define the Laplacian embedding

$$F : x \in \mathcal{M} \mapsto \begin{pmatrix} u_1(x) \\ \vdots \\ u_N(x) \end{pmatrix} \in \mathbb{R}^N. \quad (14)$$

The push-forward of the measure ν by F , denoted $\mu := F_\# \nu$, is a measure on $\mathbb{R}^N$ which describes the distribution of points originally in $\mathcal{M}$ after transformation by the Laplacian map F . Our first main result asserts that if the mixture model (1) is well-separated, then the measure μ has an orthogonal cone structure; this is a geometric notion originally introduced for point clouds in Schiebinger et al. (2015) and extended here for arbitrary probability distributions.

Definition 8 *Let $\sigma \in (0, \pi/4)$, $\delta \in [0, 1)$, and $r > 0$. We say that the probability measure $\mu \in \mathcal{P}(\mathbb{R}^k)$ has an orthogonal cone structure with parameters (σ, δ, r) if there exists an orthonormal basis for $\mathbb{R}^k$, $e_1, \dots, e_k$, such that*

$$\mu \left(\bigcup_{j=1}^k C(e_j, \sigma, r) \right) \geq 1 - \delta,$$

where $C(e_j, \sigma, r)$ is the set

$$C(e_j, \sigma, r) := \left\{ z \in \mathbb{R}^k : \frac{z \cdot e_j}{|z|} > \cos(\sigma), \quad |z| > r \right\}.$$

In simple words, a measure $\mu \in \mathcal{P}(\mathbb{R}^k)$ has an orthogonal cone structure if it is concentrated on cones centered at vectors which form an orthonormal basis for $\mathbb{R}^k$. The parameter σ is the opening angle of the cones measured from the axis e_j , r is the radius of the ball centered at the origin that we remove from the cones to form the sets $C(e_j, \sigma, r)$, and δ is the amount of mass that is allowed to lie outside the cones. The condition $\sigma < \pi/4$ ensures that the sets $C(e_j, \sigma, r)$ do not overlap.

Remark 9 *Instead of the above definition, one could also consider a component-wise cone structure, namely the existence of parameters $\{(\sigma_j, \delta_j, r_j)\}_{j=1}^k$ such that $\sum_{j=1}^k \delta_j \in [0, 1)$ and $\mu(C(e_j, \sigma_j, r_j)) \geq w_j - \delta_j$ for all $j \in \{1, \dots, k\}$. This notion of orthogonal cone structure implies the weaker notion in Definition 8. We do not introduce this component-wise*

orthogonal cone structure here as our assumption are not strong enough to conclude such a structure holds; in particular, one would need a more detailed, pointwise notion of the overlapping parameter $\mathcal{S}$. However, we believe that this component-wise cone structure can lead to stronger versions of our main theorems with individual bounds on the probability masses lying within each cone as opposed to the total mass outside of the union of the cones.

We are now ready to state our first main result.

Theorem 10 *Let ν be a mixture model with density ρ of the form (1) which we assume satisfies Assumptions 7. For $\sigma \in (0, \pi/4)$, suppose*

$$\mathcal{S} < \frac{w_{\min}(1 - \cos^2(\sigma))}{w_{\max} \cos^2(\sigma) N^2}.$$

Define

$$\tau := 4 \left(\sqrt{\frac{\Theta(1 - N\mathcal{S})}{\mathcal{C}}} - \frac{\sqrt{N\mathcal{S}}}{(1 - \mathcal{S})} \right)^{-1} + \sqrt{\mathcal{S}}.$$

and

$$\delta^* := \frac{w_{\max} \cos^2(\sigma) N^2 \mathcal{S}}{w_{\min}(1 - \cos^2(\sigma))}.$$

Here, $w_{\max} := \max_{i=1, \dots, N} w_i$ and $w_{\min} := \min_{i=1, \dots, N} w_i$.

Suppose that

$$\tau - \sqrt{\mathcal{S}} > 0, \quad \tau N < 1, \tag{15}$$

and $s, t > 0$ satisfy

$$\frac{t \sin(s)}{\sqrt{N w_{\max}}} \geq \sqrt{N \left(\frac{\tau - \sqrt{\mathcal{S}}}{2} \right)^2 + 4N^{3/2} \left(\frac{1}{\sqrt{1 - N\tau}} - 1 \right)}, \quad s + \sigma < \frac{\pi}{4}. \tag{16}$$

Then, the probability measure $\mu := F_{\sharp} \nu$ has an orthogonal cone structure with parameters $(\sigma + s, \delta + t^2, \frac{1 - \sin(s)}{\sqrt{w_{\max}}})$ for any $\delta \in [\delta^*, 1)$. Here, F is the continuum Laplacian embedding defined in (14).

Notice firstly that the smaller the quantities $\mathcal{S}$ and $\mathcal{C}/\Theta$ are (i.e. the better separated the mixture model is), the smaller τ is. As a consequence the right-hand side of (16) is also small. This allows a choice of parameters s, t close to zero. Secondly, for small $\mathcal{S}$, we can choose σ to be small, and if $\mathcal{S}$ is small compared to $\frac{w_{\min}(1 - \cos^2(\sigma))}{w_{\max} \cos^2(\sigma) N^2}$, then δ^* will be small also. Following Definition 8, small choices of δ and t mean that a bigger proportion of the mass of the measure $F_{\sharp} \nu$ can be found inside the cones, whereas if σ and s are small, then the cones are more concentrated around orthogonal vectors. The bottom line is that the smaller $\mathcal{S}$ and $\mathcal{C}/\Theta$ are, the more concentrated around orthogonal vectors the measure $F_{\sharp} \nu$ is. Note also that keeping all other parameters fixed, smaller choices of σ will force a bigger choice of δ . This reflects a trade-off between the total mass of the measure $F_{\sharp} \nu$ inside the cones, and the opening angle of the cones.

To see that the measure $F_{\sharp}\nu$ has an orthogonal cone structure under the assumption of a well-separated mixture model, we first study the geometry of a related measure $F_{\sharp}^Q\nu$ where F^Q is the map

$$F^Q : x \in \mathcal{M} \mapsto \begin{pmatrix} \frac{q_1(x)}{\sqrt{w_1}} \\ \vdots \\ \frac{q_N(x)}{\sqrt{w_N}} \end{pmatrix} \in \mathbb{R}^N, \quad (17)$$

where we recall the functions q_k are as in (8). Note that the functions q_k are normalized in the sense that $\sum_{k=1}^N |q_k(x)|^2 = 1$ for all $x \in \mathcal{M}$, whereas the entries of F^Q are normalized in $L^2(\rho)$ since

$$\|F_k^Q\|_{\rho}^2 = \int_{\mathcal{M}} \frac{\rho_k(x)}{\rho(x)} \rho(x) dx = 1 \quad \forall k \in \{1, \dots, N\}.$$

We will show that $F_{\sharp}^Q\nu$ has an orthogonal cone structure provided the overlapping parameter of the mixture model $\mathcal{S}$ is small enough (see Proposition 29). In Section A.1 we show that orthogonal cone structures are stable under Wasserstein perturbations. Theorem 10 is then proved by showing that the measures $F_{\sharp}^Q\nu$ and $F_{\sharp}\nu$ are close to each other in the Wasserstein sense modulo an orthogonal transformation. To achieve this we will essentially show that $q_1, \dots, q_N$ are close to spanning the subspace generated by the first N eigenfunctions of Δ_{ρ} . This statement is trivially true in the extreme case when $\mathcal{M}$ has N connected components and the $\rho_k dx$ are the uniform distributions of the connected components. Indeed, in that case, the first N eigenvalues of Δ_{ρ} are zero, and the first N eigenfunctions of Δ_{ρ} are rescaled versions of the indicator functions of the components. In other words, the span of the ρ_k is exactly equal to the span of the first N eigenfunctions of Δ_{ρ} , and in particular up to a rotation we have $F^Q = F$.

3.2 The discrete setting

Having discussed spectral clustering at the continuum level, we now move to the discrete setting. Let $\mathcal{M}_n := \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ be i.i.d. samples from the probability distribution $d\nu = \rho(x)dx$ and let Δ_n be the discrete graph Laplacian operator on $L^2(\nu_n)$ as in (5). The spectrum of Δ_n induces a graph Laplacian embedding

$$F_n : \mathcal{M}_n \rightarrow \mathbb{R}^N$$

$$F_n : \mathbf{x}_i \mapsto \begin{pmatrix} u_{n,1}(\mathbf{x}_i) \\ \vdots \\ u_{n,N}(\mathbf{x}_i) \end{pmatrix}, \quad (18)$$

where $u_{n,1}, \dots, u_{n,N}$ are the first N eigenvectors of Δ_n . Our second main result states that, with very high-probability, the measure $F_{n\sharp}\nu_n$ has an orthogonal cone structure if the mixture model ν is well-separated, the number of data points n is sufficiently large, and the connectivity parameter ε is sufficiently small. In order to state and prove this result we make the following assumption.

Assumption 11 *We assume that $\mathcal{M}$ is a smooth, compact, orientable, connected m -dimensional manifold without a boundary. Furthermore, we assume that $\rho \in C^1(\mathcal{M})$, and that there exist positive numbers $0 < \rho^- < \rho^+$ such that*

$$\rho^- \leq \rho(x) \leq \rho^+, \quad \forall x \in \mathcal{M}.$$

We denote by $C_{\text{Lip}} > 0$ the Lipschitz constant of ρ .

We also assume that n is large enough and ε is small enough, so that

$$\varepsilon \left(1 + \sqrt{\lambda_N} \right) + \frac{\log(n)^{p_m}}{n^{1/m} \varepsilon} \leq \min \left\{ C, \frac{1}{2} (\lambda_{N+1} - \lambda_N) \right\}, \quad (19)$$

where $p_m = 3/4$ if $m = 2$ and $p_m = 1/m$ when $m \geq 3$, $C > 0$ is a finite constant that depends only on $\mathcal{M}$, and λ_N and λ_{N+1} are the N -th and $(N+1)$ -th eigenvalues of Δ_ρ in increasing order.

Remark 12 *Several theoretical results deducing convergence rates for the spectra of graph Laplacians towards their continuum counterparts make similar assumptions to those in Assumption 11 (e.g. García Trillos et al. (2020); García Trillos and Slepčev (2018)), and for the moment being we will only be able to state our second main result under the assumption that $\mathcal{M}$ is compact (see Section 4.2.1 where we discuss further extensions and generalizations). Regarding the assumptions on n and ε , we remark that requiring the quantity on the left hand side to be less than a constant C is the same condition used in the literature to quantify rates of convergence for the spectra of graph Laplacians (see for example Burago et al. (2014); García Trillos et al. (2020)). On the other hand, requiring the left hand side of (19) to also be smaller than $\frac{1}{2}(\lambda_{N+1} - \lambda_N)$ is a condition that is easier to satisfy when the model is well-separated. Indeed, we will later see that it is possible to find an upper bound for λ_N (Lemma 38) and a lower bound for λ_{N+1} (Proposition 27) in terms of the parameters of the mixture model. That is, one can restate our assumption and write an inequality in terms of the overlapping, indivisibility and coupling parameters. However, at this point we do not believe there is any benefit in developing this discussion any further, and we remark that these are simply concrete and technical conditions that guarantee that we have entered the regime where our theorems hold.*

Remark 13 *Notice that when the manifold $\mathcal{M}$ is smooth and compact as in Assumption 11, Assumption 7 is automatically satisfied.*

Remark 14 *Notice that the fact that ρ is bounded away from zero does not imply that the individual components ρ_k are as well. In particular, an individual component ρ_k may be arbitrarily close to zero (or even equal to zero) in certain region of $\mathcal{M}$ without contradicting Assumption 11. As a matter of fact, one precisely needs this to be the case in order for a model to be well-separated (for otherwise $\mathcal{S}$ would not typically be small).*

In Appendix A.3, we make the connection between the spectrum of the graph Laplacian Δ_n and the spectrum of the continuous operator Δ_ρ . The proof of the main technical result, Theorem 30, relies on a comparison between the Dirichlet forms associated to the operators Δ_n and Δ_ρ by way of careful interpolation and discretization of discrete and continuum

functions. Only small modifications to the proofs in Burago et al. (2014); García Trillos et al. (2020) are necessary and Section A.3 should be understood as a short guide on how to adapt the machinery developed there to the context of this paper.

Finally, we leverage Theorems 10 and 30 to establish our second main result, which asserts that under Assumption 11 and provided the mixture model is well-separated, the measure $F_{n\sharp}\nu_n$ also has an orthogonal cone structure with very high probability. More precisely we have the following.

Theorem 15 *Let $N \geq 2$ and $\beta > 1$. Suppose that $\mathcal{M}$ and ρ satisfy Assumption 11 for some $\varepsilon > 0$. Assume also that the parameters $\mathcal{S}, \mathcal{C}, \Theta$ of the mixture model ρ are such that the quantities τ, δ^* and σ from Theorem 10 satisfy (15) and (16). In addition, assume $\mathcal{S} < N^{-2}$. Let $\mathbf{x}_1, \dots, \mathbf{x}_n$ be i.i.d. samples from the measure $d\nu = \rho dx$ with associated empirical measure ν_n , and let F_n be the Laplacian embedding defined in (18).*

Then, there exists a constant $C_\beta > 0$ depending only on β , so that with probability at least $1 - C_\beta n^{-\beta}$, the probability measure $\mu_n := F_{n\sharp}\nu_n$ has an orthogonal cone structure with parameters

$$\left(\sigma + s, \delta + t^2, \frac{1 - \sin(s)}{\sqrt{w_{\max}}} \right),$$

for any $\delta \in [\delta^*, 1)$, and $s, t > 0$ satisfying

$$\sigma + s < \frac{\pi}{4},$$

$$\begin{aligned} \frac{t \sin(s)}{\sqrt{N w_{\max}}} \geq & \sqrt{N \left(\frac{\tau - \sqrt{\mathcal{S}}}{2} \right)^2 + 4N^{3/2} \left(\frac{1}{\sqrt{1 - N\tau}} - 1 \right)} \\ & + \sqrt{N \phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m)}, \end{aligned} \quad (20)$$

where

$$\begin{aligned} \phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m) = & c_{\mathcal{M}} \left(\frac{N}{1 - N\mathcal{S}^{1/2}} \right) \left(\varepsilon + \frac{\log(n)^{p_m}}{\varepsilon n^{1/m}} \right) \\ & \cdot \left(\left(\sqrt{\frac{\Theta}{\mathcal{C}}(1 - N\mathcal{S})} - \frac{\sqrt{N\mathcal{S}}}{1 - \mathcal{S}} \right)^2 - \frac{N}{1 - N\mathcal{S}^{1/2}} \right)^{-1}, \end{aligned}$$

and $c_{\mathcal{M}} > 0$ is a constant depending on $\mathcal{M}, \rho^\pm, C_{Lip}, \eta$ and N . In the above we require $\mathcal{S}$ and $\mathcal{C}/\Theta$ small enough so that all the inner terms in the definition of ϕ make sense and are positive.

Notice that the only extra term in inequality (20) that does not appear in (16) is the term $\phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m)$. As we will see in Appendix A.4 (following the estimates in Appendix A.3) this term is an upper bound for the Wasserstein distance between the measures $F_{\sharp}\nu$ and $F_{n\sharp}\nu_n$ (up to an orthogonal transformation). The assumptions in Theorem 15 imply that $F_{\sharp}\nu$ has an orthogonal cone structure as described in Theorem 10, and hence in order to deduce that $F_{n\sharp}\nu_n$ has an orthogonal cone structure (as it is stated in Theorem 15) it is

enough to combine Proposition 22 with the estimates on the Wasserstein distance between $F_{n\sharp}\nu_n$ and $F_{\sharp}\nu$.

We observe that asymptotically, if the parameter $\varepsilon > 0$ scales with n like

$$\frac{\log(n)^{p_m}}{n^{1/m}} \ll \varepsilon \ll 1,$$

then $\phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m)$ gets smaller as $n \rightarrow \infty$, and we deduce that the geometry of the point cloud $F_n(\mathcal{M}_n)$ converges to that of $F_{\sharp}\nu$. Nevertheless, it is worth emphasizing that our results are meaningful in the finite (but possibly large) n case. As we discussed in Section 1.2 these rates can be improved with some straightforward but tedious computations following the ideas in the recent work of Calder and García Trillos (2019). We opted for concreteness and simplicity in our estimates (as much as possible) in order to avoid detracting from the main analytical idea that we explore in this paper: study properties of continuum objects using PDE techniques and then connect to the finite data objects using *some* available approximation result.

4. Discussion

The purpose of this section is to provide a more detailed account of the notions introduced above, as well as to discuss our main results and provide several remarks on extensions and generalizations.

In Subsection 4.1 we discuss the overlapping, coupling, and indivisibility parameters and the notion of a well-separated mixture model from various view points and in the context of two concrete examples. Subsection 4.2 is then dedicated to various extensions and generalizations of our results to unbounded domains and k -NN graphs as well as different normalizations of the graph Laplacian.

4.1 Overlapping, coupling and indivisibility parameters

We now give a more detailed description of the parameters $\mathcal{S}$, $\mathcal{C}$ and Θ introduced in Section 2.2. The overlapping parameter $\mathcal{S}$ is small if the components ρ_k have well-separated mass in the sense that if at some point $x \in \mathcal{M}$ one component $\rho_k(x)$ is large, then all other components $\rho_i(x)$, $i \neq k$, are small at that point. In this case, at $x \in \mathcal{M}$, we have $\rho(x) \approx w_k \rho_k(x)$, and so

$$\frac{\rho_k(x)\rho_i(x)}{\rho(x)} \approx \frac{\rho_i(x)}{w_k},$$

which is small. If ρ_k is small at $x \in \mathcal{M}$ on the other hand, then

$$\frac{\rho_k(x)\rho_i(x)}{\rho(x)} \leq \frac{\rho_k(x)}{w_i}$$

is also small. In other words, the parameter $\mathcal{S}$ being small requires the components not to *overlap* too much (hence the name given to it).

If the indivisibility parameter Θ is small, this is an indication that our choice of mixture model $\rho(x) = \sum_k w_k \rho_k(x)$ for a given density ρ is not optimal in the sense that it could

be split into smaller meaningful components. This property is captured by the spectrum of the operators Δ_{ρ_k} . For any $k \in \{1, \dots, N\}$, and $u, v \in C^1(\mathcal{M})$,

$$\langle \Delta_{\rho_k} u, v \rangle_{\rho_k} = \int_{\mathcal{M}} \nabla u \cdot \nabla v \rho_k dx = \langle \nabla u, \nabla v \rangle_{\rho_k},$$

and so the first eigenvalue is zero with constant eigenfunction $\mathbb{1}$. For the sake of illustration, suppose that the support of ρ_k has two disconnected components. In that case, one can construct an eigenfunction orthogonal to $\mathbb{1}$ that takes different constant values on the two different components, and so the second eigenvalue is zero. In light of the max-min theorem (see for example (Attouch et al., 2014, Thm. 8.4.2), also known as the Courant–Fisher theorem) and definition (13), we observe that Θ_k is exactly the second eigenvalue of the operator Δ_{ρ_k} .

It is well-known, in both the spectral geometry and machine learning literature, that the first non-trivial eigenvalue of the Laplacian associated to a geometric object (discrete or continuum) is intimately related to the problem of two-way clustering. Having a small first non-trivial eigenvalue indicates the presence of two clusters in the data set in the discrete setting. This is essentially because the eigenvalue problem can be interpreted as a relaxation of balanced cut minimization problems (see von Luxburg (2007) and also Cheeger’s inequality for graphs (Chung, 1997, Sec. 2.3)). The more ρ_k concentrates on two separate components, the smaller the second eigenvalue will become (see Hoffmann et al. (2019) for precise estimates). This is why the second eigenfunction can be interpreted as an approximate identifier for clusters in the problem of two-way clustering, and it is usually referred to as the *Fiedler vector*. Intuitively, a low value of Θ_k indicates that the component ρ_k can be split into further significant components. For this reason, we require that Θ , the indivisibility parameter for the model to be large enough, guaranteeing that the components in the model cannot be divided into further meaningful components. This follows from the observation that if there exists a $k \in \{1, \dots, N\}$ such that ρ_k can be split into separate components, then the second eigenvalue of Δ_{ρ_k} will be small resulting in Θ being small as well.

The coupling parameter $\mathcal{C}$ being small can be interpreted as a metastability condition on the relative entropy of measures $\rho_k dx$ with respect to the measure ρdx . We define the relative entropy of a probability measure ϱdx with respect to ρdx by

$$H(\varrho|\rho) := \int_{\mathcal{M}} \left(\frac{\varrho}{\rho} \right) \log \left(\frac{\varrho}{\rho} \right) \rho dx = \int_{\mathcal{M}} \varrho (\log \varrho - \log \rho) dx.$$

Given a probability density ρ , consider $\varrho(t, x)$ varying with time that satisfies the evolution equation

$$\begin{aligned} \partial_t \varrho &= \Delta \varrho + \operatorname{div}(\varrho \nabla \log \rho) \\ &= \operatorname{div}(\varrho \nabla (\log \varrho - \log \rho)), \end{aligned} \tag{21}$$

with initial condition $\varrho(0, x) = \rho_k(x)$ for a fixed $k \in \{1, \dots, N\}$. Model (21) is a linear Fokker-Planck equation that is driven by the competition between linear diffusion and a confining drift term given by the potential $\log \rho$. Observe that ρ is a stationary state of the

system, and the quantity $H(\varrho(t)|\rho)$ gives a sense of how far $\varrho(t, x)$ is from the stationary state $\rho(x)$ at time $t \geq 0$. In this context, the coupling parameter $\mathcal{C}_k$ can be interpreted as an initial rate of decay for the entropy. More precisely, following Boltzmann's H -theorem, the *relative Fisher Information* $I(\varrho|\rho)$ is defined as the entropy dissipation along solutions to (21),

$$\begin{aligned} \frac{d}{dt}H(\varrho(t)|\rho) &= \int_{\mathcal{M}} \partial_t \varrho (\log \varrho - \log \rho) \, dx \\ &= - \int_{\mathcal{M}} \varrho |\nabla (\log \varrho - \log \rho)|^2 \, dx =: -I(\varrho(t)|\rho), \end{aligned}$$

and so the definition of the coupling parameter (11) can be interpreted as

$$\mathcal{C}_k = \frac{1}{4}I(\varrho(0)|\rho) = \frac{1}{4} \int_{\mathcal{M}} \left| \frac{\nabla \rho_k}{\rho_k} - \frac{\nabla \rho}{\rho} \right|^2 \rho_k \, dx.$$

Approximating the entropy for small initial times $t > 0$,

$$H(\varrho(t)|\rho) = H(\rho_k|\rho) - 4\mathcal{C}_k t + O(t^2),$$

we observe that if the coupling parameter $\mathcal{C}_k$ is small, then $H(\varrho(t)|\rho)$ only varies very slowly in a neighborhood of $t = 0$, and so we can consider $H(\rho_k|\rho)$ to be in a metastable state. For a well-separated mixture model, we expect the initial entropy $H(\rho_k|\rho)$ to be large enough as small entropy would indicate that ρ_k takes values close to ρ (note that $H(\varrho|\rho) = 0$ if and only if $\varrho = \rho$ almost everywhere). In other words, having a small overlapping parameter $\mathcal{S}$ indicates that the initial entropy $H(\rho_k|\rho)$ cannot be too small. For well-separatedness, however, we require that $\mathcal{C}_k$ is small for all $k \in \{1, \dots, N\}$, that is to say that the Fisher Information is small initially when starting the evolution process (21) at any of the components ρ_k . In this sense, we require the mixture model to be in a metastable state.

Having discussed the parameters $\mathcal{S}$, Θ and $\mathcal{C}$, we now give some intuition on the notion of a well-separated mixture as introduced in Definition 4. Roughly speaking, a mixture model is well-separated if it has small coupling and overlapping parameters $\mathcal{C}$ and $\mathcal{S}$, but a large indivisibility parameter Θ . A natural question is whether it is necessary to require both $\mathcal{C}$ and $\mathcal{S}$ to be small. Intuitively, $\mathcal{S}$ is small when the components ρ_j are concentrated on disjoint sets that are far apart. In such a setting one expects the coupling parameter $\mathcal{C}$ to also be small. However, we present two examples where $\mathcal{C}$ and $\mathcal{S}$ are not necessarily simultaneously small, demonstrating that it is not enough to only impose a small $\mathcal{C}$ (Example 1) or a small $\mathcal{S}$ (Example 2). The examples were chosen to illustrate how to estimate the parameters $\mathcal{S}$, $\mathcal{C}$ and Θ in some concrete settings, and to see that our mathematical definitions do capture our intuitive interpretation of what well-separated components should look like.

Example 1 (Mixture of two Gaussians) *Consider a mixture of two standard Gaussian densities on $\mathbb{R}^2$ obtained by shifting the two densities. More precisely, let $\rho = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$, where*




Figure 1: Cross section of a 2D mixture model with Gaussian components.

$$\rho_1(x) := \frac{1}{\sqrt{2\pi}} e^{-|x|^2/2}, \quad \rho_2(x) := \frac{1}{\sqrt{2\pi}} e^{-|x-\gamma|^2/2}$$

for a fixed vector $\gamma \in \mathbb{R}^2$. A cross section of the mixture model along the vector γ is illustrated in Figure 1. We estimate how the overlapping and coupling parameters of the model scale with the size of the off-set γ . Intuitively, the components become better separated as $|\gamma|$ increases and we will show that both $\mathcal{C}$ and $\mathcal{S}$ become small in this regime. However, as $|\gamma|$ becomes small the coupling parameter $\mathcal{C}$ decreases once more while the overlapping parameter $\mathcal{S}$ increases. This example demonstrates why we require both $\mathcal{C}$ and $\mathcal{S}$ to be small.

Observe that the indivisibility parameter is unaffected by shifts in space of the components, i.e., Θ does not depend on γ . On the other hand, straightforward calculations show

$$\frac{\nabla \rho_1(x)}{\rho_1(x)} = -x, \quad \frac{\nabla \rho_2(x)}{\rho_2(x)} = -(x - \gamma),$$

and that

$$\frac{\nabla \rho(x)}{\rho(x)} = -x + \frac{\gamma \rho_2(x)}{2 \rho(x)} = -(x - \gamma) - \frac{\gamma \rho_1(x)}{2 \rho(x)}.$$

In particular,

$$\frac{\nabla \rho(x)}{\rho(x)} - \frac{\nabla \rho_1(x)}{\rho_1(x)} = \frac{\gamma \rho_2(x)}{2 \rho(x)}, \quad \frac{\nabla \rho(x)}{\rho(x)} - \frac{\nabla \rho_2(x)}{\rho_2(x)} = -\frac{\gamma \rho_1(x)}{2 \rho(x)}.$$

It then follows from (11) that

$$\mathcal{C} \leq \mathcal{C}_1 + \mathcal{C}_2 = \frac{|\gamma|^2}{16} \int_{\mathbb{R}^2} \frac{\rho_1 \rho_2}{\rho^2} (\rho_1 + \rho_2) dx = \frac{|\gamma|^2}{8} \int_{\mathbb{R}^2} \frac{\rho_1 \rho_2}{\rho} dx = \frac{|\gamma|^2}{8} \mathcal{S}.$$

Thus, the coupling parameter $\mathcal{C}$ is controlled by the overlapping parameter $\mathcal{S}$ and $|\gamma|^2$. It follows that when $|\gamma|$ is close to zero, the overlapping parameter is close to one since both ρ_1 and ρ_2 are roughly equal to ρ , and that the coupling parameter is small of order $|\gamma|^2$. On the other hand, for large $|\gamma|$, the overlapping parameter decays exponentially fast and the coupling parameter decays as well. From an intuitive point of view it is to be expected that for large $|\gamma|$ the model is well separated.

It is straightforward to verify that, for large values of $|\gamma|$ the parameter $\mathcal{S} \approx \alpha_1 \exp(-\beta_1 |\gamma|^2)$ and it follows from the above arguments that $\mathcal{C}/\Theta \approx \alpha_2 \exp(-\beta_2 |\gamma|^2)$ for constants $\alpha_i, \beta_i > 0$, $i = 1, 2$. Now consider the bound (16) and fix $s, \sigma > 0$. Then a formal calculation yields

$t \geq \alpha_3 \exp(-\beta_3 |\gamma|^2)$ provided that $|\gamma|$ is sufficiently large. Let F denote the Laplacian embedding obtained from the first two eigenfunctions of Δ_ρ and let $\mu = F_\# \nu$. Since δ^* vanishes like $\mathcal{S}$ as well we infer the probability mass of μ outside of its orthogonal cones should vanish like $\alpha_4 \exp(-\beta_4 |\gamma|^2)$.

We verify the sharpness of the above formal calculation, up to the constants α_4, β_4 , for a mixture of Gaussian distributions in two dimensions. Let $\rho = \frac{1}{2}\rho_1 + \frac{1}{2}\rho_2$ where $\rho_1 = N(0, 0.04I)$ and $\rho_2 = N(\gamma, 0.04I)$ with I denoting the identity matrix in $\mathbb{R}^{2 \times 2}$. We then generate a random graph by drawing $n = 2^{11}$ vertices from ρ and construct the graph Laplacian matrix Δ_n as in (5) by setting $\varepsilon = 2 \frac{(\log n)^{3/4}}{n^{1/2}} \approx 0.2$. We then construct the discrete Laplacian embedding F_n using the first two eigenvectors of Δ_n and compute the number n_{outside} of embedded points that fall outside the cones $C((1, -1)^T, \pi/4 - 1/4, 0)$ and $C((1, 1)^T, \pi/4 - 1/4, 0)$. We choose $|\gamma| \in \{0.7, 0.8, 0.9, 1, 1.1, 1.2\}$ and repeat the experiment over 20 trials for each value of $|\gamma|$ where the vertices are redrawn from the mixture. We report the averaged values of n_{outside}/n , as an empirical approximation to the probability mass outside the cones, versus $|\gamma|^2$ in Figure 2(a) indicating that indeed $-\log(n_{\text{outside}}/n) = \mathcal{O}(|\gamma|^2)$ as expected.




Figure 2: Empirical approximation of the probability mass of $\mu = F_\# \nu$ lying outside of the orthogonal cones $C((1, -1)^T, \pi/4 - 1/4, 0)$ and $C((1, 1)^T, \pi/4 - 1/4, 0)$. ν is a mixture model of two components and F is the Laplacian embedding. Weighted graphs with $n = 2^{11}$ vertices were constructed with the vertices distributed according to ν and F was approximated using the eigenvectors of the graph Laplacian matrix Δ_n constructed using as kernel $\eta(r) = \chi_{B_1}(r)/\pi$, the indicator function on the unit ball. (a) Here ν is a mixture of two Gaussians as in Example 1 and the figure shows the fraction n_{outside}/n of embedded vertices that fall outside of the cones versus $|\gamma|^2$, where $|\gamma|$ denotes the distance between the component means. (b) ν is the uniform measure on a dumbbell shaped domain as in Example 2 with components as in Figure 3(a) and the figure shows the fraction n_{outside}/n of embedded vertices that fall outside of the cones versus ϑ the width of the thin channel between the two components of the dumbbell.

Example 2 (Partitioning a dumbbell shaped domain) We consider a probability measure ρ on a dumbbell shaped domain $\mathcal{M} \subset \mathbb{R}^2$. More precisely define the sets

$$\begin{aligned}\mathcal{M}_0 &= \{(x_1, x_2) : -\ell/2 < x_1 < \ell/2, \quad |x_2| < \vartheta\}, \\ \mathcal{M}_1 &= \{(x_1, x_2) : \ell/2 < x_1 < \ell, \quad |x_2| < \ell/2\}, \\ \mathcal{M}_2 &= \{(x_1, x_2) : -\ell < x_1 < -\ell/2, \quad |x_2| < \ell/2\},\end{aligned}$$

for $\ell, \vartheta > 0$ and define $\mathcal{M} := \bigcup_{j=0}^2 \mathcal{M}_j \subset \mathbb{R}^2$. We have in mind $\vartheta \ll \ell$ so that $\mathcal{M}$ is a dumbbell shaped domain consisting of two rectangles connected by a thin rectangle.

We assume $|\mathcal{M}| = 1$, i.e. $\ell^2 + 2\ell\vartheta = 1$. We take $\rho = 1$ on $\mathcal{M}$ and then define the mixture elements ρ_k by partitioning $\mathcal{M}$ into subsets and taking the ρ_k to be a partition of unity for $\mathcal{M}$ with appropriate normalization. We present two choices of partitions and demonstrate that for small ϑ , one choice results in well-separated densities ρ_k while the other partition does not.

Our first partition aims at cutting the dumbbell perpendicularly to its axis (see Figure 3(a)). For $0 < \varepsilon < \ell/2$ define the function $\psi : \mathbb{R} \rightarrow \mathbb{R}$

$$\psi(t) := \begin{cases} 1 & \text{if } t \leq 0 \\ 1 - \frac{2}{\varepsilon^2} t^2 & \text{if } t \in [0, \varepsilon/2] \\ \frac{2(t-\varepsilon)^2}{\varepsilon^2} & \text{if } t \in [\varepsilon/2, \varepsilon] \\ 0 & \text{if } t \geq \varepsilon \end{cases}.$$

We let $\tilde{\rho}_1, \tilde{\rho}_2 : \mathcal{M} \rightarrow \mathbb{R}$,

$$\tilde{\rho}_1(x) := \psi(x_1), \quad \tilde{\rho}_2(x) := \psi(-x_1), \quad x = (x_1, x_2) \in \mathcal{M}.$$

The densities ρ_1 and ρ_2 are then defined by

$$\rho_1(x) := \frac{2\tilde{\rho}_1(x)}{\tilde{\rho}_1(x) + \tilde{\rho}_2(x)}, \quad \rho_2(x) := \frac{2\tilde{\rho}_2(x)}{\tilde{\rho}_1(x) + \tilde{\rho}_2(x)}, \quad (22)$$

from where we see that $\rho(x) = \frac{1}{2}\rho_1(x) + \frac{1}{2}\rho_2(x)$ with weights $w_1 = w_2 = \frac{1}{2}$ in the notation of (1).

We now show that the above mixture model is well-separated according to Definition 4 for an appropriate choice of $\varepsilon(\vartheta)$ as $\vartheta \rightarrow 0$. Let us begin by estimating the parameters $\mathcal{S}, \mathcal{C}$ and Θ . Following (9) we have

$$\begin{aligned}\mathcal{S} &= \int_{\mathcal{M}} \rho_1(x) \rho_2(x) dx \\ &\leq 4 \int_{\mathcal{M}} \tilde{\rho}_1(x) \tilde{\rho}_2(x) dx \\ &\leq 16\vartheta\varepsilon\end{aligned} \quad (23)$$

where in the first inequality we have used the fact that

$$1 \leq \tilde{\rho}_1(x) + \tilde{\rho}_2(x) \leq 2.$$

Thus $\mathcal{S}$ can be made arbitrarily small by either taking ϑ or ε to zero. Furthermore, by (11) one can see that

$$\mathcal{C} = \max_k \frac{1}{4} \int_{\mathcal{M}} \left| \frac{\nabla \rho_k(x)}{\rho_k(x)} \right|^2 \rho_k(x) dx \leq 4\vartheta \int_0^\varepsilon \left| \frac{\psi'(t)}{\psi(t)} \right|^2 \psi(t) dt \leq \frac{4\vartheta}{\varepsilon}. \quad (24)$$

This implies that if we pick

$$\vartheta \ll \varepsilon,$$

then we can make $\mathcal{C}$ as small as we want. It remains to show that the indivisibility parameter Θ remains bounded away from zero for the right choice of parameters. In order to find a lower bound for Θ it will be convenient to use Cheeger's inequality, which gives us a lower bound for Θ in terms of a geometric quantity that is much easier to understand and estimate. Namely, by Theorem 43,

$$\Theta_k = \min_{u \perp \mathbb{1}} \frac{\int_{\mathcal{M}} |\nabla u|^2 \rho_k dx}{\int_{\mathcal{M}} u^2 \rho_k dx} \geq \frac{(h(\mathcal{M}, \rho_k))^2}{4}, \quad (25)$$

where in the above $\perp$ means orthogonal with respect to the inner product $\langle \cdot, \cdot \rangle_{\rho_k}$ and

$$h(\mathcal{M}, \rho_k) := \min_{A \subseteq \mathcal{M}} \frac{\int_{\partial A \cap \mathcal{M}} \rho_k(x) dS(x)}{\min \left\{ \int_A \rho_k dx, \int_{\mathcal{M} \setminus A} \rho_k dx \right\}}.$$

The inequality is due to Cheeger (1969), and for the convenience of the reader we provide in the Appendix a sketch of the proof that in particular highlights the specific structure of the Raleigh quotient associated to the operators Δ_{ρ_k} and Δ_ρ that makes the inequality work (see Section B.2). In particular we emphasize that for Cheeger's inequality to hold, both the numerator and denominator of the Raleigh quotient defining Θ_k must be weighted with the same measure (in our case both integrals are weighted by ρ_k).

Due to symmetry we have $\Theta = \Theta_1 = \Theta_2$. Thanks to Cheeger's inequality we can focus on obtaining a lower bound for $h(\mathcal{M}, \rho_1)$. Now, from the fact that $2 \geq \tilde{\rho}_1(x) + \tilde{\rho}_2(x) \geq 1$ we see that in order to get a lower bound for $h(\mathcal{M}, \rho_1)$ we can alternatively get a lower bound for $h(\mathcal{M}, \tilde{\rho}_1)$. After a moment of reflection, we can conclude that if we can find $C(\vartheta, \varepsilon)$ such that

$$\text{Cut}(A_t) := \frac{\int_{\partial A_t \cap \mathcal{M}} \tilde{\rho}_1(x) dS(x)}{\min \left\{ \int_{A_t} \tilde{\rho}_1 dx, \int_{\mathcal{M} \setminus A_t} \tilde{\rho}_1 dx \right\}} \geq C(\vartheta, \varepsilon), \quad \forall t \in [-l/2, \varepsilon],$$

where

$$A_t := \{x \in \mathcal{M} : x_1 \leq t\},$$

then

$$h(\mathcal{M}, \rho_1) \geq \min\{c, C(\vartheta, \varepsilon)\},$$

where c is a constant that does not depend on ϑ nor ε and corresponds to the cost of a cut along the x_1 axis. This is because other than the horizontal cut along the x_1 axis, the only other natural cuts that are viable competitors for minimizing Cut (the objective function in the definition of $h(\mathcal{M}, \tilde{\rho}_1)$) are the ones perpendicular to the bottleneck of the dumbbell

and along the x_2 axis. A straight forward calculation shows that the cuts along the x_1 axis result in a value of Cut that is bounded away from zero and so we focus on finding $C(\vartheta, \varepsilon)$ by analyzing vertical cuts across the bottleneck. Notice that we do not have to consider the case $t > \varepsilon$ since the density $\tilde{\rho}_1$ is equal to zero for $x = (x_1, x_2)$ with $x_1 > \varepsilon$. In other words we just have to look at non-trivial subsets of $\{x \in \mathcal{M} : \tilde{\rho}_1(x) > 0\}$.

For all $t \in [-l/2, \varepsilon]$ we have $\partial A_t \cap \mathcal{M} = \{(x_1, x_2) \in \mathcal{M} : x_1 = t, |x_2| < \vartheta\}$, and so we can write

$$Cut(A_t) = \frac{\psi(t)}{\int_t^\varepsilon \psi(s) ds}.$$

Now, when ε is small, for every $t \in [-l/2, \varepsilon/2]$ we have

$$Cut(A_t) \geq \frac{1}{2 \int_t^\varepsilon \psi(s) ds} \geq 1.$$

On the other hand, we notice that by the definition of ψ we have

$$\psi'(t) \int_t^\varepsilon \psi(s) ds + (\psi(t))^2 = \frac{4(\varepsilon - t)^4}{3\varepsilon^4} \geq 0, \quad \forall t \in [\varepsilon/2, \varepsilon],$$

implying that the function $Cut(A_t)$ is increasing between $\varepsilon/2$ and ε . In particular, it follows that for all $t \in [\varepsilon/2, \varepsilon]$ we have

$$Cut(A_t) \geq Cut(A_{\varepsilon/2}) = \frac{1}{2 \int_{\varepsilon/2}^\varepsilon \psi(s) ds} \geq \frac{1}{\varepsilon}.$$

Hence $C(\vartheta, \varepsilon) = \min\{1, 1/\varepsilon\}$. We conclude that provided $\varepsilon \leq 1$ we have that

$$\Theta \geq c,$$

for some positive constant that does not depend on ε or ϑ . The bottom line is that the model is well-separated as long as $\vartheta \ll \varepsilon$ and ε is sufficiently small.

One can then consider a partition that aims at cutting the dumbbell along its axis (see Figure 3(b)). In particular we can introduce analogue densities ρ_1 and ρ_2 smoothening out the indicator functions for the sets in the new partitioning (in this case the smoothening happens in the x_2 coordinate). Then using similar calculations as before we can show that the overlapping parameter can be made arbitrarily small (by selecting the parameter ε to be small), however, it will not be possible to control the indivisibility parameter from below as $\vartheta \rightarrow 0$. In fact, we expect the second eigenvalue of Δ_{ρ_k} to go to zero with ϑ following Anné (1995). Moreover, the coupling parameter will not be controlled by ϑ in this case, and in fact as $\varepsilon \rightarrow 0$ it will blow up. Therefore according to our definition this mixture model is not well-separated.

Notice that this example has made explicit the fact that the notion of well separated mixture model does not depend on the density ρ , but on the actual decomposition as a mixture model. The example has also exhibited that our mathematical notion of well-separated mixture model does capture our intuition.

Let us formally verify the lower bound (16) for this example similar to our numerical experiments in Example 1. Here, we consider a mixture model with the components defined

as indicated by Figure 3(a), see (22). We let $\ell = 1$ and consider dumbbell shaped domains with parameter $\vartheta \in \{.05, .1, .15, .2, .25, .3, .4, .5\}$. We then let ρ denote the uniform measure on the dumbbell shaped domain $\mathcal{M}$ and generate random graphs with $n = 2^{11}$ vertices sampled randomly from ρ . We choose the rest of the parameters in this experiment identical to the experiment in Example 1 and report the averaged values of n_{outside}/n versus ϑ over 20 trials in Figure 2(b), indicating a linear relationship between ϑ and the probability mass outside the orthogonal cones.

Let us now return to (16) with s fixed and split ρ into a mixture of two components as in (22) so that $N = 2$ and $w_{\max} = w_{\min} = 1/2$. Assuming our bounds on $\mathcal{S}, \mathcal{C}$ in (23) and (24) are sharp we can proceed to simplify the lower bound on t up to leading order assuming that $\vartheta \ll \varepsilon < 1$. A straightforward calculation then yields that $t \geq \alpha \vartheta^{1/2}$ for some constant $\alpha > 0$ which in turn suggests that the probability mass of μ lying outside orthogonal cones of a fixed angle should vanish like $\vartheta^{1/2}$. Thus we observe a gap between the lower bound in (16) and our numerical experiments, which suggest the probability mass outside of the cones vanishes like ϑ . We conclude that either the lower bound (16) is not sharp in this case or our bounds on $\mathcal{S}, \mathcal{C}$ in (23) and (24) can be improved.




Figure 3: Two partitioning strategies for the dumbbell shaped domain of Example 2. The partition (a) leads to well-separated clusters in the limit $\vartheta \rightarrow 0$, while the partition (b) does not. Clearly, the subsets in (b) can be further split into meaningful components.

The previous two examples have illustrated some aspects of the parameters $\mathcal{S}, \Theta, \mathcal{C}$ which we now put together in a remark.

Remark 16 *The parameters $\mathcal{S}, \mathcal{C}, \Theta$, can be interpreted as quantitative measures for the presence of three different types of degeneracies in the mixture model:*

- *Small $\mathcal{S}$ guarantees that there are no repeated components in the model. Notice that in Example 1 the two components of the model are the same when $\gamma = 0$. In that case, $\mathcal{C}$ would be zero and Θ would still be bounded away from zero.*

- *Large Θ guarantees that a component in the model can not be divided into further well separated components (e.g. as it would be the case in Figure 3(b)).*
- *Small $\mathcal{C}$ guarantees that our components are not slices of what should be considered one big component. Consider for example a one-dimensional mixture model with two components given by the halves of a centered Gaussian, each component equal to the part of the Gaussian lying to either side of the origin. Then the overlap $\mathcal{S}$ is zero and Θ is bounded away from zero but the coupling $\mathcal{C}$ would be large since it is easy for entropy to flow from one component to the other.*

In summary, $\mathcal{C}$ is a measure of how well entropy flows from any one given component to the rest of the distribution, which is essentially a measure of between class separation. Conversely, Θ measures the within class separation. For a clustering method to work we intuitively expect that one should have that the within class similarity is much greater than the between class similarity: this is exactly saying that the ratio $\mathcal{C}/\Theta$ should be small. Small $\mathcal{S}$ on the other hand guarantees that with high probability, data points are allocated to the correct cluster, which is a direct consequence of small overlap between components.

4.2 Extensions and generalizations of main results

As outlined in Section A the proofs of our main results are divided into two main parts.

1. We prove that the measure $F_{\sharp}\nu$ has an orthogonal cone structure provided the model is well-separated.
2. We show that after an orthogonal transformation, the measures $F_{\sharp}\nu$ and $F_{n\sharp}\nu_n$ are close to each other in the Wasserstein sense with high probability.

We emphasize that for part (1) we do not need to assume that $\mathcal{M}$ has compact support, nor that the density ρ is bounded below and above by positive constants. We only need Assumptions 7. In contrast, for part (2), we impose the extra conditions in Assumptions 11 to guarantee that the results of Burago et al. (2014); García Trillos et al. (2020) on the spectral convergence of graph Laplacians are applicable.

4.2.1 EXTENSIONS TO UNBOUNDED DOMAINS AND k -NN GRAPHS

In this paper Theorem 30 links the graph Laplacian Δ_n to the differential operator Δ_ρ . Assumption 11 is only used when proving Theorem 30, and in particular to show the existence of transport maps T_n between ν and ν_n satisfying

$$\|T_n - Id\|_\infty := \operatorname{ess\,sup}_{x \in \mathcal{M}} |T_n(x) - x| \rightarrow 0,$$

at an explicit rate in terms of n . For an unbounded manifold $\mathcal{M}$, *any* transport map T_n pushing forward ν into ν_n , must transport mass over arbitrary long distances (even a tiny bit of mass transported over a large distance already results in a large ∞ -OT distance). It is therefore of relevance to investigate other forms of establishing the discrete-continuum link that do not involve using ∞ -OT maps; this is the topic of current investigation.

It is also of relevance to extend our analysis to settings where different graph constructions are used to extract the coarse structure of the point cloud. In this paper we have focused on proximity graphs constructed using a kernel η and a connectivity parameter $\varepsilon > 0$, but other graph constructions are perhaps more popular among practitioners and enjoy higher regularity properties; such is the case of k nearest neighbors (k -NN) graphs. Results quantifying spectral convergence rates for k -NN graph Laplacians are largely missing in the literature, and it is interesting to extend the analysis showed in this paper to the k -NN setting. In García Trillos (2019), variational techniques have been used to study the statistical consistency of optimization problems defined on k -NN graphs.

4.2.2 DIFFERENT WEIGHTED VERSIONS OF THE LAPLACIAN

The kernelized Laplacian considered in this paper is just one of many operators that can be used to construct the embedding F . In light of the notion of *diffusion maps* introduced in Coifman and Lafon (2006), it seems reasonable to ask what is the effect of different graph Laplacian normalizations on the geometry of the embedded data. For example, one may consider a family of renormalizations, where for a given $\gamma \in [0, 1]$, we let

$$(\mathbf{W}_n)_{ij} := \frac{\eta_\varepsilon(|\mathbf{x}_i - \mathbf{x}_j|)}{\tilde{d}_i^\gamma \tilde{d}_j^\gamma}, \quad \tilde{d}_i := \sum_{j=1}^n \eta_\varepsilon(|\mathbf{x}_i - \mathbf{x}_j|).^1$$

The weights $(\mathbf{W}_n)_{ij}$ induce the family of normalized graph Laplacians

$$\Delta_n^\gamma := I - \mathbf{D}_n^{-1} \mathbf{W}_n$$

where $\mathbf{D}_n = \text{diag}(d_i)$, $d_i = \sum_j (\mathbf{W}_n)_{ij}$, denotes the degree matrix of the weight matrix $\mathbf{W}_n$.

Intuitively, the parameter γ controls the effect of the density ρ on the clustering: for γ close to one, the effect of ρ is minimal, and the limiting clustering is influenced only by the geometry of $\mathcal{M}$; on the other hand, when γ is close to zero, the effect of the ground-truth density ρ is maximal (see Coifman and Lafon (2006)). This observation can be made rigorous using the ideas presented in the work García Trillos and Slepčev (2018) which when applied would show that after appropriate rescaling, the spectrum of the graph Laplacian Δ_n^γ approximates that of the formally defined differential operator

$$\Delta_\rho^\gamma u := -\frac{1}{\rho^{2(1-\gamma)}} \text{div}(\rho^{2(1-\gamma)} \nabla u).$$

Furthermore, following the discussion in Section A.3 one can bootstrap the ideas presented in García Trillos et al. (2020) and obtain quantitative error estimates relating the spectra of Δ_n^γ and Δ_ρ^γ analogous to the ones that appear in Theorem 30. We notice that the limiting differential operator Δ_ρ^γ is of the same type as the one we have considered in this paper, but where now we think of the underlying density function as being proportional to $\rho^{2(1-\gamma)}$. That is, we can introduce the density

$$\tilde{\rho}(x) := \frac{\rho^\theta(x)}{\int_{\mathcal{M}} \rho^\theta(y) dy},$$

1. Formally, one could take any $\gamma \in \mathbb{R}$, but we follow here the choice of γ as in Coifman and Lafon (2006).

and notice that

$$\Delta_\rho^\gamma = \Delta_{\tilde{\rho}},$$

for $\theta = 2(1-\gamma)$. Based on the analysis presented in this paper, it is then possible to describe the geometry of the graph Laplacian embeddings constructed from Δ_n^γ , assuming that the modified density $\tilde{\rho}$ admits a representation as a well-separated mixture model.

An immediate question that arises is to understand whether there is an interesting criterion under which one can choose the value of γ that produces the “best” clustering for a given data set. Alternatively, instead of looking at a single value of γ , it maybe worth working with the full ensemble of Laplacian embeddings (for all values of γ), and determine how to use it as best as possible. The formulation and analysis of these questions are the topic of current investigation.

Remark 17 *Note that the so called random walk graph Laplacian Shi and Malik (2000) coincides with Δ_n^0 (i.e. $\gamma = 0$). The limiting differential operator then takes the form*

$$\Delta_\rho^0 = -\frac{1}{\rho^2} \operatorname{div}(\rho^2 \nabla u).$$

Remark 18 *Asymptotically, the kernelized graph Laplacian Δ_n studied here coincides with $\Delta_n^{1/2}$ (i.e. $\gamma = 1/2$). In particular, there is no reason to work with $\Delta_n^{1/2}$ (i.e. renormalizing twice). In addition, the limiting differential operator $\Delta_\rho^{1/2} = \Delta_\rho$ has the following properties. First,*

$$\langle \Delta_\rho u, v \rangle_\rho = \int_{\mathcal{M}} \nabla u \cdot \nabla v \rho dx,$$

is linear in ρ . Secondly,

$$\langle u, v \rangle_\rho = \int_{\mathcal{M}} uv \rho dx,$$

is also linear in ρ , and Δ_ρ is self-adjoint with respect to $\langle \cdot, \cdot \rangle_\rho$. Thirdly, under fairly general conditions, such as when ρ is a well-separated mixture, Δ_ρ has a spectral gap after the N -th eigenvalue.

Acknowledgments

The authors would like to thank Ulrike von Luxburg for pointing them to the paper Schiebinger et al. (2015) which was the starting point of this work. NGT was supported by NSF grant DMS 1912802. FH was partially supported by Caltech’s von Karman post-doctoral instructorship. BH is supported in part by a postdoctoral fellowship granted by Natural Sciences and Engineering Research Council of Canada.

References

Colette Anné. A note on the generalized dumbbell problem. *Proceedings of the American Mathematical Society*, 123(8):2595–2599, 1995. doi: 10.1090/S0002-9939-1995-1257096-5.

- Hedy Attouch, Giuseppe Buttazzo, and Gérard Michaille. *Variational analysis in Sobolev and BV spaces: Applications to PDEs and optimization*. SIAM, Philadelphia, 2014. doi: 10.1137/1.9781611973488.
- Mikhail Belkin and Partha Niyogi. Towards a theoretical foundation for Laplacian-based manifold methods. *Journal of Computer and System Sciences*, 74(8):1289–1308, 2008. doi: 10.1016/j.jcss.2007.08.006.
- Vladimir I. Bogachev. *Measure Theory*, volume 1. Springer, New York, 2007. doi: 10.1007/978-3-540-34514-5.
- Dmitri Burago, Sergei Ivanov, and Yaroslav Kurylev. A graph discretization of the Laplace-Beltrami operator. *Journal of Spectral Theory*, 4(4):675–714, 2014. doi: 10.4171/JST/83.
- Jeff Calder and Nicolás García Trillos. Improved spectral convergence rates for graph Laplacians on epsilon-graphs and k-NN graphs. *arXiv preprint:1910.13476*, 2019. URL <https://arxiv.org/abs/1910.13476>.
- Kamalika Chaudhuri, Fan Chung, and Alexander Tsiatas. Spectral clustering of graphs with general degrees in the extended planted partition model. In *Proceedings of the 25th Annual Conference on Learning Theory*, pages 35.1–35.23, 2012. URL <http://proceedings.mlr.press/v23/chaudhuri12.html>.
- Jeff Cheeger. A lower bound for the smallest eigenvalue of the laplacian. In *Proceedings of the Princeton conference in honor of Professor S. Bochner*, pages 195–199, 1969.
- Fan R. K. Chung. *Spectral graph theory*. Number 92 in CBMS: Regional Conference Series in Mathematics. American Mathematical Society, Providence, 1997. doi: 10.1090/cbms/092.
- Ronald R. Coifman and Stéphane Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. doi: 10.1016/j.acha.2006.04.006.
- Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.
- Matthew M Dunlop, Dejan Slepčev, Andrew M Stuart, and Matthew Thorpe. Large data and zero noise limits of graph-based semi-supervised learning algorithms. *Applied and Computational Harmonic Analysis*, 2019. doi: 10.1016/j.acha.2019.03.005.
- Nicolás García Trillos. Variational limits of k-NN graph-based functionals on data clouds. *SIAM Journal on Mathematics of Data Science*, 1(1):93–120, 2019. doi: 10.1137/18M1188999.
- Nicolás García Trillos and Ryan Murray. A new analytical approach to consistency and overfitting in regularized empirical risk minimization. *European Journal of Applied Mathematics*, 28(6):886–921, 2017. doi: 10.1017/S0956792517000201.
- Nicolás García Trillos and Daniel Sanz-Alonso. Continuum limits of posteriors in graph Bayesian inverse problems. *SIAM Journal on Mathematical Analysis*, 50(4):4020–4040, 2018. doi: 10.1137/17M1138005.

- Nicolás García Trillos and Dejan Slepčev. Continuum limit of total variation on point clouds. *Archive for Rational Mechanics and Analysis*, pages 1–49, 2015. doi: 10.1007/s00205-015-0929-z.
- Nicolás García Trillos and Dejan Slepčev. A variational approach to the consistency of spectral clustering. *Applied and Computational Harmonic Analysis*, 45(2):239–281, 2018. doi: 10.1016/j.acha.2016.09.003.
- Nicolás García Trillos, Moritz Gerlach, Matthias Hein, and Dejan Slepčev. Error estimates for spectral convergence of the graph Laplacian on random geometric graphs toward the Laplace-Beltrami operator. *Foundations of Computational Mathematics*, 20(4):827–887, 2020. doi: 10.1007/s10208-019-09436-w.
- Matthias Hein, Jean-Yves Audibert, and Ulrike von Luxburg. From graphs to manifolds—weak and strong pointwise consistency of graph Laplacians. In *Learning theory*, volume 3559 of *Lecture Notes in Computer Science*, pages 470–485. Springer, Berlin, 2005. doi: 10.1007/11503415_32.
- Nicholas J Higham. Matrix nearness problems and applications. In *Applications of matrix theory*, pages 1–27. Oxford University Press, 1989. doi: 10.2307/3619410.
- Franca Hoffmann, Bamdad Hosseini, Assad A Oberai, and Andrew M Stuart. Spectral analysis of weighted laplacians arising in data clustering. *arXiv preprint:1909.06389*, 2019. URL <https://arxiv.org/abs/1909.06389>.
- Shuyang Ling and Thomas Strohmer. Certifying global optimality of graph cuts via semidefinite relaxation: A performance guarantee for spectral clustering. *Foundations of Computational Mathematics*, pages 1–55, 2019. doi: 10.1007/s10208-019-09421-3.
- Anna V Little, Mauro Maggioni, and James M Murphy. Path-based spectral clustering: Guarantees, robustness to outliers, and fast algorithms. *Journal of Machine Learning Research*, 21(6):1–66, 2020.
- Karl Rohe, Sourav Chatterjee, and Bin Yu. Spectral clustering and the high-dimensional stochastic blockmodel. *Annals of Statistics*, 39(4):1878–1915, 08 2011. doi: 10.1214/11-AOS887.
- Geoffrey Schiebinger, Martin J. Wainwright, and Bin Yu. The geometry of kernelized spectral clustering. *Annals of Statistics*, 43(2):819–846, 04 2015. doi: 10.1214/14-AOS1283.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, August 2000. doi: 10.1109/34.868688.
- Dejan Slepčev and Matthew Thorpe. Analysis of p-laplacian regularization in semisupervised learning. *SIAM Journal on Mathematical Analysis*, 51(3):2085–2120, 2019. doi: 10.1137/17M115222X.

Wenqi Tao and Zuoqiang Shi. Convergence of Laplacian spectra from random samples. *Journal of Computational Mathematics*, 38(6):952–984, 2020. ISSN 1991-7139. doi: 10.4208/jcm.2008-m2018-0232.

Cédric Villani. *Topics in optimal transportation*. American Mathematical Society, Providence, RI, 2003. doi: 10.1007/b12016.

Ulrike von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. doi: 10.1007/s11222-007-9033-z.

Ulrike von Luxburg, Mikhail Belkin, and Olivier Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, 36(2):555–586, 2008. doi: 10.1214/009053607000000640.

Appendix A. Proofs of main results

Here we collect the detailed proofs of the main theoretical results of this article. In Appendix A.1 we show stability of orthogonal cone structures under perturbations in the Wasserstein distance, proving that nearby probability measures have similar orthogonal cone structures.

Appendix A.2 is dedicated to the proof of our first main result Theorem 10 stating that the pushforward $F_{\sharp}\nu$ has an orthogonal cone structure when ν is well-separated. In proving this theorem we also obtain some auxiliary results regarding the spectrum of Δ_{ρ} that are interesting by themselves.

In Appendix A.3 we collect convergence results on the spectrum of Δ_n and Δ_{ρ} and prove that as $n \rightarrow \infty$ and $\varepsilon \rightarrow 0$ in an appropriate regime then the spectrum of Δ_n converges to that of Δ_{ρ} and obtain quantitative rates.

We then combine the above results in Appendix A.4 to prove Theorem 15, showing that the point cloud $F_n(\mathcal{M}_n)$ has an orthogonal cone structure when ν is well-separated.

A.1 Stability of orthogonal cone structures

In this section we prove some basic results on the stability of orthogonal cone structures. The first result follows immediately from the characterization of weak convergence of probability measures via Portmanteau’s theorem (Bogachev, 2007, Theorem 8.4.7).

Proposition 19 *Let $\{\mu_n\}_{n \in \mathbb{N}}$ be a sequence of Borel probability measures on $\mathbb{R}^k$ that converge weakly to a measure $\mu \in \mathcal{P}(\mathbb{R}^k)$. If μ has an orthogonal cone structure with parameters (σ, δ, r) , then for every $\varepsilon \in (0, 1 - \delta)$, there exists $K \in \mathbb{N}$ such that for all $n \geq K$, μ_n has a orthogonal cone structure with parameters $(\sigma, \delta + \varepsilon, r)$.*

Proof Consider the open set $U := \cup_{j=1}^k C(e_j, \sigma, r)$, where the sets $C(e_j, \sigma, r)$ are as in Definition 8 on the orthogonal cone structure property for μ . Since $\mu_n \xrightarrow{w} \mu$, Portmanteau’s theorem implies that

$$\liminf_{n \rightarrow \infty} \mu_n(U) \geq \mu(U).$$

In particular, this means that for every $\varepsilon \in (0, 1 - \delta)$, there exists a K such that if $n \geq K$ then,

$$\mu_n(U) \geq \mu(U) - \varepsilon \geq 1 - \delta - \varepsilon.$$

■

An immediate corollary of the previous proposition is the following.

Corollary 20 *If a probability measure μ has an orthogonal cone structure with parameters (σ, δ, r) and μ_n is the empirical measure associated to n i.i.d. samples from μ , then $\forall \varepsilon \in (0, 1 - \delta)$, with probability one there exists $K = K(\varepsilon, \delta) \in \mathbb{N}$ such that for every $n \geq K$, μ_n has a $(\sigma, \delta + \varepsilon, r)$ -cone structure.*

The idea is now to improve the previous asymptotic result and instead study the stability of orthogonal cone structures under small perturbations of a measure; we use the Wasserstein distance to measure the distance between different probability measures. We recall its definition.

Definition 21 *Let μ_1, μ_2 be two probability measures on $\mathbb{R}^k$ with finite second moments. We define their Wasserstein distance by*

$$(W_2(\mu_1, \mu_2))^2 := \min_{\pi \in \Gamma(\mu_1, \mu_2)} \int_{\mathbb{R}^k \times \mathbb{R}^k} |x - y|^2 d\pi(x, y),$$

where $\Gamma(\mu_1, \mu_2)$ stands for the set of transportation plans between μ_1 and μ_2 , that is, the set of probability measures in $\mathcal{P}(\mathbb{R}^k \times \mathbb{R}^k)$ with first and second marginals equal to μ_1 and μ_2 respectively.

Proposition 22 *Let $\mu_1, \mu_2 \in \mathcal{P}(\mathbb{R}^k)$ and suppose that μ_1 has an orthogonal cone structure with parameters (σ, δ, r) , where $\sigma < \pi/4$. Let $s, t > 0$ be such that*

$$\frac{rt \sin(s)}{\sqrt{k}} \geq W_2(\mu_1, \mu_2), \quad (26)$$

and such that $\sigma + s < \pi/4$. Then, μ_2 has an orthogonal cone structure with parameters $(\sigma + s, \delta + t^2, r(1 - \sin(s)))$.

Proof Denote by $C_1, \dots, C_k$ the “cones” associated to μ_1 and let $e_1, \dots, e_k$ be the orthonormal vectors they are centered at. For $s, t > 0$ satisfying (26) and such that $s + \sigma < \pi/4$, define

$$\tilde{C}_j := \left\{ z \in \mathbb{R}^k : \frac{z \cdot e_j}{|z|} > \cos(\tilde{\sigma}), \quad |z| > \tilde{r} \right\},$$

where $\tilde{\sigma} := \sigma + s$ and $\tilde{r} := r(1 - \sin(s))$. Let $\pi \in \Gamma(\mu_2, \mu_1)$ be an optimal transportation plan between μ_2 and μ_1 (which exists by (Villani, 2003, Proposition 2.1)). That is,

$$\int_{\mathbb{R}^k \times \mathbb{R}^k} |x - y|^2 d\pi(x, y) = (W_2(\mu_1, \mu_2))^2.$$

First, observe that from the definition of $\tilde{r}$, simple trigonometry shows that

$$\min_{v \in C_k, w \in \partial \tilde{C}_k} |v - w| = r \sin(s).$$

From this, it follows that

$$r^2 \sin^2(s) \pi(C_j, \mathbb{R}^k \setminus \tilde{C}_j) \leq \int_{C_j \times (\mathbb{R}^k \setminus \tilde{C}_j)} |x - y|^2 d\pi(x, y) \leq (W_2(\mu_1, \mu_2))^2.$$

Hence, for every $j = 1, \dots, k$,

$$\begin{aligned} \mu_2(\tilde{C}_j) &\geq \pi(C_j, \tilde{C}_j) \\ &= \pi(C_j, \mathbb{R}^k) - \pi(C_j, \mathbb{R}^k \setminus \tilde{C}_j) \\ &\geq \mu_1(C_j) - \frac{(W_2(\mu_1, \mu_2))^2}{r^2 \sin^2(s)} \\ &\geq \mu_1(C_j) - \frac{t^2}{k}. \end{aligned}$$

The assumption $\sigma < \pi/4$ and $\sigma + s < \pi/4$ implies that the sets $C_1, \dots, C_k$ and respectively $\tilde{C}_1, \dots, \tilde{C}_k$ are pairwise disjoint and thus

$$\mu_2\left(\bigcup_{j=1}^k \tilde{C}_j\right) \geq \mu_1\left(\bigcup_{j=1}^k C_j\right) - t^2 \geq 1 - (\delta + t^2).$$

■

We finish this section by recalling that a map $T : \mathbb{R}^k \rightarrow \mathbb{R}^k$ with $\mu_2 = T_{\#}\mu_1$ (i.e. μ_2 is the pushforward of μ_1 by T) is called a *transportation map* between μ_1 and μ_2 . Such a map induces a transportation plan $\pi_T \in \Gamma(\mu_1, \mu_2)$ given by

$$\pi_T := (Id \times T)_{\#}\mu_1,$$

where the map $Id \times T$ is given by $x \in \mathbb{R}^k \mapsto (x, T(x)) \in \mathbb{R}^k \times \mathbb{R}^k$. In particular, for any such T , we have

$$(W_2(\mu_1, \mu_2))^2 \leq \int_{\mathbb{R}^k \times \mathbb{R}^k} |x - y|^2 d\pi_T(x, y) = \int_{\mathbb{R}^k} |x - T(x)|^2 d\mu_1(x). \quad (27)$$

A.2 Proof of Theorem 10

In what follows we use $u_1, \dots, u_N$ to represent the first N eigenfunctions of Δ_ρ and denote by U their span. Normalized in $L^2(\rho)$, the functions $u_1, \dots, u_N$ form an orthonormal basis for U . We use $\Pi_N : L^2(\rho) \rightarrow U$ to denote the projection onto U . Denote by Q the span of the functions $q_1, \dots, q_N$. Our first proposition quantifies how close are the functions q_k to their projections onto U .

Proposition 23 *For every $k = 1, \dots, N$ we have*

$$\frac{1}{w_k} \|q_k - \Pi_N(q_k)\|_\rho^2 \leq \frac{\mathcal{C}}{\lambda_{N+1}}, \quad (28)$$

where Π_N stands for the projection onto U the span of the first N eigenfunctions of Δ_ρ and λ_{N+1} is the $(N+1)$ -st eigenvalue of Δ_ρ , where the eigenvalues are assumed to be indexed in increasing order.

Proof A direct computation using the definition of q_k shows that

$$\langle \Delta_\rho q_k, q_k \rangle_\rho = \int_{\mathcal{M}} |\nabla q_k(x)|^2 \rho(x) dx = w_k \mathcal{C}_k,$$

where $\mathcal{C}_k$ is as defined in (11). On the other hand, we can write q_k in the orthonormal basis of eigenfunctions $\{u_1, u_2, \dots\}$ of Δ_ρ as

$$q_k = \sum_{l=1}^{\infty} a_{lk} u_l,$$

for some coefficients $\{a_{lk}\}_{l \in \mathbb{N}}$. Using this representation we can alternatively write $\langle \Delta_\rho q_k, q_k \rangle_\rho$ as

$$\langle \Delta_\rho q_k, q_k \rangle_\rho = \sum_{l=1}^N a_{lk}^2 \lambda_l + \sum_{l=N+1}^{\infty} a_{lk}^2 \lambda_l.$$

We deduce that

$$w_k \mathcal{C} \geq w_k \mathcal{C}_k \geq \lambda_{N+1} \sum_{l=N+1}^{\infty} a_{lk}^2 = \lambda_{N+1} \|q_k - \Pi_N(q_k)\|_\rho^2.$$

■

We now focus on getting lower bounds for λ_{N+1} in terms of the parameters of the mixture model. In particular, we show that when $\mathcal{S}$ and $\mathcal{C}$ are small in relation to Θ then λ_{N+1} is large. We start with two preliminary results.

Lemma 24 *For every $j \in \{1, \dots, N\}$,*

$$1 - \mathcal{S} \leq \langle q_j, q_j \rangle_{\rho_j} \leq 1 + \mathcal{S}.$$

Remark 25 *It follows from the above lemma that $\|q_j\|_{\rho_j}$ converges to one as $\mathcal{S} \rightarrow 0$. In general however, $\mathcal{S}$ may be bigger than one, in which case the lower bound is trivial.*

Proof From the definition of q_j we see that

$$\langle q_j, q_j \rangle_{\rho_j} = 1 + \int_{\mathcal{M}} \left(\frac{w_j \rho_j}{\rho} - 1 \right) \rho_j dx.$$

On the other hand,

$$\begin{aligned} \left| \int_{\mathcal{M}} \left(\frac{w_j \rho_j}{\rho} - 1 \right) \rho_j dx \right| &\leq \int_{\mathcal{M}} |w_j \rho_j - \rho| \frac{\rho_j}{\rho} dx \\ &= \int_{\mathcal{M}} \sum_{k \neq j} w_k \frac{\rho_k \rho_j}{\rho} dx \\ &= \sum_{k \neq j} w_k \int_{\mathcal{M}} \frac{\rho_k \rho_j}{\rho} dx \\ &\leq \mathcal{S}. \end{aligned}$$

From the above computations we obtain the desired inequality. $\blacksquare$

Lemma 26 *For every $j \in \{1, \dots, N\}$,*

$$\inf_{\langle v, q_j \rangle_{\rho_j} = 0} \frac{\int_{\mathcal{M}} |\nabla v|^2 \rho_j dx}{\langle v, v \rangle_{\rho_j}} \geq \Theta(1 - \mathcal{S}), \quad (29)$$

Proof Let us fix $j \in \{1, \dots, N\}$ and pick $v \in H^1(\mathcal{M}, \rho_j)$, such that $\langle v, q_j \rangle_{\rho_j} = 0$, and $\langle v, v \rangle_{\rho_j} = 1$, where

$$H^1(\mathcal{M}, \rho_j) := \left\{ u \in L^2(\mathcal{M}, \rho_j) \mid \int_{\mathcal{M}} (|\nabla u|^2 + |u|^2) \rho_j dx < +\infty \right\}. \quad (30)$$

Observe that

$$\int_{\mathcal{M}} |\nabla v|^2 \rho_j dx = \sum_{k=1}^{\infty} \langle v, e_{j,k} \rangle_{\rho_j}^2 \lambda_{j,k},$$

where in the above $\{\lambda_{j,k}, e_{j,k}\}$ are the orthonormal (w.r.t. $\langle \cdot, \cdot \rangle_{\rho_j}$) eigenpairs of Δ_{ρ_j} . Since $\lambda_{j,1} = 0$, the first eigenvector is given by $e_{j,1} \equiv \mathbb{1}$, the function which is identically equal to one. Using the above equality and the Pythagorean theorem we conclude that

$$\begin{aligned} \int_{\mathcal{M}} |\nabla v|^2 \rho_j dx &\geq \lambda_{j,2} \sum_{k=2}^{\infty} \langle v, e_{j,k} \rangle_{\rho_j}^2 = \lambda_{j,2} \left(\langle v, v \rangle_{\rho_j}^2 - \langle v, e_{j,1} \rangle_{\rho_j}^2 \right) \\ &= \Theta_j \left(1 - \langle v, e_{j,1} \rangle_{\rho_j}^2 \right), \end{aligned} \quad (31)$$

where we recall Θ_j was defined in (13), and corresponds to the second eigenvalue of the operator Δ_{ρ_j} by the max-min formula (Attouch et al., 2014, Thm. 8.4.2).

We now find an upper bound for $\langle v, e_{j,1} \rangle_{\rho_j}^2$. Notice that

$$\langle v, e_{j,1} \rangle_{\rho_j} = \int_{\mathcal{M}} v \rho_j dx = \int_{\mathcal{M}} v(1 - q_j) \rho_j dx,$$

where the second equality follows from the fact that $\langle v, q_j \rangle_{\rho_j} = 0$. The Cauchy-Schwartz inequality implies that,

$$\langle v, e_{j,1} \rangle_{\rho_j}^2 \leq \|v\|_{\rho_j}^2 \int_{\mathcal{M}} (1 - q_j)^2 \rho_j dx = \int_{\mathcal{M}} (1 - q_j)^2 \rho_j dx = 1 + \int_{\mathcal{M}} (q_j^2 - 2q_j) \rho_j dx.$$

From the fact that $0 \leq q_j \leq 1$, we conclude that $(q_j^2 - 2q_j) \leq -q_j^2$, and therefore

$$\langle v, e_{j,1} \rangle_{\rho_j}^2 \leq 1 - \int_{\mathcal{M}} q_j^2 \rho_j dx = 1 - \langle q_j, q_j \rangle_{\rho_j} \leq 1 - (1 - \mathcal{S}) = \mathcal{S},$$

where the second inequality follows from Lemma 24. To conclude, the estimate (29) is a consequence of the above inequality and (31). $\blacksquare$

We are now in a position to find a lower bound for the $(N+1)$ -th eigenvalue λ_{N+1} of Δ_{ρ} .

Proposition 27 (Lower bound for λ_{N+1}) *Suppose that $N\mathcal{S} < 1$. Then,*

$$\left(\sqrt{\Theta(1 - N\mathcal{S})} - \frac{\sqrt{CN\mathcal{S}}}{(1 - \mathcal{S})} \right)^2 \leq \lambda_{N+1}.$$

Proof Let us consider $u \in H^1(\mathcal{M}, \rho)$ satisfying $\langle u, u \rangle_\rho = 1$, and $u \in Q^\perp$, i.e. u is orthogonal to the span of $\{q_1, \dots, q_N\}$ in terms of the inner product $\langle \cdot, \cdot \rangle_\rho$. From the fact that $\langle u, q_j \rangle_\rho = 0$ for all $j \in \{1, \dots, N\}$, we deduce that for every $j = 1, \dots, N$,

$$\langle u, q_j \rangle_{\rho_j} = \int_{\mathcal{M}} u q_j \rho_j dx = \frac{1}{w_j} \int_{\mathcal{M}} u q_j (w_j \rho_j - \rho) dx = -\frac{1}{w_j} \sum_{k \neq j} w_k \int_{\mathcal{M}} u q_j \rho_k dx.$$

Hence,

$$\begin{aligned} |\langle u, q_j \rangle_{\rho_j}| &\leq \frac{1}{w_j} \sum_{k \neq j} w_k \int_{\mathcal{M}} |u| q_j \rho_k dx \\ &\leq \frac{1}{w_j} \sum_{k \neq j} w_k \left(\int_{\mathcal{M}} u^2 \rho_k dx \right)^{1/2} \left(\int_{\mathcal{M}} q_j^2 \rho_k dx \right)^{1/2} \\ &\leq \sqrt{\frac{\mathcal{S}}{w_j}} \sum_{k \neq j} w_k \left(\int_{\mathcal{M}} u^2 \rho_k dx \right)^{1/2} \\ &\leq \sqrt{\frac{(1 - w_j)\mathcal{S}}{w_j}} \left(\sum_{k \neq j} w_k \int_{\mathcal{M}} u^2 \rho_k dx \right)^{1/2} \leq \sqrt{\frac{(1 - w_j)\mathcal{S}}{w_j}}, \end{aligned} \tag{32}$$

where in the second inequality we have used the Cauchy-Schwartz inequality, in the fourth inequality we have used Jensen's inequality, and in the last inequality the fact that $\langle u, u \rangle_\rho = 1$. Define now the functions

$$v_j := u - \left(\frac{\langle u, q_j \rangle_{\rho_j}}{\langle q_j, q_j \rangle_{\rho_j}} \right) q_j, \quad j = 1, \dots, N.$$

The function v_j is orthogonal to q_j with respect to the inner product $\langle \cdot, \cdot \rangle_{\rho_j}$. In addition,

$$\begin{aligned} \int_{\mathcal{M}} |\nabla v_j|^2 \rho_j dx &= \int_{\mathcal{M}} |\nabla u|^2 \rho_j dx - 2 \frac{\langle u, q_j \rangle_{\rho_j}}{\langle q_j, q_j \rangle_{\rho_j}} \int_{\mathcal{M}} \nabla q_j \cdot \nabla u \rho_j dx \\ &\quad + \left(\frac{\langle u, q_j \rangle_{\rho_j}}{\langle q_j, q_j \rangle_{\rho_j}} \right)^2 \int_{\mathcal{M}} |\nabla q_j|^2 \rho_j dx. \end{aligned}$$

Notice that,

$$\int_{\mathcal{M}} |\nabla q_j|^2 \rho_j dx = \frac{w_j}{4} \int_{\mathcal{M}} |\nabla(\rho_j/\rho)|^2 \rho dx = \frac{1}{4} \int_{\mathcal{M}} q_j^2 \left| \frac{\nabla \rho_j}{\rho_j} - \frac{\nabla \rho}{\rho} \right|^2 \rho_j dx \leq \mathcal{C}_j,$$

where in the last step we used the fact that $0 \leq q_j \leq 1$. From the above inequality, Lemma 24, (32) and Cauchy-Schwartz inequality, it follows that

$$\int_{\mathcal{M}} |\nabla v_j|^2 \rho_j dx \leq \int_{\mathcal{M}} |\nabla u|^2 \rho_j dx + 2 \left(\sqrt{\frac{1 - w_j}{w_j}} \right) \left(\frac{\sqrt{CN\mathcal{S}}}{(1 - \mathcal{S})} \right) \left(\int_{\mathcal{M}} |\nabla u|^2 \rho_j dx \right)^{1/2}$$

$$+ \frac{(1 - w_j)\mathcal{CS}}{w_j(1 - \mathcal{S})^2},$$

According to Lemma 26, the left hand side of the above expression can be bounded from below by $\Theta(1 - \mathcal{S})\langle v_j, v_j \rangle_{\rho_j}$ (since $\langle v_j, q_j \rangle_{\rho_j} = 0$). On the other hand,

$$\langle v_j, v_j \rangle_{\rho_j} = \langle u, u \rangle_{\rho_j} - \frac{\langle u, q_j \rangle_{\rho_j}^2}{\langle q_j, q_j \rangle_{\rho_j}} \geq \langle u, u \rangle_{\rho_j} - \frac{(1 - w_j)\mathcal{S}}{w_j(1 - \mathcal{S})},$$

as it follows from (32) and Lemma 24. Therefore,

$$\begin{aligned} & \Theta(1 - \mathcal{S})\langle u, u \rangle_{\rho_j} - \frac{(1 - w_j)\Theta\mathcal{S}}{w_j} \\ & \leq \int_{\mathcal{M}} |\nabla u|^2 \rho_j dx + 2 \left(\sqrt{\frac{1 - w_j}{w_j}} \right) \left(\frac{\sqrt{\mathcal{CS}}}{(1 - \mathcal{S})} \right) \left(\int_{\mathcal{M}} |\nabla u|^2 \rho_j dx \right)^{1/2} \\ & \quad + \frac{(1 - w_j)\mathcal{CS}}{w_j(1 - \mathcal{S})^2}. \end{aligned}$$

Multiplying both sides of the above inequality by w_j and adding over j we deduce

$$\Theta(1 - N\mathcal{S}) \leq \int_{\mathcal{M}} |\nabla u|^2 \rho dx + 2 \left(\frac{\sqrt{\mathcal{CS}}}{(1 - \mathcal{S})} \right) \sum_{j=1}^N \left(\int_{\mathcal{M}} |\nabla u|^2 w_j \rho_j dx \right)^{1/2} + \frac{\mathcal{CNS}}{(1 - \mathcal{S})^2}.$$

Applying Jensen's inequality and using the fact that $\langle u, u \rangle_{\rho} = 1$ we have

$$\begin{aligned} \Theta(1 - N\mathcal{S}) & \leq \int_{\mathcal{M}} |\nabla u|^2 \rho dx + 2 \frac{\sqrt{\mathcal{CNS}}}{(1 - \mathcal{S})} \left(\int_{\mathcal{M}} |\nabla u|^2 \rho dx \right)^{1/2} + \frac{\mathcal{CNS}}{(1 - \mathcal{S})^2} \\ & = \left(\|\nabla u\|_{\rho} + \frac{\sqrt{\mathcal{CNS}}}{(1 - \mathcal{S})} \right)^2. \end{aligned}$$

That is,

$$\sqrt{\Theta(1 - N\mathcal{S})} - \frac{\sqrt{\mathcal{CNS}}}{(1 - \mathcal{S})} \leq \|\nabla u\|_{\rho}.$$

Given that the above inequality holds for every $u \in H^1(\mathcal{M}, \rho)$ with $u \in Q^{\perp}$ and $\langle u, u \rangle_{\rho} = 1$, we conclude from the fact that Q has dimension N and from the max-min formula (Attouch et al., 2014, Theorem 8.4.2),

$$\left(\sqrt{\Theta(1 - N\mathcal{S})} - \frac{\sqrt{\mathcal{CNS}}}{(1 - \mathcal{S})} \right)^2 \leq \min_{u \in Q^{\perp}} \frac{\int_{\mathcal{M}} |\nabla u|^2 \rho dx}{\int_{\mathcal{M}} u^2 \rho dx} \leq \lambda_{N+1}$$

■

Combining the previous proposition with Proposition 23 we obtain the following result.

Corollary 28 *For every $k = 1, \dots, N$ we have*

$$\frac{1}{w_k} \|q_k - \Pi_N(q_k)\|_\rho^2 \leq \left(\sqrt{\frac{\Theta(1 - N\mathcal{S})}{\mathcal{C}}} - \frac{\sqrt{N\mathcal{S}}}{(1 - \mathcal{S})} \right)^{-2}. \quad (33)$$

The above corollary shows that for a well-separated mixture model, that is, for $\mathcal{S}$ and $\mathcal{C}/\Theta$ small, the right-hand side in (33) is small also, and so in that case, the functions $q_1, \dots, q_N$ are close to their projections onto the first N eigenfunctions of Δ_ρ . We will be able to conclude that $F_\# \nu$ has an orthogonal cone structure provided we can show $F_\#^Q \nu$ has one.

Proposition 29 *The probability measure $\mu^Q := F_\#^Q \nu$ with F^Q defined in (17) has an orthogonal cone structure with parameters (σ, δ, r) for any $\sigma \in (0, \pi/4)$, $\delta^* \leq \delta < 1$ and $r = w_{\max}^{-1/2}$ where*

$$\delta^* := \frac{w_{\max} \cos^2(\sigma) N^2 \mathcal{S}}{w_{\min} (1 - \cos^2(\sigma))}.$$

Here, $w_{\max} := \max_{i=1, \dots, N} w_i$ and $w_{\min} := \min_{i=1, \dots, N} w_i$.

Proof For each $k = 1, \dots, N$, let

$$C_k := \left\{ z \in \mathbb{R}^N : \frac{z_k}{|z|} > \cos(\sigma), \quad |z| \geq r \right\},$$

with $r := \frac{1}{\sqrt{w_{\max}}}$ and fixed $\sigma \in (0, \pi/4)$. It follows that

$$\mu^Q(C_k) = F_\#^Q \nu(C_k) = \nu(A_k),$$

where A_k is the preimage of C_k through F^Q , i.e.

$$\begin{aligned} A_k &:= (F^Q)^{-1}(C_k) \\ &= \left\{ x \in \mathcal{M} : \sqrt{\frac{\rho_k(x)}{\rho(x)}} > \cos(\sigma) \left(\sum_{j=1}^N \frac{\rho_j(x)}{\rho(x)} \right)^{1/2}, \quad \left(\sum_{j=1}^N \frac{\rho_j(x)}{\rho(x)} \right)^{1/2} \geq r \right\}. \end{aligned}$$

From the definition of r we see that the condition

$$\left(\sum_{j=1}^N \frac{\rho_j(x)}{\rho(x)} \right)^{1/2} \geq r$$

is redundant and so we can write

$$A_k = \left\{ x \in \mathcal{M} : \rho_k(x) > \cos^2(\sigma) \sum_{j=1}^N \rho_j(x) \right\}.$$

Now, for an arbitrary $x_0 \in \cap_{j=1}^N A_j^c \subseteq \mathcal{M}$ we have

$$\rho_k(x_0) \leq \cos^2(\sigma) \sum_{j=1}^N \rho_j(x_0), \quad \forall k = 1, \dots, N,$$

or equivalently,

$$(1 - \cos^2(\sigma))\rho_k(x_0) \leq \cos^2(\sigma) \sum_{j \neq k} \rho_j(x_0), \quad \forall k = 1, \dots, N.$$

From the fact that,

$$\sum_{k=1}^N \frac{w_k \rho_k(x_0)}{\rho(x_0)} = 1,$$

we know there exists a $\hat{k} \in \{1, \dots, N\}$ (depending on x_0) for which

$$\frac{w_{\hat{k}} \rho_{\hat{k}}(x_0)}{\rho(x_0)} \geq \frac{1}{N}.$$

Hence,

$$\begin{aligned} \frac{1 - \cos^2(\sigma)}{N^2} &\leq (1 - \cos^2(\sigma)) \left(\frac{w_{\hat{k}} \rho_{\hat{k}}(x_0)}{\rho(x_0)} \right)^2 \leq \cos^2(\sigma) \frac{w_{\max}}{w_{\min}} \sum_{j \neq \hat{k}} \frac{w_j \rho_j(x_0)}{\rho(x_0)} \frac{w_{\hat{k}} \rho_{\hat{k}}(x_0)}{\rho(x_0)} \\ &\leq \cos^2(\sigma) \frac{w_{\max}}{w_{\min}} \sum_k \sum_{j \neq k} \frac{w_j \rho_j(x_0)}{\rho(x_0)} \frac{w_k \rho_k(x_0)}{\rho(x_0)}. \end{aligned}$$

Since this is true for every $x_0 \in \cap_{k=1}^N A_k^c$, we conclude that

$$\begin{aligned} \frac{1 - \cos^2(\sigma)}{N^2} \nu \left(\cap_{k=1}^N A_k^c \right) &\leq \cos^2(\sigma) \frac{w_{\max}}{w_{\min}} \int_{\cap_{k=1}^N A_k^c} \left(\sum_{l=1}^N \sum_{j \neq l} \frac{w_l \rho_l}{\rho} \frac{w_j \rho_j}{\rho} \right) \rho dx \\ &\leq \cos^2(\sigma) \frac{w_{\max}}{w_{\min}} \sum_{l=1}^N \sum_{j \neq l} w_l w_j \int_{\mathcal{M}} \frac{\rho_l \rho_j}{\rho} dx \leq \cos^2(\sigma) \frac{w_{\max}}{w_{\min}} \mathcal{S}. \end{aligned}$$

That is,

$$\nu \left(\cap_{k=1}^N \tilde{A}_k^c \right) \leq \frac{w_{\max} \cos^2(\sigma) N^2 \mathcal{S}}{w_{\min} (1 - \cos^2(\sigma))},$$

and so we deduce that

$$\mu^Q \left(\cup_{k=1}^N C_k \right) \geq 1 - \frac{w_{\max} \cos^2(\sigma) N^2 \mathcal{S}}{w_{\min} (1 - \cos^2(\sigma))}.$$

This concludes the proof. ■

Proof [Proof of Theorem 10] In order to show that the measure $\mu := F_{\#} \nu$ has an orthogonal cone structure, it is enough to show that the measure $(OF)_{\#} \nu$ has an orthogonal cone structure where OF is the map:

$$OF : x \in \mathcal{M} \mapsto OF(x) \in \mathbb{R}^N,$$

and where O is a conveniently chosen $N \times N$ orthogonal matrix. We will construct O in such a way that the measures $OF_{\#} \nu$ and $F_{\#}^Q \nu$ are close to each other in the Wasserstein

sense. From this, Proposition 22 and Proposition 29 we will be able to conclude that $OF_{\sharp}\nu$ has an orthogonal cone structure.

Let us introduce

$$v_i := \frac{\Pi_N(q_i)}{\|\Pi_N(q_i)\|_\rho}, \quad i = 1, \dots, N,$$

where we recall $\Pi_N : L^2(\rho) \rightarrow U$ is the orthogonal projection onto U , the span of the first N eigenfunctions of Δ_ρ . From Corollary 28 it follows that

$$\begin{aligned} \left\| \frac{q_i}{\sqrt{w_i}} - v_i \right\|_\rho &\leq \left\| \frac{q_i}{\sqrt{w_i}} - \frac{\Pi_N(q_i)}{\sqrt{w_i}} \right\|_\rho + \left\| \frac{\Pi_N(q_i)}{\sqrt{w_i}} - \frac{\Pi_N(q_i)}{\|\Pi_N(q_i)\|_\rho} \right\|_\rho \\ &= \left\| \frac{q_i}{\sqrt{w_i}} - \frac{\Pi_N(q_i)}{\sqrt{w_i}} \right\|_\rho + \frac{1}{\sqrt{w_i}} \left| \|\Pi_N(q_i)\|_\rho - \sqrt{w_i} \right| \\ &= \left\| \frac{q_i}{\sqrt{w_i}} - \frac{\Pi_N(q_i)}{\sqrt{w_i}} \right\|_\rho + \frac{1}{\sqrt{w_i}} \left| \|\Pi_N(q_i)\|_\rho - \|q_i\|_\rho \right| \\ &\leq 2 \left\| \frac{q_i}{\sqrt{w_i}} - \frac{\Pi_N(q_i)}{\sqrt{w_i}} \right\|_\rho \\ &\leq 2 \left(\sqrt{\frac{\Theta(1 - N\mathcal{S})}{\mathcal{C}}} - \frac{\sqrt{N\mathcal{S}}}{(1 - \mathcal{S})} \right)^{-1}. \end{aligned} \tag{34}$$

In particular, for $i \neq j$

$$\begin{aligned} |\langle v_i, v_j \rangle_\rho| &= \left| \langle v_i - \frac{q_i}{\sqrt{w_i}}, v_j \rangle_\rho + \langle \frac{q_i}{\sqrt{w_i}}, v_j - \frac{q_j}{\sqrt{w_j}} \rangle_\rho + \langle \frac{q_i}{\sqrt{w_i}}, \frac{q_j}{\sqrt{w_j}} \rangle_\rho \right| \\ &\leq \left\| \frac{q_i}{\sqrt{w_i}} - v_i \right\|_\rho + \left\| \frac{q_j}{\sqrt{w_j}} - v_j \right\|_\rho + \mathcal{S}^{1/2} \\ &\leq 4 \left(\sqrt{\frac{\Theta(1 - N\mathcal{S})}{\mathcal{C}}} - \frac{\sqrt{N\mathcal{S}}}{(1 - \mathcal{S})} \right)^{-1} + \mathcal{S}^{1/2} = \tau. \end{aligned}$$

The first inequality follows from an application of the Cauchy-Schwartz inequality for the first two terms (since $\|v_i\|_\rho = 1$ by definition) and Jensen's inequality for the last term. By assumption (15), we can then use Lemma 41 to conclude that there exists $\tilde{v}_1, \dots, \tilde{v}_N$, an orthonormal basis for $(U, \langle \cdot, \cdot \rangle_\rho)$ for which

$$\|v_i - \tilde{v}_i\|_\rho^2 \leq N \left(\frac{1}{\sqrt{1 - N\tau}} - 1 \right)^2, \quad i = 1, \dots, N.$$

Therefore, for any $i = 1, \dots, N$,

$$\begin{aligned} \left\| \frac{q_i}{\sqrt{w_i}} - \tilde{v}_i \right\|_\rho^2 &= \left\| \frac{q_i}{\sqrt{w_i}} - v_i \right\|_\rho^2 + 2 \langle v_i - \tilde{v}_i, \frac{q_i}{\sqrt{w_i}} \rangle_\rho - \langle v_i + \tilde{v}_i, v_i - \tilde{v}_i \rangle_\rho \\ &\leq \left\| \frac{q_i}{\sqrt{w_i}} - v_i \right\|_\rho^2 + 4 \|v_i - \tilde{v}_i\|_\rho \end{aligned}$$

$$\begin{aligned}
&\leq 4 \left(\sqrt{\frac{\Theta(1-N\mathcal{S})}{\mathcal{C}}} - \frac{\sqrt{N\mathcal{S}}}{(1-\mathcal{S})} \right)^{-2} + 4\sqrt{N} \left(\frac{1}{\sqrt{1-N\tau}} - 1 \right) \\
&= \left(\frac{\tau - \sqrt{\mathcal{S}}}{2} \right)^2 + 4\sqrt{N} \left(\frac{1}{\sqrt{1-N\tau}} - 1 \right). \tag{35}
\end{aligned}$$

Let $\tilde{F} : \mathcal{M} \mapsto \mathbb{R}^N$ be the map $\tilde{F}(x) = \sum_{j=1}^N \tilde{v}_j(x) e_j$. From the fact that $\{\tilde{v}_1, \dots, \tilde{v}_N\}$ and $\{u_1, \dots, u_N\}$ are both orthonormal bases for $(U, \langle \cdot, \cdot \rangle_\rho)$ we can conclude there exists an orthogonal matrix $O \in \mathbb{R}^N \times \mathbb{R}^N$ for which

$$OF = \tilde{F}.$$

Moreover, let $\pi := (F^Q \times \tilde{F})_\# \nu$. We notice that π is a coupling between $F_\#^Q \nu$ and $\tilde{F}_\# \nu$. From (35) we see that

$$\begin{aligned}
W_2^2(F_\#^Q \nu, \tilde{F}_\# \nu) &\leq \int_{\mathbb{R}^N} \int_{\mathbb{R}^N} |z - \tilde{z}|^2 d\pi(z, \tilde{z}) \\
&= \int_{\mathcal{M}} |F^Q(x) - \tilde{F}(x)|^2 d\nu(x) \\
&= \sum_{i=1}^N \left\| \frac{q_i}{\sqrt{w_i}} - \tilde{v}_i \right\|_\rho^2 \\
&\leq N \left(\frac{\tau - \sqrt{\mathcal{S}}}{2} \right)^2 + 4N^{3/2} \left(\frac{1}{\sqrt{1-N\tau}} - 1 \right). \tag{36}
\end{aligned}$$

It then follows from Proposition 22 and Proposition 29 that μ has an orthogonal cone structure with parameters $\left(\sigma + s, \delta + t^2, \frac{1 - \sin(s)}{\sqrt{w_{\max}}} \right)$ with $\delta \in [\delta^*, 1)$ as given in Proposition 29 for any $s, t > 0$ satisfying assumption (16). $\blacksquare$

A.3 From discrete to continuum: convergence results for the spectrum of Δ_n

In this section we connect the spectrum of the graph Laplacian Δ_n with the spectrum of Δ_ρ . For this purpose, we modify our notation slightly and write $L^2(\nu)$ for $L^2(\rho)$ to allow for comparison with $L^2(\nu_n)$ (recall (2)). Furthermore, we write $\langle u, v \rangle_{L^2(\mu)} := \int_{\mathcal{M}} uv \, d\mu(x)$ for $\mu = \nu$ or ν_n .

Theorem 30 *Let $\mathcal{M}$, ρ , ε and n be as in Assumptions 11. Let $u_{n,1}, \dots, u_{n,N}$ be the first N eigenvectors of the kernelized graph Laplacian Δ_n (i.e. the graph with weights defined in (4)). Then, for every $\beta > 1$ there exists a constant $C_\beta > 0$ such that with probability at least $1 - C_\beta n^{-\beta}$, there exists a map $T_n : \mathcal{M} \rightarrow \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ (see Remark 31) satisfying*

$$\nu(T_n^{-1}(\{\mathbf{x}_i\})) = \frac{1}{n}, \quad \forall i = 1, \dots, n, \tag{37}$$

and

$$\|g_j - u_{n,j} \circ T_n\|_{L^2(\nu)}^2 \leq c_{\mathcal{M}} \left(\frac{\lambda_N}{\lambda_{N+1} - \lambda_N} \right) \left(\varepsilon + \frac{\log(n)^{p_m}}{\varepsilon n^{1/m}} \right), \tag{38}$$

for some orthonormal functions $g_1, \dots, g_N \in L^2(\nu)$ belonging to U , the span of the first N eigenfunctions of Δ_ρ with respect to $\langle \cdot, \cdot \rangle_\rho$, and a constant $c_{\mathcal{M}} > 0$ depending only on $\mathcal{M}$, N , $\rho^\pm$, C_{Lip} and η . In the above, λ_N and λ_{N+1} denote the N -th and the $(N+1)$ -th eigenvalues of Δ_ρ .

Remark 31 Relation (37) is simply saying that the map T_n is a transportation map from ν to ν_n . T_n here is as defined by García Trillos et al. (2020), i.e. the ∞ -optimal transport (∞ -OT) map between ν and ν_n . It was proved by García Trillos et al. (2020) that the ∞ -OT cost scales like $\frac{\log(n)^{pm}}{n^{1/m}}$. Notice that this is the only term on the right hand side of (38) that explicitly depends on n .

The proof of Theorem 30 relies on a comparison between the Dirichlet forms associated to the operators Δ_n and Δ_ρ by way of careful interpolation and discretization of discrete and continuum functions. Only small modifications to the proofs in Burago et al. (2014); García Trillos et al. (2020) are necessary and in what follows we give a thorough explanation of the steps that need to be adjusted. For simplicity, we largely follow the notation in García Trillos et al. (2020).

Let us summarize our strategy for the proof. The Dirichlet forms associated to Δ_n and Δ_ρ are respectively defined as

$$b_n(u_n) := \frac{1}{2n} \sum_{i,j} \mathbf{W}_{ij} |u_n(\mathbf{x}_i) - u_n(\mathbf{x}_j)|^2, \quad u_n \in L^2(\nu_n), \quad (39)$$

and

$$D(u) := \frac{1}{2} \int_{\mathcal{M}} |\nabla u|^2 \rho(x) dx, \quad u \in H^1(\mathcal{M}, \rho), \quad (40)$$

where we recall $\mathbf{W}_{ij}$ are the weights defined in (4). We will show that

$$\begin{aligned} b_n(Pf) &\leq (1 + er(n, \varepsilon)) D(f), \quad \forall f \in L^2(\nu), \\ D(I f_n) &\leq (1 + er(n, \varepsilon)) b_n(f_n), \quad \forall f_n \in L^2(\nu_n). \end{aligned}$$

for appropriately defined *discretization* and *interpolation* maps P and I ; the $er(n, \varepsilon)$ term does not depend on f or f_n , and is small for n large, and ε small but larger than a certain quantity that depends on n that we will introduce below. Together with the properties of the maps P and I (see Lemma 33), these inequalities allow us to compare the spectra of the discrete and continuum operators.

To define the maps P and I we first introduce the transport maps T_n constructed in (García Trillos et al., 2020, Theorem 2).

Proposition 32 For a given $\beta > 1$, there exists a constant $C_\beta > 0$ depending only on β so that with probability at least $1 - C_\beta n^{-\beta}$ there exists a map $T_n : \mathcal{M} \rightarrow \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ for which:

1. $\nu(T_n^{-1}(\mathbf{x}_i)) = 1/n$ for all $i = 1, \dots, n$,
2. $\delta_n := \text{ess sup}_{x \in \mathcal{M}} d_{\mathcal{M}}(T_n(x), x) \leq C \frac{\log(n)^{pm}}{n^{1/m}}$,

where $C = C(\mathcal{M}, \rho^-, \rho^+, \beta) > 0$, $p_m = 3/4$ for $m = 2$ and $p_m = 1/m$ for $m \geq 3$, and $d_{\mathcal{M}}$ represents the geodesic distance in $\mathcal{M}$.

Given the map T_n , the discretization map

$$P : L^2(\nu) \rightarrow L^2(\nu_n),$$

is defined as the transformation that takes a given $f \in L^2(\nu)$ and maps it into the discrete function

$$Pf(\mathbf{x}_i) := n \int_{T_n^{-1}(\{\mathbf{x}_i\})} f(x) \rho(x) dx, \quad i = 1, \dots, n.$$

For the interpolation map I , we first let P^* be the adjoint of P with respect to the $L^2(\nu_n)$ inner product, i.e. the map that satisfies

$$\langle Pg, f_n \rangle_{L^2(\nu_n)} = \langle g, P^* f_n \rangle_{L^2(\nu)}, \quad \forall g \in L^2(\nu), \quad \forall f_n \in L^2(\nu_n),$$

which takes an arbitrary discrete function $f_n \in L^2(\nu_n)$ and returns a function $P^* f_n \in L^2(\nu)$ defined by

$$P^* f_n(x) = f_n \circ T_n(x), \quad x \in \mathcal{M}.$$

We also introduce a smoothing operator

$$\Lambda_{\varepsilon, n, 0} f(x) := \int_{\mathcal{M}} \frac{1}{(\varepsilon - 2\delta_n)^m} \psi\left(\frac{d_{\mathcal{M}}(x, y)}{\varepsilon - 2\delta_n}\right) f(y) dy, \quad x \in \mathcal{M}, \quad f \in L^2(\nu),$$

which is a convolution operator with radial kernel

$$\psi(t) := \frac{1}{\alpha_\eta} \int_t^\infty \eta(s) s ds,$$

where we recall $\alpha_\eta := \int_0^\infty \eta(r) r^{m+1} dr$ was defined in (3), and where $d_{\mathcal{M}}(x, y)$ is the geodesic distance between points $x, y \in \mathcal{M}$. We normalize the operator $\Lambda_{\varepsilon, n, 0}$ to produce

$$\Lambda_{\varepsilon, n} f := \frac{1}{\Lambda_{\varepsilon, n, 0} \mathbf{1}} \Lambda_{\varepsilon, n, 0} f, \quad f \in L^2(\nu),$$

where in the above $\Lambda_{\varepsilon, n, 0} \mathbf{1}$ denotes the application of $\Lambda_{\varepsilon, n, 0}$ to the function that is identically equal to one; the normalization is introduced so as to ensure that $\Lambda_{\varepsilon, n}$ leaves constant functions unchanged. Finally, I is defined as the composition of P^* with $\Lambda_{\varepsilon, n}$. Namely,

$$Iu_n := \Lambda_{\varepsilon, n} \circ P^* u_n, \quad u_n \in L^2(\nu_n).$$

Note that for any $f_n \in L^2(\nu)$ it is guaranteed that $If_n \in H^1(\mathcal{M}, \rho)$, given that η was assumed Lipschitz (and so in particular ψ is C^1).

Lemma 33 (Discretization and interpolation errors) *Under the same assumptions on $\mathcal{M}$ and ρ as in Theorem 30 it follows that:*

1. For every $f \in H^1(\mathcal{M}, \rho)$,

$$\left| \|Pf\|_{L^2(\nu_n)}^2 - \|f\|_{L^2(\nu)}^2 \right| \leq C_1 \left[\delta_n \|f\|_{L^2(\nu)} D(f)^{1/2} + \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right) \|f\|_{L^2(\nu)}^2 \right].$$

2. For every $f \in H^1(\mathcal{M}, \rho)$,

$$b_n(Pf) \leq \left[1 + C_2 \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right) \right] D(f).$$

3. For every $f_n \in L^2(\nu_n)$ we have,

$$\left| \|If_n\|_{L^2(\nu)}^2 - \|f_n\|_{L^2(\nu_n)}^2 \right| \leq C_3 \left[\varepsilon \|f_n\|_{L^2(\nu_n)} b_n(f_n)^{1/2} + \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right) \|f_n\|_{L^2(\nu_n)}^2 \right].$$

4. For every $f_n \in L^2(\nu_n)$ such that $If_n \in H^1(\mathcal{M}, \rho)$,

$$D(If_n) \leq \left[1 + C_4 \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right) \right] b_n(f_n).$$

Here the constants $C_1, C_2, C_3, C_4 > 0$ only depend on lower and upper bounds for ρ and its Lipschitz constant, and on $\mathcal{M}$ and η . The term δ_n was defined in Proposition 32.

Before we prove Lemma 33 let us introduce an intermediate, non-local continuum Dirichlet energy

$$E_r(f) := \frac{2}{\alpha_\eta} \int_{\mathcal{M}} \int_{\mathcal{M}} \eta \left(\frac{|x-y|}{r} \right) |f(x) - f(y)|^2 \sqrt{\rho(x)} \sqrt{\rho(y)} dx dy, \quad f \in L^2(\nu),$$

where $r > 0$ is a length scale to be chosen later on.

Remark 34 Observe that the Dirichlet energy D and the non-local Dirichlet energy E_r can be written as

$$D(f) = \int_{\mathcal{M}} |\nabla f|^2 \tilde{\rho}^2 dx, \quad f \in H^1(\mathcal{M}, \rho),$$

$$E_r(f) := \frac{2}{\alpha_\eta} \int_{\mathcal{M}} \int_{\mathcal{M}} \eta \left(\frac{|x-y|}{r} \right) |f(x) - f(y)|^2 \tilde{\rho}(x) \tilde{\rho}(y) dx dy, \quad f \in L^2(\nu),$$

where $\tilde{\rho} = \sqrt{\rho}$. We notice that $\tilde{\rho} \in C^1(\mathcal{M})$ since it satisfies the same regularity condition as ρ given that ρ is bounded away from zero. Hence, our continuum Dirichlet energies D and E_r take the same form as in García Trillos et al. (2020). As a consequence we may use all of the bounds relating D and E_r that were proven there.

Remark 35 The only difference between the discrete Dirichlet form b_n in (39) and the one introduced in García Trillos et al. (2020) is the normalization by the terms $\sqrt{d_\varepsilon(\mathbf{x}_i)/n}$ that appear in the definition of $\mathbf{W}$, i.e., García Trillos et al. (2020) uses a different normalization of the graph Laplacian. We notice that the term $\sqrt{d_\varepsilon(\cdot)/n}$ uniformly approximates the density ρ . More precisely, given that the kernel η is assumed to be normalized, it follows from (García Trillos et al., 2020, Lemma 18) that

$$\max_{i=1, \dots, n} \left| \frac{1}{n} d_\varepsilon(\mathbf{x}_i) - \rho(\mathbf{x}_i) \right| \leq C \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right),$$

where we recall δ_n is the ∞ -OT distance between ν_n and ν given in Proposition 32, and the constant $C > 0$ depends only on ρ and geometric quantities associated to $\mathcal{M}$. We highlight that the above is a non-optimal estimate on the error of approximation of a kernel density estimator, but has the advantage of only depending on the ∞ -OT distance between empirical and ground-truth measures.

Remark 36 The maps P, P^*, Λ_r and I are defined identically to the way they were defined in García Trillos et al. (2020) and hence we may use all properties proved there.

Proof [Proof of Lemma 33] Throughout the proof we use C to denote a finite positive constant that depends only on the lower and upper bounds on ρ , the Lipschitz constant of ρ , and geometric quantities of $\mathcal{M}$. This constant may change value from one instance to the next, unless its dependence is explicitly stated.

The first statement follows directly from (García Trillos et al., 2020, Lemma 13(i)) after noticing that $\rho \leq C\sqrt{\rho}$, given that ρ is bounded away from zero. The second statement follows almost exactly as in (García Trillos et al., 2020, Lemma 13(ii)), after making some small modifications that we now describe. First, by Remark 35 and the assumptions on ρ it follows that for all $i, j \in \{1, \dots, n\}$ and every $x \in T_n^{-1}(\{\mathbf{x}_i\})$ and $y \in T_n^{-1}(\{\mathbf{x}_j\})$ we have

$$\frac{n}{\sqrt{d_\varepsilon(\mathbf{x}_i)d_\varepsilon(\mathbf{x}_j)}} \leq (1 + C(\varepsilon + \delta_n/\varepsilon)) \frac{1}{\sqrt{\rho(x)\rho(y)}}.$$

Using the definition of Pf and Proposition 32 (1) it follows that

$$\begin{aligned} b_n(Pf) &= \frac{2}{\alpha_\eta n \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \frac{\eta \left(\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\varepsilon} \right)}{\sqrt{d_\varepsilon(\mathbf{x}_i)} \sqrt{d_\varepsilon(\mathbf{x}_j)}} \\ &\quad \cdot \left| n^2 \int_{T_n^{-1}(\{\mathbf{x}_i\})} \int_{T_n^{-1}(\{\mathbf{x}_j\})} (f(x) - f(y)) \rho(x) \rho(y) dx dy \right|^2 \\ &\leq \frac{2n}{\alpha_\eta \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \frac{\eta \left(\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\varepsilon} \right)}{\sqrt{d_\varepsilon(\mathbf{x}_i)} \sqrt{d_\varepsilon(\mathbf{x}_j)}} \\ &\quad \cdot \int_{T_n^{-1}(\{\mathbf{x}_i\})} \int_{T_n^{-1}(\{\mathbf{x}_j\})} (f(x) - f(y))^2 \rho(x) \rho(y) dx dy \\ &\leq \frac{2(1 + C(\varepsilon + \delta_n/\varepsilon))}{\alpha_\eta \varepsilon^{m+2}} \sum_{i=1}^n \sum_{j=1}^n \eta \left(\frac{|\mathbf{x}_i - \mathbf{x}_j|}{\varepsilon} \right) \\ &\quad \int_{T_n^{-1}(\{\mathbf{x}_i\})} \int_{T_n^{-1}(\{\mathbf{x}_j\})} (f(x) - f(y))^2 \sqrt{\rho(x)} \sqrt{\rho(y)} dx dy, \end{aligned}$$

where the first inequality follows by Jensen's inequality. Using the above bound together with Remark 34 and proceeding identically to the proof of (García Trillos et al., 2020, Lemma 13(ii)) we can bound b_n in terms of E_r . More precisely,

$$b_n(Pf) \leq \frac{1 + C(\varepsilon + \delta_n/\varepsilon)}{\varepsilon^{m+2}} \left(E_{\varepsilon' + 2\delta_n}(f) + \frac{\tilde{C}\delta_n}{\varepsilon} E_{2(\varepsilon' + 2\delta_n)}(f) \right),$$

where $\varepsilon' = (1 + C'\varepsilon^2)\varepsilon$, C' is a constant intrinsic to the manifold $\mathcal{M}$, and $\tilde{C} > 0$ is a constant depending only on η . At this point we can follow the rest of the proof of (García Trillos et al., 2020, Lemma 13(ii)) to obtain the desired result after bounding E_r in terms of D :

$$\begin{aligned} \frac{1}{\varepsilon^{m+2}} E_{\varepsilon'+2\delta_n}(f) &\leq \left(1 + C \left(\varepsilon + \frac{\delta_n}{\varepsilon}\right)\right) D(f), \\ \frac{1}{\varepsilon^{m+2}} \frac{\delta_n}{\varepsilon} E_{2(\varepsilon'+2\delta_n)}(f) &\leq C \frac{\delta_n}{\varepsilon} \left(1 + \varepsilon + \frac{\delta_n}{\varepsilon}\right) D(f). \end{aligned}$$

The third statement follows directly from (García Trillos et al., 2020, Lemma 14 (i)) after noticing that thanks to Remark 35 we have $\sqrt{d_\varepsilon(\mathbf{x}_i)d_\varepsilon(\mathbf{x}_j)}/n \leq C$, where C depends only on the upper bound for ρ . Indeed, this inequality allows us to upper bound the discrete Dirichlet energy introduced in García Trillos et al. (2020) with a constant multiple of our discrete Dirichlet energy b_n .

For the fourth and final statement we proceed as in (García Trillos et al., 2020, Lemma 14(ii)), but making a small modification. By Remark 35 we conclude that for all $i, j \in \{1, \dots, n\}$ and every $x \in T_n^{-1}(\{\mathbf{x}_i\})$ and $y \in T_n^{-1}(\{\mathbf{x}_j\})$ we have

$$\frac{1}{\sqrt{\rho(x)\rho(y)}} \leq \frac{n(1 + C(\varepsilon + \delta_n/\varepsilon))}{\sqrt{d_\varepsilon(\mathbf{x}_i)d_\varepsilon(\mathbf{x}_j)}}.$$

The above inequality is used to replace the degree term d_ε/n with the density ρ , allowing us to reverse the bound in Lemma 33(2). Having in mind the inequality, the proof of our statement is exactly as in García Trillos et al. (2020). Indeed, we can follow (García Trillos et al., 2020, Lemma 14(i)) and show that

$$E_{\varepsilon-2\delta_n}(P^*f_n) \leq \left(1 + C \left(\varepsilon + \frac{\delta_n}{\varepsilon}\right)\right) \varepsilon^{m+2} b_n(f_n). \quad (41)$$

In turn, (García Trillos et al., 2020, Lemma 14(ii)) (recall remark 34) gives

$$D(I f_n) \leq \left[1 + C \left(\varepsilon + \frac{\delta_n}{\varepsilon}\right)\right] \frac{1}{\varepsilon^{m+2}} E_{\varepsilon-2\delta_n}(P^*f_n),$$

for all $f_n \in L^2(\nu_n)$. Combining (41) with the above inequality we deduce the desired result. $\blacksquare$

With Lemma 33 at hand, we can obtain a precise relationship between eigenvalues of Δ_n and Δ_ρ .

Lemma 37 (Convergence rate for eigenvalues) *Let λ_i be the i -th eigenvalue of Δ_ρ and let $\lambda_{n,i}$ be the i -th eigenvalue of Δ_n . Let $\beta > 1$. Then there exist constants $C, C_\beta > 0$ such that for sufficiently large n , with probability at least $1 - C_\beta n^{-\beta}$, we have*

$$|\lambda_{n,i} - \lambda_i| \leq C \left(\varepsilon + \frac{\delta_n}{\varepsilon} + \delta_n \sqrt{\lambda_i}\right) \lambda_i, \quad i = 1, \dots, N,$$

where $C > 0$ depends only on $\mathcal{M}$, η and ρ and $C_\beta > 0$ depends only on β .

Proof Using Lemma 33, we can follow exactly the proof of (García Trillos et al., 2020, Theorem 4). $\blacksquare$

Now, we are ready to prove Theorem 30.

Proof [Proof of Theorem 30] Throughout this proof we use C to denote a finite positive constant that depends only on $\mathcal{M}$, N , $\rho^\pm$, C_{Lip} and η . This constant may change from one instance to the next.

Let $u_{n,1}, \dots, u_{n,N}$ be the first N eigenfunctions of Δ_n with unit norm. We can assume they form an orthonormal basis with respect to $\langle \cdot, \cdot \rangle_{\nu_n}$, with corresponding eigenvalues

$$\lambda_{n,1} \leq \dots \leq \lambda_{n,N}.$$

We now put (Burago et al., 2014, Lemma 7.3) together with Lemmata 33 and 37, as well as Remark 36 to conclude that for every $j = 1, \dots, N$ we have

$$\|Iu_{n,j} - \Pi_N(Iu_{n,j})\|_{L^2(\nu)}^2 \leq \frac{C_{\mathcal{M},N}\lambda_N}{\lambda_{N+1} - \lambda_N} \left(\varepsilon + \frac{\delta_n}{\varepsilon} \right) =: \gamma_0^2, \quad (42)$$

where Π_N denotes the projection onto U , the span of the first N eigenfunctions of Δ_ρ and $C_{\mathcal{M},N} > 0$ is a constant depending on $\mathcal{M}$ and N only.

We need to show that the functions $\Pi_N(Iu_{n,1}), \dots, \Pi_N(Iu_{n,N})$ can be modified slightly to form an orthonormal basis for U . First, by (42) we have

$$\|Iu_{n,j}\|_{L^2(\nu)} - \gamma_0 \leq \|\Pi_N(Iu_{n,j})\|_{L^2(\nu)} \leq \|Iu_{n,j}\|_{L^2(\nu)} + \gamma_0.$$

Next, we find a bound on $\|Iu_{n,j}\|_{L^2(\nu)}$ using Lemma 33(3). To control the $b_n(u_{n,j})$ term on the right hand side of Lemma 33(3) we make use of the bound on the eigenvalues provided in Lemma 37,

$$\begin{aligned} b_n(u_{n,j}) &= \langle u_{n,j}, \Delta_n u_{n,j} \rangle_{L^2(\nu_n)} = \lambda_{n,j} \\ &\leq C \left(1 + \varepsilon + \frac{\delta_n}{\varepsilon} + \delta_n \sqrt{\lambda_N} \right) \lambda_N \\ &\leq C \lambda_N, \end{aligned} \quad (43)$$

and so, since the $u_{n,j}$ are normalized, we obtain

$$\left| \|Iu_{n,j}\|_{L^2(\nu)}^2 - 1 \right| \leq C \left(\varepsilon \sqrt{\lambda_N} + \varepsilon + \frac{\delta_n}{\varepsilon} \right). \quad (44)$$

Combining with the estimate for $\|\Pi_N Iu_{n,j}\|_{L^2(\nu)}$ yields (for ε and δ_n/ε small enough and n large enough so that the right-hand side in the last estimate is less or equal to 1)

$$1 - \gamma_1 \leq \|\Pi_N(Iu_{n,j})\|_{L^2(\nu)} \leq 1 + \gamma_1 \quad \forall j = 1, \dots, N, \quad (45)$$

where

$$\gamma_1 := C \left[\varepsilon \sqrt{\lambda_N} + \varepsilon + \frac{\delta_n}{\varepsilon} \right]^{1/2} + \gamma_0.$$

Lemma 33(3) also allows us to bound the difference of inner products. For $i \neq j$ we have

$$\begin{aligned}\langle Iu_{n,j}, Iu_{n,i} \rangle_{L^2(\nu)} &= \frac{1}{2} \left(\|Iu_{n,j}\|_{L^2(\nu)}^2 + \|Iu_{n,i}\|_{L^2(\nu)}^2 - \|Iu_{n,j} - Iu_{n,i}\|_{L^2(\nu)}^2 \right), \\ 0 = \langle u_{n,j}, u_{n,i} \rangle_{L^2(\nu_n)} &= \frac{1}{2} \left(\|u_{n,j}\|_{L^2(\nu_n)}^2 + \|u_{n,i}\|_{L^2(\nu_n)}^2 - \|u_{n,j} - u_{n,i}\|_{L^2(\nu_n)}^2 \right).\end{aligned}$$

Subtracting the above identities and using Lemma 33(3) we obtain

$$|\langle Iu_{n,j}, Iu_{n,i} \rangle_{L^2(\nu)}| \leq C \left(\varepsilon \sqrt{\lambda_N} + \varepsilon + \frac{\delta_n}{\varepsilon} \right).$$

From this, and the fact that both $\|Iu_{n,j} - \Pi_N Iu_{n,j}\|_{L^2(\nu)}$ and $\|Iu_{n,i} - \Pi_N Iu_{n,i}\|_{L^2(\nu)}$ are smaller than γ_0 , we deduce that

$$\begin{aligned}|\langle \Pi_N Iu_{n,i}, \Pi_N Iu_{n,j} \rangle_{L^2(\nu)}| &\leq |\langle Iu_{n,i}, Iu_{n,j} \rangle_{L^2(\nu)}| + |\langle Iu_{n,j}, Iu_{n,i} - \Pi_N Iu_{n,i} \rangle_{L^2(\nu)}| \\ &\quad + |\langle \Pi_N Iu_{n,i}, Iu_{n,j} - \Pi_N Iu_{n,j} \rangle_{L^2(\nu)}| \\ &\leq C \left(\varepsilon \sqrt{\lambda_N} + \varepsilon + \frac{\delta_n}{\varepsilon} \right) + C\gamma_0 =: \gamma_2, \quad \forall i \neq j,\end{aligned}\tag{46}$$

where the second inequality follows from Cauchy-Schwartz, the fact that $\|\Pi_N u\|_{L^2(\nu)} \leq \|u\|_{L^2(\nu)}$ for any $u \in L^2(\nu)$, and from $\|Iu_{n,i}\|_{L^2(\nu)} \leq C$ by (44).

With (45), (46) and Assumptions 11 (guaranteeing the smallness of γ_0 , γ_1 and γ_2) we can use Lemma 41 and Remark 42, to deduce the existence of an orthonormal system $g_1, \dots, g_N$ for U satisfying:

$$\|\Pi_N Iu_{n,j} - g_j\|_{L^2(\nu)} \leq \sqrt{N} \left(\frac{1}{\sqrt{1 - N\gamma_2}} - 1 \right), \quad \forall j = 1, \dots, N.$$

Using (42) and expanding in $N\gamma_2$ we deduce

$$\begin{aligned}\|Iu_{n,j} - g_j\|_{L^2(\nu)}^2 &\leq (\|Iu_{n,j} - \Pi_N Iu_{n,j}\|_{L^2(\nu)} + \|\Pi_N Iu_{n,j} - g_j\|_{L^2(\nu)})^2 \\ &\leq \left(\gamma_0 + \sqrt{N} \left(\frac{1}{\sqrt{1 - N\gamma_2}} - 1 \right) \right)^2 \\ &\leq C \left(\gamma_0 + \frac{N^{3/2}}{2} \gamma_2 + \frac{3N^{5/2}}{8} \gamma_2^2 \right)^2 =: \gamma_3, \quad \forall j = 1, \dots, N.\end{aligned}$$

Now,

$$\begin{aligned}\|Iu_{n,j} - u_{n,j} \circ T_n\|_{L^2(\nu)}^2 &= \|\Lambda_{\varepsilon - 2\delta_n} P^* u_{n,j} - P^* u_{n,j}\|_{L^2(\nu)}^2 \\ &\leq \frac{C\varepsilon^2}{(\varepsilon - 2\delta_n)^{m+2}} E_{\varepsilon - 2\delta_n}(P^* u_{n,j}) \\ &\leq C\varepsilon^2 b_n(u_{n,j}) \\ &\leq C\varepsilon^2 \lambda_N =: \gamma_4,\end{aligned}$$

where we have used (García Trillos et al., 2020, Lemma 8), assuming $\delta_n < \varepsilon/(m+5)$ to obtain the first inequality and where the second and third inequalities follow from (41) and (43) respectively. Using triangle inequality and the above estimates we obtain

$$\|u_{n,j} \circ T_n - g_j\|_{L^2(\nu)}^2 \leq 2(\gamma_3 + \gamma_4).$$

Plugging in the definitions of γ_3 and γ_4 and using Assumptions 11, we may collect the highest order terms and deduce that for all $j = 1, \dots, N$ we have

$$\|u_{n,j} \circ T_n - g_j\|_{L^2(\nu)}^2 \leq C(1 + N^{3/2} + N^3)C_{\mathcal{M},N} \left(\frac{\lambda_N}{\lambda_{N+1} - \lambda_N} \right) (\varepsilon + \delta_n/\varepsilon).$$

where $C_{\mathcal{M},N}$ is the constant in (42). Proposition 32(2) implies that

$$\|u_{n,j} \circ T_n - g_j\|_{L^2(\nu)}^2 \leq c_{\mathcal{M}} \left(\frac{\lambda_N}{\lambda_{N+1} - \lambda_N} \right) \left(\varepsilon + \frac{\log(n)^{p_m}}{\varepsilon n^{1/m}} \right),$$

where $c_{\mathcal{M}}$ is a constant proportional to $C_{\mathcal{M},N}$ and C given in Proposition 32(2). This concludes the proof of the theorem. $\blacksquare$

Our goal is now to replace the terms λ_N and λ_{N+1} in (38) with quantities that only depend on the parameters of the mixture model. A lower bound for λ_{N+1} was already obtained in Proposition 27, and now we focus on obtaining an upper bound for λ_N .

Proposition 38 (Upper bound for λ_N) *Suppose that $\mathcal{S}^{1/2}N < 1$ and let λ_N be the N -th eigenvalue of Δ_ρ . Then,*

$$\lambda_N \leq \frac{NC}{1 - N\mathcal{S}^{1/2}}.$$

Proof To get an upper bound for λ_N it is enough to use the min-max formula (Attouch et al., 2014, Theorem 8.4.2) for λ_N . In particular, since $Q := \text{span}\{q_1, \dots, q_N\}$ is N -dimensional it follows that

$$\lambda_N \leq \max_{u \in Q} \frac{\langle \Delta_\rho u, u \rangle_\rho}{\langle u, u \rangle_\rho}.$$

Let $u \in Q$ for which $\langle u, u \rangle_\rho = 1$. Then,

$$u = \sum_{k=1}^N a_k q_k$$

for some scalars a_k satisfying

$$1 = \sum_{k=1}^N a_k^2 \|q_k\|_\rho^2 + \sum_{k=1}^N \sum_{j \neq k}^N a_k a_j \langle q_k, q_j \rangle_\rho = \sum_{k=1}^N a_k^2 w_k + \sum_{k=1}^N \sum_{j \neq k}^N a_k a_j \langle q_k, q_j \rangle_\rho.$$

Now,

$$\left| \sum_{k=1}^N \sum_{j \neq k}^N a_k a_j \langle q_k, q_j \rangle_\rho \right| \leq \sum_{k=1}^N \sum_{j \neq k}^N |a_k| |a_j| \sqrt{w_k} \sqrt{w_j} \mathcal{S}^{1/2} \leq N \mathcal{S}^{1/2} \sum_{k=1}^N a_k^2 w_k.$$

Thus,

$$\sum_{k=1}^N a_k^2 w_k \leq \frac{1}{1 - N\mathcal{S}^{1/2}}.$$

In addition,

$$\langle \Delta_\rho u, u \rangle_\rho = \sum_{k=1}^N \sum_{j=1}^N a_k a_j \int_{\mathcal{M}} \nabla q_k \cdot \nabla q_j \rho dx.$$

Recall that $\|\nabla q_k\|_\rho^2 = w_k \mathcal{C}_k$. Using Hölder's inequality, we obtain

$$\begin{aligned} \langle \Delta_\rho u, u \rangle_\rho &\leq \sum_{k=1}^N \sum_{j=1}^N |a_k a_j| \left(\int_{\mathcal{M}} |\nabla q_k|^2 \rho dx \right)^{1/2} \left(\int_{\mathcal{M}} |\nabla q_j|^2 \rho dx \right)^{1/2} \\ &\leq \mathcal{C} N \sum_{k=1}^N a_k^2 w_k \leq \frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}}. \end{aligned}$$

We conclude that

$$\lambda_N \leq \frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}}.$$

■

Remark 39 *In the previous result we used the trace of the operator Δ_ρ restricted to Q to bound Θ_N . This is not necessarily an optimal bound, but we remark that in the case when $\mathcal{C}/\Theta$ is small enough (in particular smaller than $1/N$), one can use this estimate to get a meaningful lower bound for the spectral gap $\lambda_{N+1} - \lambda_N$. More precisely, we obtain*

$$\lambda_{N+1} - \lambda_N \geq \left(\sqrt{\Theta(1 - N\mathcal{S})} - \frac{\sqrt{CN\mathcal{S}}}{(1 - \mathcal{S})} \right)^2 - \frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}}.$$

While in general we think of N as a relatively small number (representing the number of meaningful components in a data set) we would like to remark that N in the estimates

$$\lambda_N \leq \frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}},$$

can be replaced with N_{eff} where N_{eff} is the effective number of components any given component intersects. That is,

$$N_{eff} = \max_{k=1, \dots, N} \#\{i \neq k \text{ s.t. } \nu(\text{supp}(\rho_k) \cap \text{supp}(\rho_i)) > 0\}.$$

Estimates with better dependence on N may be obtained in terms of a quantity analogue to $\mathcal{C}$ of the form:

$$\max_{k=1, \dots, N} \int_{\mathcal{M}} |\nabla \log(\rho_k / \rho)|^3 \rho_k dx.$$

Combining the lower bound for λ_{N+1} from Proposition 27 and the upper bound for λ_N we immediately deduce the following.

Corollary 40 *In Theorem 30, inequality (38) can be replaced with:*

$$\int_{\mathcal{M}} |g_j(x) - u_{n,j} \circ T_n(x)|^2 d\nu(x) = \|g_j - u_{n,j} \circ T_n\|_{L^2(\nu)}^2 \leq \phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m), \quad (47)$$

for all $j = 1, \dots, N$, where we recall

$$\begin{aligned} \phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m) &:= c_{\mathcal{M}} \left(\frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}} \right) \left(\varepsilon + \frac{\log(n)^{p_m}}{\varepsilon n^{1/m}} \right) \\ &\quad \cdot \left(\left(\sqrt{\Theta(1 - N\mathcal{S})} - \frac{\sqrt{\mathcal{C}N\mathcal{S}}}{1 - \mathcal{S}} \right)^2 - \frac{N\mathcal{C}}{1 - N\mathcal{S}^{1/2}} \right)^{-1}, \end{aligned} \quad (48)$$

and $c_{\mathcal{M}}$ is a constant depending on $\mathcal{M}$, $\rho^{\pm}$, C_{Lip} , η and N .

A.4 Proof of Theorem 15

Proof [Proof of Theorem 15] Let $\beta > 1$. From Corollary 40, we know that with probability greater than $1 - C_{\beta} n^{-\beta}$, there exist a transportation map $T_n : \mathcal{M} \rightarrow \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ pushing forward ν into ν_n and an orthonormal set of functions $g_1, \dots, g_n$ in U (the space generated by the first N eigenfunctions of Δ_{ρ}) satisfying

1. $\sup_{x \in \mathcal{M}} d_{\mathcal{M}}(x, T_n(x)) \leq c_{\mathcal{M}} \frac{\log(n)^{p_m}}{n^{1/m}}.$
2. $\int_{\mathcal{M}} |g_i(x) - u_{n,i}(T_n(x))|^2 d\nu(x) \leq \phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m),$

where the function ϕ is defined in (48) and used throughout this proof for convenience of notation. Let $G(x) := (g_1(x), \dots, g_N(x))$ for $x \in \mathcal{M}$. We deduce that

$$\begin{aligned} \int_{\mathcal{M}} |F_n \circ T_n(x) - G(x)|^2 d\nu(x) &= \sum_{i=1}^N \int_{\mathcal{M}} |g_i(x) - u_{n,i}(T_n(x))|^2 d\nu(x) \\ &\leq N\phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m). \end{aligned}$$

We claim that the integral on the left hand side of the previous expression can be written as:

$$\int_{\mathcal{M}} |F_n \circ T_n(x) - G(x)|^2 d\nu(x) = \int_{\mathbb{R}^N \times \mathbb{R}^N} |x - y|^2 d\pi_n(x, y),$$

for a transportation plan $\pi_n \in \mathcal{P}(\mathbb{R}^N \times \mathbb{R}^N)$ between the measures $G_{\#}\nu$ and $F_{n\#}\nu_n$. To see this, let

$$\pi_n := (G \times F_n \circ T_n)_{\#}\nu.$$

It is straightforward to check that π_n is a transportation plan between $G_{\#}\nu$ and $F_{n\#}\nu_n$, and moreover, by the change of variables formula

$$\int_{\mathbb{R}^N \times \mathbb{R}^N} |x - y|^2 d\pi_n(x, y) = \int_{\mathcal{M}} |G(x) - F_n \circ T_n(x)|^2 d\nu(x).$$

Therefore,

$$\begin{aligned} (W_2(G_{\#}\nu, F_{n\#}\nu_n))^2 &\leq \int_{\mathcal{M}} |F_n \circ T_n(x) - G(x)|^2 d\nu(x) \\ &\leq N\phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m). \end{aligned} \quad (49)$$

Since $g_1, \dots, g_N$ is an orthonormal basis for U , we conclude that there exists an orthogonal matrix R such that for every $x \in \mathcal{M}$ we have

$$G(x) = RF(x),$$

where F is the continuum spectral embedding as defined in (14).

In order to show that the measure $\mu_n := F_{n\#}\nu_n$ has an orthogonal cone structure, it is enough to show that the pushforward of ν_n through any orthogonal transformation of F_n has an orthogonal cone structure. In particular, we will show that the measure $(RO^{-1}F_n)_{\#}\nu_n$ has an orthogonal cone structure, where R is defined as above, and O is the orthogonal transformation chosen in the proof of Theorem 10:

$$OF(x) = \tilde{F}(x) = \sum_{j=1}^N \frac{\Pi_N(q_j)(x)}{\|\Pi_N(q_j)\|_{\rho}} e_j.$$

We can now show that the measures $OR^{-1}F_{n\#}\nu$ and $F_{n\#}^Q\nu$ are close to each other in the Wasserstein sense. Since $G = RO^{-1}\tilde{F}$,

$$\begin{aligned} W_2(OR^{-1}F_{n\#}\nu_n, F_{n\#}^Q\nu) &\leq W_2(OR^{-1}F_{n\#}\nu_n, \tilde{F}_{n\#}\nu) + W_2(\tilde{F}_{n\#}\nu, F_{n\#}^Q\nu) \\ &= W_2(F_{n\#}\nu_n, RO^{-1}\tilde{F}_{n\#}\nu) + W_2(\tilde{F}_{n\#}\nu, F_{n\#}^Q\nu) \\ &= W_2(F_{n\#}\nu_n, G_{n\#}\nu) + W_2(\tilde{F}_{n\#}\nu, F_{n\#}^Q\nu) \\ &\leq \sqrt{N\phi(\mathcal{S}, \mathcal{C}, \Theta, N, \varepsilon, n, m)} \\ &\quad + \sqrt{N \left(\frac{\tau - \sqrt{\mathcal{S}}}{2} \right)^2 + 4N^{3/2} \left(\frac{1}{\sqrt{1 - N\tau}} - 1 \right)}, \end{aligned}$$

where the last inequality follows from (49) and (36). Thanks to the Assumption (20) we can apply Proposition 29 and Proposition 22 to conclude that $OR^{-1}F_{n\#}\nu_n$ has an orthogonal cone structure with the parameters as stated in Theorem 15, and hence so does $F_{n\#}\nu_n$. ■

Appendix B. Supplementary results

B.1 Near-orthogonal vectors

Lemma 41 *Let V be a vector space of dimension N and let $\langle \cdot, \cdot \rangle$ be an inner product on V with associated norm $\|\cdot\|$. Suppose that $v_1, \dots, v_N$ are linearly independent unit vectors in V such that*

$$|\langle v_j, v_l \rangle| \leq \delta, \quad \forall j \neq l$$

for $\delta > 0$ satisfying

$$N\delta < 1.$$

Then, there exists an orthonormal basis for V , $\{\tilde{v}_1, \dots, \tilde{v}_N\}$, such that for every $j = 1, \dots, N$

$$\|v_j - \tilde{v}_j\| \leq \tilde{\phi}(N, \delta), \quad (50)$$

where

$$\tilde{\phi}(N, \delta) := \sqrt{N} \left[\frac{1}{\sqrt{1 - N\delta}} - 1 \right].$$

Proof Without loss of generality assume V is the Euclidean space $\mathbb{R}^N$ with the usual inner product. Define the square matrix

$$\mathbf{V} = [v_1 | v_2 | \dots | v_N].$$

and the residual matrix $\mathbf{R} = \mathbf{V}^T \mathbf{V} - I$. Since $v_1, \dots, v_N$ are linearly independent, the matrix $\mathbf{V}^T \mathbf{V}$ is invertible, and following Higham (1989) we have that the nearest orthonormal approximation to $\mathbf{V}$ in the Frobenius norm is the matrix

$$\tilde{\mathbf{V}} = \mathbf{V}(\mathbf{V}^T \mathbf{V})^{-\frac{1}{2}} = \mathbf{V}(I + \mathbf{R})^{-\frac{1}{2}}.$$

Using the series expansion of the square root gives

$$\mathbf{V} - \tilde{\mathbf{V}} = \mathbf{V} \left(\frac{1}{2} \mathbf{R} - \frac{3}{8} \mathbf{R}^2 + \frac{5}{16} \mathbf{R}^3 - \dots \right)$$

Note that $|\mathbf{R}_{ij}| \leq \delta$ following our assumptions on the v_j . Furthermore, a straightforward calculation shows that $|(\mathbf{R}^k)_{ij}| \leq N^{k-1} \delta^k$. Thus,

$$\begin{aligned} |(\mathbf{V} - \tilde{\mathbf{V}})_{ij}| &\leq N \left(\sup_{ij} |\mathbf{V}_{ij}| \right) \left(\frac{1}{2} \delta + \frac{3}{8} N \delta^2 + \frac{5}{16} N^2 \delta^3 + \dots \right) \\ &= \left(\sup_{ij} |\mathbf{V}_{ij}| \right) \left[(1 - N\delta)^{-\frac{1}{2}} - 1 \right] \\ &\leq \left[(1 - N\delta)^{-\frac{1}{2}} - 1 \right] = \frac{\tilde{\phi}(N, \delta)}{\sqrt{N}}. \end{aligned}$$

It follows that

$$\|v_j - \tilde{v}_j\| = \left(\sum_{i=1}^N |(\mathbf{V} - \tilde{\mathbf{V}})_{ij}|^2 \right)^{1/2} \leq \left(\sum_{i=1}^N \frac{\tilde{\phi}(N, \delta)^2}{N} \right)^{1/2} = \tilde{\phi}(N, \delta),$$

which proves inequality (50). ■

Remark 42 *With the same notation as in the above lemma, suppose that the unit vectors $v_1, \dots, v_N$ satisfy*

$$|\langle v_i, v_j \rangle| \leq \delta, \quad \forall i \neq j,$$

where δ satisfies the slightly stronger condition

$$2N\delta < 1.$$

We claim that in that case the vectors are linearly independent. Indeed, for the sake of contradiction suppose that they are not. Then we can find numbers $a_1, \dots, a_N$ not all equal to zero, for which

$$0 = \sum_{i=1}^N a_i v_i.$$

Normalizing the a_i we can further assume that

$$\sum_{i=1}^N a_i^2 = 1.$$

It follows from Jensen's inequality that

$$0 = \left\| \sum_{i=1}^N a_i v_i \right\|^2 = \sum_{i=1}^N a_i^2 + 2 \sum_{i=1}^N \sum_{j \neq i} a_i a_j \langle v_i, v_j \rangle \geq 1 - 2\delta \left(\sum_{i=1}^N |a_i| \right)^2 \geq 1 - 2N\delta,$$

which would contradict the hypothesis on δ .

B.2 Cheeger's inequality

Cheeger's inequality in manifolds was introduced in Cheeger (1969). For the convenience of the reader here we present a proof of Cheeger's inequality for sufficiently smooth and bounded functions using techniques that are developed in spectral geometry and spectral graph theory literature. We omit some technical details and highlight the specific structure of the operator Δ_ρ which allows us to deduce the inequality. As we will see below not every normalization of Laplacian operator will produce a similar result. We also note that our result can be generalized to functions in $L^2(\mathcal{M}, \rho)$ via a density argument.

Theorem 43 *Let $u \in L^2(\mathcal{M}, \rho)$ sufficiently smooth and bounded with*

$$\int_{\mathcal{M}} u \rho dx = 0.$$

Then

$$\frac{\langle \Delta_\rho u, u \rangle_\rho}{\langle u, u \rangle_\rho} \geq \frac{1}{4} h(\mathcal{M}, \rho)^2,$$

where

$$h(\mathcal{M}, \rho) := \min_{A \subseteq \mathcal{M}} \frac{\int_{\partial A \cap \mathcal{M}} \rho(x) dS(x)}{\min \left\{ \int_A \rho dx, \int_{\mathcal{M} \setminus A} \rho dx \right\}}.$$

Proof We show that for any non-constant function $u \in L^2(\mathcal{M}, \rho)$ sufficiently smooth and bounded with

$$\int_{\mathcal{M}} u \rho dx = 0,$$

we can find a set $A \subseteq \mathcal{M}$ such that

$$R(u) := \frac{\langle \Delta_\rho u, u \rangle_\rho}{\langle u, u \rangle_\rho} = \frac{\int_{\mathcal{M}} |\nabla u|^2 \rho dx}{\int_{\mathcal{M}} u^2 \rho dx} \geq \frac{1}{4} \left(\frac{\int_{\partial A \cap \mathcal{M}} \rho dS(x)}{\min \left\{ \int_A \rho dx, \int_{\mathcal{M} \setminus A} \rho dx \right\}} \right)^2 \geq \frac{1}{4} (h(\mathcal{M}, \rho))^2.$$

To see this, we start by letting r be the smallest number for which

$$\int_{\mathcal{M}} \mathbb{1}_{\{u \leq r\}} \rho dx \geq 1/2,$$

and we define

$$z(x) := u(x) - r, \quad x \in \mathcal{M}.$$

In other words, r is the median of $u(x)$. Let m, M be the infimum and supremum of z . Then either $m < 0$ or $M > 0$. Notice that since z and u differ only by a constant we have

$$\langle \Delta_\rho u, u \rangle_\rho = \langle \Delta_\rho z, z \rangle_\rho,$$

and from the fact that $\int_{\mathcal{M}} u \rho dx = 0$, we also have

$$\langle z, z \rangle_\rho \geq \langle u, u \rangle_\rho$$

In particular,

$$R(z) \leq R(u).$$

Now, the coarea formula states that for any $L^1(dx)$ function $g : \mathcal{M} \rightarrow \mathbb{R}$ we have

$$\int_{\mathcal{M}} g(x) |\nabla z(x)| dx = \int_{-\infty}^{\infty} \left(\int_{\partial A_t \cap \mathcal{M}} g(x) dS(x) \right) dt,$$

where A_t is the level set:

$$A_t := \{x \in \mathcal{M} : z(x) \leq t\}.$$

Taking g to be the function

$$g(x) = 2\rho(x)|z(x)|, \quad x \in \mathcal{M},$$

we deduce that

$$\begin{aligned} \int_{\mathcal{M}} 2|z(x)| |\nabla z(x)| \rho(x) dx &= \int_{-\infty}^{\infty} \left(\int_{\partial A_t \cap \mathcal{M}} \rho(x) dS(x) \right) 2|t| dt \\ &= \int_m^M \left(\int_{\partial A_t \cap \mathcal{M}} \rho(x) dS(x) \right) 2|t| dt, \end{aligned} \tag{51}$$

where in the first equality we have used the fact that at the boundary ∂A_t the function z is equal to t . On the other hand,

$$\begin{aligned}
 & \int_m^M \min \left\{ \int_{A_t} \rho(x) dx, \int_{\mathcal{M} \setminus A_t} \rho(x) dx \right\} 2|t| dt \\
 &= \int_m^0 \left(\int_{A_t} \rho(x) dx \right) 2|t| dt + \int_0^M \left(\int_{\mathcal{M} \setminus A_t} \rho(x) dx \right) 2|t| dt \\
 &= - \int_m^0 2 \left(\int_{A_t} \rho(x) dx \right) t dt + \int_0^M \left(\int_{\mathcal{M} \setminus A_t} \rho(x) dx \right) 2t dt \\
 &= \int_{\mathcal{M}} \int_m^0 \mathbb{1}_{\{z(x) \leq t\}}(x) (-2t) dt \rho(x) dx + \int_{\mathcal{M}} \int_0^M \mathbb{1}_{\{z(x) > t\}}(x) (2t) dt \rho(x) dx \\
 &= \int_{\mathcal{M}} \int_{z(x)}^0 \mathbb{1}_{\{z(x) \leq 0\}}(x) (-2t) dt \rho(x) dx + \int_{\mathcal{M}} \int_0^{z(x)} \mathbb{1}_{\{z(x) > 0\}}(x) (2t) dt \rho(x) dx \\
 &= \int_{\mathcal{M}} \mathbb{1}_{\{z(x) \leq 0\}}(x) z^2 \rho(x) dx + \int_{\mathcal{M}} \mathbb{1}_{\{z(x) > 0\}}(x) z^2 \rho(x) dx = \int_{\mathcal{M}} z(x)^2 \rho(x) dx,
 \end{aligned}$$

where in the third equality we used the Fubini-Tonelli theorem to switch the order of integrals. From the previous identity and (51) we notice that

$$\int_m^M \left(\int_{\partial A_t \cap \mathcal{M}} \rho(x) dS(x) \right) 2|t| dt = K(z) \int_m^M \min \left\{ \int_{A_t} \rho dx, \int_{\mathcal{M} \setminus A_t} \rho dx \right\} 2|t| dt$$

where

$$K(z) := \frac{2 \int_{\mathcal{M}} |z(x)| |\nabla z(x)| \rho(x) dx}{\langle z, z \rangle_{\rho}}.$$

Now, Cauchy Schwartz-inequality shows that

$$K(z) \leq 2\sqrt{R(z)}.$$

This is the point where it is important to have both numerator and denominator in the Raleigh quotient R to be weighted by ρ . We notice that with a different weighting we would have not gotten the square root of the Raleigh quotient in this last step.

It follows that

$$0 \leq \int_m^M 2|t| \left(2\sqrt{R(z)} \min \left\{ \int_{A_t} \rho dx, \int_{\mathcal{M} \setminus A_t} \rho dx \right\} - \int_{\partial A_t \cap \mathcal{M}} \rho(x) dS(x) \right) dt,$$

from where we can see that there must exist some $t \in (m, M)$ for which

$$\frac{\int_{\partial A_t \cap \mathcal{M}} \rho dS(x)}{\min \left\{ \int_{A_t} \rho dx, \int_{\mathcal{M} \setminus A_t} \rho dx \right\}} \leq 2\sqrt{R(z)}.$$

■