

Weighted Cheeger-Buser Inequalities, with Applications to Cutting Probability Densities - as Easy as 1,2,3

Timothy Chu
CMU
timothyzchu@gmail.com

Gary L. Miller*
CMU
glmiller@cs.cmu.edu

Noel J. Walkington†
CMU
noelw@cmu.edu

Alex L. Wang
CMU
alw1@cs.cmu.edu

May 7, 2020

Abstract

In this paper, we show how sparse or isoperimetric cuts of a probability density function relate to Cheeger cuts of its principal eigenfunction, for appropriate definitions of ‘sparse cut’ and ‘principal eigenfunction’.

We construct these appropriate definitions of sparse cut and principal eigenfunction in the probability density setting. Then, we prove Cheeger and Buser type inequalities similar to those for the normalized graph Laplacian of Alon-Milman. We demonstrate that no such inequalities hold for most prior definitions of sparse cut and principal eigenfunction. We apply this result to generate novel algorithms for cutting probability densities and clustering data, including a principled variant of spectral clustering.

*Supported in part by National Science Foundation Grant CCF-1637523.

†Supported in part by National Science Foundation Grant DMS-1729478.

1 Introduction

Clustering, the task of partitioning data into groups, is one of the fundamental tasks in machine learning [Bis06, NJW01, EKX96, KK10, SW18]. The data that practitioners seek to cluster is commonly modeled as i.i.d. samples from a probability density function, an assumption foundational in theory, statistics, and AI [Bis06, GTS16, BC17, MV10, GLR18]. A clustering algorithm on data drawn from a probability density function should ideally converge to a partitioning of the probability density function, as the number of i.i.d. samples grows large [VLBB08, TS15, GTS16].

In this paper, we outline a new strategy for clustering, and make progress on implementing it. We propose the following two questions to help generate new clustering algorithms:

1. How can we partition probability density functions? How can we do this so that two data points drawn from the same part of the partition are likely to be similar, and two data points drawn from different parts of the partition are likely to be dissimilar?
2. What clustering algorithms converge to such a partition, as the number of samples from the density function grows large?

In this paper, we address the first point, and make partial progress on the second. We focus on the special case of 2-way partitioning, which can be seen as finding a good cut on the probability density function. First, we propose a new notion of sparse (or isoperimetric) cuts on density functions. We call this an (α, β) -sparse cut, for real parameters α and β . Next, we propose a new notion of spectral sweep cuts on probability densities, called a (α, γ) -spectral sweep cut, for real parameters α and γ . We show that a (α, γ) -spectral sweep cut provably approximates an (α, β) -sparse cut when $\beta = \alpha + 1$ and $\gamma = \alpha + 2$. In particular, $\alpha = 1, \beta = 2, \gamma = 3$ is such a setting. Our result holds for any L -Lipschitz probability density function on \mathbb{R}^d , for any d . Based on past success applying sparse graph cuts to dividing data into two pieces in the machine learning setting [GS06, BN04], we believe our similarly defined $(\alpha = 1, \beta = 2)$ -sparse cuts may have similar ideal behavior when it comes to partitioning our density function.

To our knowledge, this is the first spectral method of cutting probability density functions that has any theoretical guarantee on the cut quality. The key mathematical contribution of this paper is a new Cheeger and Buser inequality for probability density functions, which we use to prove that (α, γ) -spectral sweep cuts approximate (α, β) -sparse cuts on probability density functions for the aforementioned settings of α, β, γ . These inequalities are inspired by the Cheeger and Buser inequalities on graphs and manifolds [AM84, Che70, Bus82], which have received considerable attention in graph algorithms and machine learning [Chu97, ST04, OSVV08, OV11, KW16, BN04]. These new inequalities do not directly follow from either the graph or manifold Cheeger inequalities, something we detail in Section 1.4. We note that our Cheeger and Buser inequalities for probability density functions require a careful definition of eigenvalue and sparse/isoperimetric cut: existing definitions lead to false inequalities.

Finally, our paper will present a discrete 2-way clustering algorithm that we suspect converges to the $(\alpha = 1, \gamma = 3)$ -spectral sweep cut as the number of data points grows large. Our algorithm bears similarity to classic spectral clustering methods, although we note that classical spectral clustering does not have any theoretical guarantees on the cluster quality. We note that we do not prove convergence of our discrete clustering method to the $(\alpha = 1, \gamma = 3)$ -spectral sweep cut, and leave this for future work.

1.1 Definitions

In this subsection, we define (α, β) sparsity, (α, γ) eigenvalues/Rayleigh quotients, and (α, γ) sweep cuts.

Definition 1.1. *Let ρ be a probability density function with domain \mathbb{R}^d , and let A be a subset of \mathbb{R}^d .*

The (α, β) -*sparsity* of the cut defined by A is the $(d - 1)$ dimensional integral of ρ^β on the cut, divided by the d dimensional integral of ρ^α on the side of the cut where this integral is smaller.

Definition 1.2. The (α, γ) -*Rayleigh quotient* of u with respect to $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is:

$$R_{\alpha, \gamma}(u) := \frac{\int_{\mathbb{R}^d} \rho^\gamma |\nabla u|^2}{\int_{\mathbb{R}^d} \rho^\alpha |u|^2}$$

A (α, γ) -*principal eigenvalue* of ρ is λ_2 , where:

$$\lambda_2 := \inf_{\int \rho^\alpha u = 0} R_{\alpha, \gamma}(u).$$

Define a (α, γ) -*principal eigenfunction* of ρ to be a function u such that $R_{\alpha, \gamma}(u) = \lambda_2$.

Now we define a sweep cut for a given function with respect to a positive valued function supported on \mathbb{R}^d :

Definition 1.3. Let α, β be two real numbers, and ρ be any function from \mathbb{R}^d to $\mathbb{R}_{\geq 0}$. Let u be any function from $\mathbb{R}^d \rightarrow \mathbb{R}$, and let C_t be the cut defined by the set $\{s \in \mathbb{R}^d | u(s) > t\}$.

The *sweep-cut* algorithm for u with respect to ρ returns the cut C_t of minimum (α, β) sparsity, where this sparsity is measured with respect to ρ .

When u is a (α, γ) -principal eigenfunction, the sweep cut is called a (α, γ) -*spectral sweep cut* of ρ .

Additional Definitions:

A function $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is *L-Lipschitz* if $|\rho(x) - \rho(y)| \leq L|x - y|_2$ for all $x, y \in \mathbb{R}^d$.

A function is $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ is α -*integrable* if $\int_{\mathbb{R}^d} \rho^\alpha$ is well defined and finite. Throughout this paper, we assume ρ is always α -integrable.

1.2 Past Work

Spectral Clustering and Sweep Cut Algorithms on Data

The spectral clustering algorithms of Shi and Malik [SM97] and those of Ng, Jordan, and Weiss [NJW01] are some of the most popular clustering algorithms on data (over 10,000 citations). If we want to split data points into two clusters, their algorithm works as follows: for n data points, compute an $n \times n$ matrix M on the data, and compute the principal eigenvector e of the matrix. Then, find a threshold value t such that all points p where $e(p) \leq t$ are considered to be on one side of the cut, and all other points where $e(p) > t$ are on the other. Often, the matrix is a Laplacian matrix of some graph built from the data [VL07].

Von Luxburg, Belkin, and Bosquet [VLBB08] proved that if the data is modeled as n i.i.d samples from a probability density ρ , the matrix M is a Laplacian matrix with certain structural assumptions, and certain regularity assumptions on ρ hold, then classical spectral clustering algorithms converge to a $(\alpha = 1, \gamma = 2)$ -spectral sweep cut on ρ [1]. These results were refined in [RBV10, GTS15, TS15]. We note that there are no sparsity guarantees known for a $(\alpha = 1, \gamma = 2)$ -spectral sweep cut, and we show 1-Lipschitz examples of ρ where this spectral sweep cut leads to undesirable behavior.

Cheeger Inequality and Sparse Cuts in Graphs

In 1984, Alon and Milman discovered a graph Cheeger inequality [AM84], which showed that a graph spectral sweep cut is approximately sparse. For a formal definition of sparse (or isoperimetric) cuts in a

¹We note that these authors used different terminology to describe this result, as their papers did not define (α, γ) -spectral sweep cuts.

graph theoretic setting, see [AM84]. The graph Cheeger inequality has guided decades of algorithmic and theoretical research on graph partitioning, random walks, and spectral graph theory in general [Chu97][KW16][OSVV08][OV11][ST04].

The graph Cheeger inequality implies that partitioning a graph based on the principal graph eigenvector (via spectral-sweep methods) will find a provably sparse cut [Chu97][OV11][OSVV08][LRTV12][LGT14b]. It also implies that a slowly mixing random walk is likely to yield good information about a sparse graph cut; this intuition was leveraged in Spielman and Teng’s seminal nearly-linear time algorithm on graph partitioning in [ST04]. Cheeger’s inequality for graphs has been an inspiration for decades of spectral graph theory research (for more information, see [Chu97]).

Sparse cuts have been researched extensively in graph theory [LR88][ARV04][CGR05][AP09][Mad10][LRTV12][KW16]. Sparse cuts, and the related notion of balanced cuts, have been used to generate fast multicommodity flow algorithms [LR88][ARV04]. These cuts are deeply related to expander decomposition [ST04][WN17][SW19], which in turn has proven immensely useful in algorithm design. For a brief overview of the application of sparse and balanced cuts in computer science theory, see the introduction of [Mad10]. For a survey of uses of expander decomposition, see the introduction of [SW19].

We note that expander decompositions are based on the idea that sparse cuts form natural partitions of the graph into clustered components. This intuition is leveraged in machine learning [BN04]. **Cheeger’s Inequality and Sparse Cuts for Manifolds**

The Cheeger inequality on a manifold states that the fundamental eigenvalue λ_2 of the Beltrami-Laplace operator is bounded below by the square of the sparse or isoperimetric cut, divided by 4 (See [Che70] for details). In 1982, Buser proved an upper bound for λ_2 provided that the manifold has lower bounded Ricci curvature [Bus82]. In contrast to the graph case, this inequality is false without the curvature assumption. This inequality has been widely used in the theory of differential geometry [VSCC08][HK00]. Intuition based on manifold Cheeger and manifold Buser has been used to great effect in semi-supervised learning and image processing [BN04][GS06].

For a formal definition of sparse or isoperimetric cut on manifolds, see [Che70] or [Bus82].

The Cheeger-Buser Inequality and Sparse Cuts for Convex Bodies and Density Functions

The Cheeger-Buser inequality, and the related notion of sparse cuts, have seen renewed interest in the computer science literature [LS90][LV18b][Mil09]. Past researchers have used ideas inspired by sparse cuts and the Cheeger-Buser inequality to recover fast mixing time lemmas for random walks for log-concave density functions supported on convex polytopes [LS90][DFK91][KLS95][LV18b][GM83]. These lemmas have in turn been used to find fast sampling algorithms from such density functions [LV18b], and more.

One prominent use of sparse cuts in convex geometry is the celebrated Kannan-Lovasz-Simonovits Conjecture (KLS Conjecture). This conjecture asserts that, for a convex body, the sparsity of the sparsest hyperplane cut is a dimension-less constant approximation of the sparsity of the optimal sparse cut [KLS95][LV18a]. Generally, these results all implicitly use the settings ($\alpha = 1, \beta = 1, \gamma = 1$) for their definition of sparse cuts, eigenvalues, and eigenvectors. They operate in the setting where there is significant structure on the probability, such as being a log-concave distribution [LV18b]. We note that the Cheeger-Buser inequality fails for this setting of α, β, γ when the probability density is Lipschitz but not log-concave. For a survey on uses of Cheeger’s inequality in convex polytope theory and log-concave density function sampling, see [LV18a].

Discrete Machine Learning from Continuous Methods

Our overall approach to clustering follows the line of work generating discrete machine learning methods by analyzing its behavior in the limit. This approach was used fruitfully by Su, Boyd, and Candes [SBC14] and Wibisono, Wilson, and Jordan [WWJ16] to generate faster gradient descent variants based on continuous time differential equations. For more information, refer to their respective papers.

1.3 Contributions

Our paper has three core contributions:

1. A natural method for cutting probability density functions, based on a new notion of sparse cuts on density functions.
2. A Cheeger and Buser inequality for Lipschitz probability density functions, and
3. A clustering algorithm operating on samples, which heuristically approximates a spectral sweep cut on the density function when the number of samples grows large.

We emphasize that our primary contributions are points 1 and 2, which are formally stated in Theorems [1.4](#) and [1.5](#) respectively. Our clustering algorithm on samples, which is designed to approximate the (α, γ) -spectral sweep cut on the density function as the number of samples grows large, is of secondary importance.

We now state our two main theorems.

Theorem 1.4. *Spectral Sweep Cuts give Sparse Cuts:*

Let $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be an L -Lipschitz probability density function, and let $\alpha = \beta - 1 = \gamma - 2$.

The (α, γ) -spectral sweep cut of ρ has (α, β) sparsity Φ satisfying:

$$\Phi_{OPT} \leq \Phi \leq O(\sqrt{dL\Phi_{OPT}}).$$

Here, Φ_{OPT} refers to the optimal (α, β) sparsity of a cut on ρ .

In words, the spectral sweep cut of the (α, γ) eigenvector gives a provably good approximation to the sparsest (α, β) cut, as long as $\beta = \alpha + 1$ and $\gamma = \alpha + 2$. Proving this result is a straightforward application of two new inequalities we present, which we will refer to as the Cheeger and Buser inequalities for probability density functions.

Theorem 1.5. *Probability Density Cheeger and Buser:*

Let $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be an L -Lipschitz density function. Let $\alpha = \beta - 1 = \gamma - 2$.

Let Φ be the infimum (α, β) -sparsity of a cut through ρ , and let λ_2 be the (α, γ) -principal eigenvalue of ρ . Then:

$$\Phi^2/4 \leq \lambda_2$$

and

$$\lambda_2 \leq O_{\alpha, \beta}(d \max(L\Phi, \Phi^2)).$$

The first inequality is **Probability Density Cheeger**, and the second inequality is **Probability Density Buser**.

Note that we don't need ρ to have a total mass of 1 for any of our proofs. The overall probability mass of ρ can be arbitrary.

We note that Theorem [1.5](#) has a partial converse:

Lemma 1.6. *If $\alpha + \gamma > 2\beta$ or $\gamma - 1 < \beta$, then the Cheeger-Buser inequality in Theorem [1.5](#) does not hold.*

In particular, if $\alpha = \gamma = 1$ or $\alpha = 1, \gamma = 2$, no Cheeger-Buser inequality can hold for any β . These settings of α and γ encompass most past work on spectral cuts and Cheeger-Buser inequalities on probability density functions.

Finally, we give a discrete algorithm 1,3-SPECTRALCLUSTERING for clustering data points into two-clusters. We conjecture, but do not prove, that 1,3-SPECTRALCLUSTERING converges to the $(\alpha = 1, \gamma = 3)$ -spectral sweep cut of the probability density function ρ as the number of samples grows large.

We note that this resembles the unnormalized spectral clustering based on the work of Shi and Malik [SM97](#) and Ng, Weiss, and Jordan [NJW01](#), [TS15](#). The major difference is that we build our Laplacian from

1,3-SPECTRALCLUSTERING

Input: Point $s_1, \dots, s_n \in \mathbb{R}^d$, and similarity measure $K : \mathbb{R}^d, \mathbb{R}^d \rightarrow \mathbb{R}$.

1. Form the affinity matrix $A \in \mathbb{R}^{n \times n}$, where $A_{ij} = e^{-n^{2/d} \|s_i - s_j\|^2}$ for $i \neq j$ and $A_{ii} = 0$ for all i .
2. Define D to be the diagonal matrix whose (i, i) element is the sum of A 's i^{th} row. Let L be the Laplacian formed from the adjacency matrix $D^{1/2}AD^{1/2}$.
3. Let u be the principal eigenvector of L . Find the value t where $t := \arg \min_s \Phi_{\{u(v) > t\}}$, where Φ_S is the graph conductance of the cut defined by set S .

Output: Clusters $G_1 = \{v : u(v) > t\}, G_2 = \{v : u(v) \leq t\}$.

the matrix $D^{1/2}AD^{1/2}$. Past work on spectral clustering builds the Laplacian from the matrices A or $D^{-1/2}AD^{-1/2}$ [VL07, TS15].

1.4 Differences between our work and past work

Our work differs from past work in the following key ways:

1. Our work differs from past practical work on spectral sweep cuts [GTS15, TS15, SM97, NJW01], as those methods perform what we call a $(\alpha = 1, \gamma = 2)$ -sweep cut. These sweep cuts have no theoretical guarantees, much less a guarantee on their (α, β) sparsity. Lemma 1.6 shows that no Cheeger and Buser inequality can simultaneously hold for any setting of β when $\alpha = 1, \beta = 2$.

We will further show that using a $(\alpha = 1, \gamma = 2)$ -sweep cut can lead to undesirable cuts of 1-Lipschitz probability densities, with poor sparsity guarantees.

2. We note that probability density Cheeger-Buser is not easily implied by graph or manifold Cheeger-Buser. For a lengthier discussion on this, see Appendix A.
3. We do not require any assumptions on our probability density except that it is Lipschitz. Past work on Cheeger-Buser inequality for densities focused on log-concave distributions, or mixtures thereof [LV18a, GLR18].
4. For our work, the probability density ρ is not required to be bounded away from 0. This is a sharp departure from many existing results: past results on partitioning probability densities required a positive lower bound on ρ [VLBB08, GTS15]. The strongest results in fields like linear elliptic partial differential equations depend on ρ being bounded away from 0 [Whi17].

Our work is the first spectral sweep cut algorithm that guarantees a sparse cut on Lipschitz densities ρ , without requiring strong parametric assumptions on ρ .

1.4.1 Technical Contribution

The key technical contribution of our proof is proving Buser's inequality on Lipschitz probability densities via mollification [Sob38, Fri44] with disks of varying radius. This paper is the first time mollification with disks of varying radius have been used. We emphasize that the most difficult part of our paper is proving the Buser inequality.

Mollification has a long history in mathematics dating back to Sergei Sobolev's celebrated proof of the Sobolev embedding theorem [Sob38]. It is one of the key tools in numerical analysis, partial differential equations, fluid mechanics, and functional analysis [Fri44, LW01, SS09, Mon03], and analogs of mollification have been used in computational complexity settings [DKN10]. Informally speaking, mollification is used to create a series of smooth functions approximating a non-smooth function, by convolving the original function with a smooth function supported on a disk. Notably, an approach using convolution is used by

Buser in [Bus82] to prove the original Buser’s inequality, albeit with an intricate pre-processing step on any given cut.

To prove Buser’s inequality on Lipschitz probability density functions ρ , we will show that given a cut C with low (α, β) -sparsity, we can find a function u with low (α, γ) -Rayleigh quotient. We build u by starting with the indicator function I_C for cut C (which is 1 on one side of the cut and 0 on the other). Next, we mollify this function with disks of varying radius. In particular, for each point r in the domain of ρ , we spread out the point mass $I_C(r)$ over a disk of radius proportional to $\rho(r)L$, where L is the Lipschitz constant of ρ . The resulting function u obtained by ‘spreading out’ I_C will have low (α, γ) -Rayleigh quotient.

For all past uses of mollification, the disks on which the smooth convolving function is supported (we call this the mollification disk) have the same radius throughout the manifold. The use of a uniform radius disk is critical for most uses and proofs in mollification. Our contribution is to allow the disks to vary in radius across our density. This variation in radius allow us to deal with functions that approach 0 (and explains the importance of the density being Lipschitz). No mollification disks centered anywhere in our probability density will intersect the 0-set of the density. This overcomes significant hurdles in many results for functional analysis and PDEs, as many past significant results related to partial differential equations rely on having a positive lower bound on the density [Whi17] [GTS15].

Proving our Buser inequality using mollification by disks of various radius requires a fairly delicate proof with many pages of calculus. Our key technical lemma is a bound on how the l_1 norm of a mollified function when the mollification disks have various radius, which can be found in Section 3.3

2 Paper Organization

We prove the Buser inequality in Section 3 via a rather extensive series of calculus computations. Our proof relies on a key technical lemma, which is presented in Section 3.3. This is by far the most difficult part of our proof.

We prove the Cheeger inequality in Section 4. The proof in this section implies that the $(\alpha, \alpha + 1)$ sparsity of the $(\alpha, \alpha + 2)$ spectral sweep cut of a probability density function ρ is provably close to the $(\alpha, \alpha + 2)$ principal eigenvalue of ρ . We note that this inequality does not depend on the Lipschitz nature of the probability density function.

In Section 5, we prove Theorem 1.4 which shows that a $(\alpha, \alpha + 2)$ spectral sweep cut has $(\alpha, \alpha + 1)$ sparsity which provably approximates the optimal $(\alpha, \alpha + 1)$ sparsity.

In Section 6, we go over example 1-D distributions that show that either Cheeger or Buser inequality must fail for past definitions of sparsity and eigenfunctions. We will prove Lemma 1.6 in this section.

In Section 7 we show an example Lipschitz probability density where the $(\alpha = 1, \gamma = 2)$ spectral sweep cut has bad $(1, \beta)$ sparsity for any $\beta < 10$, and will lead to an undesirable cut (from a clustering point of view) on this density function. This is important since the spectral clustering algorithm of Ng et al [NJV01] is known to converge to a $(\alpha = 1, \gamma = 2)$ spectral sweep cut on the underlying probability density function, as the number of samples grows large [TS15].

Finally, we state conclusions and open problems in Section 8.

In the appendix, we note that the Cheeger and Buser inequalities for probability densities are not easily implied by graph or manifold Cheeger-Buser. We also provide a simplified version of Cheeger’s and Buser’s inequality for probability densities, in the 1-dimensional case. This may make easier reading for those unfamiliar with technical multivariable mollification.

3 Buser Inequality for Probability Density Functions

The key idea to proving Buser's inequality is as follows: given $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, and a cut u where $u \subset \mathbb{R}^d$, we will build a function u_θ whose (α, γ) -Rayleigh quotient is close to the (α, β) sparsity of u .

Roughly speaking, u_θ is built by convolving u at point x with a unit-weight disk with radius proportional to $\rho(x)$. Thus, we are convolving u with a disk, where the radius of the disk varies between points $x \in \mathbb{R}^d$, and the radius is directly proportional to ρ .

3.1 Weighted Buser-type Inequality

We now prove our weighted Buser-type inequality, from Theorem [1.5](#). We state our result in terms of general (α, β, γ) .

Theorem 3.1. *Let $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be an L -Lipschitz function, λ_2 be a (α, γ) -principal eigenvalue, and Φ the (α, β) isoperimetric cut.*

Then:

$$\lambda_2 \leq 3 \cdot 2^{\beta+1} d \|\rho^{\gamma-\beta-1}\|_{L^\infty} \max(L\Phi(A), 2^{\beta+1} \|\rho^{\alpha+1-\beta}\|_{L^\infty} \Phi(A)^2)$$

We note that when setting $(\alpha, \beta, \gamma) = (1, 2, 3)$, the above expression simplifies into:

$$\lambda_2 \leq 24d \max(L\Phi(A), 8\Phi(A)^2).$$

3.2 Proof Strategy: Mollification by Disks of Radius Proportional to ρ

To prove Theorem [3.1](#) we construct an approximation u_θ of u for which the numerator and denominator of the Raleigh quotient, $R(u_\theta)$, approximate respectively the numerator and denominator of this expression. Specifically, u_θ will be constructed as a mollification of u . Recall the following two equivalent definitions of a mollification. They are equivalent by the change of variables $z = x - \theta\rho(x)y$.

$$u_\theta(x) := \int_{B(0,1)} u(x - \theta\rho(x)y) \phi(y) dy = \int u(z) \phi_{\theta\rho(x)}(x - z) dz, \quad \text{where} \quad \phi_\eta(z) = \frac{1}{\eta^d} \phi\left(\frac{z}{\eta}\right), \quad (1)$$

with $\theta > 0$ a parameter to be chosen and $\phi : \mathbb{R}^d \rightarrow [0, \infty)$ a smooth radially symmetric function supported in the unit open ball $B(0, 1) = \{x \in \mathbb{R}^d \mid |x| < 1\}$ with unit mass $\int_{\mathbb{R}^d} \phi = 1$. When ρ is constant it follows from the Tonelli theorem that $\|u_\theta\|_{L^1} = \|u\|_{L^1}$; when ρ is not constant the following lemma shows that the latter still bounds the former.

3.3 Key Technical Lemma: Bounding L_1 norm of a function with the L_1 norm of its mollification

The following is our primary technical lemma, which roughly bounds the L_1 norm of a mollified function f by the L_1 norm of the original f . Here, the mollification radius is determined by a function $\delta(x)$.

Lemma 3.2. *Let $\delta : \mathbb{R}^d \rightarrow \mathbb{R}$ be Lipschitz continuous with Lipschitz constant $|\nabla\delta(x)| \leq c < 1$ for almost every $x \in \mathbb{R}^d$. Let $\phi : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be smooth, $\int_{\mathbb{R}^d} \phi = 1$, and $\text{supp}(\phi) \subseteq B(0, 1)$. Then*

$$\frac{1}{1+c} \|f\|_{L^1} \leq \int_{\mathbb{R}^d} \int_{B(0,1)} |f(x - \delta(x)y)| \phi(y) dy dx \leq \frac{1}{1-c} \|f\|_{L^1}, \quad f \in L^1(\mathbb{R}^d).$$

Proof. (of Lemma 3.2) An application of Tonelli's theorem shows

$$\int_{\mathbb{R}^d} \int_{B(0,1)} |f(x - \delta(x)y)| \phi(y) dy dx = \int_{B(0,1)} \phi(y) \int_{\mathbb{R}^d} |f(x - \delta(x)y)| dx dy. \quad (2)$$

Fix $y \in B(0,1)$ and consider the change of variables $z = x - \delta(x)y$. The Jacobian of this mapping is $I - y \otimes \nabla \delta(x)$ which by Sylvester's determinant theorem has determinant $1 - y \cdot \nabla \delta(x) > 0$. It follows that

$$\int_{\mathbb{R}^d} \int_{B(0,1)} |f(x - \delta(x)y)| \phi(y) dy dx = \int_{B(0,1)} \phi(y) \int_{\mathbb{R}^d} \frac{|f(z)|}{1 - y \cdot \nabla \delta(x)} dx dy, \quad (3)$$

and the lemma follows since $1 - c \leq 1 - y \cdot \nabla \delta(x) \leq 1 + c$. \square

(Here, $a \cdot b$ denotes the dot product between a and b .)

We present the following simple corollaries, which is the primary way our proof makes use of Lemma 3.2

Corollary 3.3. *For any Lipschitz continuous function $\rho : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq 0}$ with Lipschitz constant L and any θ with $0 < \theta L < 1$, we have:*

$$\begin{aligned} & \frac{1}{1 + \theta L} \|\rho^\beta(x) \nabla u(x)\|_{L^1} \\ & < \int_{\mathbb{R}^d} \int_{B(0,1)} \rho^\beta(x - \theta \rho(x)y) |\nabla u(x - \theta \rho(x)y)| \phi(y) dy dx \\ & < \frac{1}{1 - \theta L} \|\rho^\beta(x) \nabla u(x)\|_{L^1} \end{aligned}$$

Proof. (of Corollary 3.3) Apply Lemma 3.2 with $\delta(x) = \theta \rho(x)$, and $f(x) = \rho^\beta(x) \nabla u(x)$. \square

This corollary will be used to bound the numerator of our Rayleigh quotient. Note that the expression

$$\int_{\mathbb{R}^d} \int_{B(0,1)} \rho^\beta(x - \theta \rho(x)y) |\nabla u(x - \theta \rho(x)y)| \phi(y) dy dx$$

is close to $\int_{\mathbb{R}^d} \rho^\beta(x) \nabla |u_\theta(x)| dx$ when $\theta \leq \frac{1}{2L}$. This is the guiding intuition behind how Corollary 3.3 and Lemma 3.2 will be used, and will be formalized later in our proof of Theorem 3.1

We present another simple corollary whose proof is equally straightforward. This corollary will be used to bound the denominator, and is a small generalization of Corollary 3.3. We write down both corollaries anyhow, since this will make it easier to interpret our bounds on the Rayleigh quotient.

Corollary 3.4. *For any Lipschitz continuous function $\rho : \mathbb{R}^d \rightarrow \mathbb{R}^{\geq 0}$ with Lipschitz constant L , any $0 < t < 1$, and any θ with $0 < \theta L < 1$, we have:*

$$\begin{aligned} & \frac{1}{1 + \theta L} \|\rho^\beta(x) \nabla u(x)\|_{L^1} \\ & < \int_{\mathbb{R}^d} \int_{B(0,1)} \rho^\beta(x - \theta t \rho(x)y) |\nabla u(x - \theta t \rho(x)y)| \phi(y) dy dx \\ & < \frac{1}{1 - \theta L} \|\rho^\beta(x) \nabla u(x)\|_{L^1} \end{aligned}$$

Proof. (of Corollary 3.4) Apply Lemma 3.2 with $\delta(x) = \theta t \rho(x)$, and $f(x) = \rho^\beta(x) \nabla u(x)$. \square

Now we are ready to prove our main Theorem, which is the Buser inequality for probability densities stated in Theorem 3.1

Proof. (of Theorem [3.1](#))

Fix $A \subset \mathbb{R}^d$ with $|A|_\alpha \leq |1|_\alpha/2$ and let $u(x) = \chi_A(x)$ be the characteristic function of A . Setting \bar{u} to be the weighted average of u ,

$$\bar{u} = \frac{\int \rho^\alpha u}{\int \rho^\alpha} = \frac{\int_A \rho^\alpha}{\int \rho^\alpha} = \frac{|A|_\alpha}{|1|_\alpha} \in [0, 1/2], \quad \text{then} \quad \int \rho^\alpha (u - \bar{u}) = 0,$$

and

$$\|\rho^\alpha (u - \bar{u})\|_{L^1} = \int \rho^\alpha |u - \bar{u}| = 2|A|_\alpha (1 - \bar{u}) = 2 \int \rho^\alpha |u - \bar{u}|^2. \quad (4)$$

Since $|A|_\alpha = \|u\|_{L^1_\alpha}$ and $1 - \bar{u} \in [0, 1/2]$ it follows that

$$(1/2) \frac{\|\rho^\beta \nabla u\|_{L^1}}{\|\rho^\alpha (u - \bar{u})\|_{L^1}} \leq \Phi(A) = \frac{\|\rho^\beta \nabla u\|_{L^1}}{\|\rho^\alpha u\|_{L^1}} \leq \frac{\|\rho^\beta \nabla u\|_{L^1}}{\|\rho^\alpha (u - \bar{u})\|_{L^1}}. \quad (5)$$

In the calculations below we omit the limiting argument with smooth approximations of u outlined at the beginning of this section which justify formula involving ∇u , and for readability frequently write ρ and $\nabla \rho$ for $\rho(x)$ and $\nabla \rho(x)$.

Next, let u_θ be the mollification of (an extension of) u given by equation [\(1\)](#). Then $u_\theta(x)$ is a local average average of u so $u_\theta(x) \geq 0$, $\|u_\theta\|_{L^\infty} \leq 1$ and $\|u - u_\theta\|_{L^\infty} \leq 1$. Letting L denote the Lipschitz constant of ρ , the parameter θ will to be chosen less than $1/(2L)$ so that that Lemma [3.2](#) is applicable with constant $c = 1/2$.

The remainder of the proof constructs an upper bound on the numerator $\int_{\mathbb{R}^d} \rho^\gamma |\nabla u_\theta|^2$ of the Raleigh quotient for $u_\theta - \bar{u}_\theta$ by $\|\rho^\beta \nabla u\|_{L^1}$ and to lower bound the denominator $\int_{\mathbb{R}^d} \rho^\alpha (u_\theta - \bar{u}_\theta)^2$ by $\|\rho^\alpha (u - \bar{u})\|_{L^1}$. The conclusion of the theorem then follows from equation [\(5\)](#).

3.4 Upper Bounding the Numerator

To bound the L^2 norm in the numerator of the Raleigh quotient by the L^1 norm in the numerator of the expression for $\Phi(A)$ it is necessary to obtain uniform bound on $\rho(x)\nabla u_\theta(x)$.

Lemma 3.5. *Let u be any function, and let u_θ be defined as in Equation [1](#). Let $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be an L -Lipschitz function.*

$$\|\rho(x)\nabla u_\theta(x)\|_{L^\infty} \leq \|u\|_{L^\infty} \frac{d(2+3L)}{\theta} \quad (6)$$

Proof. In order to prove this lemma, we first need to get a handle on $\nabla u_\theta(x)$, which is the gradient of u after mollification by θ .

We take the the second representation of u_θ in equation [\(1\)](#) to get

$$\nabla u_\theta(x) = \int_{\mathbb{R}^d} u(z) \left\{ \frac{-d}{\theta \rho(x)} \phi_{\theta \rho}(x-z) \nabla \rho + \frac{1}{(\theta \rho(x))^{d+1}} \left(I + \nabla \rho(x) \otimes \frac{x-z}{\theta \rho(x)} \right) \nabla \phi \left(\frac{x-z}{\theta \rho(x)} \right) \right\} dz, \quad (7)$$

which is a consequence of the multivariable chain rule. Here, $v \otimes u$ refers to the outer product of v and u .

Multiplying by ρ gives:

$$\rho(x)\nabla u_\theta(x) = \int_{\mathbb{R}^d} u(z) \left\{ \frac{-d}{\theta} \phi_{\theta \rho}(x-z) \nabla \rho(x) + \frac{1}{(\theta^{d+1} \rho(x)^d)} \left(I + \nabla \rho(x) \otimes \frac{x-z}{\theta \rho(x)} \right) \nabla \phi \left(\frac{x-z}{\theta \rho(x)} \right) \right\} dz. \quad (8)$$

Now, we can bound the above equation by carefully bounding each part. We note:

$$\int_{\mathbb{R}^d} \frac{1}{(\theta^{d+1}\rho(x)^d)} \nabla \phi \left(\frac{x-z}{\theta\rho(x)} \right) dz \quad (9)$$

$$= \int_{\mathbb{R}^d} \frac{1}{(\theta^{d+1}\rho(x)^d)} \nabla \phi \left(\frac{-z}{\theta\rho(x)} \right) dz \quad (10)$$

$$= \frac{1}{\theta} \int_{\mathbb{R}^d} \nabla \phi(-y) dy \quad (11)$$

where the last step follows by a simple change of variable. Here, we note that $\nabla \phi(y)$ is a vector, and the integral is over \mathbb{R}^d , which is how we eliminated $\frac{1}{(\theta\rho(x))^d}$ from the expression.

Next, we examine the term:

$$I + \nabla \rho(x) \otimes \frac{x-z}{\theta\rho(x)} \quad (12)$$

Here, we aim to bound the operator norm of this matrix. Here, we note that

$$|x-z| \leq \theta\rho(x)$$

when

$$\nabla \phi \left(\frac{x-z}{\theta\rho(x)} \right) \neq 0$$

and thus, when the latter equation holds, we can say:

$$\left| \frac{x-z}{\theta\rho(x)} \right| < 1.$$

Since $|\nabla \rho(x)| < L$, we now have:

$$|I + \nabla \rho(x) \otimes \frac{x-z}{\theta\rho(x)}|_2 < 3/2 \quad (13)$$

Combining Equation [13](#) Equation [9](#) to show:

$$\left| \int_{\mathbb{R}^d} \frac{1}{(\theta^{d+1}\phi(x)^d)} \left(I + \nabla \rho(x) \otimes \frac{x-z}{\theta\rho(x)} \right) \nabla \phi \left(\frac{x-z}{\theta\rho(x)} \right) dz \right| \quad (14)$$

$$\leq \frac{(1+L)}{\theta} \int_{\mathbb{R}^d} |\nabla \phi(y)| dy, \quad (15)$$

where $L = \|\nabla \rho\|_{L^\infty}$ is the Lipschitz constant for ρ . We note that Section [3.7](#) shows that

$$\int_{\mathbb{R}^d} |\nabla \phi(y)| dy \leq 2d. \quad (16)$$

and therefore:

$$\left| \int_{\mathbb{R}^d} \frac{1}{(\theta^{d+1}\phi(x)^d)} \left(I + \nabla \rho(x) \otimes \frac{x-z}{\theta\rho(x)} \right) \nabla \phi \left(\frac{x-z}{\theta\rho(x)} \right) dz \right| \quad (17)$$

$$\leq \frac{2d(1+L)}{\theta} \quad (18)$$

Now we turn our attention to the first term, which is:

$$\int_{\mathbb{R}^d} \frac{-d}{\theta} \phi_{\theta\rho(x)}(x-z) \nabla \rho(x) dz \quad (19)$$

We note that

$$\int_{\mathbb{R}^d} |\phi_{\theta\rho(x)}(x-z)| dz = 1$$

by our definition of ϕ (which was defined when we defined u_θ). Combining this with $|\nabla\rho(x)| < L$, we get:

$$\int_{\mathbb{R}^d} \left| \frac{-d}{\theta} \phi_{\theta\rho(x)}(x-z) \nabla\rho(x) \right| dz \tag{20}$$

$$< \frac{dL}{\theta} \tag{21}$$

Therefore,

$$\begin{aligned} & \left| \int_{\mathbb{R}^d} \frac{-d}{\theta\rho(x)} \phi_{\theta\rho(x)}(x-z) \nabla\rho + \frac{1}{(\theta\rho(x))^{d+1}} \left(I + \nabla\rho(x) \otimes \frac{x-z}{\theta\rho(x)} \right) \nabla\phi \left(\frac{x-z}{\theta\rho(x)} \right) dz \right| \\ & \leq \frac{d}{\theta} (L + 2(1+L)) \\ & = \frac{d(2+3L)}{\theta}. \end{aligned} \tag{22}$$

where the first inequality comes from combining Equations [18](#) and [21](#)

This allows us to bound $\|\rho(x)\nabla u_\theta(x)\|_{L^\infty}$:

$$\|\rho(x)\nabla u_\theta(x)\|_{L^\infty} \tag{23}$$

$$= \left\| \int_{\mathbb{R}^d} u(z) \left\{ \frac{-d}{\theta} \phi_{\theta\rho(x)}(x-z) \nabla\rho(x) + \frac{1}{(\theta\rho(x))^{d+1}} \left(I + \nabla\rho(x) \otimes \frac{x-z}{\theta\rho(x)} \right) \nabla\phi \left(\frac{x-z}{\theta\rho(x)} \right) \right\} dz \right\|_{L^\infty} \tag{24}$$

$$\leq \|u\|_{L^\infty} \left\| \int_{\mathbb{R}^d} \left| \frac{-d}{\theta} \phi_{\theta\rho(x)}(x-z) \nabla\rho(x) + \frac{1}{(\theta\rho(x))^{d+1}} \left(I + \nabla\rho(x) \otimes \frac{x-z}{\theta\rho(x)} \right) \nabla\phi \left(\frac{x-z}{\theta\rho(x)} \right) \right| dz \right\|_{L^\infty} \tag{25}$$

$$\leq \|u\|_{L^\infty} \frac{d(2+3L)}{\theta} \tag{26}$$

where we make use of the fact that $\|ab\|_{L^\infty} < \|a\|_{L^\infty} \|b\|_{L^1}$. This completes our proof. \square

Next, we want an L_1 bound on $\rho^\beta(x)\nabla u_\theta(x)$.

Lemma 3.6. *Let u be any function, and let u_θ be defined as in Equation [1](#). Let $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$ be an L -Lipschitz function, and let $\theta L < 1/2$.*

Then:

$$\|\rho^\beta(x)\nabla u_\theta(x)\|_{L^1} \leq C_\beta \|\rho^\beta(x)\nabla u(x)\|_{L^1} \tag{27}$$

Proof. First, we take the gradient first representation of u_θ in equation [1](#). Using the chain rule gives us an alternate form for $\nabla u_\theta(x)$:

$$\nabla u_\theta(x) = \int_{\mathbb{R}^d} (I - \theta\nabla\rho \otimes y) \nabla u(x - \theta\rho y) \phi(y) dy, \tag{28}$$

so

$$\rho^\beta(x)\nabla u_\theta(x) = \int_{\mathbb{R}^d} (I - \theta\nabla\rho \otimes y) \frac{\rho^\beta(x)}{\rho^\beta(x - \theta\rho y)} \rho^\beta(x - \theta\rho y) \nabla u(x - \theta\rho y) \phi(y) dy. \tag{29}$$

The ratio in the integrand is bounded using the Lipschitz assumption on ρ (and $|y| \leq 1$),

$$\frac{\rho(x)}{\rho(x - \theta\rho y)} \leq \frac{\rho(x)}{\rho(x) - L\theta\rho(x)} = \frac{1}{1 - L\theta} \leq 2, \quad \text{when } \theta < 1/(2L). \quad (30)$$

Note that

$$\|I - \theta\nabla\rho \otimes y\|_2 \leq 3/2 \quad (31)$$

where $\|M\|_2$ represents the ℓ^2 matrix norm of M . This is because $|\nabla\rho(x)| \leq L$, and $\theta L < 1/2$, and $|y| \leq 1$ every time $\phi(y) \neq 0$, and thus

$$\frac{I}{2} \preceq I - \theta\nabla\rho \otimes y \preceq \frac{3I}{2}.$$

Therefore, we can now apply Corollary [3.3](#) to Equation [\(29\)](#) to show:

$$\begin{aligned} & \|\rho^\beta(x)\nabla u_\theta(x)\|_{L^1} \\ & \leq \|I - \theta\nabla\rho(x) \otimes y\|_2 \cdot \max_x \left(\frac{\rho(x)}{\rho(x - \theta\rho y)} \right) \cdot \int_{\mathbb{R}^d} \left| \int_{\mathbb{R}^d} \rho^\beta(x - \theta\rho y) \nabla u(x - \theta\rho y) \phi(y) dy \right| \\ & \leq 3 \cdot 2^{\beta-1} \int_{\mathbb{R}^d} \rho^\beta(x - \theta\rho(x)y) \nabla u(x - \theta\rho(x)y) \\ & \leq 3 \cdot 2^\beta \|\rho^\beta \nabla u\|_{L^1}, \quad \text{when } \theta < 1/(2L). \end{aligned}$$

here, the first inequality comes from the equation $\|abc\|_{L^1} \leq \|a\|_{L^\infty} \|b\|_{L^\infty} \|c\|_{L^1}$, the second inequality comes from Equations [30](#) and [31](#) and the third inequality comes from Corollary [3.3](#) assuming $\theta L \leq 1/2$. \square

Lemma 3.7. *For any L -Lipschitz distribution ρ , any function u , and any θ such that $\theta L < 1/2$:*

$$\int_{\mathbb{R}^d} \rho^\gamma |\nabla u_\theta|^2 \leq C_\beta \|\rho^{\gamma-\beta-1}\|_{L^\infty} \frac{d(2+3L)}{\theta} \|u\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1}, \quad (32)$$

Proof. Combining the two estimates from Lemma [3.5](#) and [3.6](#) gives an upper bound for the Raleigh quotient

$$\int_{\mathbb{R}^d} \rho^\gamma |\nabla u_\theta|^2 = \int_{\mathbb{R}^d} \rho^{\gamma-\beta-1} \rho |\nabla u_\theta| \rho^\beta |\nabla u_\theta| \leq 3 \cdot 2^{\beta+1} \|\rho^{\gamma-\beta-1}\|_{L^\infty} \frac{d(2+3L)}{\theta} \|u\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1}, \quad (33)$$

\square

We note that in the case where $\gamma = \beta + 1$, and if u is a step function, the expression would simplify to:

$$\int_{\mathbb{R}^d} |\rho^\gamma |\nabla u_\theta|^2 \leq 3 \cdot 2^{\beta+1} \frac{d(2+3L)}{\theta} \|u\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1},$$

3.5 Lower Bound on the Demoninator

Let \bar{u} and \bar{u}_θ be the ρ^α -weighted averages of u and u_θ and let $\|\cdot\|_{L_\alpha^2}$ denote the L^2 space with this weight. Our core lemma is a bound on $\|u_\theta - \bar{u}_\theta\|_{L_\alpha^2}$ in terms of l_1 and weighted l_1 norms of ∇u and $u - \bar{u}$ respectively.

Lemma 3.8. *Let ρ be an L -Lipschitz function $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, and let θ be such that $\theta L < 1/2$. Let u be an indicator function of a set A with finite β -perimeter. Let \bar{u} be defined as $\bar{u}(x) := u(x) - \int u(y) dy$ u_θ be defined as in Equation [1](#) and \bar{u}_θ be defined as $\bar{u}_\theta(x) := u_\theta(x) - \int u_\theta(y) dy$. Then:*

$$\|u_\theta - \bar{u}_\theta\|_{L_\alpha^2}^2 \geq (1/4) \|u - \bar{u}\|_{L_\alpha^1} - C(\beta)\theta \|\rho^{\alpha+1-\beta}\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1}, \quad \text{when } \theta < 1/(2L). \quad (34)$$

Note that when $\alpha + 1 = \beta$, as is true when $(\alpha, \beta, \gamma) = (1, 2, 3)$, the inequality in Lemma 3.8 becomes:

$$\|u_\theta - \bar{u}_\theta\|_{L_\alpha^2}^2 \geq (1/4) \|u - \bar{u}\|_{L_\alpha^1} - C(\beta)\theta \|\rho^\beta \nabla u\|_{L^1}, \quad \text{when } \theta < 1/(2L).$$

The estimate in Lemma 3.8 will be combined with the estimate in Lemma 3.7 to prove Theorem 3.1 in Section 3.6

Proof. The key to this proof is to upper bound the quantity $\|u_\theta - \bar{u}_\theta\|_{L_\alpha^2}$ with the expression appearing in Corollary 3.4. We will do so by a series of inequalities, application of the fundamental theorem of calculus, and more.

Using the property that subtracting the average from a function reduces the L^2 norm it follows that

$$\begin{aligned} & \|u_\theta - \bar{u}_\theta\|_{L_\alpha^2} \\ & \geq \|u - \bar{u}\|_{L_\alpha^2} - \|u_\theta - u - (\bar{u}_\theta - \bar{u})\|_{L_\alpha^2} \\ & \geq \|u - \bar{u}\|_{L_\alpha^2} - \|u_\theta - u\|_{L_\alpha^2}. \end{aligned}$$

If $a \geq b - c$ then $a^2 \geq b^2/2 - c^2$, so a lower bound for the denominator of the Raleigh quotient

$$\begin{aligned} & \|u_\theta - \bar{u}_\theta\|_{L_\alpha^2}^2 \\ & \geq (1/2) \|u - \bar{u}\|_{L_\alpha^2}^2 - \|u_\theta - u\|_{L_\alpha^2}^2 \\ & \geq (1/4) \|u - \bar{u}\|_{L_\alpha^1} - \|u_\theta - u\|_{L_\alpha^1}, \end{aligned} \tag{35}$$

where the identity $\|u - \bar{u}\|_{L_\alpha^2}^2 = \|u - \bar{u}\|_{L_\alpha^1}^2 / 2$ from Equation 4 and the bound $\|u_\theta - u\|_{L^\infty} \leq 1$, were used in the last step.

It remains to estimate the difference $\|u_\theta - u\|_{L_\alpha^1}$. To do this, we use the multivariable fundamental theorem of calculus to write

$$\begin{aligned} u_\theta(x) - u(x) &= \int (u(x - \theta\rho y) - u(x))\phi(y) dy \\ &= \int \int_0^1 -\theta\rho(x)\nabla u(x - t\theta\rho(x)y) \cdot y\phi(y) dt dy \\ &= \int \int_0^1 \frac{-\theta\rho(x)}{\rho^\beta(x - t\theta\rho(x)y)} \rho^\beta(x - t\theta\rho(x)y)\nabla u(x - t\theta\rho(x)y) \cdot y\phi(y) dt dy, \end{aligned}$$

where the first and second equalities came from application of the multivariable fundamental theorem of calculus, and the last equation is straightforward. This tells us that:

$$\begin{aligned} & \rho^\alpha(x)(u_\theta\rho(x) - u(x)) \\ &= \int \int_0^1 \frac{-\theta\rho^{\alpha+1}(x)}{\rho^\beta(x - t\theta\rho y)} \rho^\beta(x - t\theta\rho(x)y)\nabla u(x - t\theta\rho(x)y) \cdot y\phi(y) dt dy. \\ &= \int \int_0^1 \frac{\rho^\beta(x)}{\rho^\beta(x - t\theta\rho(x)y)} \frac{-\theta\rho^{\alpha+1}(x)}{\rho^\beta(x)} \rho^\beta(x - t\theta\rho(x)y)\nabla u(x - t\theta\rho(x)y) \cdot y\phi(y) dt dy. \end{aligned} \tag{36}$$

Equation 30 bounds the ratio $\rho(x)/\rho(x - t\theta\rho(x)y)$ as less than 2 when $\theta L < 1/2$, so Equation 36 is always less than or equal to:

$$\int \int_0^1 2^\beta \frac{-\theta\rho^{\alpha+1}(x)}{\rho^\beta(x)} \rho^\beta(x - t\theta\rho(x)y)\nabla u(x - t\theta\rho(x)y) \cdot y\phi(y) dy dt. \tag{37}$$

An application of Corollary [3.4](#) then shows

$$\int \int_0^1 2^\beta \frac{-\theta \rho^{\alpha+1}(x)}{\rho^\beta(x)} \rho^\beta(x - t\theta\rho(x)y) \nabla u(x - t\theta\rho(x)y) \cdot y \phi(y) dy dt. \quad (38)$$

$$\leq \int \int_0^1 2^\beta \frac{-\theta \rho^{\alpha+1}(x)}{\rho^\beta(x)} \rho^\beta(x - t\theta\rho(x)y) |\nabla u(x - t\theta\rho(x)y) \phi(y)| dy dt. \quad (39)$$

$$\leq 2^{\beta+1} \|\rho^{\alpha+1-\beta}\|_{L^\infty} \theta \int \int_0^1 \rho^\beta(x - t\theta\rho(x)y) |\nabla u(x - t\theta\rho(x)y) \phi(y)| dy dt. \quad \text{when } \theta < 1/(2L). \quad (40)$$

$$\leq 2^{\beta+1} \|\rho^{\alpha+1-\beta}\|_{L^\infty} \theta \|\rho^\beta \nabla u\|_{L^1} \quad (41)$$

where the last inequality follows from Corollary [3.4](#)

Using this estimate in [\(35\)](#) gives a lower bound on the denominator of the Raleigh quotient,

$$\|u_\theta - \bar{u}_\theta\|_{L^2_\alpha}^2 \geq (1/4) \|u - \bar{u}\|_{L^1_\alpha} - 2^{\beta+1}\theta \|\rho^{\alpha+1-\beta}\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1}, \quad \text{when } \theta < 1/(2L). \quad (42)$$

as desired. \square

3.6 Bounding the Rayleigh Quotient (Proof of Theorem [3.1](#))

Combining Lemmas [3.7](#) and Lemmas [3.8](#) provides an upper bound for the Rayleigh quotient of $u_\theta - \bar{u}_\theta$,

$$\begin{aligned} \lambda_2 &\leq \frac{\int_{\mathbb{R}^d} \rho^\gamma |\nabla u_\theta|^2}{\int_{\mathbb{R}^d} \rho^\alpha (u_\theta - \bar{u}_\theta)^2} \\ &\leq \frac{d \cdot 3 \cdot 2^\beta}{\theta} \frac{\|\rho^{\gamma-\beta-1}\|_{L^\infty} (2+3L) \|\rho^\beta \nabla u\|_{L^1}}{\|u - \bar{u}\|_{L^1_\alpha} - 2^{\beta+1}\theta \|\rho^{\alpha+1-\beta}\|_{L^\infty} \|\rho^\beta \nabla u\|_{L^1}} \\ &\leq \frac{d \cdot 3 \cdot 2^\beta}{\theta} \frac{\|\rho^{\gamma-\beta-1}\|_{L^\infty} (2+3L)}{1 - 2^{\beta+1}\theta \|\rho^{\alpha+1-\beta}\|_{L^\infty}} \Phi(A). \end{aligned}$$

Selecting $\theta = (1/2) \min(1/(2^{\beta+1} \|\rho^{\alpha+1-\beta}\|_{L^\infty} \Phi(A)), 1/L)$ shows

$$\lambda_2 \leq 2d \cdot 3 \cdot 2^\beta \|\rho^{\gamma-\beta-1}\|_{L^\infty} (2+3L) \max(L\Phi(A), 2^{\beta+1} \|\rho^{\alpha+1-\beta}\|_{L^\infty} \Phi(A)^2).$$

When $\gamma = (1, 2, 3)$, this simplifies into:

$$\lambda_2 \leq 12(2+3L)d \max(L\Phi(A), 8\Phi(A)^2).$$

We note that, via the work shown in Section [3.8](#), we can strengthen our inequality to:

$$\lambda_2 \leq 24d \max(L\Phi(A), 8\Phi(A)^2).$$

\square

3.7 Gradient of Mollifier

Let ϕ be a standard mollifier i.e. $\phi \in C_c^\infty(\mathbb{R}^d)$ is a function from $\mathbb{R}^d \rightarrow [0, \infty)$ satisfying $\int_{\mathbb{R}^d} \phi dx = 1$ and $\text{supp}(\phi) \subseteq B(0, 1)$. We will define ϕ by its profile. Namely, let $\hat{\phi}(r) : [0, \infty) \rightarrow [0, 1]$ be a fixed monotone decreasing profile with $\hat{\phi}(0) = 1$, $0 < \hat{\phi}(r) < 1$ for $0 < r < 1$, and $\hat{\phi}(r) = 0$ for $r \geq 1$. Then define $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ by $\phi(x) = c\hat{\phi}(|x|)$ with $c > 0$ chosen so that $\int_{\mathbb{R}^d} \phi(x) dx = 1$; that is,

$$1 = \int_{\mathbb{R}^d} \phi(x) dx = c|S^{d-1}| \int_0^1 \hat{\phi}(r)r^{d-1} dr \quad \Rightarrow \quad c = \frac{1}{|S^{d-1}| \int_0^1 \hat{\phi}(r)r^{d-1} dr},$$

where $|S^{d-1}|$ is the $(d-1)$ -area of the unit sphere in \mathbb{R}^d . We claim the L_1 norm of the gradient of $\nabla\phi(x)$ is linear in d .

Lemma 3.9.

$$\int_{\mathbb{R}^d} |\nabla\phi(x)| dx \leq (d-1) \left(\frac{d2^d}{\hat{\phi}(1/2)} \right)^{1/(d-1)} \xrightarrow{d \rightarrow \infty} 2(d-1).$$

For the classic mollifier $\hat{\phi}(r) = \exp(-1/(1-r^2))$ we get

$$\int_{\mathbb{R}^d} |\nabla\phi(x)| dx \leq 2d.$$

From the formula $\nabla\phi(x) = c\hat{\phi}'(|x|)(x/|x|)$ we compute

$$\begin{aligned} \int_{\mathbb{R}^d} |\nabla\phi(x)| dx &= c|S^{d-1}| \int_0^1 |\hat{\phi}'(r)|r^{d-1} dr \\ &= c|S^{d-1}| \int_0^1 -\hat{\phi}'(r)r^{d-1} dr \\ &= c|S^{d-1}| \int_0^1 \hat{\phi}(r)(d-1)r^{d-2} dr \\ &= (d-1) \frac{\int_0^1 \hat{\phi}(r)r^{d-2} dr}{\int_0^1 \hat{\phi}(r)r^{d-1} dr}. \end{aligned}$$

To estimate the numerator use Holder's inequality: for $1 \leq s, s' \leq \infty$ with $1/s + 1/s' = 1$

$$\int fg \leq \left(\int |f|^s \right)^{1/s} \left(\int |g|^{s'} \right)^{1/s'}.$$

Set $s = (d-1)/(d-2)$ and $s' = d-1$ to get

$$\int_0^1 \hat{\phi}(r)r^{d-2} dr = \int_0^1 \hat{\phi}(r)^{1/s} r^{d-2} \times \hat{\phi}(r)^{1/s'} dr \leq \left(\int_0^1 \hat{\phi}(r)r^{d-1} dr \right)^{1/s} \left(\int_0^1 \hat{\phi}(r) dr \right)^{1/s'}.$$

It follows that

$$\int_{\mathbb{R}^d} |\nabla\phi(x)| dx \leq (d-1) \left(\frac{\int_0^1 \hat{\phi}(r) dr}{\int_0^1 \hat{\phi}(r)r^{d-1} dr} \right)^{1/(d-1)}. \quad (43)$$

Since $0 \leq \hat{\phi}(r) \leq 1$ we can bound the numerator by 1, and since $\hat{\phi}(r)$ is monotone decreasing we have $\hat{\phi}(r) \geq \hat{\phi}(1/2)$ on $(0, 1/2)$, so

$$\int_{\mathbb{R}^d} |\nabla\phi(x)| dx \leq (d-1) \left(\frac{1}{\hat{\phi}(1/2) \int_0^{1/2} r^{d-1} dr} \right)^{1/(d-1)} \leq (d-1) \left(\frac{d2^d}{\hat{\phi}(1/2)} \right)^{1/(d-1)} \xrightarrow{d \rightarrow \infty} 2(d-1). \quad (44)$$

It will be convenient to write equation [44](#) as a simple inequality. Observe that

$$\left(\frac{d2^d}{\hat{\phi}(1/2)} \right)^{1/(d-1)}$$

is monotone decreasing. We now pick the classic $\hat{\phi}(r) = \exp(-1/(1-r^2))$ we have $\hat{\phi}(1/2) \geq 1/4$ and if $d \geq 5$ the right hand side of equation (44) is bounded by $2d$. If $d < 5$ explicit computations of the integrals shows the right hand side of equation (43) is bounded by $2d$.

3.8 Scaling

In this section we show that if one scales the density function ρ then the isoperimetric value $\Phi(A)$ and the Raleigh quotient $R(u)$ scale nicely. More formally Let $A \subset \Omega \subseteq \mathbb{R}^d$, ρ a density function over a domain Ω , and u an arbitrary differentiable function over Ω .

Consider the transformation $\hat{x} = \ell x$ with $\ell > 0$ which maps Ω to the domain $\hat{\Omega} = \{\ell x \mid x \in \Omega\}$. Given $u : \Omega \rightarrow \mathbb{R}$, we define $\hat{u} : \hat{\Omega} \rightarrow \mathbb{R}$ by $\hat{u}(\hat{x}) = u(x)$. We will future scale ρ by $\alpha\hat{\rho}(\hat{x}) = \ell\rho(x)$ where $\alpha > 0$.

Theorem 3.10. *When scaling by α and ℓ then*

$$\Phi(A) = \alpha\hat{\Phi}(\hat{A})$$

and

$$R(u) = \alpha^2\hat{R}(\hat{u}) \quad \text{and thus} \quad \lambda_2 = \alpha^2\hat{\lambda}_2$$

.

We will use this scaling theorem to improve the bounds of theorem (3.1).

That is, if we have a density function ρ over a domain Ω the isoperimetric number that the fundamental eigenvalue only change as a function of the scaling. Thus the optimal cut and eigenvector are unchanged by scaling up to the transformation.

If u and l are as defined above then we get the simple but basic identity. Suppose that $u : \mathcal{R} \rightarrow \mathcal{R}$ then:

$$\frac{\partial u}{\partial x} = \frac{\partial \hat{u}}{\partial \hat{x}} \frac{\partial \hat{x}}{\partial x} = \frac{\partial \hat{u}}{\partial \hat{x}} \ell, \quad \text{in general we get} \quad |\nabla u(x)| = \ell |\hat{\nabla} \hat{u}(\hat{x})|.$$

In the case of $\rho : \Omega \rightarrow (0, \infty)$, where $\hat{\rho} : \hat{\Omega} \rightarrow (0, \infty)$ is defined by $\alpha\hat{\rho}(\hat{x}) = \ell\rho(x)$ we get that

$$|\nabla \rho(x)| = \alpha |\hat{\nabla} \hat{\rho}(\hat{x})|.$$

It follows that $L_{\hat{\rho}}$ and L_{ρ} , the Lipschitz constants for $\hat{\rho}$ and ρ , satisfy $L_{\hat{\rho}} = (1/\alpha)L_{\rho}$.

- Since $d\hat{x} = \ell^d dx$ we have

$$\int_{\Omega} \rho dx = \frac{\alpha}{\ell^{d+1}} \int_{\hat{\Omega}} \hat{\rho} d\hat{x},$$

- If $A \subset \Omega$ and $\hat{A} = \ell A \subset \hat{\Omega}$, let $f_A(x) = 1$ if $x \in A$ and zero otherwise, and similarly $f_{\hat{A}} = 1$ if $\hat{x} \in \hat{A}$ and zero otherwise. We next perform a set of standard integral calculations.

$$\int_{\Omega} \rho^2 |\nabla f_A| dx = \int_{\hat{\Omega}} \left(\frac{\alpha}{\ell}\right)^2 \hat{\rho}^2 \ell |\hat{\nabla} f_{\hat{A}}| \frac{1}{\ell^d} d\hat{x} \tag{45}$$

$$= \frac{\alpha^2}{\ell^{d+1}} \int_{\hat{\Omega}} \hat{\rho}^2 |\hat{\nabla} f_{\hat{A}}| d\hat{x} \tag{46}$$

Equation (45) follows by making the substitutions:

$$\rho(x) = \left(\frac{\alpha}{\ell}\right)\hat{\rho}(\hat{x}) \quad |\nabla f_A| = \ell |\hat{\nabla} f_{\hat{A}}| \quad dx = \frac{1}{\ell^d} d\hat{x}$$

Observing the $f_A(x) = f_{\hat{A}}(\hat{x})$ we get the following identity.

$$\int_{\Omega} \rho f_A dx = \int_{\hat{\Omega}} \frac{\alpha}{\ell} \hat{\rho} f_{\hat{A}} \frac{1}{\ell^d} d\hat{x} = \frac{\alpha}{\ell^{d+1}} \int_{\hat{\Omega}} \hat{\rho} f_{\hat{u}} d\hat{x} \quad (47)$$

Combining equation [46](#) and equation [47](#) we get that:

$$\Phi(A) = \alpha \hat{\Phi}(\hat{A}) \quad (48)$$

- We next do a similar calculation for the Raleigh quotient. If $u : \Omega \rightarrow \mathbb{R}$ and $\hat{u}(\hat{x}) = u(x)$, the Raleigh quotients can be computed as follows,

$$\int_{\Omega} \rho^3 |\nabla u|^2 dx = \int_{\hat{\Omega}} \left(\frac{\alpha}{\ell}\right)^3 \hat{\rho}^3 \ell^2 |\hat{\nabla} \hat{u}|^2 \frac{1}{\ell^d} d\hat{x} = \frac{\alpha^3}{\ell^{d+1}} \int_{\hat{\Omega}} \hat{\rho}^3 |\hat{\nabla} \hat{u}|^2 d\hat{x}$$

$$\int_{\Omega} \rho u^2 dx = \int_{\hat{\Omega}} \frac{\alpha}{\ell} \hat{\rho} \hat{u}^2 \frac{1}{\ell^d} dx = \frac{\alpha}{\ell^{d+1}} \int_{\hat{\Omega}} \hat{\rho} \hat{u}^2 dx$$

Thus

$$R(u) = \alpha^2 \hat{R}(\hat{u})$$

We next use our scaling result in the (1, 2, 3) case to our Buser-type bound, Theorem [3.1](#) Theorem [3.1](#) states that the following hold:

$$\lambda_2 \leq 24d(1 + L) \max(L\Phi(A), 12\Phi(A)^2). \quad (49)$$

We now make substitutions into equation [49](#) from Theorem [3.10](#) and its proof for some parameter α to be determined.

$$\begin{aligned} \lambda_2 &= \alpha^2 \hat{\lambda}_2 \leq \alpha^2 24d(1 + \hat{L}) \max(\hat{L}\hat{\Phi}(\hat{A}), 12\hat{\Phi}(\hat{A})^2) \\ &= \alpha^2 24d(1 + (L/\alpha)) \max((L/\alpha)(\Phi(A)/\alpha), 12(\Phi(A)/\alpha)^2) \\ &= 24d(1 + (L/\alpha)) \max(L\Phi(A), 12\Phi(A)^2) \\ &= 24d \max(L\Phi(A), 12\Phi(A)^2) \quad \because \alpha \rightarrow \infty \end{aligned}$$

Thus we get that λ_2 only depends linear in the dimension and the Lipschitz constant:

Corollary 3.11.

$$\lambda_2 \leq 24d \max(L\Phi, 12\Phi^2)$$

4 Cheeger Inequality for Probability Density Functions

In this section, we prove the Cheeger inequality from Theorem [1.5](#). That is a weighted Cheeger inequality in higher dimensions. This is the easier to prove than Buser's inequality, which contrasts with what happens in the graph case (the graph Buser inequality is trivial).

For a simplified proof of the Cheeger inequality for distributions in one-dimension, see Appendix [B](#)

As we will see from simple counterexamples in Section [6](#) the Cheeger-direction does not hold for all setting of (α, β, γ) . The proof we give requires fewer assumptions than the Buser inequality for probability densities. One, the Cheeger inequality is independent of the Lipschitz constant of ρ and two, the proof also holds when ρ is supported on a set $\Omega \subset \mathbb{R}^d$.

The proof is almost identical to the proof in one dimension and only a slight modification of standard proofs. The only change in the proof is replacing the change of variables formula with a co-area formula. Let $\rho : \Omega \rightarrow \mathbb{R}_>$ be an Lipschitz density function that is (α, β, γ) -integrable over an open set $\Omega \subseteq \mathbb{R}^d$. Note a stronger hypothesis on Ω is that it is the support of ρ when $\rho : \mathbb{R}^d \rightarrow \mathbb{R}_\leq$.

Theorem 4.1. *Let $\rho : \Omega \rightarrow \mathbb{R}_{>0}$ be a Lipschitz function. Then,*

$$\Phi^2 \leq 4 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_\infty^2 \lambda_2.$$

In particular, when $(\alpha, \beta, \gamma) = (1, 2, 3)$ we have

$$\Phi^2 \leq 4\lambda_2.$$

Here, Φ is the optimal (α, β) -sparsity of a cut through ρ . We note that we can say something a little stronger:

Theorem 4.2. *Let $\rho : \Omega \rightarrow \mathbb{R}_{>0}$ be a Lipschitz function. Let $\Phi_{(\alpha, \beta, \gamma)}$ be the (α, β) sparsity of the (α, γ) spectral sweep cut. If $\alpha = \beta - 1 = \gamma - 2$, then:*

$$\Phi_{(\alpha, \beta, \gamma)}^2 \leq 4\lambda_2$$

Proof. (of both theorems): Let $w \in W^{1,2}$, functions whose gradient is square integrable, nonzero with $\int_\Omega \rho^\alpha w dx = 0$. Let $v = w + a1$ where a is chosen such that $|\{v < 0\}|_\alpha = |\{v > 0\}|$. Note that

$$\begin{aligned} R(w) &= \frac{\int_\Omega \rho^\gamma |\nabla w|^2 dx}{\int_\Omega \rho^\alpha w^2 dx} \\ &\geq \frac{\int_\Omega \rho^\gamma |\nabla w|^2 dx}{\int_\Omega \rho^\alpha w^2 dx + a^2 |\Omega|_\alpha} \\ &= R(v). \end{aligned}$$

Without loss of generality, the function $u = \max(v, 0)$ satisfies $R(u) \leq R(v)$.

Let $\Omega_0 = \{v > 0\}$. Let $g = u^2$. Noting that $\nabla g = 2u\nabla u$ a.e., we can apply Cauchy-Schwarz to obtain

$$\begin{aligned} \int_{\Omega_0} \rho^\beta |\nabla g| dx &= 2 \int_{\Omega_0} \rho^\beta |u| |\nabla u| dx \\ &\leq 2 \sqrt{\int_{\Omega_0} \rho^{2\beta - \alpha} |\nabla u|^2 dx} \sqrt{\int_{\Omega_0} \rho^\alpha u^2 dx} \\ &\leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_\infty \sqrt{\int_{\Omega_0} \rho^\gamma |\nabla u|^2 dx} \sqrt{\int_{\Omega_0} \rho^\alpha u^2 dx}. \end{aligned}$$

Then, dividing by $\int_{\Omega_0} \rho^\alpha g dx$, we have

$$\frac{\int_{\Omega_0} \rho^\beta |\nabla g| dx}{\int_{\Omega_0} \rho^\alpha g dx} \leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_\infty \sqrt{R(w)}.$$

Let $A_t = \{g > t\}$. Then, by the weighted co-area formula,

$$\int_{\Omega_0} \rho^\beta |\nabla g| dx = \int_0^\infty |\partial A_t|_\beta dt.$$

Writing $g(x) = \int_0^{g(x)} 1 dt$ and applying Tonelli's theorem, we rewrite the denominator

$$\int_{\Omega_0} \rho^\alpha g dx = \int_0^\infty |A_t|_\alpha dt.$$

Thus, by averaging, there exists some t^* such that

$$\begin{aligned} \Phi &\leq \Phi(A_{t^*}) \\ &\leq \frac{\int_{\Omega_0} \rho^\beta |\nabla g| dx}{\int_{\Omega_0} \rho^\alpha g dx} \\ &\leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_\infty \sqrt{R(w)}. \end{aligned}$$

Optimizing over the set $\{w \in W^{1,2} \mid w \neq 0, \int_\Omega \rho^\alpha w dx = 0\}$ completes the proof. \square

5 Spectral Sweep Cuts have Provably Good Sparsity (proof of Theorem 1.4)

Theorem 4.2 tells us that

$$\Phi_{(1,2,3)}^2 / 4 \leq \lambda_2^{(1,3)}$$

for all 1-Lipschitz ρ whose domain is on \mathbb{R}^d . Here, $\phi_{(1,2,3)}$ is the (1, 2) sparsity of the (1, 3)-spectral sweep cut, and $\lambda_2^{(1,3)}$ is the (1, 3)-principal eigenvalue.

Next Theorem 3.1 tells us that

$$\lambda_2^{(1,3)} \leq O(d\Phi_{(1,2)}),$$

where $\Phi_{(1,2)}$ is the minimal (1, 2)-sparsity of any cut through ρ .

Therefore,

$$\Phi_{(1,2,3)}^2 \leq \Phi_{(1,2)} \leq \Phi_{(1,2,3)}^2,$$

where $\Phi_{(1,2)}$ is the minimum (1, 2)-sparsity of a cut through ρ , proving Theorem 1.4

6 Cheeger-Buser inequality fails for Bad Settings of α, β, γ : Examples

In this section, we will analyze some simple 1-Lipschitz density functions in 1 dimension, and see the requirements we need on (α, β, γ) in order to recover Cheeger and Buser type inequalities. This will prove Lemma 1.6

Specifically, we refer to an inequality of the form

$$\Phi^2 < O(\lambda_2)$$

as a Cheeger-type inequality and an inequality of the form

$$\lambda_2 < O(\max(\Phi, \Phi^2))$$

as a Buser-type inequality. The presence of the Φ^2 term in the Buser-type inequality is necessary as Φ may be larger than 1; this contrasts the normalized graph case where Φ is always bounded above by 1.

We consider two simple density functions: the first is the function ρ that's $1/n$ for a support of interval $[-n/2, n/2]$. The distribution is then made to be Lipschitz in the obvious way, by making a Lipschitz drop-off at the ends of the interval, and the interval is rescaled so that it is a probability density function. In this analysis, we assume α, β, γ to be constant. The examples are discontinuous across the boundaries, $-1, 1$, but it's straight forward to extend them to be continuous on their boundary without appreciably changing them as counter examples.

In this case, the cut size in the isoperimetric cut is within a constant approximation of $1/n^\beta$, the mass term in the isoperimetric cut is within a constant approximation of n/n^α , so the isoperimetric cut is approximately $n^{\alpha-\beta-1}$. and the Rayleigh quotient can be bounded above using the Hardy Muckenhoupt inequality [Muc72] [MWW18], which shows that it is a constant approximation of $n^{\alpha-\gamma-2}$.

The Cheeger inequality would then say:

$$n^{2(\alpha-\beta-1)} < O(n^{\alpha-\gamma-2})$$

or

$$\frac{\alpha + \gamma}{2} \leq \beta \tag{50}$$

for any Cheeger inequality to hold.

We now turn our attention to a scenario in which the Buser direction will fail for improperly set α, β, γ .

6.1 Notation

We will write $a \gtrsim b$ if $a \geq cb$ for some absolute constant $0 < c < \infty$. Similarly define $a \lesssim b$. We will write $a \asymp b$ if both relations hold.

6.2 A Lipschitz weight

Consider the density function $\rho(x) = |x| + \epsilon$ on the domain $(-1, 1)$, where $\epsilon \in (0, 1/4)$, and $\rho(x) = \max(0, 2 + \epsilon - |x|)$ for all other x .

It is clear that

$$\Phi(\Omega) = \Phi(0) \asymp \epsilon^\beta.$$

Next, we apply the Hardy-Muckenhoupt inequality to estimate λ_2 . We upper bound \mathcal{H} as

$$\begin{aligned} \mathcal{H} &\leq \mathcal{R}(1)\mathcal{M}(0) \\ &\asymp \int_0^1 \frac{1}{(x + \epsilon)^\gamma} dx \\ &\lesssim \begin{cases} 1 & \text{if } \gamma < 1 \\ \ln(1/\epsilon) & \text{if } \gamma = 1 \\ O\epsilon^{1-\gamma} & \text{if } \gamma > 1. \end{cases} \end{aligned}$$

Thus, if we want a Buser-type inequality to hold, then (α, β, γ) needs to satisfy,

$$\begin{cases} 1 \lesssim \lambda_2 \lesssim \max(\Phi, \Phi^2) \asymp \epsilon^\beta & \text{if } \gamma < 1 \\ \frac{1}{\ln(1/\epsilon)} \lesssim \lambda_2 \lesssim \max(\Phi, \Phi^2) \asymp \epsilon^\beta & \text{if } \gamma = 1 \\ \epsilon^{\gamma-1} \lesssim \lambda_2 \lesssim \max(\Phi, \Phi^2) \asymp \epsilon^\beta & \text{if } \gamma > 1. \end{cases}$$

By letting ϵ go to zero, it is clear that $\gamma - 1 \geq \beta$.

By combining the condition (50) required for a Cheeger-type inequality and the requirement $\gamma - 1 \geq \beta$ required for a Buser-type inequality, we conclude that there does not exist a β such that $(1, \beta, 2)$ satisfies both Cheeger- and Buser-type inequalities.

Thus, we note that if we want a Cheeger and Buser type inequality to hold even for 1-dimensional 1-Lipschitz functions, we at least require

$$\frac{\alpha + \gamma}{2} \leq \beta$$

and

$$\gamma - 1 \geq \beta$$

These two inequalities have solutions if and only if $\gamma - 2 \geq \alpha$, in which case β can occupy any number between $\frac{\alpha + \gamma}{2}$ and $\gamma - 1$.

Note that in the special case of $(1, 1, 1)$, we have no hope of proving any inequality of the form $\lambda_2 < O(\Phi^p)$. Similar estimates can be derived for the *smooth* weight function $\rho(x) = \sqrt{x^2 + \epsilon^2}$. The calculations are more involved, however. One can similarly show that

$$\Phi \asymp \epsilon, \quad \lambda_2 \gtrsim \frac{1}{\ln(1/\epsilon)}$$

Thus combining these bounds, we deduce

$$\frac{\lambda_2}{\Phi} > \Omega\left(\frac{1}{\ln(1/\epsilon)}\right)$$

which diverges to infinity as $\epsilon \rightarrow 0$. As before we conclude that, for the choice $(\alpha, \beta, \gamma) = (1, 1, 1)$, there is no hope of proving an inequality of the form $\lambda_2 \leq O(\Phi^p)$ for any $p > 0$ and any class of ρ containing the smooth Lipschitz weights $\sqrt{x^2 + \epsilon^2}$.

7 Problems with Existing Spectral Cut Methods

In this section, we introduce a simple Lipschitz distribution where the $(\alpha = 1, \gamma = 2)$ -spectral sweep cut fails to find a $(1, \beta)$ sparse cut for any $0 \leq \beta < 10$. Meanwhile, the $(1, 3)$ -spectral sweep cut finds a desirable cut with good $(1, 2)$ -sparsity. We note that $(\alpha = 1, \beta > 10)$ -sparse cuts are likely to find cuts where one side has extremely small probability mass, making it undesirable for machine learning.

We note that this section combined with Theorem 1.4 and Lemma 1.6 shows that no Cheeger and Buser inequality can hold when $\alpha = 1$ and $\gamma = 2$ for any β : this section combined with Theorem 1.4 will show that the Cheeger-Buser inequalities can only hold for $\beta > 10$, while Lemma 1.6 shows that they can only hold for $\beta \leq 1$. Therefore, the Cheeger-Buser inequalities cannot hold for any β , for $\alpha = 1$ and $\gamma = 2$.




Figure 1: The probability density function where $\rho(x, y) = \min(\epsilon + x, \frac{1}{n})$ for arbitrary X, Y, n . Here, $\rho(x, y)$ is plotted in the z axis, and E is at point $(0, -Y, \epsilon)$.

Theorem 7.1. ($\alpha = 1, \gamma = 2$)-*Spectral Sweep Cut Counterexample:*

For a 1-Lipschitz positive valued function ρ , let Φ be the sparsity of the $(1, 3)$ -spectral sweep cut, and let Φ_{OPT} be the cut of optimal $(1, \beta)$ sparsity for any $\beta < 10$. There there exists a 1-Lipschitz density function ρ such that:

$$\Phi > C \max(\Phi_{OPT}, \sqrt{\Phi_{OPT}})$$

for any constant C .

7.1 Our density function

We first construct our 1-Lipschitz Density function for which a $(1, 2)$ spectral cut has poor $(1, \beta)$ sparsity. Our density function has parameters X, Y, ϵ, n which we will set later.

Definition 7.2. Let $\rho : [-X, X] \times [-Y, Y] \rightarrow \mathbb{R}$ be a density function such that:

$$\rho(x, y) = \min(\epsilon + x, 1/n)$$

To turn this into a 1-Lipschitz probability density function, we simply extend it to a function $\rho' : \mathbb{R}^2 \rightarrow \mathbb{R}$ where ρ' agrees with ρ on $[-X, X] \times [-Y, Y]$, and the function goes 1-Lipschitzly to 0 outside this range.

We will set $X = \sqrt{n}/10, Y = 10\sqrt{n}$, and n large, to obtain a density function where the $(1, 2)$ spectral cut has arbitrarily bad $(1, \beta)$ sparsity for all $\beta < 10$.

7.2 Proof Overview

First, we prove theorems about the zero-set of this density's $(\alpha = 1, \gamma = 2)$ eigenfunction. In particular, the zero-set of this eigenfunction must cut from the line $x = -X$ to $x = X$. It cannot cut from the line $y = -Y$ to the line $y = Y$.

We prove that any level-set of the eigenfunction can't cut from $y = -Y$ to $y = Y$. We then show that any cut that doesn't cut from $y = -Y$ to $y = Y$ has bad $(1, \beta)$ sparsity for $\beta < 10$. This completes our proof. Moreover, any cut that doesn't cut from $y = -Y$ to $y = Y$ is intuitively a poor cut of our density function, according to standard machine learning intuition.

We note the natural cut of this distribution is the straight line cut $x = 0$, which the $(1, 3)$ -spectral sweep cut will find (this is an artifact of our proof, though we do not explicitly prove it here).

First, we prove a few lemmas on the zero-set of the $(1, 2)$ eigenfunction.

7.3 The Zero-set of a principal $(1, 2)$ eigenfunction is the line $y = 0$

Theorem 7.3. *The Zero-set of the eigenfunction for our given density function, is the line $y = 0$.*

Lemma 7.4. *Let f be any eigenfunction of our given density function, for which $f(x, y) \neq f(x, y')$ for some $x, y \neq y'$. Then*

$$\int_0^Y f(x, y) dy = 0.$$

Lemma 7.5. *There exists a principal eigenfunction f_2 of our given density function, for which*

$$f_2(x, y) = f_2(-x, y) = -f_2(x, -y)$$

Proof. This follows from a (non-trivial) symmetrization argument put forward in the graph case in Guattery and Miller [GM95]. □

Lemma 7.6. *(Nodal domains for Densities) Every principal eigenfunction f_2 of our given density function satisfies: the closure of the set $\{S = (x, y) | f_2(x, y) > 0\}$ is connected.*

Proof. This follows analogously to the proof of Fiedler's nodal domains for eigenfunctions of a graph [Fie73]. □

Lemma 7.7. *Let f be a (α, β) eigenfunction of any density function supported on a compact set $S \subset \mathbb{R}^n$ for some n . For every point in the zero-set, if any open set containing that point contains a positive element, it must also contain a negative element.*

Proof. This follows directly from the definition of eigenfunction. □

Proof. (of Theorem [7.3]): First, we note that there is a principal eigenfunction whose zero set contains $y = 0$, by Lemma [7.5]. We claim there is a principal eigenfunction for which this is the entire zero-set. This follows from Lemma [7.6] and Lemma [7.7]. □

7.4 Any spectral sweep cut has high $(1, \beta)$ sparsity

In this section, we prove that the spectral sweep cut must have high $(1, \beta)$ -sparsity for $0 < \beta < 10$, and for $\beta > 10$ the spectral sweep-cut either has high $(1, \beta)$ sparsity or else divides the probability density into two pieces, one of which has less than $\leq 1/n$ fraction of the probability mass.

Lemma 7.8. *Any spectral sweep cut (of the principal $(1, 2)$ eigenfunction whose eigenvector's zero-set is the line $y = 0$) can't cut through $y = Y$ and $y = -Y$.*

Proof. This is clear. □

Lemma 7.9. *Any cut that doesn't cut through both $y = Y$ and $y = -Y$ has poor $(1, \beta)$ sparsity for any $0 < \beta < 10$. For $\beta > 10$, a cut of good $(1, \beta)$ sparsity must have its smaller side contain $o_n(1)$ fraction of the mass.*

To be precise, if Φ_β is the optimal $(1, \beta)$ sparsity of the cut, and Φ is the $(1, \beta)$ sparsity induced by a cut that doesn't cut both $y = Y$ and $y = -Y$, then there is no constant C independent of n for which

$$\Phi^2 < C\Phi_\beta.$$

We note that Theorem [7.1](#) follows from Lemma [7.8](#) and [7.9](#). Thus, it remains to show Lemma [7.9](#).

Proof. (of Lemma [7.9](#)). We split this into two cases. Consider the side of the level set cut with smaller probability mass. The first case is when this side has at least half its probability mass outside the region $|x| < 1/n - \epsilon$. The second case is when the side has less than half its mass in this region.

In the first case, we note that we can lower bound the cut by its projection onto the x axis. A quick calculation shows that when $X = \frac{1}{10\sqrt{n}}$ and $Y = \frac{10}{\sqrt{n}}$, the $(1, \beta)$ -sparsity of this cut is within a factor of 2 of the $(1, \beta)$ -sparsity of the cut $y = 0$ through the uniform distribution of height $\frac{1}{n}$ supported on $[-X, X] \times [-Y, Y]$. This $(1, \beta)$ sparsity is

$$A := O\left(\frac{X}{n^\beta}\right) = O\left(\frac{\sqrt{n}}{n^\beta}\right)$$

When ϵ is chosen to be $\frac{1}{n^2\sqrt{n}}$, then the $(1, \beta)$ sparsity of the optimal cut is the cut $x = 0$, which has $(1, \beta)$ sparsity of:

$$B := O\left(\frac{Y}{(n\sqrt{n})^\beta}\right) = O\left(\frac{\sqrt{n}}{(n^2\sqrt{n})^\beta}\right).$$

We note that this choice of ϵ is the minimum such choice such that the principal eigenvector is not constant on the Y axis.

Now we note that A^2/B goes to infinity as n gets large, if and only if

$$n^{2\beta}\sqrt{n}^\beta n/n^{2\beta}\sqrt{n}$$

goes to infinity,

or

$$\sqrt{n}(\sqrt{n}^\beta)$$

goes to infinity. This is true for any $\beta > 0$. This proves Theorem [7.1](#) in case 1, where at least half of the probability mass is outside the region $|x| < 1/n$.

In case 2, we consider the case when the smaller side of the cut has more than half its probability mass inside the region $|x| < 1/n - \epsilon$, which we note is a very small portion of the probability mass of the overall probability density. In this case, it turns out that we need $\beta < 10$ to give isoperimetry guarantees, since for any $\beta > 10$, it turns out that even cuts containing small probability mass are considered to have good $(1, \beta)$ sparsity, since for large β , $(1, \beta)$ sparse cuts tremendously favor small cuts, even if the smaller side has negligible probability mass.

Since at least half the mass is inside the region $|x| < 1/n$, we can assume without loss of generality that the entire probability mass of the smaller side of the cut is inside this region, by simply projecting the cut onto this region (reducing its β -perimeter while decreasing probability mass by at most a factor of 2). We can again use a symmetry argument analogous to [7.5](#) to show that any level set of this principal eigenfunction is symmetric about the x axis (we note Lemma [7.9](#) is slightly stronger than this as it does not assume symmetry,

but for our purposes we can strictly deal with symmetric cuts, and the non-symmetric case follows through a similar argument).

Now given the cut is symmetric about the x axis, if the cut cuts through (x', y') , then it also cuts through $(-x', y')$, and we can lower bound the probability mass contained by the cut $y = y'$ with $x' \cdot \rho(x', y')$. A simple calculation using this estimate finishes the proof for us.

□

8 Conclusion and Future Directions

We define a new notion of spectral sweep cuts, eigenvalues, Rayleigh quotients, and sparsity for probability densities. We present the first known Cheeger and Buser inequality on probability density functions, and use this to show an $(1, 3)$ spectral sweep cut on a L -Lipschitz probability density function has provably low $(1, 2)$ -sparsity. This work is the first spectral sweep cut algorithm on Lipschitz densities with any guarantees on the cut quality.

Further, we show that existing spectral sweep cut methods (such as those implicit in spectral clustering) compute $(1, 1)$ or $(1, 2)$ spectral sweep cuts, neither of which has any sparsity guarantees. We prove that $(1, 2)$ spectral sweep cuts, which are implicitly used in traditional spectral clustering, can lead to undesirable partitions of simple 1-Lipschitz probability densities. We also show that Cheeger and Buser’s inequality for probability density relies on a careful setting of three parameters: α, β , and γ .

For future directions, we conjecture that $\beta = \alpha + 1$ and $\gamma = \alpha + 2$ is the only settings of (α, β, γ) in which both Cheeger and Buser inequalities are provable. This would be a stronger theorem than we currently have for Lemma 1.6, which shows that $\gamma - 2 \leq \alpha$ and $\frac{\alpha + \beta}{2} \leq \gamma$ is required.

In the Buser inequality, we would like to iron out the exact dimensional dependence on the dimension, d (Theorem 3.1). The authors believe that this dependence can be reduced to \sqrt{d} . It is an open question whether *any* dimension dependence is required. In particular, the latest version of Buser’s inequality for manifolds has no dimension dependence [Led04], and this can be proven through the heat kernels, the Bochner formula, and the Li-Yau inequality in manifold theory. It is an open question how to generalize their techniques into the distribution setting, as the Bochner formula does not easily generalize to distributions.

We also conjecture that our algorithm, 1,3-SPECTRALCLUSTERING, converges to the $(1, 3)$ -spectral sweep cut of ρ as the number of samples drawn from ρ grows large. This would be analogous to the results of Slepcev and Trillos on standard spectral clustering [TS15].

Another open question is whether multi-way Cheeger and Buser inequalities can be proven on distributions, mirroring the work on graphs [LRTV12, KW16, LGT14a, LGT14b]. This would allow our clustering algorithms to generalize into k -way clusterings.

Finally, we would like to know whether Buser and Cheeger inequalities may exist for L -Lipschitz probability densities supported on manifolds with bounded curvature. If true, this would fully generalize the work of Cheeger and Buser on manifolds, which may lead to deeper insight into manifold theory. Moreover, it could have foundational impact: a fundamental assumption underlying modern machine learning is that most data comes from probability density supported on a manifold, and a Cheeger and Buser inequality in this setting would give provable sparsity guarantees about spectral sweep cuts in this setting.

References

- [AM84] N. Alon and V. D. Milman. Eigenvalues, expanders and superconcentrators. In *Proceedings of the 25th Annual Symposium on Foundations of Computer Science, 1984*, FOCS '84, pages 320–322, Washington, DC, USA, 1984. IEEE Computer Society.
- [AP09] Reid Andersen and Yuval Peres. Finding sparse cuts locally using evolving sets. In *Proceedings of the forty-first annual ACM symposium on Theory of computing*, STOC '09, pages 235–244. ACM, 2009.
- [ARV04] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In László Babai, editor, *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC), Chicago, IL, USA, June 13-16, 2004*, pages 222–231. ACM, 2004.
- [BC17] Tugkan Batu and Clément L. Canonne. Generalized uniformity testing. In Chris Umans, editor, *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 880–889. IEEE Computer Society, 2017.
- [Bis06] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- [BN04] Mikhail Belkin and Partha Niyogi. Semi-supervised learning on riemannian manifolds. *Mach. Learn.*, 56(1-3):209–239, June 2004.
- [Bus82] Peter Buser. A note on the isoperimetric constant. *Annales scientifiques de l'École Normale Supérieure*, Ser. 4, 15(2):213–230, 1982.
- [CGR05] Shuchi Chawla, Anupam Gupta, and Harald Räcke. Embeddings of negative-type metrics and an improved approximation to generalized sparsest cut. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005, Vancouver, British Columbia, Canada, January 23-25, 2005*, pages 102–111. SIAM, 2005.
- [Che70] Jeff Cheeger. A lower bound for the smallest eigenvalue of the laplacian. In *Problems in Analysis: A Symposium in Honor of Salomon Bochner*, pages 195–199. Princeton Univ. Press, Princeton, N. J., 1970.
- [Chu97] F.R.K. Chung. *Spectral Graph Theory*, volume 92 of *Regional Conference Series in Mathematics*. American Mathematical Society, 1997.
- [DFK91] Martin Dyer, Alan Frieze, and Ravi Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, January 1991.
- [DKN10] Ilias Diakonikolas, Daniel M. Kane, and Jelani Nelson. Bounded independence fools degree-2 threshold functions. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23-26, 2010, Las Vegas, Nevada, USA*, pages 11–20, 2010.
- [EKSX96] Martin Ester, Hans-Peter Kriegel, Jorg Sander, and Xiaowei Xu. A density-based algorithm for discovering clusters a density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD'96*, pages 226–231. AAAI Press, 1996.
- [Fie73] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Math. J.*, 23(98):298–305, 1973.
- [Fri44] Kurt Otto Friedrichs. The identity of weak and strong extensions of differential operators. *Transactions of the American Mathematical Society*, 55:132–151, 1944.
- [GLR18] Rong Ge, Holden Lee, and Andrej Risteski. Beyond log-concavity: Provable guarantees for sampling multi-modal distributions using simulated tempering langevin monte carlo. In *Proceedings*

- of the 32Nd International Conference on Neural Information Processing Systems, NIPS'18, pages 7858–7867, USA, 2018. Curran Associates Inc.
- [GM83] Mikhail Gromov and Vitali D Milman. A topological application of the isoperimetric inequality. *American Journal of Mathematics*, 105(4):843–854, 1983.
 - [GM95] Stephen Guattery and Gary L. Miller. On the performance of the spectral graph partitioning methods. In *SODA '95*, pages 233–242. ACM-SIAM, 1995.
 - [GS06] Leo Grady and Eric L Schwartz. Isoperimetric graph partitioning for image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 28(3):469–475, 2006.
 - [GTS15] Nicolás García Trillos and Dejan Slepčev. On the rate of convergence of empirical measures in ∞ -transportation distance. *Canadian Journal of Mathematics*, 67(6):1358–1383, 2015.
 - [GTS16] Nicolás García Trillos and Dejan Slepčev. Continuum limit of total variation on point clouds. *Archive for Rational Mechanics and Analysis*, 220(1):193–241, 04 2016.
 - [HK00] Piotr Hajłasz and Pekka Koskela. *Sobolev met poincaré*, volume 688. American Mathematical Soc., 2000.
 - [KK10] A. Kumar and R. Kannan. Clustering with spectral norm and the k-means algorithm. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 299–308, 2010.
 - [KLS95] Ravi Kannan, László Lovász, and Miklós Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13:541–559, 1995.
 - [KW16] Robert Krauthgamer and Tal Wagner. Cheeger-type approximation for sparsest *st*-cut. *ACM Trans. Algorithms*, 13(1):14:1–14:21, 2016.
 - [Led04] Michel Ledoux. Spectral gap, logarithmic sobolev constant, and geometric bounds. *Surveys in differential geometry*, 9(1):219–240, 2004.
 - [LGT14a] James R. Lee, Shayan Oveis Gharan, and Luca Trevisan. Multi-way spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM*, 61(6)(37), 2014.
 - [LGT14b] James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):37, 2014.
 - [LR88] Frank Thomson Leighton and Satish Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *29th Annual Symposium on Foundations of Computer Science, White Plains, New York, USA, 24-26 October 1988*, pages 422–431. IEEE Computer Society, 1988.
 - [LRTV12] Anand Louis, Prasad Raghavendra, Prasad Tetali, and Santosh Vempala. Many sparse cuts via higher eigenvalues. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, FOCS 2012, pages 1131–1140. ACM, 2012.
 - [LS90] L. Lovasz and M. Simonovits. The mixing rate of markov chains, an isoperimetric inequality, and computing the volume. In *Proceedings of the 31st Annual Symposium on Foundations of Computer Science (FOCS)*, FOCS '90, pages 346–354 vol. 1, Washington, DC, USA, 1990. IEEE Computer Society.
 - [LV18a] Yin Tat Lee and Santosh S. Vempala. The Kannan-Lovász-Simonovits conjecture, 2018.
 - [LV18b] Yin Tat Lee and Santosh S. Vempala. Stochastic localization + stieltjes barrier = tight bound for log-sobolev. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 1122–1129, New York, NY, USA, 2018. ACM.

- [LW01] Chun Liu and Noel J Walkington. An eulerian description of fluids containing visco-elastic particles. *Archive for rational mechanics and analysis*, 159(3):229–252, 2001.
- [Mad10] Aleksander Madry. Fast approximation algorithms for cut-based problems in undirected graphs. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 245–254. IEEE, 2010.
- [Mil09] Emanuel Milman. On the role of convexity in isoperimetry, spectral gap and concentration. *Inventiones mathematicae*, 177(1):1–43, 2009.
- [Mon03] Peter Monk. *Finite Element Methods for Maxwell’s Equations (Numerical Analysis and Scientific Computation Series)*. Oxford Press, 01 2003.
- [Muc72] Benjamin Muckenhoupt. Hardy’s inequality with weights. *Studia Mathematica*, 44(1):31–38, 1972.
- [MV10] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *51th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2010, October 23–26, 2010, Las Vegas, Nevada, USA*, pages 93–102, 2010.
- [MWW18] Gary L. Miller, Noel J. Walkington, and Alex L. Wang. Hardy-muckenhoupt bounds for laplacian eigenvalues. *CoRR*, abs/1812.02841, 2018.
- [NJW01] Andrew Y. Ng, Michael I. Jordan, and Yair Weiss. On spectral clustering: Analysis and an algorithm. In *Proceedings of the 14th International Conference on Neural Information Processing Systems: Natural and Synthetic*, NIPS’01, pages 849–856, Cambridge, MA, USA, 2001. MIT Press.
- [OSVV08] Lorenzo Orecchia, Leonard J Schulman, Umesh V Vazirani, and Nisheeth K Vishnoi. On partitioning graphs via single commodity flows. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, STOC ’08, pages 461–470. ACM, 2008.
- [OV11] Lorenzo Orecchia and Nisheeth K. Vishnoi. Towards an sdp-based approach to spectral methods: A nearly-linear-time algorithm for graph partitioning and decomposition. In *Proceedings of the Twenty-second Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’11, pages 532–545, Philadelphia, PA, USA, 2011. Society for Industrial and Applied Mathematics.
- [RBV10] Lorenzo Rosasco, Mikhail Belkin, and Ernesto De Vito. On learning with integral operators. *Journal of Machine Learning Research*, 11(Feb):905–934, 2010.
- [SBC14] Weijie Su, Stephen Boyd, and Emmanuel J. Candès. A differential equation for modeling nesterov’s accelerated gradient method: Theory and insights. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, NIPS’14, page 2510–2518, Cambridge, MA, USA, 2014. MIT Press.
- [SM97] J.B. Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. on Pattern Anal. and Mach. Intell.*, 22:888–905, 01 1997.
- [Sob38] S Soboleff. Sur un théorème d’analyse fonctionnelle. *Matematicheskii Sbornik*, 46(3):471–497, 1938.
- [SS09] Elias M Stein and Rami Shakarchi. *Real analysis: measure theory, integration, and Hilbert spaces*. Princeton University Press, 2009.
- [ST04] Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*, STOC ’04, pages 81–90, New York, NY, USA, 2004. ACM.

- [ST07] Daniel A. Spielman and Shang-Hua Teng. Spectral partitioning works: Planar graphs and finite element meshes. *Linear Algebra and its Applications*, 421(2):284 – 305, 2007. Special Issue in honor of Miroslav Fiedler.
- [SW18] Christian Sohler and David P. Woodruff. Strong coresets for k-median and subspace approximation: Goodbye dimension. In *59th IEEE Annual Symposium on Foundations of Computer Science, (FOCS) 2018, Paris, France, October 7-9, 2018*, pages 802–813, 2018.
- [SW19] Thatchaphol Saranurak and Di Wang. Expander decomposition and pruning: Faster, stronger, and simpler. In Timothy M. Chan, editor, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 2616–2635. SIAM, 2019.
- [TS15] Nicolás García Trillos and Dejan Slepčev. A variational approach to the consistency of spectral clustering, 2015.
- [VL07] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.
- [VLBB08] Ulrike Von Luxburg, Mikhail Belkin, and Olivier Bousquet. Consistency of spectral clustering. *The Annals of Statistics*, pages 555–586, 2008.
- [VSCC08] Nicholas T Varopoulos, Laurent Saloff-Coste, and Thierry Coulhon. *Analysis and geometry on groups*, volume 100. Cambridge university press, 2008.
- [Whi17] Jonathan Whiteley. Linear elliptic partial differential equations. In *Finite Element Methods*, pages 119–141. Springer, 2017.
- [WN17] Christian Wulff-Nilsen. Fully-dynamic minimum spanning forest with improved worst-case update time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, STOC 2017, page 1130–1143, New York, NY, USA, 2017. Association for Computing Machinery.
- [WWJ16] Andre Wibisono, Ashia C. Wilson, and Michael I. Jordan. A variational perspective on accelerated methods in optimization. *Proceedings of the National Academy of Sciences of the United States of America*, 113 47:E7351–E7358, 2016.

Appendix A Cheeger and Buser for Density Functions does not easily follow from Graph or Manifold Cheeger and Buser

A.1 Comments on Graph Cheeger-Buser

The most natural method of proving distributional Cheeger-Buser inequality using the graph Cheeger-Buser inequality is to generate a vertex and edge weighted graph approximating the distribution, and write down graph Cheeger-Buser. Then, one would generate a sequence of graphs with an increasing number of vertices. Ideally, the graph Cheeger-Buser inequality on these graphs would converge to a Cheeger-Buser inequality on the underlying distribution. This discretization approach follows a standard paradigm of approximating distributions with graphs, present in numerical methods, finite element methods, and machine learning [GTS15](#) [TS15](#) [ST07](#).

Such an approach cannot work (no matter how the eigenvalues and isoperimetric cuts are defined for distributions). The easiest way to see this is to attempt to execute this strategy for a simple uniform distribution in 1 dimension, on the interval $[0, 1]$. One would naively approximate this distribution with a line graph with n vertices, with edge weights w_n and vertex weights m_n . Then one would take n to go to infinity.

If one writes down the Cheeger and Buser inequalities for graphs in this example, we get:

$$\frac{w_n}{m_n n^2} \leq \Phi_{OPT} \leq \frac{w_n}{m_n n}$$

No matter what m_n and w_n are, the ratio between the upper and lower bound is n , which diverges. Thus, either the Cheeger inequality or the Buser inequality becomes meaningless: either the lower bound goes to 0 or the upper bound goes to ∞ , or both, depending on how w_n and m_n are set.

Thus, even for the simple case of a uniform distribution on $[0, 1]$ the natural strategy for deriving probability density Cheeger/Buser from graph Cheeger/Buser fails.

A.2 Comments on Manifold Cheeger-Buser

Distributional Buser does not easily follow from an application of the manifold Buser inequality. We recall that manifold Buser only applies for manifolds with bounded Ricci curvature. The natural way to parlay manifold Buser into distributional Buser on \mathbb{R}^d is to change the underlying metric tensor on \mathbb{R}^d to factor in the probability density function at that point. However, the authors are unaware of any method of doing this for which one can recover a meaningful Cheeger and Buser inequality. Moreover, it is unclear how to obtain any Ricci curvature bounds when we change the metric tensor.

Most modern approaches to proving Buser's inequality for manifolds rely on the Li-Yau inequality, which in turn depends on the Bochner identity for manifolds on bounded Ricci curvature [Led04](#). The authors are unaware of a clean Bochner-like identity for distributions. Older techniques use Almgren's minimizing currents and/or Epsilon nets [Bus82](#). For the former, we do not know of any analog for distributions. For the latter, the corresponding Buser inequality has a 2^d multiplicative dependence, which is significantly worse than our d dependence.

Appendix B A weighted Cheeger inequality in one dimension

Theorem B.1. *Let $\Omega = (a, b)$ where $-\infty < a < b < \infty$. Let $\rho : (a, b) \rightarrow \mathbb{R}_{>0}$ be Lipschitz continuous. Then,*

$$\Phi(\Omega)^2 \leq 4 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_{\infty}^2 \lambda_2(\Omega).$$

In particular, when $(\alpha, \beta, \gamma) = (1, 2, 3)$, we have

$$\Phi(\Omega)^2 \leq 4\lambda_2(\Omega).$$

Proof. Let $w \in W^{1,2}(\Omega) \cap C^{\infty}(\Omega)$ be a strictly decreasing function with $\int_{\Omega} \rho^{\alpha} w \, dx = 0$. Let $v = w + a1$ where a is chosen such that $|\{v < 0\}|_{\alpha} = |\{v > 0\}|_{\alpha}$. Note that

$$\begin{aligned} R(w) &= \frac{\int_{\Omega} \rho^{\gamma} (w')^2 \, dx}{\int_{\Omega} \rho^{\alpha} w^2 \, dx} \\ &\geq \frac{\int_{\Omega} \rho^{\gamma} (w')^2 \, dx}{\int_{\Omega} \rho^{\alpha} w^2 \, dx + a^2 |\Omega|_{\alpha}} \\ &= R(v). \end{aligned}$$

Let $\hat{x} \in (a, b)$ be the unique value such that $v(\hat{x}) = 0$. Without loss of generality, the function $u = \max(v, 0)$ satisfies $R(u) \leq R(v)$ and has $u(a) = 1$.

Let $g = u^2$. Noting that $g' = 2uu'$ a.e., we can apply Cauchy-Schwarz to obtain

$$\begin{aligned} \int_a^{\hat{x}} \rho^{\beta} |g'| \, dx &= 2 \int_a^{\hat{x}} \rho^{\beta} |u| |u'| \, dx \\ &\leq 2 \sqrt{\int_a^{\hat{x}} \rho^{2\beta - \alpha} (u')^2 \, dx} \sqrt{\int_a^{\hat{x}} \rho^{\alpha} u^2 \, dx} \\ &\leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_{\infty} \sqrt{\int_a^{\hat{x}} \rho^{\gamma} (u')^2 \, dx} \sqrt{\int_a^{\hat{x}} \rho^{\alpha} u^2 \, dx}. \end{aligned}$$

Then, dividing by $\int_a^{\hat{x}} \rho^{\alpha} g \, dx$, we have

$$\frac{\int_a^{\hat{x}} \rho^{\beta} |g'| \, dx}{\int_a^{\hat{x}} \rho^{\alpha} g \, dx} \leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_{\infty} \sqrt{R(w)}.$$

By change of variables,

$$\int_a^{\hat{x}} \rho^{\beta} |g'| \, dx = \int_0^1 \rho^{\beta} (g^{-1}(t)) \, dt.$$

Writing $g(x) = \int_0^{g(x)} 1 \, dt$ and applying Tonelli's theorem, we rewrite the denominator

$$\int_a^{\hat{x}} \rho^{\alpha} g \, dx = \int_0^1 |(a, g^{-1}(t))|_{\alpha} \, dt.$$

Thus, by averaging, there exists some t^* such that,

$$\Phi(\Omega) \leq \frac{\rho^{\beta}(t^*)}{|(a, t^*)|_{\alpha}} \leq \frac{\int_a^{\hat{x}} \rho^{\beta} |g'| \, dx}{\int_a^{\hat{x}} \rho^{\alpha} g \, dx} \leq 2 \left\| \rho^{\beta - \frac{\alpha + \gamma}{2}} \right\|_{\infty} \sqrt{R(w)}.$$

g

□

Theorem B.2. Let $\Omega = (a, b)$ where $-\infty < a < b < \infty$. Let $\rho : (a, b) \rightarrow \mathbb{R}_{>0}$ be Lipschitz continuous with Lipschitz constant L . Then,

$$\lambda_2(\Omega) \leq 8 \cdot (3/2)^{\gamma/\alpha} \|\rho^{\gamma-1-\beta}\|_\infty \max \left(4 \|\rho^{\alpha+1-\beta}\|_\infty \Phi^2(\Omega), \frac{\alpha}{\ln(3/2)} L\Phi(\Omega) \right).$$

In particular, when $(\alpha, \beta, \gamma) = (1, 2, 3)$, we have

$$\lambda_2(\Omega) \leq O \left(\max \left(\Phi^2(\Omega), L\Phi(\Omega) \right) \right).$$

Proof. Let $\hat{x} \in (a, b)$. We will show that there exists a $u \in W^{1,2}(\Omega)$ with small Rayleigh quotient compared to $\Phi(\hat{x})$. Let $A = (a, \hat{x})$ and $B = (\hat{x}, b)$. Without loss of generality $|A|_\alpha \leq |B|_\alpha$ and hence $\Phi(\hat{x}) = \frac{\rho^\beta(\hat{x})}{|A|_\alpha}$. For notational convenience, we will write $\Phi = \Phi(\hat{x})$ in this proof.

Let

$$u(x) = \begin{cases} |A|_\alpha & a \leq x \leq \hat{x} \\ -|B|_\alpha & \hat{x} < x \leq b. \end{cases}$$

Let $\delta = \theta\rho(\hat{x})$ where $\theta > 0$ will be picked later. Define the continuous function

$$u_\delta(x) = \begin{cases} |A|_\alpha & a \leq x \leq x_1 \\ \text{linear with slope } \frac{-|A|_\alpha}{\delta} & x_1 \leq x \leq x_2 \\ -|B|_\alpha & x_2 \leq x \leq b \end{cases}$$

where $a \leq x_1 < \hat{x} < x_2 \leq b$ are picked such that $\int_a^b \rho^\alpha u_\delta dx = 0$. Note $x_2 - x_1 \leq \delta$.

We bound the numerator in $R(u_\delta)$ using the mean value theorem.

$$\begin{aligned} \int_a^b \rho^\gamma (u'_\delta)^2 dx &= \frac{|\Omega|_\alpha^2}{\delta^2} \int_{x_1}^{x_2} \rho^\gamma dx \\ &\leq \frac{|\Omega|_\alpha^2}{\delta} \rho^\gamma(\tilde{x}) \quad \text{for some } \tilde{x} \in [x_1, x_2] \\ &\leq |\Omega|_\alpha^2 \rho^{\gamma-1}(\hat{x})(1 + L\theta)^\gamma/\theta \end{aligned}$$

In the third line we used the Lipschitz estimate $\rho(\tilde{x}) \leq \rho(\hat{x})(1 + L\theta)$. We lower bound the denominator in $R(u_\delta)$ using the mean value theorem and the same Lipschitz estimate. We will also recall that $\Phi = \rho^\beta(\hat{x})/|A|_\alpha$.

$$\begin{aligned} \int_a^b \rho^\alpha u_\delta^2 dx &\geq \int_a^b \rho^\alpha u^2 dx - \int_{x_1}^{x_2} \rho^\alpha u^2 dx \\ &\geq |A|_\alpha |B|_\alpha |\Omega|_\alpha - \delta \rho^\alpha(\tilde{x}) |B|_\alpha^2 \quad \text{for some } \tilde{x} \in [x_1, x_2] \\ &\geq |A|_\alpha |B|_\alpha |\Omega|_\alpha - \rho^{\alpha+1}(\hat{x}) |B|_\alpha^2 (1 + L\theta)^\alpha \theta \\ &\geq |\Omega|_\alpha^2 (|A|_\alpha/2 - \rho^{\alpha+1}(\hat{x})(1 + L\theta)^\alpha \theta) \\ &\geq |\Omega|_\alpha^2 |A|_\alpha (1/2 - \|\rho^{\alpha+1-\beta}\|_\infty \Phi (1 + L\theta)^\alpha \theta) \end{aligned}$$

The parameter θ will be chosen such that the estimate of the denominator is positive. We combine the two

bounds above.

$$\begin{aligned}
R(u_\delta) &\leq \frac{|\Omega|_\alpha^2 \rho^{\gamma-1}(\hat{x})(1+L\theta)^\gamma/\theta}{|\Omega|_\alpha^2 |A|_\alpha (1/2 - \|\rho^{\alpha+1-\beta}\|_\infty \Phi(1+L\theta)^\alpha\theta)} \\
&= \frac{\rho^{\gamma-1-\beta}(\hat{x})\Phi(1+L\theta)^\gamma/\theta}{1/2 - \|\rho^{\alpha+1-\beta}\|_\infty \Phi(1+L\theta)^\alpha\theta} \\
&\leq \frac{\|\rho^{\gamma-1-\beta}\|_\infty \Phi(1+L\theta)^\gamma/\theta}{1/2 - \|\rho^{\alpha+1-\beta}\|_\infty \Phi(1+L\theta)^\alpha\theta}.
\end{aligned}$$

We make the following choice of $\theta > 0$,

$$\theta = \min\left(\frac{1}{4\Phi\|\rho^{\alpha+1-\beta}\|_\infty}, \frac{\ln(3/2)}{\alpha L}\right).$$

Then, $(1+L\theta) \leq (3/2)^{1/\alpha}$ and $\Phi\theta \leq \frac{1}{4\|\rho^{\alpha+1-\beta}\|_\infty}$. Thus,

$$\begin{aligned}
\lambda_2 &\leq R(u_\delta) \\
&\leq 8 \cdot (3/2)^{\gamma/\alpha} \|\rho^{\gamma-1-\beta}\|_\infty \frac{\Phi}{\theta} \\
&= 8 \cdot (3/2)^{\gamma/\alpha} \|\rho^{\gamma-1-\beta}\|_\infty \max\left(4\|\rho^{\alpha+1-\beta}\|_\infty \Phi^2, \frac{\alpha}{\ln(3/2)} L\Phi\right).
\end{aligned}$$

Finally, picking \hat{x} such that $\Phi(\hat{x}) \rightarrow \Phi(\Omega)$ completes the proof. \square

Remark B.3. Recall the example presented in Section [6.2](#), i.e. $\Omega = (-1, 1)$, $\rho = |x| + \epsilon$. For the choice $(\alpha, \beta, \gamma) = (1, 1, 1)$, it was shown that $\frac{\lambda_2(\Omega)}{\Phi(\Omega)^p}$ diverges to infinity as $\epsilon \rightarrow 0$ for any $p > 0$. This does not contradict our Theorem [B.2](#) which only asserts that

$$\lambda_2(\Omega) \lesssim \frac{1}{\epsilon} \max(\Phi^2(\Omega), \Phi(\Omega)).$$