UNIVERSITY OF MINNESOTA

This is to certify that I have examined
this copy of a doctoral dissertation by

Sabyasachi Bera

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Snigdhansu Chatterjee

Name of Faculty Adviser

Signature of Faculty Adviser

Date

GRADUATE SCHOOL

Inference using Geometry and Density Information in Manifold Data

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Sabyasachi Bera

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Snigdhansu Chatterjee, Adviser

August 2023

Acknowledgements

I would like to thank my advisor, Prof. Snigdhansu Chatterjee for everything he has done for me throughout last 10 years. This dissertation would not have been possible without the very fundamental ideas on clustering that came out of our discussions during summer of 2019. Ansu da gave me the freedom to pursue those ideas and shape them as I wanted. I am thankful for the professional support he provided, as well as for being patient with me during difficult years.

I would like to thank the ex-DGS Prof. Yuhong Yang and the current DGS Prof. Adam Rothman for their patience with me throughout my tenure as a PhD student. Thanks to Prof. Yang for teaching me the fundamentals of applied statistics, for being a reviewer of this dissertation and being on the final exam committee as well. I would like to thank Profs. Qian Qin and Wei-Kuo Chen for reviewing the thesis and being on the final exam committee. I would also like to thank other professors in the department, who have always been helpful in answering questions, and the staff in the department office for their support regarding official matters. Thanks to other students in the department for being there with coursework, feedback and discussions: especially Riddhiman, Sijia, Yu, Sakshi, Adam, Nandy da, Abhirup da and Subho da.

I would like to thank Prof. John Fieberg from Department of Fisheries, Wildlife and Conservation Biology with whom I collaborated during last 2 years of my PhD. He has been an amazing mentor, and discussions with him has helped me learn a lot about applied statistics. I would like to thank International Student and Scholar Services and Boyton Health at University of Minnesota for providing with invaluable support and resources throughout my PhD. In particular, I would like to thank my

therapist Joey Lusvardi at Boynton for helping me navigate through the most difficult phase of my PhD journey.

Coming to personal life, I am thankful to my old friends from Indian Statistical Institute who have been always there for me and have been fundamental to my growth as a person. I am thankful to my new friends in Minnesota for teaching me the American way of life and for not being "Minnesota nice" to me (for the most part). I am thankful to the creative community in Minnesota and beyond for helping me become a professional photographer.

Finally, I would like to thank my parents, my elder sister and everyone I have ever loved for their kindness, love and support.

Dedication

This dissertation is dedicated to my parents, Malina Bera and Mathur Mohan Bera, my sister Aparna Bera and my Mathematics teacher from high school Pulak Kundu.

Abstract

Clustering is the task of grouping a dataset so that data in the same group (called a cluster) are more similar in some sense to each other than to those in other groups. While different notions of clustering exist in literature, it is commonly understood that data which are "close" to each other (geometric proximity) should be in the same cluster and clusters should capture the concentration pattern (high density regions) in the data. In many applications, especially when the data is from a topological manifold, we are required to capture both geometry and density information from the data simultaneously in order to cluster them in a meaningful way. We introduce g-distance, a data driven density sensitive distance, and explore its theoretical properties, geometry and usefulness in clustering applications under several data generating models. We derive the convergence limit of longest leg path distance (LLPD), a purely density based limiting form of g-distance. We compare several distances, for example, Euclidean distance, g-distance, LLPD, in clustering and manifold learning applications under several data generating models. Finally, as an application of high-dimensional learning and manifold learning, we develop a technique for record linkage on high-dimensional data using sparse principal components.

Contents

List of Figures

List of Symbols

Notation	Meaning
1_A	Indicator function on set A
$\mathbb{E}$	Expectation of a random variable
Var	Variance of a random variable
Cor	Correlation coefficient between two random variables
$\|.\|$	Euclidean Norm
$\|.\|_2$	Spectral Norm
$\mathbb{P}$	Probability Measure
χ_n	Data $\{X_1, X_2, \ldots, X_n\}$
χ_N	Data of size N after de-noising procedure in chapter 5
M	Riemannian Manifold
∂M	Boundary of M
f	A density function on M
f_0	$\inf_M f(x)$
f^0	$\sup_M f(x)$
$\mathbb{R}^d$	d dimensional Euclidean space
d	Intrinsic/ lower dimension
D	Ambient dimension
$L_g(x, y, \chi_n)$	g-distance between x and y given χ_n

$L_g^N(x, y, \chi_n)$	Normalized g-distance between x and y given χ_n
k or k^*	Lowest power of g-distance
k^{**}	Second lowest power of g-distance
K	Number of clusters in chapter 5
β	$\frac{k-1}{d} \geq 0$ in chapter 4
H_λ	Poisson process in $\mathbb{R}^d$ with rate λ in chapter 2
$LLPD(x, y, \chi_n)$	Longest leg path distance between x and y
$f_{x,y}$	Bottleneck value between x and y
W	Similarity matrix in chapter 5
L	Un-normalized graph Laplacian in chapter 5
L_{sym}	Normalized graph Laplacian in chapter 5
(λ_i, ϕ_i)	i-th smallest eigenvalue λ_i and corresponding eigenvector ϕ_i in chapter 5
N_x	Intrinsic dimensional neighborhood around x
M_λ	Super-level set at level λ, defined by $\{x \in M : f(x) \geq \lambda\}$

Chapter 1

Introduction

1.1 Background

Research on manifold estimation and clustering has made remarkable progress in the last 20 years. Some of the main challenges researchers face in manifold estimation and clustering include extracting the density and local geometry information from the data in a meaningful way, finding a data representation that reveals intrinsic lower dimensional structure (clusters) in the high dimensional data, finding a clustering procedure which is robust towards minor mis-specifications in the underlying data generating model and robust against both background and systematic noise.

Since early 2000s, density (generically denoted by f in this thesis) dependent population level distance measures have been proposed (Orlitsky et al., 2005; Vincent and Bengio, 2003; Bijral et al., 2012) for clustering purposes instead of working with Euclidean/default geodesic measure which do not depend on f and is given by

$$dist_1(x, y) \coloneqq \inf_{\gamma} \int_0^1 |\gamma_t^{'}| dt \tag{1.1}$$

with $\gamma_0 = x$ and $\gamma_1 = y$ and γ_t being a piece-wise continuous curve from x to y in a D dimensional Riemannian manifold.

For example,

$$d_h(x, y) \coloneqq \inf_{\gamma} \int_0^1 h(f(\gamma_t))|\gamma_t^{'}|dt \tag{1.2}$$

has been proposed by Bijral et al. (2012); Orlitsky et al. (2005) for clustering purposes where f is the underlying density on the manifold M, $h : \mathbb{R}^+ \to \mathbb{R}^+$ is a positive and monotonically decreasing function, γ_t is a curve from x to y satisfying $\gamma_0 = x$ and $\gamma_1 = y$.

This distance measure puts more penalty when γ_t travels through low density region (since h is a decreasing function) thus encouraging γ_t to stay in the high density region while connecting x and y. Since in density based clustering methodologies, clusters are loosely defined as connected high density regions of M, γ_t between points from the same clusters under $d_h(.,.)$ are more likely to remain within that cluster/high density region for the most part, while γ_t between points from different clusters are forced to pass through a lower density region thus making within cluster $d_h(.,.)$'s significantly lower than between cluster $d_h(.,.)$'s leading to near perfect clustering.

Although conceptually this looks promising, since we don't know f we need to estimate it from data $\chi_n \coloneqq \{x_1, .., x_n\}$ in order to use the sample-version of (1.2). Estimation of f is usually done by kernel methods (Vincent and Bengio, 2003; Orlitsky et al., 2005) or nearest neighbor estimate (Bijral et al., 2012). However, using these estimation methods often involve a smoothing parameter such as the bandwidth in case of kernel based methods and number of neighbors in case of nearest neighbor based method. This influences the number of clusters and it is also sub-optimal in the sense we do not need to estimate f at every point in the manifold to know the underlying cluster structure. Similarly mode based clustering methods such as mean shift, quick shift (Chen et al., 2016; Vedaldi and Soatto, 2008) etc also suffer from the same problem of requiring an estimate of f or gradient vector for each data point and

requiring a smoothing parameter which indirectly can influence the elicited cluster structure.

Sample level graph-based density sensitive distances have been proposed for clustering since early 2000s as well. In 2004, Fischer et al. (2004) and Vincent and Bengio (2003) introduced longest leg path distance (LLPD) as an alternative to the Euclidean distance for clustering purposes which is defined as

$$LLPD(x, y, \chi_n) \coloneqq \min_{p} \max_{i} dist_1(x_i, x_{i+1}) \tag{1.3}$$

where p is a path from $x = x_0 \to x_1 \to x_2 .. \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary. The intuition for proposing this distance is as follows: if we consider any graph-based path joining x and y, at some point the path has to "make a jump" from one cluster to another if x and y are from different clusters. Since this jump will happen in a low density region the jump length (LLPD distance between x and y) will be large. However, if x and y are from the same cluster, it's possible to go from x to y staying in a high-density region of f implying LLPD between x and y in this case will be small.

Since it's proposal, many papers (Chang and Yeung, 2008; Fischer et al., 2004; Maggioni and Murphy, 2018) have considered the use of LLPD for classification and clustering methods (K-means and spectral clustering, SVM's). In particular Maggioni and Murphy (2018) recently have provided theoretical justification of using LLPD over Euclidean distance in spectral clustering and also provided approximate theoretical conditions for performance guarantees of spectral clustering with LLPD for low dimensional large noise (LDLN) type (Fefferman et al., 2016) data model. But from the simulation studies done in these papers, it is also well known that although able to extract density information, LLPD is highly unstable and is influenced by (specially lower di-

mensional) outliers. So as a trade-off between LLPD and Euclidean/Default geodesic distance, power-weighted distances (L_{power}) and exponentially-weighted (L_{exp}) distances (which are interpolations between Euclidean distance and LLPD in some sense) have been proposed by other authors and are defined as

$$L_{power}(x, y, \chi_n) \coloneqq \min_{p} \Big(\sum_{i=0}^{r} dist_1(x_i, x_{i+1})^k \Big)^{\frac{1}{k}}$$

where p is a path from $x = x_0 \to x_1 \to x_2 .. \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary and $1 < k < \infty$

and

$$L_{exp}(x, y, \chi_n) \coloneqq \min_{p} \frac{1}{k} ln \Big(1 + \sum_{i=0}^{r} (exp(k\, dist_1(x_i, x_{i+1})) - 1) \Big)$$

where p is a path from $x = x_0 \to x_1 \to x_2 .. \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary and $0 < k < \infty$, respectively.

In particular, power-weighted distance for clustering purposes was first proposed by Vincent and Bengio (2003) and later used for K-means and spectral clustering instead of Euclidean distance by Orlitsky et al. (2005); Bijral et al. (2012), whereas exponential distance was proposed by Wang et al. (2006) and has been used by Bijral et al. (2012); Tao et al. (2019). But until very recently, the theoretical properties of these distances remained unknown and although several applied works and simulation studies have been performed over last 10 years, no to little work has been done towards establishing the underlying distributional properties of these distances and theoretical guarantees of their performance in clustering.

However, recently Hwang et al. (2016) established that population and sample level density sensitive distances are in-fact related. Formally, using first passage percolation theory they showed that suitably scaled sample level power-weighted distance

$n^{\frac{k-1}{D}} L^k_{power}$ converge to the population level density dependent distance

$$dist_k(x, y) \coloneqq \inf_{\gamma} \int_0^1 f^{\frac{1-k}{D}}(\gamma_t)|\gamma_t^{'}|dt$$

up to a constant almost surely. This means that the power weighted distances can be used as an approximation to $dist_k(x, y)$, a special form of (1.2), without needing to estimate f as long as sample size n is sufficiently large (Mckenzie and Damelin, 2019).

1.2 Motivation

Although a connection between sample level path based distance L_{power} and its population counterpart $dist_k$ was made, many interesting research problems in this area of research still remain. We list some of these research problems below which are addressed in this thesis.

1. In last decade (apart from L_{power}, L_{exp} or LLPD), several other sample level density sensitive distance measures have been used in many papers for clustering purposes as well. For example, Lan et al. (2011) proposed to use

$$\min_p \sum_{i=0}^{r} \{(dist_1(x_i, x_{i+1}) + \sigma)^k - \sigma^k\}$$

for some offset factor $\sigma > 0$ instead of using $\min_p \sum_{i=0}^{r} dist_1(x_i, x_{i+1})^k$. Similarly

$$\min_p \sum_{i=0}^{r} \{a\,cosh\, dist_1(x_i, x_{i+1}) - 1\}$$

for $a > 0$,

$$\min_p \sum_{i=0}^{r} \{a\, dist_1(x_i, x_{i+1}) + b\, dist_1(x_i, x_{i+1})^2\}$$

for $a, b > 0$ have been proposed as well. It turns out that all these proposed sample level density sensitive distances after standardization are of the general form

$$L_g^N(x, y, \chi_n) \coloneqq g^{-1}(\min_p \sum_{i=0}^{r} g(dist_1(x_i, x_{i+1})). \tag{1.4}$$

for some continuous and strictly super-additive function $g : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ with $g(0) = 0$. A natural research problem is to study the properties of L_g^N as a distance measure and derive convergence limit and rate of L_g^N.

2. Although LLPD (which can be viewed a limiting form of L_g^N as well) has been used for clustering applications in many papers over last 15 years, convergence limit as well as rate of convergence of LLPD remain unknown.

3. As discussed above, the primary motivation for using both sample and population level density sensitive distances for clustering purposes is that they bend the geodesic joining 2 points towards higher density region of M which lead to near perfect clustering. But how to quantify this bending phenomenon? From a physics point of view bending is caused by a field, in this case the gradient field of f, ∇f. When ∇f is absent and $\boldsymbol{x}(t)$ denotes the trajectory of the default geodesic, then it is well known that the trajectory of $\boldsymbol{x}(t)$ is governed by Einstein's geodesic equation

$$\boldsymbol{x}^{''} + \Gamma = 0$$

where $\boldsymbol{x}'' := \begin{pmatrix} x_1'' \\ . \\ . \\ x_D'' \end{pmatrix}$ and $\Gamma = \begin{pmatrix} \sum_{j=1}^{D}\sum_{k=1}^{D}\Gamma_{jk}^{1}x_j'x_k' \\ . \\ . \\ \sum_{j=1}^{D}\sum_{k=1}^{D}\Gamma_{jk}^{D}x_j'x_k' \end{pmatrix}$, $\{\Gamma_{ij}^{k}\}_{i,j,k=1(1)D}$'s being the Christoffel symbols.

An interesting research problem is to quantify how ∇f changes the default geodesic equation $\boldsymbol{x}'' + \Gamma = 0$ for L_g^N or LLPD and what type of geometry does these distances induce?

4. In many clustering applications (specially in image and mircrobiome data) at stage 1 we obtain *islands* as clusters which are connected components of super-level set $M_\lambda := \{x \in M : f(x) \geq \lambda\}$ where λ denote the *sea-level*. But often there are too many *islands* and at stage 2 we want to merge these *islands* in order to obtain *continents* for meaningful interpretation. In order to solve the ***continent problem*** the factors often in consideration are *size*, *topographic isolation* (these 2 factors are geometric characteristics of islands) and *topographic prominence* (this is a density based characteristic of islands) of *islands*. However, it turns out that in clustering literature this problem is not yet well-formulated and in many papers these are addressed in an ad hoc fashion.

5. As discussed above although LLPD and different variants of L_g^N have been used for spectral clustering purposes over two decades, very little work has been done so far to derive theoretical conditions for performance guarantees of these algorithms or to compare different distances (LLPD or default geodesic distance or L_g^N distance? What type of g is best for cluster purposes?) under various data generating models and sample size regime. In particular, the data generating models we have considered in this thesis are

   - Data generated from a density f in a D dimensional Riemannian manifold.

- Data generated from multiple disjoint lower dimensional Riemannian manifold with or without D-dimensional ambient noise, known as LDLN model in literature.
- The previous settings with the additional requirement of being able to detect intersections between manifolds, known as multi-manifold clustering (MMC) in literature.

1.3 Organization of work

- In **Chapter 2**, we have introduced g-distance as a generalization of power (which corresponds to $g(t) = t^p$, $p > 1$) weighted distances and other sample level density sensitive distances proposed by several authors over last two decades. We have shown that normalized g-distance can be viewed as an interpolation between Euclidean distance and LLPD. Then asymptotic limits of g-distances have been derived by partly following Hwang et al. (2016)'s proof technique. As a corollary we have shown that in small/medium sample scenario, a non-trivial g-distance behaves like a density based distance in low density region while behaving like default geodesic distance in high density region of data. Thus g-distance can effectively extract information from available data by prioritizing density information in low-density region of data (where density information is pricey) while prioritizing geometry information in high-density region (where density information is cheap) of data.

- In **Chapter 3**, almost sure limit of LLPD and it's convergence rate has been derived.

- In **Chapter 4** we show that sample g distance arc between x and y converge to the geodesic of a density dependent metric as $n \to \infty$. This geodesic is an

interpolation between *straight line* (i.e. default geodesic) and *field lines* induced by gradient field of f joining x and y and is able to capture both the density as well as geometry information present. This property makes suitably scaled g distance ideal for clustering on manifolds. In particular we have derived underlying differential equation governing g-distance geodesic using calculus of variation followed by introducing several geometrical concepts such as *basin of attractions* and *bottleneck values* and then exploring the geometry of g-distance through some Gaussian mixture model (GMM) examples.

- In **Chapter 5** we consider the problem of spectral clustering (SC) with g-distance and the longest-leg path distance (LLPD) metric, which are able to capture both the geometry and the density information from the data. We show that SC with g-distance is an interpolation between traditional K-means or centroid based clustering and mode or topological prominence based clustering. We prove guarantees on the performance of SC with respect to these metrics when random samples are drawn from different data generating models such as a d dimensional Riemannian manifold for $d > 1$, multiple lower dimensional Riemannian manifolds embedded inside ambient D dimensional space with or without the presence of background noise, multiple self-intersecting Riemannian manifolds etc. In particular, we provide conditions under which the Laplacian eigen-gap statistic correctly determines the number of clusters under these data generating models and prove guarantees on the labeling accuracy of the proposed algorithms while simultaneously comparing the performance of different g distances and LLPD.

- In **Chapter 6** we develop a technique for record linkage on high dimensional data, where the two datasets may not have any common variable, and there may be no training set available. Our methodology is based on sparse, high di-

mensional principal components. Since large and high dimensional datasets are often prone to outliers and aberrant observations, we propose a technique for estimating robust, high dimensional principal components. We present theoretical results validating the robust, high dimensional principal component estimation steps, and justifying their use for record linkage. Some numeric results and remarks are also presented.

- In **Chapter 7** we expound on future directions of research.

Chapter 2

The g Distance: a general density sensitive distance for clustering

As discussed in chapter 1, we have introduced g-distance as a generalization of power (which corresponds to $g(t) = t^p$, $p > 1$) weighted distances and other sample level density sensitive distances proposed by several authors over last 15 years. In section 2.1 we provide a brief motivation and state the main results. In section 2.2 we provide some properties of g distance in relation to default geodesic distance and the longest leg path distance. In section 2.3 we provide the proof of convergence limit and rate of g-distance under some mild assumptions on g. In section 2.4 we discuss what happens when those assumptions are removed and provide some examples.

2.1 Problem description, notations and main results

Let (M, g_1) be a smooth d-dimensional ($d > 1$) Riemannian manifold without boundary equipped with Riemannian metric tensor g_1. Recall that a Riemannian metric tensor, often simply called a Riemannian metric, on a manifold is a family of positive definite inner products on the tangent spaces of the manifold. When $M = \mathbb{R}^d$ or a convex subset of $\mathbb{R}^d$, g_1 is simply the standard Euclidean inner product. This metric

tensor induces the usual distance, which we shall call by $dist_1(.,.)$. Formally,

$$dist_1(x,y) = \inf_{\gamma} \int_0^1 \sqrt{g_1(\gamma_t', \gamma_t')}dt = \inf_{\gamma} \int_0^1 |\gamma_t'|dt \tag{2.1}$$

where the infimum is taken over all piecewise smooth curves $\gamma : [0,1] \to M$ such that $\gamma(0) = x$ and $\gamma(1) = y$. We use the notation $\overline{\gamma_1(x,y)}$ or simply $\overline{xy}$ to be the g_1-geodesic from x to y, i.e. path traced-out by γ_t which minimizes $dist_1(x,y)$ subject to $\gamma(0) = x$ and $\gamma(1) = y$.

Consider a probability space $(M, \mathbb{B}, P)$ where P is a probability distribution over Borel subsets $\mathbb{B}$ of the sample space M. Assume that the distribution has a smooth (i.e. $f \in C^\infty$) Lebesgue probability density function f with respect to g_1. Let $\chi_n \coloneqq \{X_1, X_2, .., X_n\}$ be a set of n sample points in M drawn from f. The sequence χ_n will be associated with the nodes in an undirected simple graph whose edges have weight equal to the $w(x,y) = g(dist_1(x,y))$ between pairs of nodes x and y where $g(.) : [0,\infty) \to \mathbb{R}$ is a function such that $g^{-1}(.)$ exists. We will also sometimes use indexing by n of a generic non-random point $x_n \in M$. This point is not related in any way to the random variable X_n or the data χ_n. For realizations, we will use the notation $X_n(\omega)$ where ω is an elementary outcome in the sample space. We assume $\inf_M f(x) = f_0 > 0$ (unless otherwise is mentioned) and $\sup_M f(x) = f^0 < \infty$.

Also, following Hwang et al. (2016), we define a new metric tensor $g_k = f^{2(1-k)/d} g_1$ for $k \geq 1$ i.e. if Z_x and W_x are two tangent vectors at a point $x \in M$, then $g_k(Z_x, W_x) = f(x)^{2(1-k)/d} g_1(Z_x, W_x)$.

The deformed metric tensor g_k is well-defined for every x with $f(x) > 0$, and g_k is a Riemannian metric tensor when $f \in C^\infty$. The corresponding induced distance is

$$dist_k(x,y) = \inf_{\gamma} \int_0^1 f(\gamma_t)^{\frac{1-k}{d}} |\gamma_t'|dt \tag{2.2}$$

where the infimum is taken over all piecewise smooth curves $\gamma : [0,1] \to M$ such that

$\gamma(0) = x$ and $\gamma(1) = y$. We denote $\gamma_\beta(x, y)$ to be the g_k-geodesic from x to y, i.e. path traced-out by γ_t which minimizes $dist_k(x, y)$ subject to $\gamma(0) = x$ and $\gamma(1) = y$, where $\beta \coloneqq \frac{k-1}{d}$.

2.1.1 A general metric

Consider a function $g : [0, \infty) \to [0, \infty)$, having the following properties:

G1: $g(\cdot)$ is a continuous function.

G2: $g(t) = 0$ if and only if $t = 0$.

G3: $g(\cdot)$ is a strictly increasing function.

G4: $g(\cdot)$ is a strictly super-additive function, that is,

$$g(a + b) > g(a) + g(b), \quad \text{for any } a, b \in [0, \infty).$$

There is some degree of redundancy in the above conditions: this is for simplicity and clarity. Let x, y be two points in M. Define $L_g\big(x, y, \chi_n\big)$ as

$$L_g\big(x, y, \chi_n\big) \coloneqq \min_p \sum_{i=0}^{r} g\big(dist_1(x_i, x_{i+1})\big), \tag{2.3}$$

where p is any path through points $x_0 = x \to x_1 \to x_2 \to \ldots \to x_r \to x_{r+1} = y$, $x_i \in \chi_n$ for $i = 1(1)r$, r is free to vary. Thus, $L_g\big(x, y, \chi_n\big)$ is the shortest g-path distance from x to y through χ_n. We define $\overline{p_g(x, y, \chi_n)}$ as the corresponding almost surely unique shortest g-path.

The normalized g-distance $L_g^N\big(x, y, \chi_n\big)$ between the two points x and y in M is defined by

$$L_g^N\big(x, y, \chi_n\big) \coloneqq g^{-1}\big(L_g(x, y, \chi_n)\big). \tag{2.4}$$

In this context, also recall that longest leg path distance (LLPD) is defined as

$$LLPD\big(x, y, \chi_n\big) = \min_{p} \max_{i} dist_1(x_i, x_{i+1}), \tag{2.5}$$

where as usual p is a path from $x_0 = x \to x_1 \to x_2 \to \ldots \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary.

We now have the the following important result:

Proposition 2.1

Under the set of conditions **G1**-**G4** on $g(\cdot)$, both $L_g\big(x, y, \chi_n\big)$ and $L_g^N\big(x, y, \chi_n\big)$ are continuous metrics on M. Additionally, $L_g^N\big(x, y, \chi_n\big)$ is a non-trivial metric on M.$\square$

Proof 2.1

First we show that L_g is a continuous metric. Non-negativity of L_g follows from **G2** and **G3**, identity of indiscernible follow from property **G2**, $L_g(x, y, \chi_n)$ is symmetric because $dist_1(x, y)$ is symmetric and finally $L_g(x, y, \chi_n) \leq L_g(x, z, \chi_n) + L_g(y, z, \chi_n)$ follows from the definition of shortest path (triangle Inequality). Additionally, note that L_g is continuous because of **G1**.

Next we show that L_g^N is a continuous and non-trivial metric. Note that because of **G3**, g^{-1} exists. Non-negativity of L_g^N follows from **G2** and **G3**, identity of indiscernible follow from property **G2**, $L_g^N(x, y, \chi_n)$ is symmetric because $dist_1(x, y)$ is symmetric. Triangle inequality follows since g^{-1} is sub-additive. Indeed,

$$L_g^N(x, y, \chi_n) = g^{-1}[L_g(x, y, \chi_n)] \leq g^{-1}[L_g(x, z, \chi_n) + L_g(y, z, \chi_n)]$$

$$\leq g^{-1}[L_g(x, z, \chi_n)] + g^{-1}[L_g(y, z, \chi_n)] = L_g^N(x, z, \chi_n) + L_g^N(y, z, \chi_n).$$

Also note that L_g^N is continuous because of **G1**.

Finally we show that g needs to be strictly super-additive for L_g^N to be a non-trivial metric. Suppose g is sub-additive. Consider any path p, $x_0 = x \to x_1 \to ... \to$

$x_{m+1} = y$. Then,

$$\sum_{i=0}^{m} g(dist_1(x_i, x_{i+1})) \geq g(\sum_{i=0}^{m} dist_1(x_i, x_{i+1})) \geq g(dist_1(x, y))$$

i.e.

$$L_g(x, y, \chi_n) = \min_p \sum_{i=0}^{m} g(dist_1(x_i, x_{i+1})) = g(dist_1(x, y))$$

$$\Rightarrow L_g^N(x, y, \chi_n) = dist_1(x, y)$$

for all x and y. i.e. L_g^N becomes trivial. □

We formally define

$$\mathbb{G} \coloneqq \{g : [0, \infty) \to [0, \infty) \,|\, g(0) = 0,\ g \text{ is continuous and strictly super-additive}\}$$

as the collection of all g distance functions satisfying **G1-G4**. The assumptions **G1-G4** on g are necessary for ensuring L_g and L_g^N are metrics on M.

We now list one more assumption on g that facilitate easier analysis using these metrics. For some $\delta > 0$ and $0 < t < \delta$ we assume that

G5:

$$g(t) = \sum_{i=1}^{m} a_i t^{b_i}, : \quad m \in \bar{\mathbb{N}},\ a_i \in \mathbb{R}^+,\ b_1 < b_2 < ..$$

Note that from the necessary assumptions **G1-G4** alone, when additionally **G5** is satisfied one can derive $a_1 > 0$ and $b_1 \geq 1$. However, it's not true that a_2, a_3.. etc also need to be > 0. However, we assume so for the ease of our analysis.

Proposition 2.2

Under **G5**, $b_1 \geq 1$. □

Proof 2.2

Suppose that $b_1 < 1$. Then for $r, s > 0$ we have

$$g(r+s) = a_1(r+s)^{b_1}[1+o(1)] \tag{2.6}$$
$$< a_1(r^{b_1} + s^{b_1})[1+o(1)] \tag{2.7}$$
$$= g(r) + g(s), \tag{2.8}$$

and thus **G4** is violated. □

Additionally it is to be noted that, under **G1-G5**, when $b_1 = 1$, $b_2 > 1$ must exist because of our requirement of strict super-additivity. We formally define

$$\mathbb{G}_0 \coloneqq \{g = \sum_{i=1}^{m} a_i t^{b_i} : [0,\infty) \to [0,\infty) \,|m \in \bar{\mathbb{N}},\, a_i \in \mathbb{R}^+, b_i \in [1,\infty),\, 1 \leq b_1 < b_2 < ..\} \tag{2.9}$$

as the collection of all g distance function satisfying **G1-G5**.

2.1.2 Examples of *g*-functions

One of the simplest way of obtaining a function satisfying assumptions **G1-G4** is to consider a *strictly increasing convex* function on $[0,\infty)$ that satisfies the condition $g(0) = 0$. The continuity and strict super-additivity of $g(\cdot)$ automatically follows.

Example 2.1

Here we provide some examples of g for which all the assumptions **G1-G5** are satisfied i.e. $g \in \mathbb{G}_0$.

1. *Power-weighted distance family* is defined by $g_p(t) = t^p$ for $p > 1$. In this case $m = 1$, $b_1 = p$, $a_1 = 1$.

2. *Exponential-weighted distance family* is defined by $g_a(t) = e^{at} - 1$ for $a > 0$. In this case $m = \infty$, $b_k = k$ and $a_k = \frac{a^k}{k!}$ for $k \in \mathbb{N}$.

3. For $g_a(t) = \frac{1}{a^2}e^{at} - \frac{t}{a} - \frac{1}{a^2}$ with $a > 0$, we have $m = \infty$, $b_k = k+1$ and $a_k = \frac{a^{k-1}}{k+1!}$ for $k \in \mathbb{N}$.

4. For $g_{p,q,r,s}(t) = e^{p+e^{qt^r+s}} - e^{p+e^s}$ with $p, s \in \mathbb{R}$ and $q > 0$, $r \geq 1$ we have

$$g_{p,q,r,s}(t) = e^p[(e^{e^s})^{e^{qt^r}} - e^{e^s}]$$

$$= X[Y^{Z^{t^c}} - Y]$$

where $X = e^p > 0$, $Y = e^{e^s} > 1$, $Z = e^q > 1$. Note that

$$Z^{t^r} = e^{ln\, Z\, t^r} = 1 + q\, t^r + \frac{q^2}{2}t^{2r} + ..$$

. Hence,

$$g_{p,q,r,s}(t) = XY\big(Y^{(qt^r+\frac{q^2}{2}t^{2r}+..)} - 1\big) = XY\big((e^{qln(Y)t^r}e^{\frac{q^2}{2}ln(Y)t^{2r}}..) - 1\big)$$

$$= XY\big(qln(Y)t^r + O(t^{2r}) + ..\big)$$

Thus in this case we have $m = \infty$, $b_k = kr$ for $k \in \mathbb{N}$ and $a_1 = qe^{s+p+e^s}$. □

However, note that **G5** do not follow from **G1**-**G4**. Below we provide some examples of g where $g \in \mathbb{G}$ but $g \notin \mathbb{G}_0$.

Example 2.2

Consider $g(t) = e^{-\frac{1}{t}}$ in $(0, \delta)$ with $g(0) = 0$ and δ sufficiently small so that g is strictly super-additive. Since $g''(t) = \frac{e^{-\frac{1}{t}}}{t^4}(1 - 2t)$, $\delta < \frac{1}{2}$ suffices. Thus g satisfies **G1**-**G4**. However it can be easily shown that $\lim_{t\downarrow 0} \frac{g(t)}{t^k} = 0$ for all $k \in (0, \infty)$. Thus **G5** is not satisfied. □

Example 2.3

Consider

$$g(t) = t\int_0^t (1+h(x))dx,\ 0 \le t \le 1$$

where

$$h(x) = \begin{cases} \sum_{i=0}^{\infty} 1_{x\in[2^{-2i},2^{-(2i+1)})} & x \in (0,1] \\ 0 & x = 0. \end{cases}$$

Clearly g is continuous, $g(0) = 0$. Also, for $r, s > 0$

$$g(r+s) = (r+s)\int_0^{r+s} (1+h(x))dx$$

$$= r\int_0^{r+s} (1+h(x))dx + s\int_0^{r+s} (1+h(x))dx$$

$$> r\int_0^{r} (1+h(x))dx + s\int_0^{s} (1+h(x))dx = g(r) + g(s)$$

i.e. g is strictly super-additive. Hence g satisfies **G1-G4**. Note that for $n \in \mathbb{N}$,

$$\int_0^t h(x)dx = \begin{cases} \sum_{i=0}^{\infty} 2^{-(2n+1+2i)} = \frac{2}{3}2^{-(2n)} = \frac{2t}{3} & \text{for } t = 2^{-2n} \\ \sum_{i=0}^{\infty} 2^{-(2n+3+2i)} = \frac{1}{3}2^{-(2n+1)} = \frac{t}{3} & \text{for } t = 2^{-(2n+1)} \end{cases}$$

Thus, for $n \in \mathbb{N}$,

$$g(t) = \begin{cases} \frac{5}{3}t^2 & \text{for } t = 2^{-2n} \\ \frac{4}{3}t^2 & \text{for } t = 2^{-(2n+1)} \end{cases}$$

Thus clearly **G5** is not satisfied. □

Example 2.4

It is easy to verify that for $g_1(t) = t + e^{-\frac{1}{t}}$, $g_2(t) = t^2 - \frac{t^3}{10} + t^4$ and $g_3(t) = t + g(t)$ where $g(t)$ is same from Example 2.3 , **G1-G4** are satisfied, but **G5** is not satisfied.□

Remark 2.1

We shall assume $g \in \mathbb{G}_0$ in the proof of theorem 2.1 although as we shall see later in section 2.5 that this assumption is mainly technical and convergence limit and rate can be obtained in-general without this assumption. □

Apart from the ease of proof, there are primarily two other reasons for working with $g \in \mathbb{G}_0$ class instead of the bigger class $g \in \mathbb{G}$.

1. Theorem 2.1 states that for *any* $g \in \mathbb{G}_0$, in high sample size/high density regime suitably scaled g^N converge to the normalized power-weighted distance corresponding to the lowest power of g. An important side result is in low sample size/low density regime, suitably scaled g^N behaves like normalized power-weighted distance corresponding to the highest power of g if m is finite. While, same conclusion can be made for any $g \in \mathbb{G}$ in high sample size/high density regime by arguing similarly, what happens in low sample size/low density regime is in-general unclear and depends on g. Below we provide 2 examples to illustrate this.

   Example 2.5

   Consider $g(t) = \{\frac{1}{a^2}e^{-at} + \frac{t}{a} - \frac{1}{a^2}\}_{a>0}$. Since $g^{'}(t) = \frac{1}{a}(1 - e^{-at}) > 0$ for $t > 0$ and $g^{''}(t) = e^{-at} > 0$, $g \in \mathbb{G}$. Also $g(t) = \frac{t^2}{2!} - a\frac{t^3}{3!} + a^2\frac{t^4}{4!} - ..$, thus $g \notin \mathbb{G}_0$. Since $b_1 = 2$, in high sample size/density regime suitably scaled g^N converge to normalized power(p)-weighted distance with $p = 2$. However as $t \to \infty$, $g(t) = O(t)$. Thus in low sample size/low density regime suitably scaled g^N converge to Euclidean distance. □

   Example 2.6

   Consider $g(t) = \{(at + 1)\, log(at + 1) - at\}_{a>0}$. Since $g^{'}(t) = log(at + 1) > 0$ for $t > 0$ and $g^{''}(t) = \frac{a}{at+1} > 0$, $g \in \mathbb{G}$. Also $g(t) = \frac{(at)^2}{2} - \frac{(at)^3}{6} + \frac{(at)^4}{12} - ..$,

thus $g \notin \mathbb{G}_0$. Since $b_1 = 2$, in high sample size/density regime suitably scaled g^N converge to normalized power(p)-weighted distance with $p = 2$. However as $t \to \infty$, $g(t) = O(t\,logt)$. Thus in low sample size/ low density regime what happens to g^N is unclear at least from theorem 2.1 and needs additional analysis as provided in section 2.5. □

For reasons stated in the following paragraph, we would like to behave g^N more like it's highest power-weighted distance in high sample size/density regime and lowest power-weighted distance in low sample size/density regime.

2. For clustering applications, in high sample size/high density regime (when density information is easily available) g^N should prioritize geometry based information while in low sample size/low density regime (when density information is pricey) g^N should prioritize density based information. This is *exactly* what g^N with $g \in \mathbb{G}_0$ do by graduating from geometry dominant distances such as Euclidean distance/smallest normalized power-weighted distance to density dominant distances such as highest normalized power-weighted distance/LLPD as density information becomes more scarce. As we have seen in some examples above, in-general g^N with $g \in \mathbb{G} \setminus \mathbb{G}_0$ fails to do this. Thus, $\mathbb{G}_0$ is more suitable in clustering application.

2.1.3 Main result

The following two theorems are the main results of this chapter

Theorem 2.1

Assume that M is compact and f is continuous with $f_0 > 0$ and $g \in \mathbb{G}_0$. Then for sufficient small $\epsilon > 0$, there exists a constant $C(d, b_1|g) > 0$ depending on d, b_1 and g

and a constant $\theta_0 > 0$ such that whenever $dist_1(x,y) \geq b > 0$,

$$\mathbb{P}\big(\sup_{x,y}|\frac{L_g(x,y,\chi_n)}{a_1 n^{\frac{1-b_1}{d}} dist_{b_1}(x,y)} - C(d,b_1|g)| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\alpha^*}) \tag{2.10}$$

for all sufficiently large n, where the supremum is taken all over x, $y \in M$ with $dist_1(x,y) \geq b > 0$ and α^* is given by

$$\alpha^* = \begin{cases} \frac{1}{d+2b_1} & \text{if } b_1 > 1 \\ \frac{1}{d+b_2+1} & \text{if } b_1 = 1 \end{cases} \tag{2.11}$$

for some $\theta_0 > 0$ depending on d and g. Also, if $b_1 = 1$, $C(d,b_1|g) = 1$. □

For example, if we apply theorem 2.1 on the 4 g-distances listed in Example 2.1, we obtain

1. $\mathbb{P}\big(\sup_{x,y}|\frac{L_{g_p}(x,y,\chi_n)}{n^{\frac{1-p}{d}} dist_p(x,y)} - C(d,p)| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\frac{1}{d+2p}})$

2. $\mathbb{P}\big(\sup_{x,y}|\frac{L_{g_a}(x,y,\chi_n)}{a\, dist_1(x,y)} - 1| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\frac{1}{d+3}})$

3. $\mathbb{P}\big(\sup_{x,y}|\frac{2L_{g_a}(x,y,\chi_n)}{n^{-\frac{1}{d}} dist_2(x,y)} - C(d,2|g_a)| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\frac{1}{d+4}})$

4. $\begin{cases} \mathbb{P}\big(\sup_{x,y}|\frac{L_{g_{p,q,r,s}}(x,y,\chi_n)}{qe^{s+p+e^s} n^{\frac{1-r}{d}} dist_r(x,y)} - C(d,r|g_{p,q,r,s})| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\frac{1}{d+2r}}) & \text{if } r > 1 \\ \mathbb{P}\big(\sup_{x,y}|\frac{L_{g_{p,q,r,s}}(x,y,\chi_n)}{qe^{s+p+e^s} dist_1(x,y)} - 1| > \epsilon\big) \leq exp(-\theta_0\epsilon^2 n^{\frac{1}{d+3}}) & \text{if } r = 1 \end{cases}$

Corollary 2.1

An elementary Lipschitz analysis on (2.10) leads to

$$\mathbb{P}\big(\sup_{x,y}|\frac{L_g^N(x,y,\chi_n)}{g^{-1}(a_1 n^{\frac{1-b_1}{d}} dist_1(x,y))} - C(d,b_1|g)| > \epsilon\big) \leq exp(-\theta_0'\epsilon^2 n^{\alpha^*})$$

as well, for some constant $\theta_0' > 0$ depending on d and g. □

Note that when $b_1 > 1$, this further simplifies to

$$\mathbb{P}\big(\sup_{x,y}\frac{L_g^N(x,y,\chi_n)}{n^{\frac{1-b_1}{b_1 d}}\, dist_{b_1}(x,y)^{\frac{1}{b_1}}} - C(d,b_1|g)| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\alpha^*})$$

since $g^{-1}(t) \approx (a_1^{-1}t)^{\frac{1}{b_1}}$ as $t \downarrow 0$. When $b_1 = 1$, we have

$$\mathbb{P}\big(\sup_{x,y}\frac{L_g^N(x,y,\chi_n)}{g^{-1}(a_1 dist_1(x,y))} - 1| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\alpha^*})$$

for some $\theta_0^{'} > 0$ depending on d and g.

Thus if we apply corollary 2.1 on the 4 g-distances listed in Example 2.1, we obtain

1. $$\mathbb{P}\big(\sup_{x,y}\frac{L_{g_p}^N(x,y,\chi_n)}{n^{\frac{1-p}{pd}}\, dist_p(x,y)^{\frac{1}{p}}} - C(d,p)| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\frac{1}{d+2p}})$$

2. $$\mathbb{P}\big(\sup_{x,y}\frac{L_{g_a}^N(x,y,\chi_n)}{\frac{1}{a}\, ln(1+a\, dist_1(x,y))} - 1| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\frac{1}{d+3}})$$

3. $$\mathbb{P}\big(\sup_{x,y}\frac{L_{g_a}^N(x,y,\chi_n)}{n^{-\frac{1}{2d}} dist_2(x,y)^{\frac{1}{2}}} - C(d,2|g_a)| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\frac{1}{d+4}})$$

4. $$\begin{cases} \mathbb{P}\big(\sup_{x,y}\frac{L_{g_{p,q,r,s}}^N(x,y,\chi_n)}{n^{\frac{1-r}{rd}}\, dist_r(x,y)^{\frac{1}{r}}} - C(d,r|g_{p,q,r,s})| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\frac{1}{d+2r}}) & \text{if } r > 1 \\ \mathbb{P}\big(\sup_{x,y}\frac{L_{g_{p,q,r,s}}^N(x,y,\chi_n)}{[\frac{1}{q}(ln(ln(e^{p+e^s}\,(1+qe^s\, dist_1(x,y))-p)-s)]^{\frac{1}{r}}} - 1| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\frac{1}{d+3}}) & \text{if } r = 1 \end{cases}$$

since $g_{p,q,r,s}^{-1}(t) = \big(\frac{1}{q}[ln(ln(t+e^{p+e^s}) - p) - s]\big)^{\frac{1}{r}}$.

Remark 2.2

The assumption $dist_1(x,y) \geq b > 0$ (i.e. x and y are fixed non-identical points with some minimum distance between them) can be relaxed in theorem 2.1. When $b_1 > 1$, instead of looking at a fixed point x and y if we consider a sequence of points $\{x_n\}$ and $\{y_n\}$ with $dist_1(x_n,y_n) \geq r_n$ for some positive sequence $\{r_n\} \downarrow 0$, the probability upper bound $exp(-\theta_0\epsilon^2 n^{\alpha^*})$ can be re-written as $exp\big(-\theta_0^{'}\epsilon^2(nr_n^d)^{\alpha^*} + O(logn)\big)$ where $\theta_0^{'} = \theta_0 b^{-d\alpha^*}$. Therefore for the probability upper bound to be non-trivial, we require $\frac{nr_n^d}{logn} \to \infty$, i.e. $dist(x_n,y_n) > O\big((\frac{logn}{n})^{\frac{1}{d}}\big)$. □

Remark 2.3

Similarly when $b_1 = 1$, probability upper bound $exp(-\theta_0 \epsilon^2 n^{\alpha^*})$ can be re-written as $exp\big(-\theta_0^{'} \epsilon^2 (n r_n^{d+b_2-1})^{\alpha^*} + O(logn)\big)$ where $\theta_0^{'} = \theta_0 b^{-d\alpha^*}$and where we assume $dist_1(x_n, y_n) \geq r_n$ for some positive sequence $\{r_n\}$. In this case we require $\frac{n r_n^{d+b_2-1}}{logn} \to \infty$, i.e. $dist(x_n, y_n) > O\big((\frac{logn}{n})^{\frac{1}{d+b_2-1}}\big)$ for the probability upper bound to be non-trivial. In this case the effective dimension increases from d to $d + b_2 - 1$. □

A similar convergence result holds when M is complete (but not necessarily compact) giving the almost sure limit stated below

Theorem 2.2

Assume that M is complete and that f is continuous with $f(u) > 0$ for all $u \in M$. Fix $x, y \in M$. Then

$$\lim_{n\to\infty} \frac{1}{a_1} n^{\frac{b_1-1}{d}} L_g(x, y, \chi_n) = C(d, b_1|g)\, dist_{b_1}(x, y) \ \text{ almost surely.}, \tag{2.12}$$

The constant $C(d, b_1|g)$ is the same constant as in equation (2.10). □

Corollary 2.2

It follows that under the same set-up as theorem 2.2, almost surely

$$n^{\frac{b_1-1}{b_1\, d}}\, L_g^N(x, y, \chi_n) \to \begin{cases} C(d, b_1|g)^{\frac{1}{b_1}}\, dist_{b_1}(x, y)^{\frac{1}{b_1}} & b_1 > 1 \\ g^{-1}(a_1\, dist_1(x, y)) & b_1 = 1 \end{cases}$$

□

2.2 $L_g^N(x, y, \chi_n)$: an interpolation between $dist_1$ and LLPD

In this section we derive several results that shows that normalized g-distance L_g^N can be viewed as an interpolation between $dist_1$ and LLPD.

Proposition 2.3

For any $g \in \mathbb{G}$, $LLPD(x, y, \chi_n) \leq L_g^N(x, y, \chi_n) \leq dist_1(x, y)$ for any x and $y \in M$. □

Proof 2.3

One useful property in clustering and dimension reduction applications of $L_g^N(x, y, \chi_n)$ is that it is a contraction map of $dist_1(x, y)$. Indeed,

$$L_g^N(x, y, \chi_n) = g^{-1}[\min_p \sum_{i=0}^{m} g(dist_1(x_i, x_{i+1}))]$$

$$\leq \min_p \sum_{i=0}^{m} g^{-1}[g(dist_1(x_i, x_{i+1}))]$$

$$= \min_p \sum_{i=0}^{m} dist_1(x_i, x_{i+1})) = dist_1(x, y)$$

Equality holds if and only if $L_g(x, y, \chi_n) = g(dist_1(x, y))$, i.e. when shortest g-path between x and y is the direct edge connecting x and y. It is easy to check that $LLPD(x, y, \chi_n) \leq L_g^N(x, y, \chi_n)$ since for any path p,

$$\max_i g(dist_1(x_i, x_{i+1})) \leq \sum_{i=0}^{m} g(dist_1(x_i, x_{i+1}))$$

. Thus we always have,

$$LLPD(x, y, \chi_n) \leq L_g^N(x, y, \chi_n) \leq dist_1(x, y) \tag{2.13}$$

□

Now we present a sufficient condition for a family of functions $\{g_a\}$ under which $\lim_{a\to\infty} L_{g_a}^N(x,y,\chi_n) \to LLPD(x,y,\chi_n)$.

Proposition 2.4

Assume $\{g_a(t)\}_{a\in\mathbb{R}_0^+}$ are a family of functions with $g_a \in \mathbb{G}$ and $g_a(t)$ is defined for all $t \in [0,\epsilon)$ such that for any $0 < x < y < \epsilon$ (ϵ can be ∞) we have $\lim_{a\to\infty} \frac{g_a(x)}{g_a(y)} = 0$. Then for any finite n,

$$\lim_{a\to\infty} L_{g_a}^N(x,y,\chi_n) = LLPD(x,y,\chi_n)$$

for any $x, y \in M$. □

Proof 2.4

$$\lim_{a\to\infty} L_{g_a}^N(x,y,\chi_n) = \lim_{a\to\infty} \min_p g_a^{-1}\big[\sum_{i=0}^{m} g_a(dist_1(x_i,x_{i+1}))\big]$$

$$= \min_p \lim_{a\to\infty} g_a^{-1}\big[\frac{\sum_{i=0}^m g_a(dist_1(x_i,x_{i+1}))}{\max_i g_a(dist_1(x_i,x_{i+1}))} \max_i g_a(dist_1(x_i,x_{i+1}))\big]$$

$$= \min_p g_a^{-1}\big[\lim_{a\to\infty}\big[\frac{\sum_{i=0}^m g_a(dist_1(x_i,x_{i+1}))}{\max_i g_a(dist_1(x_i,x_{i+1}))}\big] \max_i g_a(dist_1(x_i,x_{i+1}))\big]$$

$$= \min_p g_a^{-1}\big[1 \times \max_i g_a(dist_1(x_i,x_{i+1}))\big] = \min_p \max_i g_a^{-1}(g_a(dist_1(x_i,x_{i+1}))$$

$$= LLPD(x,y,\chi_n)$$

Since almost surely $\max_i g_a(dist_1(x_i,x_{i+1}))$ is unique for any path p and n is finite.□

Corollary 2.3

For the families in Example 2.1, i.e. $g_p(t) = \{t^p\}_{p>1}$, $g_p(t) = \{e^{pt}-1\}_{p>0}$ and $g_p(t) = \{\frac{1}{p^2}e^{pt} - \frac{t}{p} - \frac{1}{p^2}\}_{p>0}$, we have

$$\lim_{p\to\infty} L_{g_p}^N(x,y,\chi_n) = LLPD(x,y,\chi_n)$$

□

Proposition 2.5

Suppose $g \in \mathbb{G}_0$ and $m = \infty$. Then for any finite n

$$\lim_{|x-y|\to\infty} L_g^N(x, y, \chi_n) = LLPD(x, y, \chi_n)$$

□

Proof 2.5

Consider the family $\{g_a(t) = g(at)\}_{a\in\mathbb{R}_0^+}$. Clearly for $0 < p < q$ $\lim_{a\to\infty} \frac{g_a(p)}{g_a(q)} = 0$ if $m = \infty$. This in particular implies $\lim_{|x-y|\to\infty} L_g^N(x, y, \chi_n) = LLPD(x, y, \chi_n)$ for $g(t) = e^t - 1$. □

2.3 Proof of the main result

An important property of $L_g(x, y, A)$ for $x, y \in M$ and locally finite $A \subset M$ is that if $A_0 \subset A$ then $L_g(x, y, A) \leq L(x, y, A_0)$. This property will be used in several places in the proofs. Going forward, we shall use **case 1** and **case 2** to denote $b_1 > 1$ and $b_1 = 1$ respectively.

The generic step of the proof is to first prove version of theorem for a homogeneous Poisson process H_λ on $\mathbb{R}^d$, then extend it to a uniform density on $M \subset \mathbb{R}^d$ case followed by an extension to a generic density and finally extending it for the Riemannian manifold case. We shall only prove the result for a uniform density case since the rest is very similar to the proof technique in Hwang et al. (2016).

2.3.1 Empty region graph of g distance

In this sub-section we introduce the notion of empty region graph (ERG) of g distance which will be useful in the proof of theorem 2.1. In particular, asymptotic convergence rate of g-distance is determined by the limiting volume of the *empty region* in ERG as derived in lemma 2.1 and proposition 2.12. We shall additionally use several results

such as propositions 2.7, 2.9 and 2.15 in the proof of theorem 2.1. The remaining results are of independent interest and also demonstrate that (specially lemma 2.2) the normalized g distance behaves like it's lowest power-weighted distance in high sample size/density regime and highest power-weighted distance or LLPD in low sample size/density regime.

ERG was introduced in 2009 by Cardinal et al. (2009) as a generalization of proximity/neighborhood graphs and since then it has been used in manifold estimation, clustering, pattern recognition and image processing and many other applications (Correa and Lindstrom, 2012; Devillers and Duménil, 2021; Katz and Tal, 2017). Some examples of ERG include relative neighborhood graphs (Supowit, 1983; Ichino and Sklansky, 1985), Gabriel graphs (Zhang and King, 2002), β-Skeleton graphs (Bose et al., 2002; Wang et al., 2003). Also, throughout this sub-section we include Euclidean distance and LLPD in our analysis along with normalized g distances for any $g \in \mathbb{G}_0$.

Given data $\chi_n = \{X_1, X_2, .., X_n\}$, $g \in \mathbb{G}_0$ induces an ERG. For any 2 points x and y in M, $\overline{p_g(x, y, \chi_n)} = \overline{xy}$ (i.e. shortest g-distance path between x and y is the direct edge between x and y) occurs if and only if $R_g(x, y) \cap \chi_n$ is empty where $R_g(x, y)$ (*empty region* of x and y induced by g) is defined by

$$R_g(x, y) \coloneqq \{u \in \mathbb{R}^d : g(|x - u|) + g(|y - u|) - g(|x - y|) < 0\}. \tag{2.14}$$

ERG of g, $G_g = (V, E)$, has vertex set $V = \chi_n$ and $E(X_i, X_j) = 1_{R_g(X_i, X_j) \cap \chi_n = \emptyset}$.

If $u \in R_g(x, y) \cap \chi_n$, then the path $x \to u \to v$ is preferred over $x \to y$ under g-distance by the definition of $R_g(x, y)$. So shortest g-distance path between x and y is not the direct distance between x and y. Because of our assumption of strict super-additivity on g, $R_g(x, y)) \neq \emptyset$ for any $x \neq y$.

Some well-known examples of ERGs what are also induced by g or limiting form

of g are Gabriel graphs (which corresponds to $g(t) = t^2$), which is also β-skeleton graph corresponding to $\beta = 1$, and relative neighborhood graphs (which corresponds to LLPD).

Proposition 2.6

Given data $\chi_n = \{X_1, X_2, .., X_n\}$ the shortest g distance path joining X_i and X_j, $i \neq j$ is in ERG induced by g. □

Proof 2.6

If any edge (w.l.g. joining X_k and X_l, $k \neq l$) in the candidate path joining X_i and X_j is not in ERG, then by definition of ERG, $\exists$ a path joining X_k and X_l (which is not the direct edge joining X_k and X_l) which have shorter g distance than the direct edge joining X_k and X_l. Thus the candidate path isn't the shortest g distance minimizing path joining X_i and X_j. □

In-general, $R_g(x, y)$'s shape is invariant under translation and rotation but not under uniform scaling. Note that $R_g(x, y)$'s shape is invariant under uniform scaling (i.e. $R_g(x, y) \propto R_g(ax, ay)$ for any $a > 0$) if and only if $g(t) = t^k$ for some $k > 1$. In this case $R_k(x, y)$ (we shall denote $R_g(x, y)$ by $R_k(x, y)$ for $g(t) = t^k$, $k > 1$) is given by

$$R_k(x, y) = \{u \in \mathbb{R}^d : |x - u|^k + |y - u|^k - |x - y|^k < 0\} \tag{2.15}$$

and we have $R_k(ax + b, ay + b) = b + aR_k(x, y)$. In particular, $\mu(R_k(ax + b, ay + b)) = a^d \mu(R_k(x, y))$ where $\mu(.)$ is the Lebesgue measure in dimension d. Our goal in this sub-section is to obtain $\mu(R_g(x, y))$ for $g \in \mathbb{G}_0$ and to prove several propositions related to $R_g(x, y)$ which will be useful in the proof of theorem 2.1 . We shall focus on these following propositions which will be used to obtain $\mu(R_g(x, y))$ subsequently.

Proposition 2.7

$R_g(x, y)$ is bounded and strictly convex for any $g \in \mathbb{G}_0$. □

Proof 2.7

Bounded part is obvious. W.l.g, suppose $u, v \in \partial R_g(x, y)$ where ∂A denotes the boundary of A. i.e. $g(|x-u|) + g(|y-u|) = g(|x-y|)$ and $g(|x-v|) + g(|y-v|) = g(|x-y|)$. This in particular implies $g(|x-u|) + g(|y-u|) = g(|x-v|) + g(|y-v|)$. Then for $0 < \alpha < 1$

$$g[|x - (\alpha u + (1-\alpha)v)|] + g[|y - (\alpha u + (1-\alpha)v)|]$$

$$= g[\alpha(x-u) + (1-\alpha)(x-v)] + g[\alpha(y-u) + (1-\alpha)(y-v)]$$

$$< \alpha\, g(x-u) + (1-\alpha)\, g(x-v) + \alpha\, g(y-u) + (1-\alpha)\, g(y-v)$$

$$= g(|x-y|)$$

So, $\alpha u + (1-\alpha)v \in R_g(x, y)$ since $g \in \mathbb{G}_0$ is strictly convex. □

Proposition 2.8

$R_{k_1}(x, y) \subset R_{k_2}(x, y)$ if $1 < k_1 < k_2$ where $R_k(x, y)$ is defined in (2.15). Moreover, $\lim\limits_{k\to\infty} R_k(x, y) = R_\infty(x, y) = R_{LLPD}(x, y)$ which is given by

$$R_\infty(x, y) = \{u \in \mathbb{R}^d : max\{|x-u|, |y-u|\} < |x-y|\} \tag{2.16}$$

$$= B_d^0(x, |x-y|) \cap B_d^0(y, |x-y|)$$

where $B_d^0(a, b) \coloneqq \{x \in \mathbb{R}^d : |x-a| < b\}$.

$R_\infty(x, y)$, which is a "lune", is bounded, convex and invariant under translation, rotation and uniform scaling. Also $R_\infty(x, y)$ is the empty region corresponding to LLPD. □

Proof 2.8

Since $R_g(x,y)$ is bounded and strictly convex, it is sufficient to prove if $z \in \partial R_{k_1}(x,y)$ i.e. if $|x-y|^{k_1} = |x-z|^{k_1} + |y-z|^{k_1}$ then $z \in R_{k_2}(x,y)$ i.e.

$$|x-y|^{k_2} = (|x-z|^{k_1} + |y-z|^{k_1})^{\frac{k_2}{k_1}} > |x-z|^{k_2} + |y-z|^{k_2}$$

which follows from Holder's inequality.

The second part follows after observing

$$\lim_{k\to\infty}(|x-u|^k + |y-u|^k)^{\frac{1}{k}} = max\{|x-u|, |y-u|\}.$$

It is trivial to see that $R_\infty(x,y)$ is bounded, convex and invariant under translation, rotation and uniform scaling. Note that LLPD edge between x and y is the direct edge between x and y if and only if $R_\infty(x,y)$ is empty, or in other word there is no point in χ_n which is closer to both x and y. Thus $R_\infty(x,y)$ is the empty region of LLPD. □

Remark 2.4

We shall denote

$$\mu(R_k(0,1e_1)) \coloneqq \theta_{k,d}$$

where e_1 is the unit vector along the first co-ordinate axis for all $k \geq 1$ (although $g(t) = t \notin \mathbb{G}_0$, we are extending our definition of $R_k(x,y)$ for $k = 1$ as well). Then $\{\theta_{k,d}\}_{k\geq 1}$ increases in k from $\mu(R_0(0,1e_1)) = 0$ to $\mu(R_\infty(0,1e_1)) = \theta_{\infty,d} = \frac{2\pi^{\frac{d-1}{2}}}{\Gamma(\frac{d+1}{2})}\int_0^{\pi/3} sin^d x\, dx < \infty$. □

Proposition 2.9

$R_{b_1}(x,y) \subseteq R_g(x,y) \subset R_\infty(x,y)$ for any $g \in \mathbb{G}_0$. □

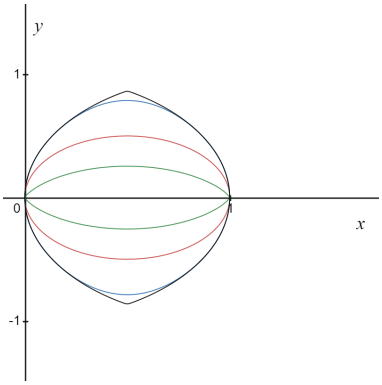

Figure 2.1: $\partial R_k(0, e_1)$ in $d = 2$ for $k = 1.2$ (green), 2 (red), 10 (blue) and 100 (black).

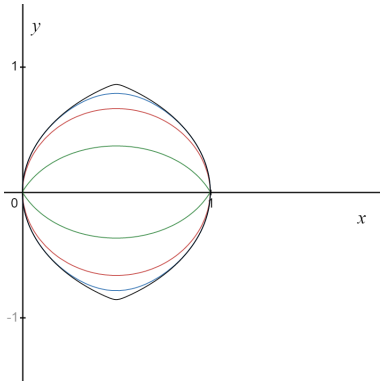

Figure 2.2: $\frac{1}{x}\,\partial R_g(0, xe_1)$ for $g(t) = e^t - 1$ in $d = 2$ for $x = 1$ (green), 4 (red), 10 (blue) and 100 (black).

Figure 2.3: As $k \uparrow$ from 1 to ∞ in the left panel, $\partial R_k(0, e_1)$ transitions from the straight line joining 0 and e_1 (corresponding to Euclidean distance) to a lune (corresponding to LLPD). Similar transition in shape can be observed in the right panel as $x \uparrow$ from 0 to ∞. This hints that g^N distance behaves like Euclidean distance and LLPD in high and low sample size/density regimes respectively.

Proof 2.9

Let $u \in \partial R_{b_1}(x, y)$, i.e. for (ignoring the trivial cases) $0 < \alpha = \frac{|x-u|}{|x-y|} < 1$ and $0 < \beta = \frac{|y-u|}{|x-y|} < 1$ we have $\alpha^{b_1} + \beta^{b_1} - 1 = 0$. Then

$$g_u(x, y) = \sum_{i=1}^{m} a_i x^{b_i} (\alpha^{b_i} + \beta^{b_i} - 1) < 0$$

For the second part, let $u \in \partial R_\infty(x, y)$. i.e. $max\{|x - u|, |y - u|\} = |x - y|$. Letting $\frac{|x-u|}{|x-y|} = \alpha \geq 0$ and $\frac{|y-u|}{|x-y|} = \beta \geq 0$, we see $u \in \partial R_\infty(x, y) \Leftrightarrow max(\alpha, \beta) = 1$.

Let

$$g_u(x, y) = g(|x - u|) + g(|y - u|) - g(|x - y|)$$

Since $\alpha^i + \beta^i - 1 > 0$, when $max(\alpha, \beta) = 1$ for any $i > 1$ we see that $g_u(x, y) > 0$ unless $u = x$ or y (which corresponds to $\alpha = 0$, $\beta = 1$ and $\alpha = 1$and $\beta = 0$ respectively) so u is strictly outside $\partial R_g(x, y)$. □

Lemma 2.1

$\mu(R_g(x,y)) = |x-y|^d \frac{\sum_{i=1}^m a_i \theta_{b_i,d} |x-y|^{b_i}}{\sum_{i=1}^m a_i |x-y|^{b_i}}$

for any $g \in \mathbb{G}_0$, m can be ∞. □

Proof 2.10

Consider $u \in R_g(x,y)$, we let $\alpha = \frac{|x-u|}{|x-y|}$ and $\beta = \frac{|y-u|}{|x-y|}$. So we have $\alpha+\beta \geq 1$, $\alpha, \beta > 0$.

Let

$$g_u(x,y) = g(|x-u|) + g(|y-u|) - g(|x-y|)$$
$$= \sum_{i=1}^m a_i |x-y|^{b_i} (\alpha^{b_i} + \beta^{b_i} - 1) = g(|x-y|)\, \mathbb{E}\left[\alpha^X + \beta^X - 1\right]$$

where X is a probability distribution on $\{b_1, .., b_m\}$ such that $\mathbb{P}[X = b_i] = \frac{a_i |x-y|^{b_i}}{g(|x-y|)}$ for $i = 1, 2, .., m$. This follows since given x and y, $(\alpha^{b_i} + \beta^{b_i} - 1)$ is independent of $|x-y|$ and purely depends on the position of u with respect to x and y.

Now note that

$$u \in R_g(x,y) \Leftrightarrow g_u(x,y) < 0 \Leftrightarrow \mathbb{E}[\alpha^X + \beta^X - 1] < 0$$

A way to compute $\mu(R_g(x,y))$ is by rejection sampling from $R_\infty(x,y)$, since we know $\mu(R_\infty(x,y)) = \theta_{\infty,d}|x-y|^d < \infty$ and $R_g(x,y) \subset R_\infty(x,y)$. Formally,

$$\mu(R_g(x,y)) = \mu(R_\infty(x,y))\, \mathbb{P}[\mathbb{E}[\alpha^X + \beta^X - 1 < 0 | u \in R_\infty(x,y)]]$$
$$= \mu(R_\infty(x,y))\, \mathbb{E}[\mathbb{P}[\alpha^X + \beta^X - 1 < 0 | u \in R_\infty(x,y)]]$$
$$= \mu(R_\infty(x,y)) \sum_{i=1}^m \mathbb{P}[\alpha^{b_i} + \beta^{b_i} - 1 < 0 | u \in R_\infty(x,y)]\, \mathbb{P}[X = b_i]$$
$$= \theta_{\infty,d}|x-y|^d \sum_{i=1}^m \frac{\theta_{b_i,d}}{\theta_{\infty,d}} \frac{a_i |x-y|^{b_i}}{g(|x-y|)}$$
$$= |x-y|^d \frac{\sum_{i=1}^m a_i \theta_{b_i,d} |x-y|^{b_i}}{\sum_{i=1}^m a_i |x-y|^{b_i}}$$

since $\mathbb{P}[\alpha^{b_i} + \beta^{b_i} - 1 < 0 \,|\, u \in R_\infty(x,y)] = \frac{\theta_{b_i,d}}{\theta_{\infty,d}}$. □

Proposition 2.10

$\mu(R_g(x,y)) \coloneqq \mu(R_g(|x-y|))$ depends on x and y only through $|x-y|$. Thus $\mu(R_g(x,y)) \coloneqq \mu(R_g(|x-y|))$ is invariant under translation and rotation. □

Proof 2.11

Follows trivially from lemma 2.1. □

Proposition 2.11

$\mu(R_g(x,y))$ is a strictly increasing function of $|x-y|$ for any $g \in \mathbb{G}_0$. □

Proof 2.12

This follows from the fact that $|x-y|^d$ is strictly increasing as well as $\frac{\sum_{i=1}^m a_i \theta_{b_i,d}|x-y|^{b_i}}{\sum_{i=1}^m a_i |x-y|^{b_i}}$ is non-decreasing since $\{\theta_{i,d}\}_{i\geq 1}$ are strictly increasing and $a_i \in \mathbb{R}^+$. □

Proposition 2.12

As $|x-y| \to 0^+$, under case 1 $R_g(x,y)$ has d-dimensional volume whereas under case 2 $R_g(x,y)$ has $d+b_2-1$ dimensional volume. □

Proof 2.13

Indeed, from it follows from lemma 2.1 that

$$\lim_{|x-y|\to 0^+} \mu(R_g(x,y)) = \begin{cases} \theta_{b_1,d}|x-y|^d & \text{if } b_1 > 1 \\ \frac{a_2\theta_{b_2,d}}{a_1}|x-y|^{d+b_2-1} & \text{if } b_1 = 1 \end{cases}$$

□

Remark 2.5

For any $g \in \mathbb{G}_0$, $\mu(R_g(x,y))$ is related to the probability that x and y are direct edge (we denote it by the notation $x \sim y$) in the ERG of $H_\lambda \cup \{x,y\}$, where H_λ is the

homogeneous Poisson process in $\mathbb{R}^d$ with rate $\lambda > 0$. Indeed, we have

$$P_g(x \sim y) = P_g(|x-y|) = e^{-\lambda\mu(R_g(x,y))} \tag{2.17}$$

Thus under case 1 $P_g(x \sim y)$ is non-trivial (i.e. strictly between 0 and 1) when $|x-y| = O(\lambda^{-\frac{1}{d}})$, which is the order of the k nearest neighbor distance in dimension d. Thus, a g-distance minimizing path between x and y such that $|x-y| = O(1)$ under case 1 will have $O(\lambda^{\frac{1}{d}})$ nodes.

On the other hand under case $P(x \sim y)$ is non-trivial when $|x-y| = O(\lambda^{-\frac{1}{d+b_2-1}})$. Thus, a g-distance minimizing path with minimal number of nodes between x and y such that $|x-y| = O(1)$ under case 2 will contain $O(\lambda^{\frac{1}{d+b_2-1}})$ nodes. Note that in this case one can always choose a g-distance minimizing path with number of nodes between $O(\lambda^{\frac{1}{d+b_2-1}})$ and $O(\lambda^{\frac{1}{d}})$. □

Proposition 2.13

$g \to P_g(|x-y|)$ map is continuous and $P_g(|x-y|)$ or equivalently $\mu(R_g(|x-y|))$ uniquely characterize $g \in \mathbb{G}_0$ up to a constant thus it uniquely characterize g^N distances. □

Proof 2.14

It's trivial to check that $g \to P_g(|x-y|)$ map is continuous.

If $g = ch$ for some $c > 0$, clearly $P_g(t) = P_h(t)$ for all $t > 0$.

Suppose, $P_g(t) = P_h(t)$ for all $t > 0$, or equivalently

$$\frac{\sum_{i=1}^{m} a_i \theta_{b_i,d} t^{b_i}}{\sum_{i=1}^{m} a_i t^{b_i}} = \frac{\sum_{i=1}^{m'} a_i' \theta_{b_i',d} t^{b_i'}}{\sum_{i=1}^{m'} a_i' t^{b_i'}}, t > 0$$

It is easy to prove that (after cross multiplying) all the powers have to be equal, i.e. $m = m'$ (both can be ∞) and $b_i = b_i'$ since $\{\theta_{i,d}\}_{i\geq 1}$'s are strictly increasing by an induction argument. Here we are using the standard real analysis fact that if

$\sum_{i=1}^{m} c_i t^{b_i} = 0$ for all $t > 0$ or in some interval (a, b) with $a < b$, $c_i = 0$ for all i.

For example, coefficient of the lowest $(b_1 + b_1^{'}\text{th})$ power (after cross-multiplying) in LHS is $a_1 a_1^{'} \theta_{b_1,d}$ while it is $a_1 a_1^{'} \theta_{b_1^{'},d}$, implying $b_1 = b_1^{'}$. Assume $b_2 > b_2^{'}$, then coefficient of the lowest $(b_1 + b_2^{'}\text{th})$ power in LHS is $a_1 a_2^{'} \theta_{b_1,d}$ while it is $a_1 a_2^{'} \theta_{b_2^{'},d}$ which is a contradiction since $\theta_{b_2^{'},d} > \theta_{b_1^{'},d} = \theta_{b_1,d}$. Thus $b_2 = b_2^{'}$ and so on..

Hence, we have

$$\sum_{i=1}^{m} \left(\frac{a_i t^{b_i}}{\sum_{i=1}^{m} a_i t^{b_i}} - \frac{a_i^{'} t^{b_i}}{\sum_{i=1}^{m} a_i^{'} t^{b_i}}\right) \theta_{b_i,d} = 0, t > 0$$

implying

$$\frac{a_i t^{b_i}}{\sum_{i=1}^{m} a_i t^{b_i}} = \frac{a_i^{'} t^{b_i}}{\sum_{i=1}^{m} a_i^{'} t^{b_i}} \text{ for all } i$$

i.e.

$$\frac{g}{h} = \frac{\sum_{i=1}^{m} a_i t^{b_i}}{\sum_{i=1}^{m} a_i^{'} t^{b_i}} = \frac{a_i}{a_i^{'}} = c \text{ for all } i$$

□

Proposition 2.14

For any $\lambda > 0$,

$$T_\lambda(g^N, h^N) = \max_{t \in (a,b)} |P_g(t) - P_h(t)|$$

is a metric for any $0 \leq a < b \leq \infty$. W.l.g. we take $a = 0$ and $b = \infty$. □

Proof 2.15

It's trivial to see that T_λ is continuous. By proposition 2.13, $T_\lambda(g^N, h^N) = 0$ if and only if g^N distance is identical to h^N distance. Also clearly T_λ is symmetric and satisfies triangle inequality. □

Lemma 2.2

If $\lambda \to \infty$ (which correspond to high sample size/high density regime), under T_λ normalized g distance converge to it's lowest power distance and if $\lambda \to 0$ (which

correspond to low sample size/low density regime) under T_λ normalized g distance converge to it's highest power distance if it exists and to LLPD if it doesn't. □

Proof 2.16

W.l.g. suppose $\lambda^d|x-y| = c = O(1)$ because otherwise the probabilities on both sides become trivial (and equal).

Then if $\lambda \to \infty$ (and $|x-y| \to 0$), when $b_1 > 1$,

$$P_g(|x-y|) = e^{-c\,\theta_{b_1,d}} = P_{h(t)=t^{b_1}}(|x-y|)$$

Under same condition, when $b_1 = 1$

$$P_g(|x-y|) = 1 = P_{h(t)=t}(|x-y|)$$

On the other hand, when $\lambda \to 0$ (and $|x-y| \to \infty$) and m finite

$$P_g(|x-y|) = e^{-c\,\theta_{b_m,d}} = P_{h(t)=t^{b_m}}(|x-y|)$$

and when m is infinite

$$P_g(|x-y|) = e^{-c\,\theta_{\infty,d}} = P_{LLPD}(|x-y|)$$ □

Proposition 2.15

For any $g \in \mathbb{G}_0$, $R_g(0, xe_1) \subseteq R_g(0, ye_1)$ for $x < y$. □

Proof 2.17

Suppose $u \in \partial R_g(0, xe_1)$. Then

$$g_u(0, xe_1) = \sum_{i=1}^{m} a_i x^{b_i}(\alpha^{b_i} + \beta^{b_i} - 1) = 0$$

where $\alpha = \frac{|u|}{x}$ and $\beta = \frac{|xe_1 - u|}{x}$, w.l.g. $\alpha, \beta \in (0,1)$, $\alpha + \beta \geq 1$.

Let $y = (1+\triangle)xe_1$ for $\triangle << 1$. Then for a fixed $0 < \beta < 1$ we have,

$$g_u(0, ye_1) = \sum_{i=1}^{m} a_i((1+\triangle)x)^{b_i}((\frac{\alpha}{1+\triangle})^{b_i} + (\frac{|(1+\triangle)xe_1 - u|}{(1+\triangle)x})^{b_i} - (1+\triangle)^{b_i})$$

$$= \sum_{i=1}^{m} a_i x^{b_i}(\alpha^{b_i} + (\frac{|(1+\triangle)xe_1 - u|}{x})^{b_i} - (1+\triangle)^{b_i})$$

$$\leq \sum_{i=1}^{m} a_i x^{b_i}(\alpha^{b_i} + (\beta+\triangle)^{b_i} - (1+\triangle)^{b_i})$$

$$\approx \sum_{i=1}^{m} a_i x^{b_i}(\alpha^{b_i} + \beta^{b_i} + \beta^{b_i-1}b_i\triangle - 1 - b_i\triangle)$$

$$= \sum_{i=1}^{m} a_i x^{b_i}(b_i\triangle(\beta^{b_i-1} - 1)) < 0$$

□

2.3.2 First passage percolation results

The goal of this sub-section is to state and prove some useful propositions which will be used in the following sub-sections. While convergence rate in case 1 reduces to Hwang et al. (2016)'s result, convergence rate in case 2 becomes non-trivial because empty region corresponding to g distance in this case has a limiting $d + b_2 - 1$ dimensional volume.

Proposition 2.16

Let p be a shortest path from x to y in a homogeneous Poisson point process H_λ. Suppose p has g-distance path length at most $c_0\lambda^{\frac{1-b_1}{d}}$ and has at least $c_1\lambda^{\frac{1}{d}}$ nodes for some positive c_0, c_1. Then $\exists$ a constant $\rho_0 > 0$ (dependent on d and g) such that if $c_1 > \rho_0 c_0$, the probability that such path p exists is exponentially small in $c_1\lambda^{\frac{1}{d}}$ as $\lambda \to \infty$. □

Proof 2.18

The proof is very similar to lemma 11 in Hwang et al. (2016) after observing that as $\lambda \to \infty$, p has g-distance path length at most $c_0\lambda^{\frac{1-b_1}{d}}$ for some $c_0 > 0$ if and only if it has b_1 power weighted path length at most $c_0^{'}\lambda^{\frac{1-b_1}{d}}$ for some $c_0^{'} > 0$. To be precise, $\exists \lambda^* > 0$ such that $0 < a_1 < \frac{g(t)}{t^{b_1}} < Ka_1$ for $t \in (0, \lambda^*)$ and by Borel Cantelli lemma, for $\lambda \geq \lambda_0 > 0$, w.p. 1 all edge lengths in a shortest paths are $\leq \lambda^*$ and Probability that in p there is an edge length $> \lambda^*$ is given by $1 - (1 - e^{-O(\lambda)\lambda^{*d}})^{O(\lambda^{\frac{1}{d}})}$ which has a higher decay rate than $O(e^{-\lambda^{\frac{1}{d}}})$. Also note that λ^* can be taken arbitrarily small which will be used in proposition 2.17. □

Proposition 2.17

Furthermore, when $b_1 = 1$, suppose p is the shortest path joining x and y having minimal number of nodes in a homogeneous Poisson point process H_λ. Suppose p has g-distance path length c_0 and has at least $c_1\lambda^{\frac{1}{d+b_2-1}}$ nodes for some positive c_0, c_1. Then $\exists$ a constant $\rho_0 > 0$ (dependent on d and g) such that if $c_1 > \rho_0 c_0$, the probability that such path p exists is exponentially small in $c_1\lambda^{\frac{1}{d+b_2-1}}$ as $\lambda \to \infty$. □

Proof 2.19

Let $H_{d,\lambda}$ denote homogeneous Poisson processes with rate $\lambda > 0$ in $\mathbb{R}^d$. For $b_1 = 1$, in ERG of g for $x_1, x_2 \in H_{d,\lambda}$, as $|x_2 - x_1| \to 0$ we have

$$P_g[x_1 \sim x_2] = e^{-\lambda|x_2-x_1|^{d+b_2-1}\frac{a_2}{a_1}\theta_{b_2,d}}$$

Consider $g^{'}(t) = t^{b_1}$ for some $b_1 > 1$ in $H_{d+b_2-1,\lambda^{'}}$ with $\lambda^{'} = c\lambda$. In this case as $|x_2 - x_1| \to 0$,

$$P_{g^{'}}[x_1 \sim x_2] = e^{-\lambda^{'}|x_2-x_1|^{d+b_2-1}\theta_{b_1,d+b_2-1}}$$

For any $\lambda > 0$ and $g, g^{'} \in \mathbb{G}_0$ with $b_1 = 1$, we can choose $c > 0$ such that $P_g[x_1 \sim x_2] = P_{g^{'}}[x_1 \sim x_2]$ as $|x_1 - x_2| \to 0$. Since $P_g(|x - y|)$ uniquely characterize g distance (upto a constant), as $\lambda \to \infty$ we can see that $g^{'N}$ in $H_{d+b_2-1,\lambda^{'}}$ is equivalent

to g^N in $H_{d,\lambda}$ with $b_1 = 1$ under T_λ. Conclusion follows after applying proposition 2.16 to $g^{'}$ in $H_{d+b_2-1,\lambda'}$ and observing as $\lambda \to \infty$, any path with g' distance $\leq c_0^{'}\lambda^{\frac{1-b_1}{d}}$ will have Euclidean distance $\leq c_0$. □

Proposition 2.18

Consider the collection of points $z \in \mathbb{R}^d$ such that $H_\lambda \cap R_g(0,z)$ is empty. Let $\xi \coloneqq sup\,|z|$. We further define the tube of radius b containing x and y by

$$T_b(x,y) \coloneqq \cup_{0\leq u\leq 1} B(ux + (1-u)y, b)$$

Similarly, consider the collection of points z such that $H_\lambda \cap (\cup_{u>0} T_b(-ue_1, ue_1)) \cap R_g(0,z)$ is empty and define $\xi_b \coloneqq sup\,|z|$ Then we have

$$\lim_{b\to\infty} \mathbb{E}[g(\xi_b)] = \mathbb{E}[g(\xi)]$$

and under case 1 when $\lambda \to \infty$ and $b \to 0$ such that $\lambda^{\frac{1}{d}} b \to \infty$,

$$\lim_{\lambda^{\frac{1}{d}} b\to\infty} \frac{\mathbb{E}[g(\xi_b)]}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})b} = 0$$

and under case 2 when $\lambda \to \infty$ and $b \to 0$ such that $\lambda^{1/(d+b_2-1)} b \to \infty$

$$\lim_{\lambda^{\frac{1}{d+b_2-1}} b\to\infty} \frac{\mathbb{E}[g(\xi_b)]}{\lambda^{\frac{1}{d+b_2-1}} g(\lambda^{-\frac{1}{d+b_2-1}})b} = 0$$

□

Proof 2.20

We have for $p \geq 1$,

$$\mathbb{E}[\xi_b^p] \leq \begin{cases} \frac{n\Gamma(1+\frac{p}{d+b_2-1})}{(\lambda A)^{\frac{p}{d+b_2-1}}} + \frac{m2^p\Gamma(1+\frac{p}{d})}{\lambda^p(\delta b)^{p(d-1)}} & \text{when } b_1 = 1 \\ \frac{n\Gamma(1+\frac{p}{d})}{(\lambda A)^{\frac{p}{d}}} + \frac{m2^p\Gamma(1+\frac{p}{d})}{\lambda^p(\delta b)^{p(d-1)}} & \text{when } b_1 > 1 \end{cases} \tag{2.18}$$

for some $m, n, \delta, A > 0$ independent of b and λ. Proof of this follows similarly from the arguments in lemma 2.5 in Hwang (2012). □

Proposition 2.19

Fix n and let $\chi_n = \{X_1, .., X_n\}$ are i.i.d. sampled from f and $0 < \alpha < 1$.

Assume $b_1 > 1$. Define the event $E_F(i,j)$ for each pair $1 \leq i, j \leq n$ such that $E_F(i,j)^c$ occur if and only if

1. both X_i and X_j are in $B(z, R_2)$
2. $|X_i - X_j| > (nf_0)^{\frac{\alpha-1}{d}}$
3. the g- shortest path from X_i to X_j contains no sample point X_k other than X_i and X_j.

Let $E_F = \cap_{i,j} E_F(i,j)$. i.e. E_F^c denotes the event that within a finite radius R_2, no direct edge in the g-distance path has edge length $\geq (nf_0)^{\frac{\alpha-1}{d}}$. Then

$$\limsup_{n\to\infty} \frac{1}{(nf_0)^\alpha} log(1 - \mathbb{P}(E_F)) \leq -\theta_{11} \tag{2.19}$$

for some constant $\theta_{11} > 0$ which depends only on d and g.

Assume $b_1 = 1$. Define the event $E_F(i,j)$ for each pair $1 \leq i, j \leq n$ such that $E_F(i,j)^c$ occur if and only if

1. both X_i and X_j are in $B(z, R_2)$
2. $|X_i - X_j| > (nf_0)^{\frac{\alpha-1}{d+b_2-1}}$
3. the g- shortest path from x_i to $x_j$contains no sample point X_k other than X_i and X_j.

Let $E_F = \cap_{i,j} E_F(i,j)$. Then

$$\limsup_{n\to\infty} \frac{1}{(nf_0)^\alpha} log(1 - \mathbb{P}(E_F)) \leq -\theta_{21}$$

for some constant $\theta_{21} > 0$ which depends only on d and g. □

Proof 2.21

For $b_1 = 1$ case, recall from lemma 2.1 that $\mu(R_g(x_i, x_j) \cap B(z, R_2))$ can be taken at-least $\theta_{21}|x_i - x_j|^{d+b_2-1}$ for some $\theta_{21} > 0$. Suppose that $E_F(1,2)^c$ occur. Then

the shortest path from X_1 to X_2 contains no sample point other than X_1 and X_2, and $R_g(X_1, X_2) \cap B(z, R_2)$ must not contain any of $X_3, X_4, .., X_n$. Since $|X_1 - X_2| > (nf_0)^{\frac{\alpha-1}{d+b_2-1}}$, probability that E_F does not occur is $\leq (1 - f_0\theta_{21}(nf_0)^{(\alpha-1)})^{n-2}$. Since there are $\frac{n(n-1)}{2}$ pair of points,

$$1 - \mathbb{P}(E_F) \leq n^2(1 - \theta_{21}\frac{(nf_0)^\alpha}{n})^{n-2} \leq exp\,(-\theta_{21}(nf_0)^\alpha)$$

If $b_1 > 1$, similarly we can take $\mu(R_g(x_i, x_j) \cap B(z, R_2)) \geq \theta_{11}|x_i - x_j|^d$ for some $\theta_{11} > 0$ and a similar calculation gives us the result. □

2.3.3 Mean convergence of g-distance for Poisson point processes

In this sub-section we prove the mean convergence of g distance for a homogeneous Poisson process H_λ on $\mathbb{R}^d$.

Lemma 2.3

For every $x, y \in \mathbb{R}^d$, $\lambda > 0$, and $g \in \mathbb{G}_0$, with probability 1 the random variable $L_g(x, y, H_\lambda)$ is well-defined. □

Proof 2.22

Note that this is not trivial since H_λ contains infinitely many points and thus g distance minimizing path from x to y might not exist. In particular, in Howard and Newman (2001) it is mentioned that it is not yet known if $LLPD(x, y, H_\lambda)$ is well-defined for all $d > 2$.

The point process H_λ is locally finite almost surely. Therefore $L_g(x, y, H_\lambda \cap T_b(x, y))$ is well-defined for any $0 < b < \infty$. Let $r_1 = L_g(x, y, H_\lambda \cap T_1(x, y))$.

Suppose a path from x to y has g-path length $r < r_1$ and it is not contained in $B(x, b) := \{z \in \mathbb{R}^d : |x - z| \leq b\}, b > 0$. Then from x, $\exists$ a path p that ends at some

point outside $B(x, b)$ and has g-path length less than r_1. Assume that p has edges with Euclidean edge lengths $c_1, c_2, .., c_n$. Thus, by definition

$$r = \sum_{i=1}^{m} \sum_{k=1}^{n} a_i c_k^{b_i}$$

First we consider the case when m is finite, i.e. let $g(t) = \sum_{i=1}^{m} a_i t^{b_i}$. The proof in this case is similar to the one presented in Howard and Newman (2001); Hwang et al. (2016) for the case $g(t) = t^p$, $p > 1$. From triangle inequality and Holder's inequality we have,

$$g(b) \leq g(\sum_{k=1}^{n} c_k) = \sum_{i=1}^{m} a_i (\sum_{k=1}^{n} c_k)^{b_i}$$

$$\leq \sum_{i=1}^{m} a_i \left[(\sum_{k=1}^{n} c_k^{b_i}) n^{b_i - 1} \right] < n^{b_m - 1} \sum_{i=1}^{m} \sum_{k=1}^{n} a_i c_k^{b_i} < n^{b_m - 1} r_1$$

Hence $n > (r_1^{-1} g(b))^{\frac{1}{b_m - 1}}$. By proposition 2.16, for b large enough, probability that such a path p exists is exponentially small in $b^{\frac{b_m}{b_m - 1}}$, since b_m is the largest power in g. By the Borel-Cantelli lemma, with probability 1 there exists $b_* > 0$ such that no candidate path from x to y has g-path length $< r_1 = L_g(x, y, H_\lambda \cap T_1(x, y))$ reaches some point outside $B(x, b_*)$. It follows that $L_g(x, y, H_\lambda)$ is determined by all the paths in $B(x, b_*)$, and number of paths in $B(x, b_*)$ is finite with probability 1.

When m is ∞ i.e. $g(t) = \sum_{i=1}^{\infty} a_i t^{b_i}$, $\exists M$ such that $|g(t) - \sum_{i=1}^{M} a_i t^{b_i}| < \epsilon$ for any $0 \leq t \leq max\{\lambda^* n^2, b\} < \infty$. Once again consider a candidate path p that ends at some point outside $B(x, b)$ and has g-path length $r < r_1$. Assume that p has edges with Euclidean edge lengths $c_1, c_2, .., c_n$. We denote $\sum_{i=1}^{M} a_i t^{b_i}$ by $g_M(t)$.

Moreover, for any $\lambda > 0$, by proposition 2.6 any candidate g distance minimizing path need to be within the ERG of g. Since

$$P(x_1 \sim x_2) = \begin{cases} O(e^{-\lambda |x_2 - x_1|^d}) & \text{when } b_1 > 1 \\ O(e^{-\lambda |x_2 - x_1|^{d + b_2 - 1}}) & \text{when } b_1 = 1, \end{cases}$$

in both cases, by Borel-Cantelli lemma w.p. 1, $\exists \lambda^* > 0$ such that $c_i \leq \lambda^*$ for all $i = 1(1)n$. Thus

$$g_M(b) + \epsilon < g(b) \leq g(\sum_{k=1}^{n} c_k) = \sum_{i=1}^{\infty} a_i \, (\sum_{k=1}^{n} c_k)^{b_i}$$

$$\leq \frac{1}{n} \sum_{i=1}^{\infty} a_i (\sum_{k=1}^{n} n c_k)^{b_i} = \frac{1}{n} [\sum_{i=1}^{M} a_i \, (\sum_{k=1}^{n} n c_k)^{b_i} + \epsilon] < n^{b_M - 1} \sum_{i=1}^{M} \sum_{k=1}^{n} a_i \, c_k^{b_i} + \frac{\epsilon}{n}$$

Thus

$$g_M(b) < 2 r_1 n^{b_M - 1}$$

And hence, $n > (2 r_1^{-1} g_M(b))^{\frac{1}{b_M - 1}}$ and again by proposition 2.16 for b large enough probability that such a path p exists is exponentially small in $b^{\frac{b_M}{b_M - 1}}$ and conclusion follows after applying Borel-Cantelli lemma. □

Lemma 2.4

For any $g \in G_0$

$$0 < C(d, g) = \lim_{t \to \infty} \frac{1}{t} \, \mathbb{E} \, L_g[0, te_1, H_1] < \infty \tag{2.20}$$

exists. Furthermore if positive $\{c_t\}_{t \in \mathbb{R}_0^+}$ satisfies $c_t \to \infty$ as $t \to \infty$, then

$$\lim_{t \to \infty} \frac{1}{t} \, \mathbb{E} \, L_g[0, te_1, H_1 \cap T_{c_t}(0, te_1)] = C(d, g) \tag{2.21}$$

□

Proof 2.23

Follows similarly from the arguments made in lemma 12 of Hwang et al. (2016) after observing that

$$f(t) = \mathbb{E} \, L_g[0, te_1, H_1 \cap T_b(0, te_1)] + E[g(2\xi_b)]$$

is a sub-additive for any $b > 0$ in t and applying Fekete's lemma. □

Now we shall utilize that H_λ in $\mathbb{R}^d$ has the following scaling property for $a > 0$,

$$aH_\lambda \overset{d}{=} H_{a^{-d}\lambda}$$

This is a consequence of $|H_\lambda \cap B| \sim Poi(\lambda\mu(B))$ and $\mu(aB) = a^d\mu(B)$, where $aB = \{ax : x \in B\}$ and $\mu(.)$ is the Lebesgue measure in dimension d.

Now, we note that if $X_0(\omega) = 0 \to X_1(\omega) .. \to X_{n+1}(\omega) = te_1$ is the realized shortest g-path in $L_g[0, te_1, H_1 \cap T_{c_t}(0, te_1)]$, then $\frac{X_0(\omega)}{\lambda^{\frac{1}{d}}} = 0 \to \frac{X_1(\omega)}{\lambda^{\frac{1}{d}}} .. \to \frac{X_{n+1}(\omega)}{\lambda^{\frac{1}{d}}} = \frac{t}{\lambda^{\frac{1}{d}}} e_1$ is the corresponding realized shortest g-path in $L_g(0, \lambda^{-\frac{1}{d}}te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}}c_t}(0, e_1))$.

The realized shortest g-path length can be written as

$$\sum_{i=1}^{m}\sum_{k=0}^{n} a_i(\lambda^{-\frac{b_i}{d}}|X_{k+1}(\omega) - X_k(\omega)|^{b_i}) = \sum_{i=1}^{m} a_i\lambda^{-\frac{b_i}{d}}(\sum_{k=0}^{n}|X_{k+1}(\omega) - X_k(\omega)|^{b_j})$$

Therefore

$$\mathbb{E}\left[L_g(0, \lambda^{-\frac{1}{d}}te_1, H_\lambda \cap T_{\lambda^{\frac{1}{d}}c_t}(0, e_1))\right] = \sum_{i=1}^{m} a_i\lambda^{-\frac{b_i}{d}} B_{b_i}(t)$$

where $B_{b_i}(t)$ is the expected b_i-th power-weighted distance evaluated on the shortest g-path between 0 and te_1 under H_1 for $i = 1, 2, .., m$.

Formally we define

$$B_j(t) \coloneqq \mathbb{E}\left[L_{j|g}(0, te_1, H_1 \cap T_{b_t}(0, te_1))\right]$$

for $j \geq 1$. The subscript $j|g$ denote we are evaluating the j-th power-weighted distance on g-distance shortest path. Therefore, for $h(\lambda) = \lambda^{\frac{1}{d}}g(\lambda^{-\frac{1}{d}}) = \sum_{i=1}^{m} a_i\lambda^{\frac{1-b_i}{d}}$,we have from lemma 2.4 (substituting $x = \lambda^{-\frac{1}{d}}t$ and $c_x = \lambda^{-\frac{1}{d}}c_t$),

$$\lim_{\lambda^{\frac{1}{d}}x\to\infty,\, \lambda^{\frac{1}{d}}c_x\to\infty} \frac{\mathbb{E}\, L_g[0, xe_1, H_\lambda \cap T_{c_x}(0, e_1)]}{h(\lambda)x} = \lim_{t\to\infty} \frac{\sum_{i=1}^{m} a_i\lambda^{-\frac{b_i}{d}} B_{b_i}(t)}{h(\lambda)\lambda^{-\frac{1}{d}}t}$$

$$= \sum_{i=1}^{m}\Big(\lim_{t\to\infty}\frac{B_{b_i}(t)}{t}\Big)\frac{a_i\lambda^{\frac{1-b_i}{d}}}{\sum_{j=1}^{m} a_j\lambda^{\frac{1-b_i}{d}}}$$

Note that from lemma 2.4, we already know $\lim_{t\to\infty}\frac{\sum_{i=1}^{m} a_i B_{b_i}(t)}{t} = C(d, g)$.

Proposition 2.20

$\lim_{t\to\infty} \frac{B_{b_i}(t)}{t} = C(d, b_i|g)$ where $C(d, b_i|g) \geq C(d, b_i)$ is a constant dependent only on d and g for $i = 1, 2, .., m$. Recall that $C(d, b_i) \coloneqq \lim_{t\to\infty} \frac{\mathbb{E}\,[L_{g(t)=t^{b_i}}(0, te_1)]}{t}$. □

Proof 2.24

Notice that $L_{g|j}(0, te_1, H_1 \cap T_b(0, te_1))$ is also nearly-subadditive like $L_g(0, te_1, H_1 \cap T_b(0, te_1))$. We have in this case

$$f_j(t) = \mathbb{E}\, L_{g|j}[0, te_1, H_1 \cap T_b(0, te_1)] + \mathbb{E}\,[g(2\xi_b)]$$

is sub-additive for any $b > 0$ and conclusion follows by applying Fekete's lemma. □

Thus we have

$$C(d, g) = \sum_{i=1}^{m} a_i C(d, b_i|g) \tag{2.22}$$

It follows that

Proposition 2.21

For a sequence of Poisson point process $\{H_{\lambda_n}\}$ and 2 points $\{x_n\}$ and $\{y_n\}$,

When $b_1 > 1$, if $\lambda_n^{\frac{1}{d}}\,|x_n - y_n| \to \infty$ and $\lambda_n^{\frac{1}{d}}\, c(|x_n - y_n|) \to \infty$ hold then

$$\lim_{n\to\infty} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{\lambda_n^{\frac{1}{d}} g(\lambda_n^{-\frac{1}{d}})|x_n - y_n|} = \lim_{n\to\infty} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n} \cap T_{c(|x_n - y_n|)}(x_n, y_n)]}{\lambda_n^{\frac{1}{d}} g(\lambda_n^{-\frac{1}{d}})|x_n - y_n|}$$

$$= \lim_{n\to\infty} \sum_{i=1}^{m} \alpha_i(\lambda_n) C(d, b_i|g) \tag{2.23}$$

where $a_i(\lambda) = \frac{a_i \lambda^{\frac{1-b_i}{d}}}{\sum_{j=1}^{m} a_j \lambda^{\frac{1-b_j}{d}}}$ for $i = 1, 2, .., m$. Clearly $\sum_{i=1}^{m} \alpha_i(\lambda) = 1$. □

Proof 2.25

Proof follows trivially from proposition 2.20 and the discussion before that. □

Corollary 2.4

When $\lambda_n \to \infty$, (2.23) reduces to $\lim_{\lambda_n \to \infty} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{a_1 \lambda_n^{\frac{1-b_1}{d}} |x_n - y_n|} \to C(d, b_1|g)$, on the other hand when $\lambda_n \to 0$, (2.23) reduces to $\lim_{\lambda_n \to 0} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{a_m \lambda_n^{\frac{1-b_m}{d}} |x_n - y_n|} \to C(d, b_m|g)$. □

Proof 2.26

Immediately follows from (2.23). □

Proposition 2.22

For a sequence of Poisson point process $\{H_{\lambda_n}\}$ and 2 points $\{x_n\}$ and $\{y_n\}$, when $b_1 = 1$, if $\lambda_n^{\frac{1}{d+b_2-1}} |x_n - y_n| \to \infty$ and $\lambda_n^{\frac{1}{d+b_2-1}} c(|x_n - y_n|) \to \infty$ hold,

then

$$\lim_{n \to \infty} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{\lambda_n^{\frac{1}{d}}\, g(\lambda_n^{-\frac{1}{d}})\, |x_n - y_n|} = \lim_{n \to \infty} \sum_{i=1}^{m} \alpha_i(\lambda_n) C(d, b_i|g)$$

. □

Proof 2.27

Immediately follows from (2.23). □

Corollary 2.5

In particular, when $\lambda_n \to \infty$,

$$\lim_{n \to \infty} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{a_1 |x_n - y_n|} = \lim_{n \to \infty} \frac{\mathbb{E} L_g[x_n, y_n, H_{\lambda_n} \cap T_{c(|x_n - y_n|)}(x_n, y_n)]}{a_1 |x_n - y_n|} = C(d, 1|g)$$

and when $\lambda_n \to 0$, $\lim_{\lambda_n \to 0} \frac{\mathbb{E}\, L_g[x_n, y_n, H_{\lambda_n}]}{a_m\, \lambda_n^{\frac{1-b_m}{d}} |x_n - y_n|} \to C(d, b_m|g)$. □

Proof 2.28

Immediately follows from (2.23). □

Remark 2.6

Thus the key takeaway of (2.23) is that g^N distance behaves like it's smallest power-weighted distance when sample size is high and density of 2 points are high and g^N behaves like it's highest power-weighted distance when sample size is low and density of 2 points are low. This adaptability of g^N according to different density and sample size regimes makes it a better candidate in cluster applications than power-weighted distances specially in small/medium sample size scenario. □

For now on we shall be focusing on $\lambda = \lambda_n \to \infty$ and $|x_n - y_n|$ is bounded (i.e. large sample/high density) scenario unless otherwise mentioned and we shall take $h(\lambda) = a_1\lambda^{\frac{1-b_1}{d}}$.

Let $z \in \mathbb{R}^d$, $0 < R_1 < R_2$. Let

$$v(u) = \begin{cases} 1 & \text{if } u \in B(z, R_2) \\ 0 & \text{if } u \notin B(z, R_2) \end{cases}$$

Let $H_{\lambda v(u)}$ denote the non-homogeneous Poisson point process with intensity λv. We shall be using $L_\lambda(x_\lambda, y_\lambda)$, $L_{\lambda v}(x_\lambda, y_\lambda)$ and $L_\lambda(x_\lambda, y_\lambda, t_\lambda)$ for $L_g(x_\lambda, y_\lambda, H_\lambda)$, $L_g(x_\lambda, y_\lambda, H_{\lambda v(u)})$ and $L_g(x_\lambda, y_\lambda, H_\lambda \cap T_{t_\lambda}(x_\lambda, y_\lambda))$ respectively for notation simplicity.

Corollary 2.6

Assume that $x_\lambda, y_\lambda \in B(z, R_1)$. If $b_1 > 1$ and $\lim\limits_{\lambda\to\infty} \lambda^{\frac{1}{d}}|x_\lambda - y_\lambda| = \infty$ then $\lim\limits_{\lambda\to\infty} \frac{\mathbb{E}[L_{\lambda v}(x_\lambda, y_\lambda)]}{a_1\lambda^{\frac{1-b_1}{d}}|x_\lambda - y_\lambda|} = C(d, b_1|g)$. Furthermore, if positive b_λ satisfies $b_\lambda \to 0$ and $\lambda^{\frac{1}{d}}b_\lambda \to \infty$ as $\lambda \to \infty$, then also $\lim\limits_{\lambda\to\infty} \frac{\mathbb{E}[L_{\lambda v}(x_\lambda, y_\lambda, b_\lambda)]}{a_1\lambda^{\frac{1-b_1}{d}}|x_\lambda - y_\lambda|} = C(d, b_1|g)$.

Similarly when $b_1 = 1$ and $\lim\limits_{\lambda\to\infty} \lambda^{\frac{1}{d+b_2-1}}|x_\lambda - y_\lambda| = \infty$ then $\lim\limits_{\lambda\to\infty} \frac{\mathbb{E}[L_{\lambda v}(x_\lambda, y_\lambda)]}{a_1|x_\lambda - y_\lambda|} = C(d, b_1|g)$. Furthermore, if b_λ satisfies $b_\lambda \to 0$ and $\lambda^{\frac{1}{d+b_2-1}}b_\lambda \to \infty$ as $\lambda \to \infty$, then $\lim\limits_{\lambda\to\infty} \frac{\mathbb{E}[L_{\lambda v}(x_\lambda, y_\lambda, b_\lambda)]}{a_1|x_\lambda - y_\lambda|} = C(d, b_1|g)$.

Proof 2.29

Similar to corollary 13 in Hwang et al. (2016) and is omitted. □

2.3.4 Concentration of g distance around mean

In previous sub-section we showed that $\lim_{\lambda\to\infty} \frac{\mathbb{E}[L_{\lambda v}(x_\lambda,y_\lambda)]}{a_1\lambda^{\frac{1-b_1}{d}}|x_\lambda-y_\lambda|} = C(d,b_1|g)$. The general goal of this and the next sub-section will be to show that as $\lambda\to\infty$, the random variable $\frac{L_{\lambda v}(x_\lambda,y_\lambda)}{a_1\lambda^{\frac{1-b_1}{d}}|x_\lambda-y_\lambda|}$ is concentrated around $C(d,b_1|g)$ as $\lambda\to\infty$.

Lemma 2.5

Let $\epsilon>0$ be fixed, and let

$$\alpha^* = \begin{cases} \frac{1}{d+2b_1-1} & \text{for case 1} \\ \frac{1}{d+b_2} & \text{for case 2} \end{cases}$$

If $\lambda\to\infty$ and $\lim_{\lambda\to\infty} inf\ \lambda^{\alpha^*}|x_\lambda-y_\lambda| = \infty$ then

$$\lim_{\lambda\to\infty} sup \frac{1}{\lambda^{\alpha^*}|x_\lambda-y_\lambda|} log\,\mathbb{P}\left[\frac{L_{\lambda v}(x_\lambda,y_\lambda)}{a_1\lambda^{\frac{1-b_1}{d}}|x_\lambda-y_\lambda|} \geq C(d,b_1|g)+\epsilon\right] \leq -min\{C_1,C_2\epsilon^2\} \tag{2.24}$$

for some $C_1,C_2>0$ dependent on d and g. □

The inequality states that the probability $\mathbb{P}\{.\}$ quantity decays no slower than $exp\,(-c\lambda^{\alpha^*}|x_\lambda-y_\lambda|)$, for some positive constant c, as λ increases to ∞. Also note that the condition is stronger than the condition in corollary 2.6 for case 1 and similar is true for case 2 as well.

Thus corollary 2.6 may be applied in the proof of lemma 2.5. The key idea in this as well as in the next sub-section is to replace the necessary condition $\lambda^{\frac{1}{d}}|x_\lambda-y_\lambda|=\infty$ for case 1 and $\lambda^{\frac{1}{d+b_2-1}}|x_\lambda-y_\lambda|=\infty$ for case 2 in corollary 2.6 , with something stronger

and in return obtain the complete convergence of the random variable $\frac{L_{\lambda v}(x_\lambda, y_\lambda)}{a_1 \lambda^{\frac{1-b_1}{d}} |x_\lambda - y_\lambda|}$ to it's mean.

Proof 2.30

Fix λ. From Corollary 2.6 , it is sufficient to prove that the bound holds for $L_\lambda(x_\lambda, y_\lambda, b_\lambda)$ instead of $L_{\lambda v}(x_\lambda, y_\lambda)$ for positive b_λ satisfying $b_\lambda \to 0$ as $\lambda \to \infty$.

Now, $L_\lambda(0, 2be_1, b)$ may be bounded above by $L_\lambda(0, be_1, b) + L_\lambda(be_1, 2be_1, b) + \sum_{i=1}^{m} a_i(2^{b_i-1}-1)(Z_1^{b_i} + Y_0^{b_i})$, where Z_k and Y_k are the first and the last link distances in $L_\lambda(kbe_1, (k+1)be_1, b)$, respectively.

Note that the shortest path for $L_\lambda(kbe_1, (k+1)be_1, b)$ is not the direct path $kbe_1 \to (k+1)be_1$ with probability same as

$$\mathbb{P}\left[|H_\lambda \cap R_g(kbe_1, (k+1)be_1)| = 0\right] = p(\lambda, b) = exp(-\lambda \mu(R_g(0, be_1)).$$

Recall that $\mu(R_g(.))$ denote the empty region in ERG of g.

If none of the shortest paths for $L_\lambda(kbe_1, (k+1)be_1, b)$ is a direct path, then this paste procedure may be repeated and we obtain

$$L_\lambda(0, rbe_1, b) \leq \sum_{k=0}^{r-1} [L_\lambda(kbe_1, (k+1)be_1, b) + \sum_{i=1}^{m} a_i(2^{b_i-1}-1)(Z_k^{b_i} + Y_k^{b_i})] \quad (2.25)$$

On the other hand, if this event does not occur then there is at least 1 direct edge between $kbe_1 \to (k+1)be_1$ for some $k = 0, 1, .., r$ which has probability at-most $rp(\lambda, b)$ where $p(\lambda, b) = exp(-\lambda \theta_{b_1,d} b^d)$ for case 1 and $exp(-\lambda \frac{a_2 \theta_{b_2,d}}{a_1} b^{d+b_2-1})$ for case 2 for $b \to 0$ as provided in proposition 2.12 .

Also, note that if k, l are integers and $l - k \geq 3$, then $L_\lambda(kb, (k+1)b, b)$ and $L_\lambda(lb, (l+1)b, b)$ become mutually independent, and so are Z_k and Z_l. Then the sum in can be split into $K \geq 3$ sums of independent variables and each sum has at most $K^{-1}r$ summands. Note that each summand is a.s. bounded since $(Z_k^{b_i} + Y_k^{b_i}) \leq L_\lambda(kb, (k+1)b, b) \leq g(b)$.

We denote

$$X_k \coloneqq L_\lambda(kbe_1, (k+1)be_1, b) + \sum_{i=1}^{m} a_i(2^{b_i-1}-1)(Z_k^{b_i} + Y_k^{b_i})$$

for $k = 0, 1, .., r-1$. As discussed earlier, X_k and X_l are independent if $|k-l| \geq 3$. Also all X_k has generic expectation denoted by $\mathbb{E}[X]$. We denote for case 1

$$\mu = \frac{\mathbb{E}[X]}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})b} = \frac{\mathbb{E}[L_\lambda(0, be_1, b)] + \mathbb{E}[\sum_{i=1}^{m} a_i(2^{b_i-1}-1)(Z_k^{b_i} + Y_k^{b_i})]}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})b}$$

Also, $X_k \leq 2^{b_j-1} g(b)$.

Thus we can apply Azuma Hoeffding's inequality for $K = 4$ separate sequences (w.l.g. let $r \equiv 1 (\mathrm{mod}\, 4)$) namely

$$\{X_0 - \mathbb{E}[X], X_4 - \mathbb{E}[X], .., X_{r-1} - \mathbb{E}[X]\}$$

,

$$\{X_1 - \mathbb{E}[X], X_5 - \mathbb{E}[X], .., X_{r-2} - \mathbb{E}[X]\}$$

,

$$\{X_2 - \mathbb{E}[X], X_6 - \mathbb{E}[X], .., X_{r-3} - \mathbb{E}[X]\}$$

and

$$\{X_3 - \mathbb{E}[X], X_7 - \mathbb{E}[X], .., X_{r-4} - \mathbb{E}[X]\}.$$

As we discussed, elements within a each sequences are independent and we have $|X_k - X_{k-4}| \leq 2^{b_j-1} g(b)$.

Now we look at the following probability for case 1 ($b_1 > 1$),

$$\mathbb{P}(A) = \mathbb{P}[\frac{L_\lambda(0, rbe_1, b)}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})rb} \geq \mu + \epsilon]$$

$$= \mathbb{P}[A|B]\,\mathbb{P}[B] + \mathbb{P}[A|B^c]\,\mathbb{P}[B^c]$$

where $B = \{L_\lambda(0, rbe_1, b) \leq \sum_{k=1}^{r-1} X_k\}$.

$$\leq \mathbb{P}[\frac{L_\lambda(0, rbe_1, b)}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}}) rb} \geq \mu + \epsilon | L_\lambda(0, rbe_1, b) \leq \sum_{k=1}^{r-1} X_k] + rp(\lambda, b)$$

$$\leq \mathbb{P}[\sum_{k=1}^{r-1} (X_k - \mathbb{E}[X_k]) > \epsilon \lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}}) rb] + rp(\lambda, b)$$

$$\leq 4\, exp(-\frac{[\frac{\epsilon \lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})(r-3)b}{4}]^2}{2^{\frac{(r-3)}{4}} [2^{b_m - 1} g(b)]^2}) + rp(\lambda, b)$$

$$= 4\, exp(-\frac{(r-3)\epsilon^2 b^2 [\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}})]^2}{2^{2b_m+1} [g(b)]^2}) + rp(\lambda, b) \tag{2.26}$$

Now we set $r = [\lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda|]$ for $0 < \alpha < 1$ where $[.]$ denote the floor function (α will be chosen later) and $rb = |x_\lambda - y_\lambda|$. Therefore we have $b = \lambda^{\frac{\alpha-1}{d}} \to 0$ but $\lambda^{\frac{1}{d}} b \to \infty$.

Similarly for case 2 ($b_1 = 1$), we define

$$\mu = \frac{\mathbb{E}[X]}{\lambda^{\frac{1}{d+b_2-1}} g(\lambda^{-\frac{1}{d+b_2-1}}) b} = \frac{\mathbb{E}[L_\lambda(0, be_1, b)] + \mathbb{E}[\sum_{i=1}^m a_i (2^{b_i - 1} - 1)(Z_k^{b_i} + Y_k^{b_i})]}{\lambda^{\frac{1}{d+b_2-1}} g(\lambda^{-\frac{1}{d+b_2-1}}) b}$$

and set $r = [\lambda^{\frac{1-\alpha}{d+b_2-1}} |x_\lambda - y_\lambda|]$ for $0 < \alpha < 1$ (α will be chosen later) and $rb = |x_\lambda - y_\lambda|$. Similarly the application of Azuma Hoeffding in second case yields

$$\mathbb{P}[\frac{L_\lambda(0, rbe_1, b)}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}}) rb} \geq \mu + \epsilon] \leq$$

$$4\, exp(-\frac{(r-3)\epsilon^2 b^2 [\lambda^{\frac{1}{d+b_2-1}} g(\lambda^{-\frac{1}{d+b_2-1}})]^2}{2^{2b_m+1} [g(b)]^2}) + rp(\lambda, b) \tag{2.27}$$

Proposition 2.23

$\mu \to C(d, b_1 | g)$ in case 1 and $C(d, 1 | g)$ in case 2. □

:

Proof 2.31

Note that

$$\mu \leq \frac{\mathbb{E}[L_\lambda(0, be_1, b)]}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}}) b} + \frac{2(2^{b_m - 1} - 1) g(\xi_b)}{\lambda^{\frac{1}{d}} b g(\lambda^{-\frac{1}{d}})}$$

We have second quantity on the RHS $\to 0$ by proposition 2.18 since $\lambda^{\frac{1}{d}} b \to \infty$. Therefore

$$\mu \to \lim_{\lambda \to \infty} \frac{\mathbb{E}[L_\lambda(0, be_1, b)]}{\lambda^{\frac{1}{d}} g(\lambda^{-\frac{1}{d}}) b} \to C(d, b_1 | g)$$

in case 1.

The case 2 follows similarly from proposition 2.18. □

Now we simplify (2.26)

for case 1, as $\lambda \to \infty$, 1st term in RHS of (2.26) is

$$4\, exp\big(-O(\frac{\epsilon^2 \lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda| a_1^2 \lambda^{2(\frac{1-b_1}{d})}}{2^{2b_m+1} a_1^2 b^{2b_1-2}})\big)$$

$$= 4\, exp\left(-O(\frac{\epsilon^2 (\lambda^{\frac{1-\alpha}{d}})(\lambda^{\frac{1-\alpha}{d}})^{2b_1-2} |x_\lambda - y_\lambda| \lambda^{2(\frac{1-b_1}{d})}}{2^{2b_m+1}})\right)$$

$$= 4\, exp\left(-C|x_\lambda - y_\lambda| \epsilon^2 \lambda^{\frac{1+\alpha-2\alpha b_1}{d}}\right)$$

Similarly for case 2, 1st term in RHS of (2.27) is $4\, exp(-C|x_\lambda - y_\lambda| \epsilon^2 \lambda^{\frac{1-\alpha}{d+b_2-1}})$. We have for case 1, the 2nd term in RHS of (2.26) is $\lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda|\, exp(-\theta_{b_1,d} \lambda^\alpha)$ and for case 2, the 2nd term in RHS of (2.27) is $\lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda|\, exp(-\frac{a_2 \theta_{b_2,d}}{a_1} \lambda^\alpha)$.

Finally we solve for α (subject to $\alpha \leq \frac{1}{d}$ in case 1 and $\alpha \leq \frac{1}{d+b_2-1}$ in case 2) so that 1st term and 2nd term in (2.26) and (2.27) respectively has common exponential decay and thus we have overall sharpest probability decay rate

So for case 1,

$$\alpha = \frac{1 + \alpha - 2\alpha b_1}{d} \Rightarrow \alpha = \frac{1}{d + 2b_1 - 1}$$

and for case 2,

$$\alpha = \frac{1-\alpha}{d+b_2-1} \Rightarrow \alpha = \frac{1}{d+b_2}$$

So we have, for case 1 as $\lambda \to \infty$

$$\mathbb{P}[\frac{L_\lambda(x_\lambda, y_\lambda, \lambda^{\frac{1-\alpha^*}{d}})}{a_1 \lambda^{\frac{1-b_1}{d}} |x_\lambda - y_\lambda|} \geq C(d, b_1|g) + 2\epsilon]$$

$$\leq \lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda| \, exp(-C_1 \lambda^{\alpha^*}) + 4 \, exp(-\epsilon^2 C_2 \lambda^{\alpha^*})$$

for constants $C_1 > 0$, $C_2 > 0$ depending only on R_1, d, g and $\alpha^* = \frac{1}{d+2b_1-1}$. Similarly for case 2 as $\lambda \to \infty$,

$$\mathbb{P}[\frac{L_\lambda(x_\lambda, y_\lambda, \lambda^{\frac{1-\alpha^*}{d}})}{a_1 |x_\lambda - y_\lambda|} \geq C(d, 1|g) + 2\epsilon]$$

$$\leq \lambda^{\frac{1-\alpha}{d}} |x_\lambda - y_\lambda| \, exp(-C_1 \lambda^{\alpha^*}) + 4 \, exp(-\epsilon^2 C_2 \lambda^{\alpha^*})$$

for constants $C_1 > 0$, $C_2 > 0$ depending only on R_1, d, g and $\alpha^* = \frac{1}{d+2b_2}$. as λ and $\lambda^{\alpha^*} |x_\lambda - y_\lambda| \to \infty$. □

2.3.5 De-Poissonization

The primary goal of this sub-section is to upgrade the proof techniques from Poisson process to uniform density (on a circle) scenario.

The generic set-up we shall assume in this sub-section is as follows: Let $z \in \mathbb{R}^d$, $0 < R_1 < R_2$. Assume that the probability density f is uniformly supported in $B(z, R_2)$. i.e. $f(u) = f_0 1_{\{u \in B(z,R_2)\}}$ and $f_0 > 0$. Let x_n and y_n be sequences of points in $B(z, R_1)$. Let $\chi_n = \{X_1, .., X_n\}$ are generated from f.

Proposition 2.24

Let α^* be same as in lemma 2.5 and $\rho_0 > 0$ be the constant introduced in proposition 2.16. Recall that $|L_g(x_n, y_n, \chi_n)|$ denote the number of nodes in the g-shortest path

from x_n to y_n with minimal number of nodes. If $C_1 > C(d, b_1|g)\rho_0$ for case 1 and $C_1 > C(d, 1|g)\rho_0$ for case 2 and $\underset{n}{lim\ inf}\,(nf_0)^{\alpha^*}|x_n - y_n| = \infty$, then

$$\underset{n\to\infty}{lim\ sup}\frac{1}{(nf_0)^{\alpha^*}|x_n - y_n|}\,log\,\mathbb{P}\{\frac{|L_g(x_n, y_n, \chi_n)|}{(nf_0)^{\frac{1}{d}}|x_n - y_n|} \geq C_1\} < 0 \tag{2.28}$$

for case 1 and

$$\underset{n\to\infty}{lim\ sup}\frac{1}{(nf_0)^{\alpha^*}|x_n - y_n|}\,log\,\mathbb{P}\,\{\frac{|L_g(x_n, y_n, \chi_n)|}{(nf_0)^{\frac{1}{d+b_2-1}}|x_n - y_n|} \geq C_1\} < 0$$

for case 2. □

Proof 2.32

Let

$$\alpha^{**} = \begin{cases} \frac{1}{d} & \text{for case 1} \\ \frac{1}{d+b_2-1} & \text{for case 2} \end{cases}$$

Consider case 1. Let N be a Poisson variable with mean na, $a > 1$. Suppose that $N \geq n$. Then a path in X_n is also a path in $X_N = H_{naf(u)} \subset H_{naf_0}$ since $f(u) \leq f_0$ for all $u \in \mathbb{R}^d$. Choose $\epsilon > 0$ sufficiently small so that $C_1 > (C(d, b_1|g) + 2\epsilon)\rho_0$. Suppose that

$$L_g(x_n, y_n, \chi_n) \leq (C(d, b_1|g) + 2\epsilon)(naf_0)^{\frac{1}{d}}g((naf_0)^{-\frac{1}{d}})|x_n - y_n| \tag{2.29}$$

and that the number of nodes in the path is at-least $C_1(naf_0)^{\frac{1}{d}}|x_n - y_n|$. Then by proposition 2.16 , the probability that such a path exists decays exponentially in $(naf_0)^{\alpha^{**}}|x_n - y_n|$ since $C_1 > (C(d, b_1|g) + 2\epsilon)\rho_0$ in case 1.

Let $N^{'}$ be a Poisson variable with mean $na^{'}$, $a^{'} < 1$. Suppose that $N^{'} \leq n$. Then $L_g(x_n, y_n, \chi_n) \leq L_N(x_n, y_n) = L_{na'f}(x_n, y_n)$, and the probability that

$$\frac{L_g(x_n, y_n, \chi_n)}{(naf_0)^{\frac{1}{d}}g[(naf_0)^{-\frac{1}{d}}]|x_n - y_n|} > C(d, b_1|g) + \epsilon \tag{2.30}$$

is exponentially small in $(na^{'}f_0)^{\alpha^*}|x_n - y_n|$ by lemma 2.5 .

Choose a and $a^{'}$so that

$$(C(d, b_1|g) + \epsilon)a^{'\frac{1}{d}}g((na^{'}f_0)^{-\frac{1}{d}}) < (C(d, b_1|g) + 2\epsilon)a^{\frac{1}{d}}g((naf_0)^{-\frac{1}{d}})$$

so that the assumption (2.29) is satisfied when (2.30) is false. In addition, both $\mathbb{P}\{N < n\}$ and $\mathbb{P}\{N^{'} > n\}$ are exponentially small in n by the Chernoff bound. To be more precise, if $X \sim Poi(\lambda)$, from Chernoff bound one can obtain

$$\mathbb{P}\{|X - \lambda| \geq x\} \leq 2e^{-\frac{x^2}{2(\lambda+x)}}; x > 0$$

Then the claim is proved since the slowest probability decay is determined by lemma 2.5 .

A proof for case 2 can be argued similarly and is omitted. □

Proposition 2.25

If $\underset{n\to\infty}{lim\ inf}\ n^{\alpha^*}|x_n - y_n| = \infty$, where α^* is same as in Lemma 2.5 then,

$$\underset{n\to\infty}{lim}\ \frac{\mathbb{E}\, L_n(x_n, y_n)}{(nf_0)^{\frac{1-b_1}{d}}a_1|x_n - y_n|} = C(d, b_1|g)$$

□

Proof 2.33

The proof (a standard de-Poissonization argument) is very similar to proposition 5 in Hwang et al. (2016) and is omitted. □

2.3.6 Complete convergence of g distance

Now we are ready to apply Talagrand's inequality to L_n. This concentration inequality provides a high-probability lower bound on L_n and gives a complete convergence result.

Lemma 2.6

Let

$$\alpha^* = \begin{cases} \frac{1}{d+2b_1} & \text{for case 1} \\ \frac{1}{d+b_2+1} & \text{for case 2} \end{cases}$$

.

Then,

$$\limsup_{n\to\infty} \frac{1}{(nf_0)^{\alpha^*}|x_n - y_n|} log\, \mathbb{P}\{|\frac{L_n(x_n, y_n)}{a_1(nf_0)^{\frac{1-b_1}{d}}|x_n - y_n|} - C(d, b_1|g)| > \epsilon\} < -C\epsilon^2$$

for some $C > 0$ depending on d and g and for ϵ sufficiently small. □

Note that this condition relating to α^* is stronger then lemma 2.5 and determines the final convergence rate of theorem 2.1.

Proof 2.34

We define 3 events which almost surely occurs (shown in Hwang et al. (2016) for power-weighted distance and can be easily generalized for g distance).

- E_S be the event that $|L_n(x_n, y_n, \chi_n)| \leq C_1(nf_0)^{\frac{1}{d}}|x_n - y_n|$ for some $C_1 > C(d, b_1|g)\rho_0$ for case 1 and $|L_n(x_n, y_n, \chi_n)| \leq C_1(nf_0)^{\frac{1}{d+b_2-1}}|x_n - y_n|$ for some $C_1 > C(d, 1|g)\rho_0$ for case 2.

- E_F be the event that all the edges of the shortest paths are at most $(nf_0)^{\frac{(\alpha^*-1)}{d}}$ for case 1 and $(nf_0)^{\frac{(\alpha^*-1)}{d+b_2-1}}$ for case 2.

- E_C be the event that at every point $u \in B(z; R_2)$, at least one of the sample points is in $B(u; (nf_0)^{\frac{\alpha^*-1}{d}})$ in case 1 and $B(u; (nf_0)^{\frac{\alpha^*-1}{d+b_2-1}})$ in case 2.

We work with case 1 first.

For every $a > 0$, define $E(a)$ to be the event that $L_n(x_n, y_n) = L_n(x_n, y_n, \chi_n) \geq a$. Let ω and η be outcomes in the probability space. Assume that $\omega \in E_S \cap E_F$, and

that $\eta \in E(a) \cap E_C$. Let $\pi * (\omega)$ be the shortest path $L_n(x_n, y_n)$ between x_n and y_n through realized sample points $X_1(\omega), ..., X_n(\omega)$. If $\pi * (\omega)$ is sequence

$$x_n = \pi_0(\omega) \to \pi_1(\omega) \to .. \to \pi_k(\omega) = y_n \tag{2.31}$$

where $k = |L_n(\omega)|$, then we build the path $\pi(\eta)$ from x_n to y_n through $X_1(\eta), \ldots, X_n(\eta)$ as follows:

For each $i \in \{1, ..., k-1\}$, let j denote the index where $X_j(\omega) = \pi_i(\omega)$. If $X_j(\omega) = X_j(\eta)$, then set $\pi_i(\eta) = \pi_i(\omega)$. Otherwise, since it was assumed that $\eta \in E_C$, $\exists$ some $X_l(\eta) \in B(\pi_i(\omega); (nf_0)^{(\alpha-1)/d})$. Set $\pi_i(\eta) = X_l(\eta)$. Then it follows that $|\pi_i(\eta) - \pi_i(\omega)| \leq (nf_0)^{(\alpha-1)/d}$ for all $i = 1, ..., k$.

At the same time $|\pi_{i+1}(\omega) - \pi_i(\omega)| \leq (nf_0)^{(\alpha-1)/d}$ for all $i = 1, .., k$ since $\omega \in E_F$.

Therefore,

$$|\pi_i(\eta) - \pi_{i+1}(\eta)|$$

$$\leq |\pi_i(\eta) - \pi_i(\omega)| + |\pi_{i+1}(\omega) - \pi_i(\omega)| + |\pi_i(\eta) - \pi_i(\omega)| \leq 3(nf_0)^{(\alpha-1)/d}$$

Let I be the set of indices i where $\pi_i(\omega) \neq \pi_i(\eta)$. Then the g-distance length of the path $\pi(\eta)$ is at most

$$L(\pi(\eta)) = \sum_{i=0}^{k} g(|\pi_i(\eta) - \pi_{i+1}(\eta)|) \leq \sum_{i=0}^{k} g(|\pi_i(\omega) - \pi_{i+1}(\omega)|)$$

$$+ \sum_{i \in I \vee i+1 \in I} g(|\pi_i(\eta) - \pi_{i+1}(\eta)|)$$

$$\leq L(\pi * (\omega)) + 2|I| \sum_{i=1}^{m} a_i [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i}$$

On the other hand, $\eta \in E(a)$. Therefore

$$L(\pi * (\omega)) = L_n(\omega) \geq a - 2|I| \sum_{i=1}^{m} a_i [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i} \tag{2.32}$$

Let $d_c(\omega; E(a)\cap E_C)$ be the Talagrand's convex distance (Section 4.1 in Talagrand (1994)) of ω to $E(a) \cap E_C$. By lemma 4.1.2 in Talagrand (1994), by choosing (see section 2.4)

$$s_i = \begin{cases} 1/\sqrt{|L_n(\omega)|} & \text{if } X_i \text{ is in the path } L_n(\omega) \\ 0 & \text{otherwise} \end{cases}$$

we can see that $\exists\, \eta \in E(a) \cap E_C$ such that

$$|I| \leq d_c(\omega; E(a) \cap E_C)\sqrt{|L_n(\omega)|}\sum_{i=1}^{m} a_i[3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i}$$

Therefore

$$L_n(\omega) \geq a - 2d_c(\omega; E(a) \cap E_C)\sqrt{|L_n(\omega)|}\sum_{i=1}^{m} a_i[3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i}. \tag{2.33}$$

In particular, if $L_n(\omega) \leq a - u$ for $u > 0$,

$$\begin{aligned} d_c(\omega; E(a) \cap E_C(a)) &\geq \frac{u}{2\sqrt{|L_n(\omega)|}(\sum_{i=1}^{m} a_i[3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})} \\ &\geq \frac{u}{2\sqrt{C_1(nf_0)^{\frac{1}{d}}|x_n - y_n|}(\sum_{i=1}^{m} a_i[3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})} \end{aligned} \tag{2.34}$$

since $\omega \in E_S$.

Let $t > 0$ denote the RHS of (2.34) and let $(E(a) \cap E_C)_t$ denote the enlargement of $(E(a) \cap E_C)$ i.e. the collection of all elementary outcomes whose convex distance from $(E(a) \cap E_C)$ is at most t. Then implies that

$$\mathbb{P}(\{L_n(\omega) \leq a - u\} \cap E_S \cap E_F) \leq 1 - \mathbb{P}((E(a) \cap E_C)_t) \tag{2.35}$$

Let M_n be a median of $L_n(\omega)$. Note that $\mathbb{P}[E(M_n) \cap E_C]$ is arbitrarily close to $\frac{1}{2}$ for sufficiently large n since $\mathbb{P}(\mathbb{E}(M_n)) = \mathbb{P}(L_n(\omega) \geq M_n) = \frac{1}{2}$ and $\mathbb{P}(E_C)$

approaches 1 as $n \to \infty$. In particular, for n sufficiently large, $\mathbb{P}[E(M_n) \cap E_C] \geq \frac{1}{3}$. Set $a = M_n$ in and apply Talagrand's inequality with $A = E(M_n) \cap E_C$ to get $(1 - \mathbb{P}((E(M_n) \cap E_C)_t) \leq 3\, exp\,(-\frac{t^2}{4})$.

Hence

$$\mathbb{P}\big(\{L_n(\omega) \leq M_n - u\} \cap E_S \cap E_F\big)$$

$$\leq\ 3\, exp\,\big(-\frac{u^2}{4^2 C_1 (nf_0)^{\frac{1}{d}} |x_n - y_n| (\sum_{i=1}^{m} a_i [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})^2}\big)$$

$$\Rightarrow \mathbb{P}\{L_n(\omega) \leq M_n - u\} \leq 3\, exp\,\big(-\frac{u^2}{4^2 C_1 (nf_0)^{\frac{1}{d}} |x_n - y_n| (\sum_{i=1}^{m} a_i [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})^2}\big)$$

$$+(1 - \mathbb{P}(E_F)) + (1 - \mathbb{P}(E_S))$$

For the upper bound part, let $a = M_n + u$. From $1 - \mathbb{P}((E(M_n) \cap E_C)_t) \geq \frac{1}{3}$ for sufficiently large n since both $\mathbb{P}(E_S)$ and $\mathbb{P}(E_F)$ converges to 1 as $n \to \infty$. Apply Talagrand's inequality again for $A = E(a) \cap E_C = E_C \cap \{L_n(\omega) \geq M_n + u\}$ to obtain

$$\mathbb{P}\{L_n(\omega) \geq M_n + u\}$$

$$\leq 3\, exp\,\big(-\frac{u^2}{4^2 C_1 (nf_0)^{\frac{1}{d}} |x_n - y_n| (\sum_{i=1}^{j} a_{b_i} [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})^2}) + (1 - \mathbb{P}(E_C)\big)$$

for sufficient large n since both $\mathbb{P}(E_F)$ and $\mathbb{P}(E_S)$ converge to 1 as $n \to \infty$.

Therefore

$$\mathbb{P}[\frac{|L_n - M_n|}{(nf_0)^{\frac{1}{d}} g[(nf_0)^{-\frac{1}{d}}] |x_n - y_n|} > u] \leq \tag{2.36}$$

$$6\, exp\,\big(-\frac{u^2 [(nf_0)^{\frac{1}{d}} g[(nf_0)^{-\frac{1}{d}}] |x_n - y_n|]^2}{4^2 C_1 (nf_0)^{\frac{1}{d}} |x_n - y_n| (\sum_{i=1}^{m} a_i [3(nf_0)^{\frac{\alpha-1}{d}}]^{b_i})^2}\big) + h_n$$

where $h_n = (1 - \mathbb{P}(E_C)) + (1 - \mathbb{P}(E_F)) + (1 - \mathbb{P}(E_S)) = O(exp(-(nf_0)^{\alpha})$.

Leading decay term of 1st term of RHS of (2.36) is

$$6\,exp[-\frac{u^2(nf_0)^{\frac{1-2b_1}{d}}|x_n-y_n|}{C_1(nf_0)^{\frac{2b_1(\alpha-1)}{d}}}]=6\,exp[-\frac{u^2|x_n-y_n|(nf_0)^{\frac{1-2b_1\alpha}{d}}}{C_1}]$$

So we again match decay of both term of (2.36) to get $\alpha=\frac{1-2b_1\alpha}{d}\Rightarrow\alpha^*=\frac{1}{d+2b_1}$. Since our new α^* is smaller than α^* in lemma 2.5 (which was $\frac{1}{d+2b_1-1}$), it can be applied here as well.

Now for case 2, we can repeat all the argument with $(nf_0)^{\frac{\alpha-1}{d}}$ replaced by $(nf_0)^{\frac{\alpha-1}{d+b_2-1}}$. Thus, in this case the leading decay term is

$$6\,exp[-\frac{u^2|x_n-y_n|}{12^2C_1\,(nf_0)^{\frac{1+2(\alpha-1)}{d+b_2-1}}}]$$

.

So we again match decay of both term to get $\alpha=\frac{1-2\alpha}{d+b_2-1}\Rightarrow\alpha^*=\frac{1}{d+b_2+1}$. Since our new α^* for case 2 is smaller than α^* in lemma 2.5 (which was $\frac{1}{d+b_2}$) as well , it can be applied here as well.

So for both case, we have

$$\mathbb{P}[\frac{|L_n-M_n|}{(nf_0)^{\frac{1}{d}}g[(nf_0)^{-\frac{1}{d}}]|x_n-y_n|}>u]\le 6\,exp(-Cu^2|x_n-y_n|(nf_0)^{\alpha^*})+h_n \quad (2.37)$$

Now we show that the median and the mean are close. By the Jensen's inequality, $|\mathbb{E}L_n-M_n|\le\mathbb{E}|L_n-M_n|$. Also note that $\mathbb{P}(|L_n-M_n|>u)=0$ if $u\ge g|x-y|$. If is integrated with respect to u, we obtain

$$\frac{|\mathbb{E}L_n-M_n|}{(nf_0)^{\frac{1}{d}}g[(nf_0)^{-\frac{1}{d}}]|x_n-y_n|}\le 6(\frac{\pi}{C|x_n-y_n|(nf_0)^{\alpha^*}})^{1/2}+h_n\frac{g|x_n-y_n|}{|x_n-y_n|(nf_0)^{\frac{1}{d}}g[(nf_0)^{-\frac{1}{d}}]}. \quad (2.38)$$

All events E_C, E_F, and E_S have their probability approach to 1 exponentially fast.

Therefore

$$\lim_{n\to\infty} \frac{|\mathbb{E}L_n - M_n|}{(nf_0)^{\frac{1}{d}} g[(nf_0)^{-\frac{1}{d}}]|x_n - y_n|} = 0$$

By proposition 2.25 if n is large enough then (for case 1)

$$\mathbb{P}\{|\frac{L_n}{(nf_0)^{\frac{1}{d}} g[(nf_0)^{-\frac{1}{d}}]|x_n - y_n|} - C(d, b_1|g)| > \epsilon\} \leq \mathbb{P}\{\frac{|L_n - M_n|}{(nf_0)^{\frac{1}{d}} g[(nf_0)^{-\frac{1}{d}}]|x_n - y_n|} > \frac{\epsilon}{2}\}$$

Then the claim follows from (2.38) and proposition 2.24 . The proof for case 2 is similar and is omitted. □

Remark 2.7

As discussed in the beginning of this chapter, the next steps in the proof of theorem 2.1 is to generalize the case for a non-uniform density essentially showing that shortest paths between 2 points stays relatively local, thus the local convergence of shortest path could be "summed up" globally followed by generalizing it to Riemannian manifold case. Note that we have already achieved our convergence rate stated in theorem 2.1 and nothing changes from this point on-wards from the arguments made in Hwang et al. (2016), and therefore is omitted. The proof of theorem 2.2 is also very similar to the arguments made in Hwang et al. (2016) and is omitted. □

2.3.7 Low sample size and density (λ or $nf \to 0$) scenario

All the steps from equation (2.23) to lemma 2.6 can be repeated, albeit with minor modifications, to argue that in low density and small-sample (λ or nf small) regime, suitably scaled normalized g-distance reduce to it's normalized highest power distance. To be precise, we have

Lemma 2.7

Let H_λ be a homogeneous Poisson process and m is finite, then if $\limsup_{\lambda\to\infty} \lambda^{\frac{1}{d+2b_m}} |x_\lambda -$

$y_\lambda| \to \infty$ we have

$$\lim_{\lambda\to 0} \frac{1}{\lambda^{\frac{1}{d+2b_m}}|x_\lambda - y_\lambda|} log\, \mathbb{P}\{|\frac{L_\lambda(x_\lambda, y_\lambda)}{a_m \lambda^{\frac{1-b_m}{d}}|x_\lambda - y_\lambda|} - C(d, b_m|g)| > \epsilon\} < -C\epsilon^2$$

for some $C > 0$ depending on d and g and for ϵ sufficiently small. □

In terms of g^N distance this translates into,

Corollary 2.7

$$\lim_{\lambda\to 0} \frac{1}{\lambda^{\frac{1}{d+2b_m}}|x_\lambda - y_\lambda|} log\, \mathbb{P}\{|\frac{L^N_\lambda(x_\lambda, y_\lambda)}{\lambda^{\frac{1-b_m}{b_m d}}|x_\lambda - y_\lambda|^{\frac{1}{b_m}}} - C(d, b_m|g)| > \epsilon\} < -C'\epsilon^2 \quad (2.39)$$

□

for some $C' > 0$.

When m is ∞, by proposition 2.5 g^N converge to LLPD and in that case the LLPD convergence asymptotic applies.

Remark 2.8

In Theorem 2.1 and 2.2 we assumed a boundary less Riemannian manifold. However, presence of boundary is very complicated. To give the simplest of examples, consider the half plane $H_+ = \{x : xe_1 \geq 0\}$ and $H_\lambda \cap H_+$. Now, we can still show that the limit $\lim_{t\to\infty} \frac{L_g(0,te_1,H_\lambda\cap H_+)}{t}$ exists (we call it $C(d,g)_+$) and obviously we have $C(d,g)_+ \geq C(d,g)$. But it's unclear if the inequality is strict or not and what happens to the asymptotics if $C(d,g)_+ > C(d,g)$. Boundary effect is in general a fascinating topic for research. □

2.4 What happens when $g \notin \mathbb{G}_0$?

So far, our primary assumption for $g(t)$ has been the assumption of $g(t) \in \mathbb{G}_0$ for some $t \in [0, \delta)$, $\delta > 0$. We recall that

$$\mathbb{G}_0 \coloneqq \Big\{ g(t) = \sum_{i=1}^{m} a_i t^{b_i} : [0,\infty) \to [0,\infty) | m \in \bar{\mathbb{N}}, a_i \in \mathbb{R}_0^+, b_i \in [1,\infty),$$

$$1 \leq b_1 < b_2 < .. \Big\}$$

As shown in the chapter, we require $g \in \mathbb{G}$ in order for g and g^N to be valid distances. However, clearly $\mathbb{G}_0 \subset \mathbb{G}$. While the justifications for why we chose $g \in \mathbb{G}_0$ were provided in the section 2.1 of this chapter, in this section we show how the proof technique in the chapter can be still be applied when $g(t) \notin \mathbb{G}_0$ by providing 2 examples.

Example 2.7

Let $g(t) = -t^k log t$ for $k > 1$ with $g(0) = 0$. It is not difficult to verify that $g(t) \notin \mathbb{G}_0$ and $g(t) \in \mathbb{G}$ for $t < e^{-\frac{2k-1}{k(k-1)}}$. Note that in this case for the Poisson process H_λ on $\mathbb{R}^d$ we have the following:

$$L_g[0, te_1, H_1 \cap T_{b_t}(0, e_1)] = -\sum_{i=0}^{n+1} |X_{i+1}(\omega) - X_i(\omega)|^k log |X_{i+1}(\omega) - X_i(\omega)|$$

On the other hand

$$L_g[0, \lambda^{-\frac{1}{d}} te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}} b_t}(0, e_1)]$$

$$= -\sum_{i=0}^{n+1} |X_{i+1}(\omega) - X_i(\omega)|^k \lambda^{-\frac{k}{d}} \{ log |X_{i+1}(\omega) - X_i(\omega)| + log \lambda^{-\frac{1}{d}} \}$$

Therefore,

$$\frac{L_g[0, \lambda^{-\frac{1}{d}} te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}} b_t}(0, e_1)]}{t}$$

$$\stackrel{d}{=} \lambda^{-\frac{k}{d}} \frac{L_g[0, te_1, H_1 \cap T_{b_t}(0, e_1)]}{t} + \lambda^{-\frac{k}{d}} log\lambda^{\frac{1}{d}} \frac{L_{k|g}[0, te_1, H_1 \cap T_{b_t}(0, e_1)]}{t}$$

This eventually generalize to (for a uniform distribution f on a disk in $\mathbb{R}^d$) as $n \to \infty$,

$$\mathbb{E}\, L_g(x, y, \chi_n) = C_1(nf)^{\frac{1-k}{d}}|x - y| + C_2 log(nf)^{\frac{1}{d}}(nf)^{\frac{1-k}{d}}|x - y|$$

for some $C_1, C_2 > 0$ dependent on d.

Therefore, following the proof technique in this chapter it follows that in this case,

$$C(d, g)\, \frac{n^{\frac{k-1}{d}}}{logn} L_g(x, y, \chi_n) \stackrel{a.s.}{\to} f^{\frac{1-k}{d}}|x - y|$$

for some $C(d, g) > 0$.

So, in it's most generalized form i.e. for a general density f on a Riemannian manifold, it can be shown that

$$C(d, g)\, \frac{n^{\frac{k-1}{d}}}{logn} L_g(x, y, \chi_n) \stackrel{a.s.}{\to} dist_k(x, y).$$

□

Example 2.8

Let $g(t) = e^{-t^{-k}}$ for $k > 0$. It is not difficult to verify that $g(t) \notin \mathbb{G}_0$ and $\exists \delta(k) > 0$ such that $g(t) \in \mathbb{G}$ for $t \in [0, \delta(k))$. Note that in this case for the Poisson process H_λ on $\mathbb{R}^d$ we have the following:

$$L_g[0, te_1, H_1 \cap T_{b_t}(0, e_1)] = \sum_{i=0}^{n+1} e^{-|X_{i+1}(\omega) - X_i(\omega)|^{-k}}.$$

On the other hand

$$L_g[0, \lambda^{-\frac{1}{d}} te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}} b_t}(0, e_1)] = \sum_{i=0}^{n+1} e^{-|X_{i+1}(\omega) - X_i(\omega)|^{-k} \lambda^{\frac{k}{d}}}.$$

Now as $\lambda \to \infty$,

$$\frac{L_g[0, \lambda^{-\frac{1}{d}} te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}} b_t}(0, e_1)]}{e^{-\max\limits_i |X_{i+1}(\omega) - X_i(\omega)|^{-k} \lambda^{\frac{k}{d}}}} = \sum_{i=0}^{n+1} \Big[\frac{e^{-|X_{i+1}(\omega) - X_i(\omega)|^{-k} \lambda^{\frac{k}{d}}}}{e^{-\max\limits_i |X_{i+1}(\omega) - X_i(\omega)|^{-k} \lambda^{\frac{k}{d}}}} \Big] \stackrel{a.s.}{\to} 1$$

Similarly,

$$\frac{L_g[0, te_1, H_1 \cap T_{b_t}(0, e_1)]}{e^{-\max\limits_i |X_{i+1}(\omega) - X_i(\omega)|^{-k}}} \overset{a.s.}{\to} 1.$$

Therefore we have,

$$\frac{L_g[0, \lambda^{-\frac{1}{d}} te_1, H_\lambda \cap T_{\lambda^{-\frac{1}{d}} b_t}(0, e_1)]^{\lambda^{-\frac{k}{d}}}}{t} \overset{d}{=} \frac{L_g[0, te_1, H_1 \cap T_{b_t}(0, e_1)}{t}$$

This eventually generalize to (for a uniform distribution f on a disk in $\mathbb{R}^d$) as $n \to \infty$, $\mathbb{E}\, L_g(x, y, \chi_n)^{(nf)^{-\frac{k}{d}}} = C_1 (nf)^{\frac{1}{d}} |x - y|$ for some $C_1 > 0$ dependent on d.

Therefore, following the proof technique in this chapter it is natural to conjecture that in this case,

$$C(d, g) L_g(x, y, \chi_n)^{\frac{n^{-\frac{k}{d}}}{logn}} \overset{a.s.}{\to} e^{f^{\frac{k}{d}}}$$

for some constant $C(d, g) > 0$. Further algebraic manipulation leads to

$$C^{'}(d, g) L_g^N(x, y, \chi_n) n^{\frac{1}{d}} (logn)^{\frac{1}{k}} \overset{a.s.}{\to} f^{-\frac{1}{d}}.$$

for some constant $C^{'}(d, g) > 0$. So, in it's most generalized form i.e. for a general density f on a Riemannian manifold, we conjecture that

$$C^{'}(d, g)\, n^{\frac{1}{d}} (logn)^{\frac{1}{k}} L_g^N(x, y, \chi_n) \overset{a.s.}{\to} f_{x,y}^{-\frac{1}{d}}.$$

where $f_{x,y} \coloneqq \max\limits_{\gamma} \min\limits_{t} f(\gamma_t)$, γ being any piecewise smooth curve with $\gamma_0 = x$ and $\gamma_1 = y$.

So, in this case L_g^N is equivalent to LLPD (LLPD's convergence limit is derived in the next chapter) but with a different scaling factor. □

Remark 2.9

As illustrated through these 2 examples, the convergence limit and rate of g^N can be different from theorem 2.1 and corollary 2.1 when $g \notin \mathbb{G}_0$ □

Chapter 3

A Strong Law for Longest Leg Path Distance

3.1 Introduction and organization

In 2003 Vincent and Bengio (2003); Fischer et al. (2004) introduced longest leg path distance (LLPD) as an alternative to the Euclidean distance for clustering purposes. Given a connected and compact Riemannian manifold (M, g_1), a density f supported on it and data $\chi_n \coloneqq \{X_1, X_2, .., X_n\}$; LLPD between 2 points x and y in M is defined as

$$LLPD(x, y, \chi_n) \coloneqq \min_{p} \max_{i} dist_1(x_i, x_{i+1}) \tag{3.1}$$

where p is a path from $x = x_0 \to x_1 \to x_2 .. \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary and $dist_1(.,.)$ is the default geodesic distance induced by g_1.

A desirable property of LLPD in many application is that it is an *ultrametric*.

Indeed, from the above definition it is not difficult to see that

$$LLPD(x, y, \chi_n) \leq max\left\{LLPD(x, z, \chi_n),\ LLPD(y, z, \chi_n)\right\}$$

for any x, y and z in M.

Since it's proposal, many authors (Fischer et al., 2004; Chang and Yeung, 2008; Little et al., 2017) have used LLPD for classification and clustering methods such as K-means, spectral clustering (SC), SVM algorithms. However, almost nothing is known about LLPD's convergence limit and rate. In this chapter we derive the convergence limit and rate of LLPD. In section 2 we state the convergence limit of LLPD and in section 3 a proof is provided. In section 4 convergence rate of LLPD and a proof is provided.

3.2 Main Result: Convergence limit of LLPD

In Little et al. (2017) sub-section 4.1.1, authors have pointed out the LLPD is intimately related to the longest edge of minimal spanning tree (or minimal bottleneck spanning tree) and also hinted that LLPD should have $O((\frac{logn}{n})^{\frac{1}{d}})$ dependency on n while providing an approximate high probability bound for LLPD in different data generating model. In this section we confirm their hypothesis and derive the exact limit and convergence rate of LLPD. In particular, under the following assumptions

A1: λ-level sets of M, $\partial M_\lambda \coloneqq \{x : f(x) = \lambda\}$ are finite union of smooth $d-1$ dimensional Riemannian manifolds for all $\lambda \in (f_0, f^0)$.

We prove the following

Theorem 3.1

Assume that (M, g_1) is connected, compact d-dimensional $(d > 1)$ Riemannian manifold without boundary and f is a continuous Lebesgue density on M with $f_0 \coloneqq$

$\inf_M f > 0$. Fix $x, y \in M$. Let V_d denote the volume of d-dimensional unit ball. Then

$$\lim_{n\to\infty} \left(\frac{nV_d}{\log n}\right)^{\frac{1}{d}} LLPD\big(x, y, \chi_n\big) \to {f_{x,y}}^{-\frac{1}{d}} \text{ almost surely.}$$

where $f_{x,y}$ is the "bottleneck value" of x and y defined as

$$f_{x,y} \coloneqq \max_{\gamma} \min_{t} f(\gamma_t)$$

where γ is any piecewise smooth curve with $\gamma_0 = x$ and $\gamma_1 = y$. □

3.3 Proof of the main result

We begin by re-stating the main result of Penrose et al. (1999) in our notation

Penrose et al. (1999) result: Suppose $\chi_n = \{x_1, x_2, ...x_n\}$ are i.i.d. random points in $\mathbb{R}^d$, $d \geq 2$, from density f, having connected support M with a smooth boundary ∂M with $f|_M$ continuous. Let M_n denote the smallest r such that the union of balls of diameter r centered at data points is connected. Let V_d denote the volume of the unit ball. Then as $n \to \infty$

$$\frac{n}{\log n} V_d M_n^d \to max\left\{(\inf_M f)^{-1}, 2(1-\frac{1}{d})(\inf_{\partial M} f)^{-1}\right\}, \quad a.s. \tag{3.2}$$

Remark 3.1

Note that it can be shown that M_n is the longest edge of the MST (Minimal spanning tree) or a MBST (Minimum bottleneck spanning tree) as well, although we shall not use this fact in our proof.

Remark 3.2

Also for our purpose we assume $\inf_M f > 0$ and $\inf_{\partial M} f > 0$ here. When one of these quantities are 0, LHS diverges to ∞. For example, in case of a d dimensional stan-

dard normal distribution on $\mathbb{R}^d$, we have $M_n = O_p(\frac{loglogn}{\sqrt{logn}})$ Penrose (1998). Thus $\frac{n}{logn} V_d M_n^d \to \infty$ a.s. in this case. □

We shall refer M_n as the "connectivity distance" henceforth. Theorem 3.1 can be extended in case of a d dimensional connected Riemannian manifold with a smooth $d-1$ dimensional boundary or without boundary as well with some minor adjustment. The technical argument for Riemannian manifold case (locally we can choose a coordinate chart so that Riemannian manifold remembles $\mathbb{R}^d$ with Euclidean distance) has been provided in lemma 5.1 and theorem 5.1 of Penrose (1999) for the largest nearest neighbor distance of data $\chi_n = \{x_1, x_2, .., x_n\}$. Note that almost identical argument can be applied for the connectivity distance as well. In fact, (suitably scaled) largest nearest neighbor distance and connectivity distance a.s. converge to the same asymptotic limit as noted in Penrose et al. (1999); Penrose (1999).

So, (3.2) holds for a d dimensional Riemannian manifold with a smooth $d-1$ dimensional boundary as well. It is to be noted that when the d dimensional smooth Riemannian manifold have no boundary (surface of a sphere or a toroid), then we have instead

$$\frac{n}{logn} V_d M_n^d \to (\inf_M f)^{-1} \quad a.s. \tag{3.3}$$

as noted in theorem 5.1 in Penrose (1999).

Henceforth, we shall only focus on the case when $M \subset \mathbb{R}^d$ is open and connected and $f_0 > 0$ since extension to Riemannian manifold case is routine as discussed before.

Definition: Given the data χ_n, we define a graph $G_n(r) = (\chi_n, E_n(r))$ by

$$E_n(X_i, X_j) = \begin{cases} 1 & \text{if } |x_i - x_j| \leq r \\ 0 & \text{if } |x_i - x_j| > r \end{cases}$$

Proposition 3.1

Under this formulation the following 2 definitions of connectivity distance,

$$M_n \coloneqq \{\inf_{r>0} : \cup_{i=1}^{n} B(x_i, \frac{r}{2}) \text{ is connected}\}$$

as provided in Penrose et al. (1999) and

$$M_n \coloneqq \{\inf_{r>0} : G_n(r) \text{ is connected}\}$$

are equivalent. We shall work with the second definition. □

Proof 3.1

Consider 2 points x, y. Now $|x-y| \leq r \Leftrightarrow B(x, \frac{r}{2}) \cap B(y, \frac{r}{2}) \neq \emptyset$ or in other word, for 2 points x and y, union of balls of diameter r centered at x and y is connected if and only if $|x-y| \leq r$. Hence proposition holds true for $n = 2$. Suppose, for $n-1$ points $X = \{x_1, .., x_{n-1}\}$ in $\mathbb{R}^d$ proposition holds. Then if we introduce a new point x_n it will become connected to remaining $n-1$ already connected points if and only if

$$B(x_n, \frac{r}{2}) \cap (\cup_{i=1}^{n-1} B(x_i, \frac{r}{2})) \neq \emptyset$$

$$\Leftrightarrow \cup_{i=1}^{n-1} (B(x_n, \frac{r}{2}) \cap B(x_i, \frac{r}{2})) \neq \emptyset$$

$\Leftrightarrow \exists j \in 1(1)n-1$ so that

$$B(x_n, \frac{r}{2}) \cap B(x_j, \frac{r}{2}) \neq \emptyset$$

$$\Leftrightarrow |x_n - x_j| \leq r$$

i.e. $G_n(r)$ is connected and proposition holds for n points. □

Lemma 3.1

Suppose $N \subset M$ is connected and open. We define

$$M_n|N \coloneqq \{\inf_{r>0} : G_n(r)|N \text{ is connected}\}$$

where $G_n(r)|N$ is the *restriction* of $G_n(r)$ to N formally defined by $G_n(r)|N \coloneqq (\chi_n \cap N, E_n(r))$. Then

$$\frac{n}{logn} V_d(M_n|N)^d \to (\underset{N}{inf} f)^{-1} \ a.s. \qquad \square$$

Proof 3.2

Let $\int_N f = q_N$, then considering f restricted to N, i.e. $f|N \coloneqq \frac{1}{q_N} f 1_N$ and $|\chi_n \cap N| \to nq_N$ and applying theorem Section 3.3 on N, we obtain

$$\frac{|\chi_n \cap N|}{log|\chi_n \cap N|} V_d(M_n|N)^d \to (\underset{N}{inf} f|N)^{-1} \ a.s.$$

$$\Rightarrow \frac{nq_N}{log(nq_N)} V_d(M_n|N)^d \to q_N(\underset{N}{inf} f)^{-1} \ a.s.$$

$$\Rightarrow \frac{n}{logn} V_d(M_n|N)^d \to (\underset{N}{inf} f)^{-1} \ a.s. \qquad \square$$

Remark 3.3

The main takeaway is scaled limit of $M_n|N$ do not depend on the size or shape of N but only on $\underset{N}{inf} f$. $\square$

Now we apply lemma 3.1 on ϵ-*variation* of a piecewise continuous path on M.

Definition: For a piecewise continuous path γ_t with $\gamma_0 = x$ and $\gamma_1 = y$, ϵ-*variation* of γ_t is defined by $\gamma_t(\epsilon) \coloneqq \cup_{t\in[0,1]} B(\gamma_t, \epsilon(t))$ for $\epsilon(t) \in [0, \epsilon_0)$ for some $\epsilon_0 > 0$ so that $\epsilon(0) = 0$ and $\epsilon(1) = 0$ for any $\epsilon \in [0, \epsilon_0)$ (i.e. end-points are fixed) and $\epsilon(t)$ for $t \in (0, 1)$ is continuous as a function of t and continuous and strictly decreasing as a function of ϵ so that $\underset{\epsilon\to 0^+}{lim} \epsilon(t) = 0$ for any $t \in (0, 1)$. Thus, $\underset{\epsilon\to 0^+}{lim} \gamma_t(\epsilon) = \gamma_t(0) = \gamma_t$.

Also we assume that $\gamma_t(\epsilon) \subset M \ \forall \epsilon \in [0, \epsilon_0)$. This can be guaranteed by choosing ϵ sufficiently small since M is open.

Corollary 3.1

As $\epsilon \to 0^+$ connectivity distance restricted to $\gamma_t(\epsilon)$, denoted by

$$M_n|\gamma_t(\epsilon) \coloneqq \{\inf_{r>0} : G_n(r)|\gamma_t(\epsilon) \text{ is connected}\}$$

reduce to the maximum edge of a path $p_n \in \gamma_t(\epsilon)$ joining x and y such that $\lim_{n\to\infty} d_F(p_n, \gamma_t) = 0$ where d_F denote the Fréchet distance and we have

$$\frac{n}{logn} V_d \, m^d(p_n) \to (\min_t f(\gamma_t))^{-1} \; a.s.$$

□

Proof 3.3

Follows from our assumption of $\lim_{\epsilon\to 0^+} \epsilon(t) = 0$ and application of lemma 3.1 and proposition 3.2 to $\gamma_t(\epsilon)$ and letting $\epsilon \to 0^+$. □

We shall show that LLPD of x and y almost surely goes to the same limit as the connectivity distance restricted to the open super-level set of f at level $f_{x,y}$ i.e. $M^0_{f_{x,y}} \coloneqq \{z : f(z) > f_{x,y}\}$ or equivalently the LLPD path joining x and y almost surely stays inside $M^0_{f_{x,y}}$ as $n \to \infty$.

We consider the connected component of super level set $M_{f_{x,y}} \coloneqq \{z : f(z) \geq f_{x,y}\}$ which contain x (and y), say $M'_{x,y}$. In-general, $M'_{x,y}$ may or may not have $d-1$ dimensional smooth boundary. However, by our assumption **A1**, $\exists k \in \mathbb{N}$ and $C_1, .., C_k$ (depending on x and y) so that $M'_{x,y} = \cup_{i=1}^k C_i$, $\partial M'_{x,y} = \cup_{i=1}^k \partial C_i$ and ∂C_is are all smooth $d-1$ dimensional boundaries of C_i's. So lemma 3.1 is applicable on individual $C_i \setminus \partial C_i$'s for $i = 1(1)k$. Clearly if $f_{x,y} = min\{f(x), f(y)\}$ then $k = 1$ and $k > 1$ case occurs only when $f_{x,y} < min\{f(x), f(y)\}$. When $k > 1$, clearly by definition for any i, $\exists j \neq i$ so that $\partial C_i \cap \partial C_j \neq \emptyset$.

Below we provide an example to illustrate our choices for C_i's.

Example 3.1

Consider

$$f \sim \frac{1}{3}\sum_{i=1}^{3} N_2(\mu_i, I_3)$$

where $\mu_1 = (0,0)$, $\mu_2 = (4,0)$ and $\mu_3 = (8,0)$ and $x = (0,0)$, $y = (4,0)$. It is not difficult to verify that $f_{x,y} = f(\tau, 0)$ for some τ close to 2 (we have $f_{x,y} < min\{f(x), f(y)\}$ here) and $M_{f_{x,y}}$ have 2 connected components, one ($M^{'}_{x,y}$) containing x and y and another having maxima at $(\tau^{'}, 0)$ for some $\tau^{'}$ close to 8 containing neither x nor y. Moreover, $M^{'}_{x,y} = \cup_{i=1}^{2} C_i$ where C_1 contains $x = (0,0)$, C_2 contains $y = (4,0)$, $C_1 \cap C_2 = (\tau, 0)$ and ∂C_1 and ∂C_2 are all smooth 1 dimensional manifolds in $\mathbb{R}^2$. □

We consider $k = 1$ case first and w.l.g. assume $f(y) \geq f(x)$ i.e. $M^{'}_{x,y} = C_1$. Note that we have $\frac{n}{logn} V_d (M_n | C_1 \setminus \partial C_1)^d \to (f(x))^{-1}$ *a.s.* by applying lemma 3.1 on $C_1 \setminus \partial C_1$.

Proposition 3.2

Under this set-up, let M_x be the distance of x from the nearest point in $\chi_n \cap (C_1 \setminus \partial C_1)$ i.e. $M_n(x) \coloneqq \underset{z \in \chi_n \cap (C_1 \setminus \partial C_1)}{min} |x - z|$ then $\frac{nV_d}{logn} M_n(x)^d \leq (f(x))^{-1}$ *a.s* □

Proof 3.4

Following Penrose (1999) let $\rho_n(t) \coloneqq (\frac{logn}{nV_d}\, t)^{\frac{1}{d}}$ and suppose $t > \frac{1}{f(x)}$. Locally C_1 is a half plane in a neighborhood of x since by **A1** C_1 has smooth $d-1$ dimensional boundary and local density restricted on the half-ball $C_1 \cap B(x, \rho_n(t))$ is $\frac{2}{q_{C_1}} f(x)$ where $q_{C_1} = \int_{C_1} f$. Thus following a standard de-Poissonization argument) we have

$$\mathbb{P}[M_n(x) \geq \rho_n(t)] = e^{-\frac{2}{q_{C_1}} f(x) n q_{C_1} \frac{1}{2} V_d \rho_n(t)^d}$$

$$= e^{-logn\, t f(x)} = n^{-t f(x)}$$

Since $tf(x) > 1$, $M_n(x) < \rho_n(t)$ *a.s.* by Borel-Cantelli lemma. □

Proposition 3.3

For $k = 1$ case, $\frac{nV_d}{logn} LLPD(x, y, \chi_n)^d \leq (f(x))^{-1}$ $a.s$ □

Proof 3.5

Since $\frac{nV_d}{logn}(M_n|C_1 \setminus \partial C_1)^d \rightarrow (f(x))^{-1}$ $a.s.$, $G_n(r_n)|C_1 \setminus \partial C_1$ is connected for $r_n = (\frac{logn}{nV_d}(f(x))^{-1})^{\frac{1}{d}}$. Also from proposition 3.2, $M_n(x) \leq (\frac{logn}{nV_d}(f(x))^{-1})^{\frac{1}{d}}$ and $M_n(y) \leq (\frac{logn}{nV_d}(f(y))^{-1})^{\frac{1}{d}}$, since by our assumption $f(y) \geq f(x)$, $\exists$ a path p_n from x to y such that maximum edge is $r_n = (\frac{logn}{nV_d}(f(x))^{-1})^{\frac{1}{d}}$ as $n \rightarrow \infty$. □

Proposition 3.4

For $k = 1$ case, $\frac{nV_d}{logn} LLPD(x, y, \chi_n)^d \geq (f(x))^{-1}$ $a.s$ □

Proof 3.6

Suppose, if possible, $\exists$ a path p_n $x = x_0 \rightarrow x_1 .. \rightarrow x_r = y$, $r \in \mathbb{N}$ and $x_1, .., x_{r-1} \in \chi_n$ such that it's maximum edge is $< (\frac{logn}{nV_d}(f(x))^{-1})^{\frac{1}{d}}$, say $= (\frac{logn}{nV_d}(f(x)+\delta)^{-1})^{\frac{1}{d}}$ for some $\delta > 0$. From proposition 3.3 we can ignore 2 edges joining x and y to a data-point and can only focus on $p_n^{'}$ $x_1 \rightarrow .. \rightarrow x_{r-1}$. From corollary 3.1 we have

$$\frac{n}{logn} V_d m^d(p_n^{'}) \rightarrow (\min_t f(\gamma_t))^{-1}$$

i.e. $\min_t f(\gamma_t) = f(x) + \delta$. However, since γ_t is piecewise continuous curve joining x and y, $\min_t f(\gamma_t) \leq f(x) < f(x) + \delta$ which is a contradiction. □

Now we consider the $k > 1$ case.

Proposition 3.5

For $k > 1$ case, $\frac{nV_d}{logn} LLPD(x, y, \chi_n)^d \leq (f_{x,y})^{-1}$ $a.s.$ □

Proof 3.7

Since $\inf_{C_i \setminus \partial C_i} f = f_{x,y}$ for $i = 1(1)k$, we have $\frac{nV_d}{logn}(M_n|C_i \setminus \partial C_i)^d \rightarrow (f_{x,y})^{-1}$ $a.s.$ by

applying lemma 3.1 on $C_i \setminus \partial C_i$ for $i = 1(1)k$. Thus $G_n(r_n)|(C_i \setminus \partial C_i)$ is connected for $r_n = (\frac{logn}{nV_d}(f_{x,y})^{-1})^{\frac{1}{d}}$ for all $i = 1(1)k$.

Suppose $x \in C_l$ and $y \in C_m$ $l \neq m$. Then $\exists\, r \in \mathbb{N}$ and all distinct $\{i_1 = l, i_2, .., i_r = m\}$ with $i_j \in 1(1)k$ so that $\exists$ a piecewise continuous path

$$p : x \to p_{i_1} \to C(C_{i_1}, C_{i_2}) \to p_{i_2} \to C(C_{i_2}, C_{i_3}).. \to p_{i_r} \to y$$

from x to y such that p_{i_j} is a path entirely within $C_{i_j} \setminus \partial C_{i_j}$ for $j = 1(1)r$ and $C(C_{i_j}, C_{i_{j+1}}) \in \partial C_{i_j} \cap \partial C_{i_{j+1}}$ for $j = 1(1)r - 1$. Clearly, maximum edge within p_{i_j} $\to$ $(\frac{logn}{nV_d}(f_{x,y})^{-1})^{\frac{1}{d}}$ for all $j = 1(1)r$ and all *connecting edges* i.e. $x \to p_{i_1}$, $p_{i_1} \to C(C_{i_1}, C_{i_2})$ etc. are $\leq (\frac{logn}{nV_d}(f_{x,y})^{-1})^{\frac{1}{d}}$.

Hence we have

$$\frac{nV_d}{logn}\, LLPD(x, C(C_{i_1}, C_{i_2}), \chi_n)^d \leq f_{x,y}^{-1},$$

$$\frac{nV_d}{logn}\, LLPD(C(C_{i_1}, C_{i_2}), C(C_{i_2}, C_{i_3}), \chi_n)^d \leq f_{x,y}^{-1},$$

. . .

$$\frac{nV_d}{logn}\, LLPD(C(C_{i_{r-1}}, C_{i_r}), y, \chi_n)^d \leq f_{x,y}^{-1}$$

almost surely. Rest follows from the ultrametric property of LLPD. □

Proposition 3.6

For $k > 1$ case, $\frac{nV_d}{logn}\, LLPD(x, y, \chi_n)^d \geq (f_{x,y})^{-1}$ *a.s* □

Proof 3.8

Suppose, if possible, $\exists$ a path p_n $x = x_0 \to x_1.. \to x_r = y$, $r \in \mathbb{N}$ and $x_1, .., x_{r-1} \in \chi_n$ such that it's maximum edge is $< (\frac{logn}{nV_d})(f_{x,y})^{-1})^{\frac{1}{d}}$, say $= (\frac{logn}{nV_d}(f_{x,y} + \delta)^{-1})^{\frac{1}{d}}$ for some $\delta > 0$. From proposition 3.2 we can only focus on $p_n^{'}$ $x_1 \to .. \to x_{r-1}$. From corollary

3.1 we have

$$\frac{n}{logn} V_d m^d(p_n^{'}) \to (\underset{t}{minf}(\gamma_t))^{-1}$$

i.e. $\underset{t}{minf}(\gamma_t) = f_{x,y} + \delta$. However, since γ_t is piecewise continuous curve joining x and y, $\underset{t}{minf}(\gamma_t) \leq f_{x,y} < f_{x,y} + \delta$ which is a contradiction. □

Example 3.2

Continuing from our previous example, we see that

$$LLPD(x, y, \chi_n) = LLPD(x, z, \chi_n) = LLPD(z, y, \chi_n)$$

for $z = (\tau, 0)$ and a candidate "LLPD path" as described in proposition 3.5 is $p : x \to p_1 \to z \to p_2$ with p_1 is a path entirely staying in $C_1 \setminus \partial C_1$ and p_2 is a path entirely staying in $C_2 \setminus \partial C_2$ almost surely as $n \to \infty$. Finally, $LLPD(x, y, \chi_n) \to (\frac{logn}{nV_2} f(z)^{-1})^{\frac{1}{2}}$. □

Remark 3.4

We have only used **A1** to prove proposition 3.2, which was crucial for our proof. However, proposition 3.2 holds even when **A1** is not true as long as $\underset{\epsilon \to 0^+}{lim} \frac{B_d(x,\epsilon) \cap (C_1 \setminus \partial C_1)}{V_d \epsilon^d} > 0$. Note that under **A1**, $\underset{\epsilon \to 0^+}{lim} \frac{B_d(x,\epsilon) \cap (C_1 \setminus \partial C_1)}{V_d \epsilon^d} = \frac{1}{2}$. Therefore, **A1** can be replaced by the following more weaker but necessary condition.

Definition: Following Arias-Castro (2011); Little et al. (2017, 2020), a d-dimensional Riemannian manifold S embedded in $\mathbb{R}^D$ is (κ, d, D) smooth if for all $x \in S$, we have

$$\underset{\epsilon \to 0^+}{lim} \frac{B_D(x, \epsilon) \cap S}{V_d \epsilon^d} > \kappa$$

. **A1*****:** λ level sets of M, $\partial M_\lambda := \{x : f(x) = \lambda\}$ is finite union of $d-1$ dimensional $(\kappa, d-1, d)$ *smooth* Riemannian manifolds for some $\kappa > 0$ and for all $\lambda \in (f_0, f^0)$.

3.4 Convergence rate of LLPD

In this section we shall show the following lower and upper tail convergence rate for $LLPD(x, y, \chi_n)$ for sufficiently small $\epsilon > 0$, η arbitrarily close to 1 and $c_1 > 0$ $c_2 > 0$ depending on M and f.

Theorem 3.2

$$\mathbb{P}\big(\sup_{x,y} \frac{LLPD(x, y, \chi_n)}{\big(\frac{logn}{n}\frac{1}{V_d} f_{x,y}{}^{-1}\big)^{\frac{1}{d}}} < 1 - \epsilon\big) \leq exp\big(-c_1 n^{\epsilon\eta d}(logn)^{-1}\big) \tag{3.4}$$

and

$$\mathbb{P}\big(\sup_{x,y} \frac{LLPD(x, y, \chi_n)}{\big(\frac{logn}{n}\frac{1}{V_d} f_{x,y}{}^{-1}\big)^{\frac{1}{d}}} > 1 + \epsilon\big) \leq c_2 n^{-\epsilon\eta d}(logn)^{-1} \tag{3.5}$$

□

3.4.1 Conjecture about LLPD's weak law

These convergence rates are very similar to the Gumbel type weak laws for connectivity distance or the longest k- nearest neighbor links of data χ_n for $k = O(1)$. Recall that standard Gumbel distribution Z is given by $\mathbb{P}(Z \leq x) = e^{-e^{-x}}\ \forall x \in \mathbb{R}$ For some specific type of compact manifolds (hyper-cubes and balls) and when f is uniformly distributed on them, the following results have been shown for the connectivity distance M_n as well as for the largest nearest neighbor links, we say $M_n^{'}$. For example, in Penrose (1997) it's shown that

$$nV_2 M_n^2 - logn \to Z \text{ for } d = 2 \text{ for both unit volume ball and square} \tag{3.6}$$

and in absence of boundary effect (i.e. working with Toroidal model or on a Riemannian manifold without boundary),

$$nV_dM_n^d - logn \to Z \text{ for unit volume cube in dimension } d. \tag{3.7}$$

Now a consequence of this weak law is the following:

$$\mathbb{P}(nV_dM_n^d - logn \leq x) \to e^{-e^{-x}}$$

$$\Rightarrow \mathbb{P}\big(\frac{M_n}{(\frac{logn}{nV_d})^{\frac{1}{d}}} \leq 1 + \frac{x}{dlogn}\big) \to e^{-e^{-x}}$$

$$\Rightarrow \mathbb{P}\big(\frac{M_n}{(\frac{logn}{nV_d})^{\frac{1}{d}}} > 1 + \epsilon\big) \to n^{-\epsilon d},\ \mathbb{P}\big(\frac{M_n}{(\frac{logn}{nV_d})^{\frac{1}{d}}} < 1 - \epsilon\big) \to e^{-n^{\epsilon d}}$$

We notice that these convergence rates are almost same as the expression in eqations (3.4) and (3.5) which we shall prove later in this sub-section. Although for a general $M \subset \mathbb{R}^d$ and f with $f_0 > 0$, the exact weak law of M_n is yet not known in literature, we can see that suitably transformed M_n is "Gumble"-like in some sense and we conjecture that

$$\alpha nM_n^d - \beta logn + \gamma loglogn - \delta \to Z$$

in-general, where α, β, γ and δ are non-negative constants dependent on d, M and f with $\frac{\beta}{\alpha} = \frac{1}{V_d f_0}$ and γ and δ possibly 0 in case M has no boundary. Furthermore, for LLPD we conjecture the following weak law

$$nf_{x,y}V_dM_n^d - logn \to Z \tag{3.8}$$

since near the boundary of $M_{x,y}^{'}$ defined earlier, locally the density is uniformly distributed with density $f_{x,y}$.

Now we shall prove the lower and upper tail convergence rate for $LLPD(x, y, \chi_n)$. For this we follow Penrose et al. (1999) in which the strong law for the connectivity

distance is derived. Since LLPD converges to the same limit as the connectivity distance restricted to $M'_{x,y}$ almost surely as $n \to \infty$, they'll have the same convergence rate as well.

3.4.2 Lower tail convergence rate

We prove the lower tail convergence rate as in equation (3.4) of theorem 3.2 below.

Proof 3.9

In Penrose (1999) from lemma 3.1, if $t < \frac{1}{f_0}$, $g_0 > f_0$ and $v > u > t$ such that $g_1 v < 1$ and $\delta > 0$ is such that $\delta^{\frac{1}{d}} + t^{\frac{1}{d}} < u^{\frac{1}{d}}$, then we have

$$\mathbb{P}\big(M_n \leq (\frac{tlogn}{nV_d})^{\frac{1}{d}}\big) \leq exp\big(-\sigma_n(\frac{f_0\delta}{2})n^{-g_0 v}\big)$$

where $\sigma_n \geq c_0(\frac{n}{logn})$ for some $c_0 > 0$, c_0 depends on M and f.

Therefore,

$$\mathbb{P}\big(M_n \leq (\frac{tlogn}{nV_d})^{\frac{1}{d}}\big) \leq exp\big(-c_0\frac{f_0\delta}{2}n^{1-g_0 v}(logn)^{-1}\big)$$

Now if we let $t = (1 - d\epsilon)\frac{1}{f_0}$, $\delta > 0$ such that $v = (1 - \eta d\epsilon)\frac{1}{g_1}$ for some $\eta < 1$, we get

$$\mathbb{P}\big(\frac{M_n}{(\frac{logn}{nV_d f_0})^{\frac{1}{d}}} \leq 1 - \epsilon\big) \leq exp(-c_0\frac{f_0\delta}{2}n^{\epsilon d\eta}(logn)^{-1}\big)$$

.

Finally, we can take η arbitrarily close to 1 while maintaining all the requirements. □

3.4.3 Upper tail convergence rate

We prove the upper tail convergence rate as mentioned in equation (3.5) of theorem 3.2 below.

Proof 3.10

From Penrose et al. (1999) Lemma 3.3 we obtain the following result. For $V_d t^d > \frac{1}{f_0}$, $q < t$, $n \approx m^{\alpha}$,$\beta = \alpha(f_0 V_d q^d - 1) > 1$,

$$\mathbb{P}\big(M_n > t(\frac{logn}{n})^{\frac{1}{d}}\big) \leq c_2 \frac{m^{-\beta}}{\alpha logm}$$

for some $c_2 > 0$ depending on M and f_0. Letting $t^d = (1 + d\epsilon)\frac{1}{V_d}\frac{1}{f_0}$ and $q^d = (1 + \eta d\epsilon)\frac{1}{V_d}\frac{1}{f_0}$ for some $\eta < 1$, we obtain

$$\mathbb{P}\big(\frac{M_n}{(\frac{logn}{nV_d f_0})^{\frac{1}{d}}} > 1 + \epsilon\big) \leq c_2 n^{-\epsilon d\eta}(logn)^{-1}$$

Again, we can choose η arbitrarily close to 1 while maintaining all the requirements.□

Chapter 4

Geometry of g-distance and LLPD

4.1 Organization and our contribution

In this chapter we analyze the clustering aspects of g-distance and LLPD by studying their geodesic and geometry those geodesics induce. In section 2, we provide the results showing the convergence of g-distance sample arc into population level geodesic. In section 3 and 4 we study the properties of this geodesic for a subset of $\mathbb{R}^d$ and for a Riemannian manifold respectively by deriving the differential equation governing it. It turns out that this geodesic captures both geometry and density information present in M and f respectively by bending towards higher density region of f while also capturing the local geometry of M.

4.2 Convergence of sample arc length to population geodesic

We provide the results showing the convergence of sample g-distance path to corresponding population geodesic γ_β , where $\beta \coloneqq \frac{b_1-1}{d}$, b_1 being the lowest power of g.

We first show that the sample geodesic arc length denoted by $|\overline{p_g(x, y, \chi_n)}|$ is

bounded almost surely. Formally if $p_g : x = x_0 \to x_1 \to x_2.. \to x_{k+1} = y$ is the path which minimize $L_g(x, y, \chi_n)$ then the sample geodesic arc length is denoted by

$$|\overline{p_g(x, y, \chi_n)}| \coloneqq \sum_{i=0}^{k} dist_1(x_i, x_{i+1}) \tag{4.1}$$

Proposition 4.1

Suppose $dist_1(x, y) \geq \delta > 0$ and $\min_M f(x) = f_0 > 0$. Then there exists positive constants l (depending on g and d) and $\theta(\delta)$ such that for all $n > n_0$

$$\begin{cases} P(|\overline{p_g(x, y, \chi_n)}| > l) \leq exp(-\theta(\delta) n^{\frac{1}{d+2b_1}}) & \text{if } b_1 > 1 \\ P(|\overline{p_g(x, y, \chi_n)}| > l) \leq exp(-\theta(\delta) n^{\frac{1}{d+b_2+1}}) & \text{if } b_1 = 1 \end{cases}$$

As a consequence we have $\limsup_{n\to\infty} |\overline{p_g(x, y, \chi_n)}| \leq l$ almost surely by Borel-Cantelli lemma. □

Proof 4.1

Follows similarly from a standard de-Poissonization argument and proposition 2.9 in Groisman et al. (2018). □

Theorem 4.1

If there exists an unique geodesic curve in M $\gamma^* \equiv \gamma_\beta(x, y)$ that minimize $dist_{b_1}(x, y)$, then $\overline{p_g(x, y, \chi_n)}$ converges uniformly to γ^* almost surely. □

Proof 4.2

Follows similarly from the arguments provided in corollary 2.10 in Groisman et al. (2018). In particular, it follows from the argument made there that

$$\mathbb{P}(d_F(\overline{p_g(x, y, \chi_n)}, \gamma_\beta(x, y)) \geq \epsilon) \leq exp(-\theta_0 \epsilon^2 n^{\alpha^*})$$

for some $\theta_0 > 0$ where $\alpha^* = \frac{1}{d+2b_1}$ or $\frac{1}{d+b_2+1}$ depending on $b_1 > 1$ or $b_1 = 1$ respectively, where d_F denote the Frechet distance between 2 curves. Thus the rate of convergence of sample arc length to population geodesic in Frechet distance is same as the rate of convergence of sample distance to the population distance. □

Proposition 4.2

If M is complete then for any x and y there exists a geodesic curve in M $\gamma^* \equiv \gamma_\beta(x, y)$ that minimize $dist_k(x, y)$. In particular, same conclusion holds when M is compact.□

Proof 4.3

Follows from Hopf-Rinow theorem. See O'neill (1983) theorem 5.21, page 138 for more details. □

However, as we shall see later, in-general the uniqueness of the geodesic $\gamma_\beta(x, y)$ is not guaranteed.

4.3 Geometry of population geodesic $\{\gamma_\beta\}_{\beta\geq 0}$ in a subset of $\mathbb{R}^d$

Let $M \subset \mathbb{R}^d$. In this section we shall discuss the geometry of the family of curves $\{\gamma_\beta(t)\}_{\beta\geq 0}$ between 2 fixed points x and y in M defined by the minimizer of

$$\int_0^1 f^{-\beta}(\gamma_t)|\gamma_t'|dt$$

where the minimum is taken over all piecewise smooth curves $\gamma : [0, 1] \to M$ such that $\gamma(0) = x$ and $\gamma(1) = y$, and derive a system of differential equations governing it.

It turns out that the geometry of $\{\gamma_\beta(t)\}_{\beta\geq 0}$ in $\mathbb{R}^d$ is closely related to the Fermat's principle of least time for ray paths. Note that we assume M to be complete and

additionally for any x and y, we need to have $\gamma_\beta(t)$ strictly inside M (no boundary effect).

4.3.1 Fermat's principle of least time

The *optical length* of a path $r(s)$ taken by a ray of light in passing from point A to B in $M \subset \mathbb{R}^d$ is defined by

$$A \coloneqq \int_A^B n(r(s))ds \tag{4.2}$$

where $n(r)$ is the index of refraction at the point $r \in \mathbb{R}^d$ and

$$ds^2 = dr(s)\,.\,dr(s) \tag{4.3}$$

ds is the element of arc length along the ray path $r(s)$ through that point.

In this context, *Fermat's principle* (1662) of least time for ray optics states that the path $r(s)$ taken by a ray of light passing from A to B in $M \subset \mathbb{R}^d$ leaves the optical length stationary under variations in a family of nearby paths $r(s,\epsilon)$ depending on a parameter ϵ. That is, the path $r(s)$ satisfies

$$\delta A = 0 \;\text{ with } \delta A \coloneqq \frac{d}{d\epsilon}|_{\epsilon=0} \int_A^B n(r(s,\epsilon))ds \tag{4.4}$$

where the deviations of the ray path $r(s,\epsilon)$ from $r(s)$ are assumed to vanish when $\epsilon = 0$, and to leave it's endpoint fixed.

Formally, consider the possible ray paths in Fermat's principle leading from point A to point B in $M \subset \mathbb{R}^d$ as belonging to a family of C^2 curves $r(s,\epsilon) \in M$ depending smoothly on a parameter ϵ in an interval that includes $\epsilon = 0$. This ϵ-family of paths $r(s,\epsilon)$ defines a set of smooth transformations of the ray path $r(s)$. These transformations are taken to satisfy $r(s,0) = r(s)$, $r(s_A,\epsilon) = r(s_A)$ and $r(s_B,\epsilon) =$

$r(s_B)$. That is, $\epsilon = 0$ is the identity transformation of the ray path and its endpoints are left fixed. An infinitesimal variation of the path $r(s)$ is denoted by $\delta r(s)$, defined by the *variational derivative*,

$$\delta r(s) \coloneqq \frac{d}{d\epsilon}|_{\epsilon=0} r(s,\epsilon)$$

where the fixed endpoint conditions imply $\delta r(s_A) = \delta r(s_B) = 0$.

4.3.2 Euler-Lagrange equation for axial ray optics in $\mathbb{R}^d$

Most optical instruments are designed to possess a line of sight (or primary direction of propagation of light) called the optical axis. Choosing coordinates so that the x_1-axis coincides with the optical axis expresses the arc-length element ds in terms of the increment along the optical axis, dx_1, as

$$ds = (\sum_{i=1}^{d}(dx_i)^2)^{1/2} = (1 + x_2'^2 + .. + x_d'^2)^{1/2} dx_1 = \boldsymbol{\gamma} dx_1, \tag{4.5}$$

where $x_i^{'} \coloneqq \frac{dx_i}{dx_1}$ for $i = 2, 3, .., d$ and $\boldsymbol{\gamma} \coloneqq \frac{ds}{dx_1}$.

For such optical instruments, the optical length in Fermat's principle can be expressed as

$$A \coloneqq \int_{x_{1_A}}^{x_{1_B}} L(x_1, .., x_d, x_2^{'}, .., x_d^{'}) dx_1 \tag{4.6}$$

Now we introduce the notion of *tangent bundle* as follows

The coordinates $(x_2, .., x_d, x_2^{'}, .., x_d^{'}) \in \mathbb{R}^{d-1} \times \mathbb{R}^{d-1}$ designate points along the ray path through the configuration space and the tangent space of possible vectors along the ray trajectory. The position on a ray passing through an image plane at a fixed value of x_1 is denoted by $(x_2, x_3, .., x_d) \in \mathbb{R}^{d-1}$. The notation $(x_2, .., x_d, x_2^{'}, .., x_d^{'}) \in T\mathbb{R}^{d-1} \backsimeq \mathbb{R}^{d-1} \times \mathbb{R}^{d-1}$ for the combined space of positions

and tangent vectors designates the *tangent bundle* of $\mathbb{R}^{d-1}$. The space $T\mathbb{R}^{d-1}$ consists of the union of all the position vectors $(x_2, x_3, .., x_d) \in \mathbb{R}^{d-1}$ and all the possible tangent vectors $(x_2^{'}, .., x_d^{'}) \in \mathbb{R}^{d-1}$ at each position $(x_2, x_3, .., x_d)$.

Thus, the integrand in the optical action A in (4.6) can be re-expressed as $L : T\mathbb{R}^{d-1} \times \mathbb{R} \to \mathbb{R}$ by

$$L(x_1, .., x_d, x_2^{'}, .., x_d^{'}) = n(x_1, .., x_n)(1 + x_2^{'2} + .. + x_d^{'2})^{1/2} = n\boldsymbol{\gamma} \tag{4.7}$$

in this tangent bundle formulation where n is the refractive index.

We may write the coordinates equivalently as $(x_1, x_2, .., x_d) = (x_1, \boldsymbol{q})$ where $\boldsymbol{q} = (x_2, .., x_d)$ is a vector with components in the hyper-plane orthogonal to the optical axis x_1. Thus, $\boldsymbol{\gamma} = (1 + |\boldsymbol{q}^{'}|^2)^{1/2}$ using this notation.

The possible ray paths from A to B in space may be parameterized for axial ray optics as a family of C^2 curves $\boldsymbol{q}(x_1, \epsilon) \in \mathbb{R}^{d-1}$ depending smoothly on a real parameter ϵ in an interval that includes $\epsilon = 0$. The ϵ-family of paths $\boldsymbol{q}(x_1, \epsilon)$ defines a set of smooth transformations of the ray path $\boldsymbol{q}(x_1)$. These transformations are taken to satisfy

$$\boldsymbol{q}(x_1, 0) = \boldsymbol{q}(x_1), \quad \boldsymbol{q}(x_{1_A}, \epsilon) = \boldsymbol{q}(x_{1_A}), \ \boldsymbol{q}(x_{1_B}, \epsilon) = \boldsymbol{q}(x_{1_B})$$

So $\epsilon = 0$ is the identity transformation of the ray path and its endpoints are left fixed. We define the variation of the optical action A by

$$\begin{aligned} \delta A &= \delta \int_{x_{1_A}}^{x_{1_B}} L(\boldsymbol{q}(x_1), \boldsymbol{q}^{'}(x_1))dx_1 \\ &\coloneqq \frac{d}{d\epsilon}|_{\epsilon=0} \int_{x_{1_A}}^{x_{1_B}} L(\boldsymbol{q}(x_1, \epsilon), \boldsymbol{q}^{'}(x_1, \epsilon))dx_1 \end{aligned} \tag{4.8}$$

We shall denote

$$\delta\boldsymbol{q}(x_1) \coloneqq \frac{d}{d\epsilon}|_{\epsilon=0}\boldsymbol{q}(x_1, \epsilon)$$

Thus, Fermat's principle implies

$$0 = \delta A = \delta \int L(x_1, \boldsymbol{q}, \boldsymbol{q}^{'}) dx_1$$

After integration by parts we obtain,

$$= \int (\frac{\partial L}{\partial \boldsymbol{q}} - \frac{d}{dx_1}\frac{\partial L}{\partial \boldsymbol{q}^{'}})\, \delta \boldsymbol{q} dx_1 + [\frac{\partial L}{\partial \boldsymbol{q}^{'}}\, \delta \boldsymbol{q}]_{x_{1_A}}^{x_{1_B}}.$$

The endpoint terms vanish because $\delta \boldsymbol{q}(x_{1_A}) = 0 = \delta \boldsymbol{q}(x_{1_B})$. That is, the variation in the ray path must vanish at A and B at x_{1_A} and x_{1_B} along the optical axis x_1. Since $\delta \boldsymbol{q}$ is otherwise arbitrary, Fermat's principle reduces to these system of $d-1$ equations known as **Euler-Lagrange equations** given by.

$$\frac{\partial L}{\partial \boldsymbol{q}} - \frac{d}{dx_1}\frac{\partial L}{\partial \boldsymbol{q}^{'}} = \boldsymbol{0} \tag{4.9}$$

After a short algebraic manipulation using the explicit form of $L = n\gamma$, and $\frac{\partial L}{\partial \boldsymbol{q}} = \gamma \frac{\partial n}{\partial \boldsymbol{q}}$ and $\frac{d}{dx_1}\frac{\partial L}{\partial \boldsymbol{q}^{'}} = \frac{d}{dx_1} n \frac{\partial \gamma}{\partial \boldsymbol{q}^{'}} = \frac{d}{dx_1} n \frac{\boldsymbol{q}^{'}}{\gamma}$ in equation (4.9), we obtain set of Eikonal equations

$$\frac{d}{dx_1}(\frac{n}{\gamma}\boldsymbol{q}^{'}) = \gamma \frac{\partial n}{\partial \boldsymbol{q}} \tag{4.10}$$

We shall use (4.10) to derive the system of differential equations for $\{\gamma_\beta(t)\}_{\beta\geq 0}$ in $\mathbb{R}^d$.

4.3.3 Differential equation for $\{\gamma_\beta(t)\}_{\beta\geq 0}$ in $\mathbb{R}^d$

Theorem 4.2

$\gamma_\beta(t)$ satisfies the system of $d-1$ ordinary differential equations(ODEs)

$$\boldsymbol{q}^{''} = \beta\gamma^2 \frac{\partial g}{\partial x_1}\left(\boldsymbol{q}^{'} - \frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}\right) \tag{4.11}$$

with the boundary condition of $\gamma_\beta(0) = x$ and $\gamma_\beta(1) = y$ where $g = log\, f$. □

Proof 4.4

In this case the refractive index n is given by $f^{-\beta}$. i.e. $L = f^{-\beta}\gamma$.

Now we derive the differential equations governing $\{\gamma_\beta(t)\}_{\beta\geq 0}$ in $\mathbb{R}^d$ using Eikonal equations

$$\frac{d}{dx_1}(\frac{n}{\gamma}x_j^{'}) = \gamma\frac{\partial n}{\partial x_j} \quad \text{for } j = 2, 3, ..d.$$

We have,

$$\gamma\frac{\partial n}{\partial x_j} = \gamma\frac{\partial f^{-\beta}}{\partial x_j} = \gamma(-\beta)\frac{f^{-\beta}}{f}\frac{\partial f}{\partial x_j} = -\beta\frac{L}{f}\frac{\partial f}{\partial x_j}.$$

And

$$\frac{d}{dx_1}(\frac{n}{\gamma}x_j^{'}) = x_j^{''}\frac{L}{\gamma^2} + x_j^{'}\frac{d}{dx_1}\frac{n}{\gamma}.$$

Now,

$$\frac{d}{dx_1}\frac{f^{-\beta}}{\gamma} = \frac{-\beta\frac{f^{-\beta}}{f}\frac{df}{dx_1}\gamma - f^{-\beta}\frac{\sum_{i=2}^d x_i^{'}x_i^{''}}{\gamma}}{\gamma^2}$$

$$= -\frac{\beta\frac{L}{f}\frac{df}{dx_1} + \frac{L}{\gamma^2}\sum_{i=2}^d x_i^{'}x_i^{''}}{\gamma^2}$$

So finally we obtain

$$-\beta\frac{L}{f}\frac{\partial f}{\partial x_j} = x_j^{''}\frac{L}{\gamma^2} - x_j^{'}(\frac{\beta\frac{L}{f}\frac{df}{dx_1} + \frac{L}{\gamma^2}\sum_{i=2}^d x_i^{'}x_i^{''}}{\gamma^2})$$

which after some rearranging and division by L yields

$$\Rightarrow \frac{x_j^{''}}{\gamma^2} - x_j^{'}\frac{1}{\gamma^4}\sum_{i=2}^d x_i^{'}x_i^{''} = -\frac{\beta}{f}(\frac{\partial f}{\partial x_j} - \frac{1}{\gamma^2}x_j^{'}(\frac{\partial f}{\partial x_1} + \sum_{i=2}^d \frac{\partial f}{\partial x_i}x_i^{'}))$$

The left hand side is

$$\frac{1}{\gamma^4}[x_j^{''}(\gamma^2 - x_j^{'2}) + \sum_{i\neq j} x_i^{''}(-x_i^{'}x_j^{'})]$$

whereas the right hand side is

$$\frac{\beta}{f\gamma^2}[x_j^{'}(\frac{\partial f}{\partial x_1}+\sum_{i=2}^{d}\frac{\partial f}{\partial x_i}x_i^{'})-\gamma^2\frac{\partial f}{\partial x_j}]$$

$$=-\frac{\beta}{f\gamma^2}[\frac{\partial f}{\partial x_1}(-x_j^{'})+\frac{\partial f}{\partial x_j}(\gamma^2-x_j^{'2})+\sum_{i=2,i\neq j}^{d}\frac{\partial f}{\partial x_i}(-x_i^{'}x_j^{'})]$$

We define the operator

$$D\coloneqq\gamma^2 I-\boldsymbol{q}^{'}\boldsymbol{q}^{'T}=\begin{bmatrix}\gamma^2-x_2^{'2} & -x_2^{'}x_3^{'} & . & . & -x_2^{'}x_d^{'}\\ -x_2^{'}x_3^{'} & \gamma^2-x_3^{'2} & . & . & -x_3^{'}x_d^{'}\\ . & . & . & . & .\\ . & . & . & . & .\\ -x_2^{'}x_d^{'} & -x_3^{'}x_d^{'} & . & . & \gamma^2-x_d^{'2}\end{bmatrix}$$

It can be verified that D is non-singular, and it's inverse operator is

$$D^{-1}=\frac{1}{\gamma^2}(I+\boldsymbol{q}^{'}\boldsymbol{q}^{'T})=\frac{1}{\gamma^2}\begin{bmatrix}1+x_2^{'2} & x_2^{'}x_3^{'} & . & . & x_2^{'}x_d^{'}\\ x_2^{'}x_3^{'} & 1+x_3^{'2} & . & . & x_3^{'}x_d^{'}\\ . & . & . & . & .\\ . & . & . & . & .\\ x_2^{'}x_d^{'} & x_3^{'}x_d^{'} & . & . & 1+x_d^{'2}\end{bmatrix}$$

Then, the system of $d-1$ equations can be rewritten is the following compact form

$$\frac{1}{\gamma^4}D\boldsymbol{q}^{''}=-\frac{\beta}{f\gamma^2}[D\frac{\partial f}{\partial \boldsymbol{q}}-\frac{\partial f}{\partial x_1}\boldsymbol{q}^{'}]$$

$$\Rightarrow \boldsymbol{q}^{''}=\frac{\beta\gamma^2}{f}[\frac{\partial f}{\partial x_1}D^{-1}\boldsymbol{q}^{'}-\frac{\partial f}{\partial \boldsymbol{q}}]$$

$$\Rightarrow \boldsymbol{q}^{''}=\beta\gamma^2\frac{\partial g}{\partial x_1}[\boldsymbol{q}^{'}-\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}] \text{ where } g=log\, f$$

since $D^{-1}\boldsymbol{q}^{'}=\boldsymbol{q}^{'}$.

Thus, $\gamma_\beta\coloneqq(x_1,x_2(x_1),..,x_d(x_1))$ satisfy the system of ODEs $\boldsymbol{q}^{''}=\beta\gamma^2\frac{\partial g}{\partial x_1}[\boldsymbol{q}^{'}-\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}]$ with the boundary condition of starting at x and ending at y. □

4.3.4 γ_β: An interpolation between straight-line and field-line

Now we introduce the concept of flow-lines/field-lines for a vector field which describe the motion of particles under the influence of a vector field.

Formally, a vector field on $\mathbb{R}^d$ is a function $F : M \subset \mathbb{R}^d \to \mathbb{R}^d$ that assigns to each point x in its domain M an d-dimensional vector $F(x)$. Let $x(t) : \mathbb{R} \to \mathbb{R}^d$ be a path indexed by time t. Since we want to capture the motion of a particle under the influence of the vector field, we require that the velocity of that particle is proportional to the vector field at any time-point i.e. **field-lines** are defined by the following equation

$$x^{'}(t) \coloneqq \alpha F(x(t)) \;\; \alpha \neq 0$$

Thus, for the vector fields induced by the gradient of a function f, we have

$$\frac{dx_i(t)}{dt} = \alpha \frac{\partial f(x(t))}{\partial x_i(t)} \;\; \text{for } \alpha \neq 0 \text{ and } i = 1, 2, ..d$$

For $\alpha > 0$, (without loss of generality we can take $\alpha = 1$) we call the flow-lines as *gradient ascent flow* and for $\alpha < 0$ (without loss of generality we can take $\alpha = -1$) we call them *gradient descent flow.*

or equivalently, in our notation

$$x_i^{'} = \frac{\frac{\partial f}{\partial x_i}}{\frac{\partial f}{\partial x_1}} \;\; \text{for } i = 2, 3, ..d$$

$$\Rightarrow \boldsymbol{q}^{'} = \frac{\frac{\partial f}{\partial \boldsymbol{q}}}{\frac{\partial f}{\partial x_1}}$$

Thus the set of differential equations

$$\boldsymbol{q}^{'} - \frac{\frac{\partial f}{\partial \boldsymbol{q}}}{\frac{\partial f}{\partial x_1}} = \mathbf{0} \tag{4.12}$$

uniquely represent *any* gradient ascent/descent flow-line.

Thus (4.11) can be interpreted as follows. Here is a discrete version of (4.11)

$$\frac{\boldsymbol{q}'_{n+1} - \boldsymbol{q}'_{n}}{\Delta x_1} = \beta(1 + |\boldsymbol{q}'_{n}|^2)\frac{\partial g}{\partial x_1}\,[\boldsymbol{q_n}' - \frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}|_{\boldsymbol{q}=\boldsymbol{q_n}}]$$

$$\Rightarrow \boldsymbol{q}'_{n+1} = \boldsymbol{q}'_{n} + \beta\,\Delta x_1(1 + |\boldsymbol{q}'_{n}|^2)\frac{\partial g}{\partial x_1}\,[\boldsymbol{q}'_n - \frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}|_{\boldsymbol{q}=\boldsymbol{q_n}}] \tag{4.13}$$

Note that we get the new velocity $\boldsymbol{q}'_{n+1}$ from the old velocity $\boldsymbol{q}'_n$ by drifting towards the velocity induced by the gradient of g given by $\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}$. Amount of drift is controlled by β. If β is small, $\boldsymbol{q}'_{n+1} \approx \boldsymbol{q}'_{n}$ (or $\boldsymbol{q}'' \approx \boldsymbol{0}$). Note that $\boldsymbol{q}'' = \boldsymbol{0}$ ($\beta = 0$ case) represent the straight line joining x and y in $\mathbb{R}^d$. While when β is large, $\boldsymbol{q}'$ is dominated by the drift generated by gradient of g. Finally $\beta \to \infty$ case reduces to $\boldsymbol{q}' = \frac{\frac{\partial f}{\partial \boldsymbol{q}}}{\frac{\partial f}{\partial x_1}}$ i.e. equation of a field line/union of field lines joining x and y.

4.3.5 γ_∞ and LLPD

There is ambiguity regarding *which* field line/union of field lines joining x and y does $\beta \to \infty$ case correspond to, since in-general there can be more than one field line/union of field lines joining 2 points in $\mathbb{R}^d$. It is very important to understand the γ_∞ case since as we have seen in the previous sub-section, γ_β is an interpolation between γ_0 (straight line) and γ_∞. In this sub-section we describe the geometry of γ_∞ and discuss it's connection with LLPD.

Following Chen et al. (2017); Gerber et al. (2013), under the assumption that

A1: f is a Morse-Smale function,

except for a lower-dimensional set L, $M \subset \mathbb{R}^d$ is the union of finitely many ***basin of attractions*** (Loosely speaking basin of attraction A_i corresponding to a maxima m_i of f is a d dimensional descending manifold which consist of all points whose ascending flow lines converge to m_i.) say $A_1, A_2, .., A_K$ corresponding to local maximas

$m_1, m_2, .., m_K$ of f, i.e.

$$M = (\cup_{i=1}^{K} A_i) \cup L \tag{4.14}$$

where L consists union of lower (from 0 to $d-1$) dimensional descending manifolds and forms the boundary of these k basin of attractions. See section 2 of Chen et al. (2017) for definitions of ascending and descending manifolds, basin of attractions and in-general for more details on Morse theory. In chapter 5 these definitions are provided as well.

Under the additional assumption that

A2: $\int_L f = 0$.

we can focus on 2 distinct cases, namely

1. x and y are from same basin of attraction.
2. x and y are from different basin of attractions.

Remark 4.1

Note that by Kupka-Smale theorem (see e.g. theorem 6.6 in Banyaga and Hurtubise (2013)) Morse-Smale functions are dense in the collection of smooth functions and for any continuous f, $\int_L f = 0$. Thus both **A1** and **A2** are very reasonable assumptions.□

Proposition 4.3

Under **A1, A2** all basin of attractions are open sets. □

Proof 4.5

Consider $x \in A_i$, an open set $U \subset A_i$ around m_i and a sequence of points $\{a_i\}$ along the ascending field line of x with $a_0 = x$, $k(a_i) = a_{i+1}$ converging to m_i. Thus for

sufficiently large N, $k^N(x) \in U$. Conclusion follows after observing that under **A2** pre-image of U under k^N, V is also an open set containing x and clearly is in A_i since $k^N(V) = U$. □

Proposition 4.4

If both x and y are from same basin of attraction, say A_i; then there is a unique field line joining x and m_i, and another unique field line joining y and m_i. In this case γ_∞ joining x and y is simply the union of the field line joining x and m_i and the field line joining m_i and y (unless x and y happen to be on the same field line or one of them is m_i, in which case γ_∞ joining x and y is simply the part of the field line between x and y joining x/y and m_i). □

Proof 4.6

Uniqueness of the (ascending) field line joining x (or y) and m_i follows from the property of field lines. Now, recall that γ_β minimizes $(\int_a^b f^{-\beta}(\gamma_t)|\gamma_t^{'}|dt)^{\frac{1}{\beta}}$ for any $0 \leq a < b \leq 1$ subject to $\gamma_0 = x$, $\gamma_1 = y$. So $\beta \to \infty$ case would correspond to union of field lines γ_t such that $\gamma_0 = x$, $\gamma_1 = y$ and it maximizes $\min\limits_t f(\gamma_t)$ for any $0 \leq a < t < b \leq 1$ since

$$\lim_{\beta\to\infty} \min_{\gamma} \Big(\int_a^b f^{-\beta}(\gamma_t)|\gamma_t^{'}|dt\Big)^{\frac{1}{\beta}} = \min_{\gamma} \max_{a\leq t\leq b} f^{-1}(\gamma_t) = \Big(\max_{\gamma} \min_{a\leq t\leq b} f(\gamma_t)\Big)^{-1}$$

W.l.g. we assume that $x \neq y \neq m_i$ and x or y do not lie on the same field line joining one of them and m_i. Then the field line starting at x need to go towards the maxima m_i (the descending field line from x may go towards some points in L or continue indefinitely.) i.e. in the ascending direction because otherwise it won't maximize $\min\limits_t f(\gamma_t)$ for any $0 \leq t < b$ and similarly the field line starting at y also need to go towards maxima m_i. We denote the field line joining x, m_i and y by

$$x \to m_i \to y$$

□

To understand what happens to γ_∞ when x and y are from different basin of attractions, say $x \in A_i$ and $y \in A_j$ $i \neq j$, we need to introduce some definitions.

We define for 2 connected basin of attractions (2 basin of attractions A_i and A_j are *connected* if they share common boundary, i.e. $\partial A_i \cap \partial A_j \neq \emptyset$.) their ***conductance value*** by

$$\text{Conductance } (A_i, A_j) \coloneqq \max_{\gamma} \min_{t} f(\gamma_t)$$

such that $\gamma_0 = m_i$ and $\gamma_1 = m_j$, γ is union of field lines and stays entirely within $(A_i \cup A_j \cup (\partial A_i \cap \partial A_j)$. Any point on γ where this value is attained is called a ***conductance point*** correspond to A_i and A_j, in our notation $C(A_i, A_j)$.

Proposition 4.5

For any 2 connected basin of attractions A_i and A_j $i \neq j$, there is at least one conductance point and all conductance points are in $\partial A_i \cap \partial A_j$, in particular are in L. □

Proof 4.7

W.l.g. suppose a conductance point c is strictly inside A_i. Then we can pick a point d on the descending field line starting from c and strictly inside A_i (since A_i's are open sets) such that $f(d) < f(c)$- which is a contradiction on the assumption that c is a conductance point. Thus, all conductance points between A_i and A_j, if they exist, must be on $\partial A_i \cap \partial A_j$. Since $\partial A_i \cap \partial A_j$ is a closed set, the conductance value will be attained somewhere on it. □

To generalize the notion of conductance value and conductance point for any 2 basin of attractions (which may or may not be connected), we define the ***bottleneck value***

for any 2 different basin of attractions A_i and A_j by

$$\text{Bottleneck } (A_i, A_j) \coloneqq \max_{\gamma} \min_{t} f(\gamma_t) \tag{4.15}$$

such that γ is union of field lines, $\gamma_0 = m_i$ and $\gamma_1 = m_j$, and any point on γ joining m_i and m_j where this value is attained is called a ***bottleneck point*** correspond to A_i and A_j.

Note that, instead of defining conductance and bottleneck value/points for 2 different basin of attractions A_i and A_j, we can define them for corresponding maximas m_i and m_j as well. In chapter 5, we shall use Bottleneck (A_i, A_j) and Bottleneck (m_i, m_j) interchangeably.

Proposition 4.6

For any 2 basin of attractions A_i and A_j $i \neq j$, there is at least one bottleneck point and all bottleneck points are in L. □

Proof 4.8

We construct a weighted graph $G = (V, E)$ with the vertices as the basin of attractions $V = \{A_1, .., A_k\}$ and edge between A_i and A_j if and only if A_i and A_j are connected, in which case, the edge weight is given by

$$E(A_i, A_j) = \text{Conductance } (A_i, A_j)$$

Note that under this formulation, Bottleneck $(A_i, A_j) = \max_{p} \min_{k} E(A_k, A_{k+1})$ where p is any path of the form $A_i = A_0 \to A_1 .. \to A_{r-1} \to A_r = A_j$ in G, $r \in \mathbb{N}$ can vary. Indeed, consider a union of field lines from m_i to m_j such that bottleneck value between A_i and A_j is attained. From proposition 4.5 it follows that any union of field lines that directly goes from m_k to m_l via the common boundary of A_k and A_l for some $k \neq l$ will have f value at most Conductance (A_k, A_l) at the boundary.

Thus, it is sufficient to consider those union of field lines only which attain respective conductance values at each direct transition between basin of attractions. Each of these candidate union of field lines can be represented as a path in G from A_i to A_j and by definition of conductance value, the minimum edge of such a path in G correspond to the minimum f value of the union of field line joining m_i and m_j. Thus, from definition of bottleneck value, our conclusion follows. Since for any 2 different basin of attractions, at least one conductance point exists and is in L, at least one bottleneck point exists as well between A_i and A_j and is in L.

Proposition 4.7

Suppose in G, $A_i \to A_{i_1}.. \to A_{i_r} \to A_j$ is any path whose minimum edge is the bottleneck value between A_i and A_j. Then assuming $x \neq m_i$, $y \neq m_j$, x not belonging in the field line joining m_i and $C(A_{i_1}, A_i)$ and y not belonging in the field line joining m_j and $C(A_{i_r}, A_j)$, in our notation as introduced in proposition 4.4, the trajectory of $\lim_{\beta\to\infty} \gamma_\beta$ can be described as

$$x \to m_i \to C(A_i, A_{i_1}) \to m_{i_1} \to C(A_{i_1}, A_{i_2}) \to m_{i_2}.. \to C(A_{i_r}, A_j) \to m_j \to y \quad \square$$

where all arrows indicate the unique field line joining the point before and after them.

Proof 4.9

Follows from proposition 4.4 and 4.6. $\square$

Now we provide a side result relating the arc lengths of γ_0, γ_β and γ_∞ which will be used in chapter 5 and is of independent interest.

Lemma 4.1

Let $|\gamma_\beta(x, y)|$ denote the arc length of γ_β joining x and y for $0 \leq \beta < \infty$. Then

$$|\gamma_0(x, y)| = |x - y| \leq |\gamma_\beta(x, y)| \leq |\gamma_\infty(x, y)|$$

where $|\gamma_\infty(x,y)|$ denote the arc length of almost surely unique field line/union of field lines joining x and y as in proposition 4.4 and 4.7. □

Proof 4.10

It is sufficient to prove this for a small change in x_1 coordinate, by $\triangle x_1$ say. w.l.g. we can also assume $\boldsymbol{q}^{'}{}_0 = \mathbf{0}$ by suitably transforming our coordinates. Then from (4.13) we obtain

$$\boldsymbol{q}_1^{'} = -\beta \triangle x_1 \frac{\partial g}{\partial x_1} [\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}|_{\boldsymbol{q}=\boldsymbol{q}_0}]$$

$$\Rightarrow 0 \leq |\boldsymbol{q}_1^{'}| \leq |\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}|_{\boldsymbol{q}=\boldsymbol{q}_0}|$$

since for a fixed β one can choose $\triangle x_1$ sufficiently small so that $|\beta \triangle x_1 \frac{\partial g}{\partial x_1}| < 1$. Conclusion follows by observing $|\gamma_0| = \triangle x_1$, $|\gamma_\infty| = \triangle x_1 \, (1 + |\frac{\frac{\partial g}{\partial \boldsymbol{q}}}{\frac{\partial g}{\partial x_1}}|_{\boldsymbol{q}=\boldsymbol{q}_0}|^2)^{1/2}$ and $|\gamma_\beta| = \triangle x_1 \, (1 + |\boldsymbol{q}_1^{'}|)^{1/2}$.

This, in particular, shows that if x and y are from same basin of attraction A_i,

$$|\gamma_\beta(x,y)| \leq |\gamma_\infty(x,m_i)| + |\gamma_\infty(y,m_i)| \leq |\gamma_\infty(A_i)|$$ □

where $|\gamma_\infty(A_i)| \coloneqq \underset{x,y\in A_i}{max} |\gamma_\infty(x,y)|$.

4.3.6 Some examples: Gaussian mixture models in $\mathbb{R}^2$

Now we explicitly write down (4.11) for Gaussian Mixture models (GMM) in $\mathbb{R}^2$ as follows

Let

$$f \sim \sum_{i=1}^{n} w_i N_2(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)$$

where $w_i \geq 0$ and $\sum_{i=1}^{n} w_i = 1$.

i.e.

$$f(x_1,x_2)=f(\boldsymbol{x})=\sum_{i=1}^{n}w_if_i(\boldsymbol{x})=\sum_{i=1}^{n}w_i\,\frac{1}{2\pi|\boldsymbol{\Sigma}_i|^{1/2}}exp\big(-\frac{1}{2}(x-\boldsymbol{\mu}_i)^T\boldsymbol{\Sigma}_i^{-1}(x-\boldsymbol{\mu}_i)\big)$$

where $\boldsymbol{\mu}_i=\begin{pmatrix}\mu_{1i}\\ \mu_{2i}\end{pmatrix}$ and $\boldsymbol{\Sigma}_i=\begin{pmatrix}\sigma_{1i}^2 & \rho_i\sigma_{1i}\sigma_{2i}\\ \rho_i\sigma_{1i}\sigma_{2i} & \sigma_{2i}^2\end{pmatrix}$ for $i=1(1)n$.
Then (4.11) reduce to

$$x_2^{''}=\beta(1+x_2^{'2})\sum_{i=1}^{n}u_iv_i$$

where

$$u_i(x_1,x_2)=\frac{w_if_i}{f}$$

and

$$v_i(x_1,x_2,x_2^{'})=\frac{1}{1-\rho_i^2}[(x_2-\mu_{2i})(\frac{1}{\sigma_{2i}^2}+\frac{\rho_ix_2^{'}}{\sigma_{1i}\sigma_{2i}})-(x_1-\mu_{1i})(\frac{\rho_i}{\sigma_{1i}\sigma_{2i}}+\frac{x_2^{'}}{\sigma_{1i}^2})].$$

Note that u_i represent the relative weight of f_i at a specific location and v_i is the equation of f_i's field line.

Now we provide some examples showing γ_β for different β using GMMs in $\mathbb{R}^2$.

Example 4.1

Let

$$f\sim N((0,0),\boldsymbol{I}_2)$$

in $\mathbb{R}^2$. The differential equation for γ_β here is

$$x_2^{''}=\beta(1+x_2^{'2})(x_2-x_1x_2^{'}).$$

Let $x=(-1,1)$ and $y=(1,1)$. Clearly in this case the field lines are given by $\{x_2=mx_1:m\in\mathbb{R}\}$. Thus as $\beta\to\infty$, γ_β approaches the union of straight lines

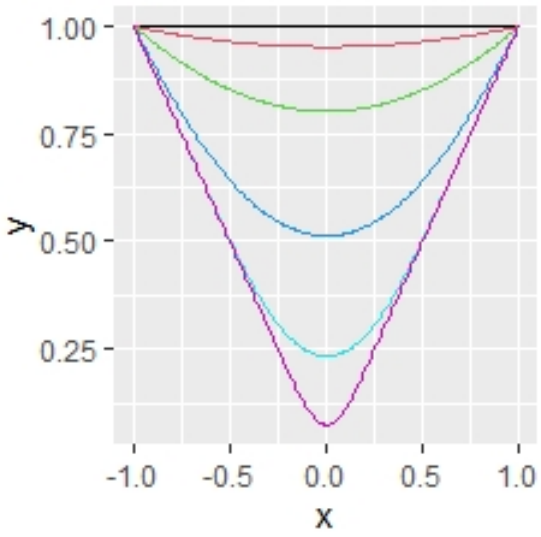

Figure 4.1: γ_β joining $(-1,1)$ and $(1,1)$ for $\beta = 0,\ .1,\ .5,\ 2,\ 10$ and 100 in example 4.1. Notice that as β becomes larger γ_β approaches the straight line joining $(-1,1) \rightarrow (0,0) \rightarrow (1,1)$.

joining $(-1,1) \rightarrow (0,0) \rightarrow (1,1)$ (in this case there's only 1 basin of attraction, namely entire $\mathbb{R}^2$). On the other hand, if $x = (-1,-1)$ and $y = (1,1)$ then γ_β will be the straight line joining them for any $\beta \geq 0$ since in this case γ_0 and γ_∞ coincide. □

Example 4.2

Let

$$f \sim \frac{1}{2}N((0,0), \boldsymbol{I}_2) + \frac{1}{2}N((3,0), \boldsymbol{I}_2)$$

in $\mathbb{R}^2$. The differential equation here is

$$x_2'' = \beta(1 + x_2'^2)\big[(x_2 - x_1 x_2') \frac{exp - (\frac{1}{2}(x_1^2 + x_2^2))}{exp - (\frac{1}{2}(x_1^2 + x_2^2)) + exp - (\frac{1}{2}((x_1 - 3)^2 + x_2^2))} +$$

$$(x_2 - 3 - x_1 x_2') \frac{exp - (\frac{1}{2}((x_1 - 3)^2 + x_2^2))}{exp - (\frac{1}{2}(x_1^2 + x_2^2)) + exp - (\frac{1}{2}((x_1 - 3)^2 + x_2^2))}\big]$$ □

Let $x = (-1,1)$ and $y = (4,1)$. In this case 2 basin of attractions are $A_1 = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 < 1.5\}$ and $A_2 = \{(x_1, x_2) \in \mathbb{R}^2 : x_1 > 1.5\}$ and the bottle-

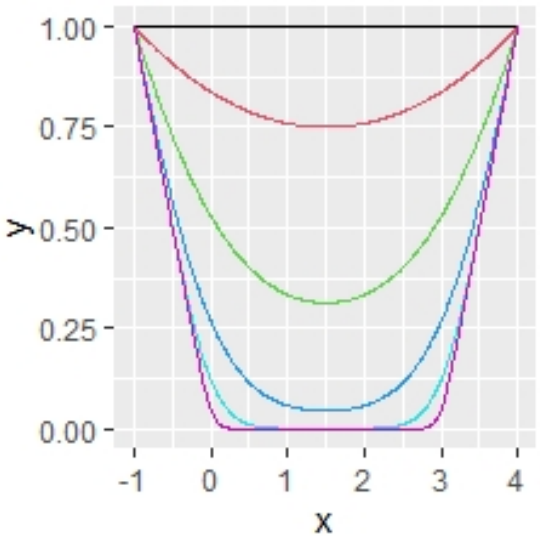

Figure 4.2: γ_β joining $(-1,1)$ and $(4,1)$ for $\beta = 0$, .1, .5, 2, 10 and 100 in example 4.2.

neck/conductance point between them is $(1.5,0)$. So in this case as $\beta \to \infty$, γ_β would correspond to the field line joining $(-1,1)$ and m_1, which is close to $(0,0)$, followed by the field line (straight line) joining m_1 and $(1.5,0)$ followed by the field line (straight line) joining $(1.5,0)$ and m_2, which is close to $(4,1)$, followed by the field line joining m_2 and $(4,1)$.

Example 4.3

Phase transition: Let

$$f \sim \frac{1}{3}N((0,0),\boldsymbol{I}_2) + \frac{1}{3}N((4,0),\boldsymbol{I}_2) + \frac{1}{3}N((2,\alpha),\boldsymbol{I}_2)$$

, $2 \leq \alpha \leq 5$, with $x = (-1,-1)$ and $y = (5,-1)$.

In this specified range of α even for $\alpha = 2$, 3 means $\mu_1 = (0,0)$, $\mu_2 = (4,0)$ and $\mu_3 = (2,\alpha)$ are far apart from one another (compared to each mixture's variance) so that there are 3 maximas m_1, m_2 and m_3 (and therefore 3 corresponding basin of

attractions A_1, A_2, A_3) close to the respective means. Note that for $\alpha = 2\sqrt{3} \approx 3.464$, μ_i's form vertices of an equilateral triangle and by symmetry, we can easily see that in this special case the 3 basin of attractions are given by the voronoi partition of $\mathbb{R}^2$ w.r.t. the μ_i's, namely 3 regions between the bisectors of sides of the triangle fromed by μ_i's from it's centroid. In-general, for other values of α in this region, boundary between A_1, A_3 and A_2, A_3 will not be straight lines, but in any case 3 basin of attractions are given by the 3 regions between 3 curves emiting from a common point $(2, r_\alpha)$ for some $0 < r_\alpha < \alpha$ by symmetry.

Moreover, by symmetry of the location of μ_3 for all choice of α, the boundary betweenA_1 and A_2 will be given by part of the straight line, $B_{1,2} \coloneqq \{(2, s) : s \leq r_\alpha\}$. We shall similarly call the boundary between A_1 and A_3 as $B_{1,3}$ and boundary between A_2 and A_3 as $B_{2,3}$. $B_{1,3}$ and $B_{2,3}$ are not straight lines in-general. Also note that, clearly $B_{1,3}$ and $B_{2,3}$ have symmetry w.r.t the line $\{(2, s) : s \in \mathbb{R}\}$. Also, $x \in A_1$ and $y \in A_2$. For smaller values of β (≈ 2), along $B_{1,2}$ as $s \uparrow$ to r_α, f increases monotonically (also clear from the contour plot), thus in this case the conductance point between A_1 and A_2 is $(2, r_\alpha)$. On the other hand, for higher values of α ($\alpha \approx 5$), there is a maxima on f along $B_{1,2}$ at $(2, \tau_\alpha)$ for some $0 \lessapprox \tau_\alpha < r_\alpha$, which is the conductance point between A_1 and A_2. Additionally, it is clear from the contour plots that for smaller value of α the conductance value between A_1 and A_3 (which is same as the conductance value between A_2 and A_3 by symmetry) is more than $f(2, r_\alpha)$, the conductance value between A_1 and A_2. Thus, for smaller values of α the bottleneck value between A_1 and A_2 will be the conductance value between A_1 and A_3 and thus in this case, from our discussion, $\lim\limits_{\beta\to\infty} \gamma_\beta(x, y)$ would correspond to the union of field lines

$$x \to m_1 \to C(A_1, A_3) \to m_3 \to C(A_2, A_3) \to m_2 \to y$$

For larger values of α, on the other hand conductance value between A_1 and A_3 (which is same as the conductance value between A_2 and A_3 by symmetry) is less

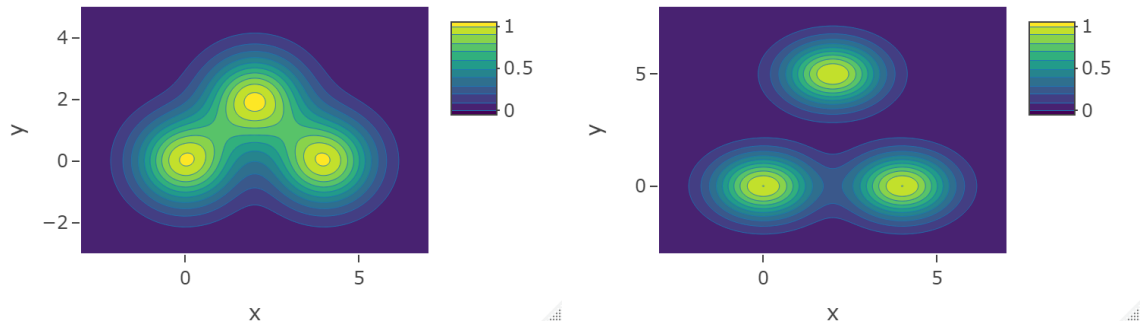

Figure 4.3: Contour plots for $\alpha = 2$ and $\alpha = 5$ in example 4.3.

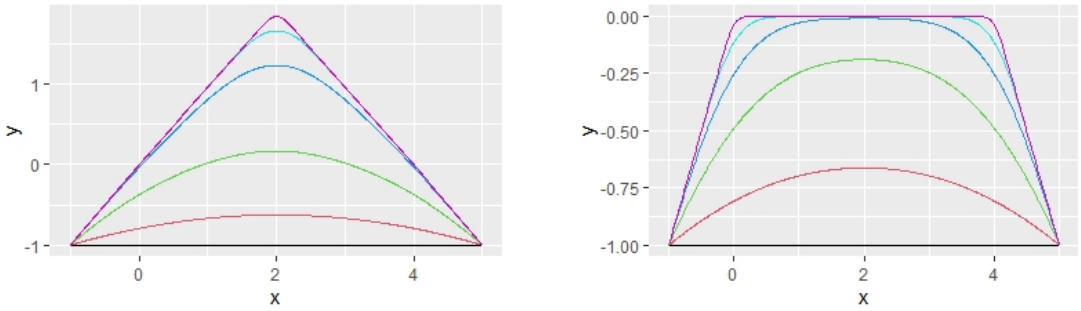

Figure 4.4: γ_β joining $(-1,-1)$ and $(5,-1)$ for $\beta = 0,\ .1,\ .5,\ 2,\ 10$ and 100 and $\alpha = 2$ and 5 in example 4.3.

than $f(2,\tau_\alpha)$, the conductance value between A_1 and A_2 because μ_3 is too far away. Thus in this case, from our discussion, $\lim_{\beta\to\infty}\gamma_\beta(x,y)$ would correspond to the union of field lines

$$x \to m_1 \to C(A_1, A_2) \to m_2 \to y$$

and in particular, $\lim_{\beta\to\infty}\gamma_\beta(x,y)$ won't enter A_3 at all, i.e. there is a phase transition happening somewhere between $2 < \alpha < 5$. □

4.4 Generalization to Riemannian Manifold

In this section we generalize theorem 4.2 for a Riemannian Manifold.

4.4.1 Differential equation for $\{\gamma_\beta(t)\}_{\beta\geq 0}$ in Riemannian manifold

Let (M, g_1) be a boundary-less d-dimensional ($d \geq 2$) Riemannian manifold. The tangent bundle (introduced in $\mathbb{R}^d$ case) of M assigns to each point p of M a vector space T_pM called the tangent space of M at p. A Riemannian metric assigns to each p a positive-definite inner product $g_p : T_pM \times T_pM \to \mathbb{R}$.

When given a system of smooth local coordinate chart (U, x) on $U \subset M$ given by d real-valued functions $x \coloneqq (x_1, .., x_d) : U \to \mathbb{R}^d$, the vectors

$$\{\frac{\partial}{\partial x_1}|_p, ..., \frac{\partial}{\partial x_d}|_p\}$$

form a basis of T_pM, for any $p \in U$. Relative to this basis, one can define metric tensor components at each point p by

$$g_{ij}|_p \coloneqq g_p(\frac{\partial}{\partial x_i}|_p, \frac{\partial}{\partial x_j}|_p)\ i, j = 1(1)d.$$

Thus we have

$$\gamma^{'}(t)|_p \coloneqq \sum_{i=1}^{d} \frac{\partial \gamma}{\partial x_i}|_p x_i^{'}$$

We denote $G_{d\times d}|p \coloneqq ((g_{ij}|_p))$, which is a positive definite matrix. Note that when $M = \mathbb{R}^d$ and g is the Euclidean metric, we have $G|_p = I$ for all $p \in M$.

Let $\gamma(t) \coloneqq x(t) = (x_1(t), .., x_d(t)) : [0, 1] \to M$ be a piecewise continuous curve with $\gamma(0) = a$ and $\gamma(1) = b$. Note that, unlike the case when $M = \mathbb{R}^d$ as discussed in the previous section, we do not assume here an "optical axis", i.e. a parametrization of $x_2, .., x_d$ in terms of x_1 since such an optical axis in-general do not exist on a general M.

Also, we denote $x_i^{'} \coloneqq \frac{\partial x_i}{\partial t}$ and $x_i^{''} \coloneqq \frac{\partial x_i^{'}}{\partial t}$ for $i = 1(1)n$.

Thus in this case the arc-length functional $L(t, x_1, .., x_d, x_1^{'}, .., x_d^{'})$ has the following expression

$$L \coloneqq \int_0^1 f^{-\beta}(\gamma_t)|\gamma_t^{'}|dt$$

$$= \int_0^1 \sum_{i=1}^{d} \sum_{j=1}^{d} f^{-\beta}(x_1, .., x_d)[g_{ij}(x_1, .., x_d)x_i^{'}x_j^{'}]^{\frac{1}{2}} dt$$

since

$$|\gamma_t^{'}|^2 \coloneqq g_p(\sum_{i=1}^{d} \frac{\partial \gamma}{\partial x_i} x_i^{'}, \sum_{i=1}^{d} \frac{\partial \gamma}{\partial x_i} x_i^{'}) = \sum_{i=1}^{d} \sum_{j=1}^{d} g_{ij}(x_1, .., x_d)x_i^{'}x_j^{'}$$

Theorem 4.3

$\{\gamma_\beta(t)\}_{\beta \geq 0}$ satisfy the system of differential equations

$$\beta \left[(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}^{'})\boldsymbol{x}^{'} - (\boldsymbol{x}^{'T} G \boldsymbol{x}^{'}) G^{-1} \frac{\partial h}{\partial \boldsymbol{x}} \right] = \boldsymbol{x}^{''} + \Gamma \tag{4.16}$$

where $h = logf$, $\frac{\partial h}{\partial \boldsymbol{x}} = \begin{pmatrix} \frac{\partial h}{\partial x_1} \\ . \\ . \\ \frac{\partial h}{\partial x_d} \end{pmatrix}$, $\boldsymbol{x}^{'} = \begin{pmatrix} x_1^{'} \\ . \\ . \\ x_d^{'} \end{pmatrix}$, $\boldsymbol{x}^{''} = \begin{pmatrix} x_1^{''} \\ . \\ . \\ x_d^{''} \end{pmatrix}$, $\Gamma = \begin{pmatrix} \sum_{j=1}^{d} \sum_{k=1}^{d} \Gamma_{jk}^1 x_j^{'} x_k^{'} \\ . \\ . \\ \sum_{j=1}^{d} \sum_{k=1}^{d} \Gamma_{jk}^d x_j^{'} x_k^{'} \end{pmatrix}$

and $\{\Gamma_{ij}^k\}_{i,j,k=1(1)d}$'s are the Christoffel symbols of the second kind (Herman, 2008) defined by

$$\Gamma_{ij}^k \coloneqq \frac{1}{2} \sum_{l=1}^{d} (g_{il,j} + g_{jl,i} - g_{ij,l}) g^{lk}$$

where $g^{ij} \coloneqq ((G^{-1}))_{ij}$. □

Proof 4.11

We begin by noting that the geodesic γ_t need to have constant speed, i.e.

$$\frac{d}{dt}(|\gamma_t^{'}|^2) = \frac{d}{dt}(\sum_{i=1}^{d} \sum_{j=1}^{d} g_{ij}(x_1, .., x_d)x_i^{'}x_j^{'}) = \frac{1}{2} < \gamma_t^{'}, \gamma_t^{''} >= 0$$

Inner product between $\gamma_t^{'}$ and $\gamma_t^{''}$ is 0 since any geodesic need to have it's acceleration normal to the surface, i.e. in particular it is orthogonal to any vector in T_pM.

From Eular-Lagrange equation we have

$$\frac{d}{dt}[\frac{\partial Q}{\partial x_i^{'}}] = \frac{\partial Q}{\partial x_i}$$

for $i = 1(1)d$ where $Q = \sum_{i=1}^{d}\sum_{j=1}^{d} f^{-\beta}(x_1, .., x_d)[g_{ij}(x_1, .., x_d)x_i^{'}x_j^{'}]^{\frac{1}{2}}$.

Now,

$$\frac{d}{dt}[\frac{\partial Q}{\partial x_i^{'}}]$$

$$= \frac{d}{dt}[f^{-\beta}\frac{\sum_{k=1}^{d} g_{ik}x_k^{'}}{(\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'})^{\frac{1}{2}}}]$$

$$= \frac{1}{(\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'})^{\frac{1}{2}}}[\frac{d}{dt}[f^{-\beta}]\sum_{k=1}^{d} g_{ik}x_k^{'} + f^{-\beta}\frac{d}{dt}(\sum_{k=1}^{d} g_{ik}x_k^{'})]$$

$$= \frac{1}{(\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'})^{\frac{1}{2}}} f^{-\beta}[-\beta(\sum_{k=1}^{d}\frac{\partial h}{\partial x_k}x_k^{'})(\sum_{k=1}^{d} g_{ik}x_k^{'})$$

$$+(\sum_{k=1}^{d} g_{ik}x_k^{''} + \sum_{k=1}^{d}\sum_{j=1}^{d} g_{ik,j}x_j^{'}x_k^{'})]$$

where $g_{ik,l} \coloneqq \frac{\partial g_{ik}}{\partial x_l}$ and $h = logf$. Similarly,

$$\frac{\partial Q}{\partial x_i}$$

$$= f^{-\beta}[\frac{1}{2}\frac{\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk,i}x_j^{'}x_k^{'}}{(\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'})^{\frac{1}{2}}} + (-\beta)\frac{\partial h}{\partial x_i}(\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'})^{\frac{1}{2}}]$$

Thus,

$$\frac{d}{dt}[\frac{\partial Q}{\partial x_i^{'}}] - \frac{\partial Q}{\partial x_i} = 0$$

$$\Rightarrow \beta[(\sum_{k=1}^{d}\frac{\partial h}{\partial x_k}x_k^{'})(\sum_{k=1}^{d} g_{ik}x_k^{'}) - \frac{\partial h}{\partial x_i}\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk}x_j^{'}x_k^{'}]$$

$$= \sum_{k=1}^{d} g_{ik} x_k^{''} + \sum_{j=1}^{d}\sum_{k=1}^{d} g_{ik,j} x_j^{'} x_k^{'} - \frac{1}{2}\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk,i} x_j^{'} x_k^{'}$$

$$\Rightarrow \beta[(\sum_{k=1}^{d} \frac{\partial h}{\partial x_k} x_k^{'})(\sum_{k=1}^{d} g_{ik} x_k^{'}) - \frac{\partial h}{\partial x_i}\sum_{j=1}^{d}\sum_{k=1}^{d} g_{jk} x_j^{'} x_k^{'}]$$

$$= \sum_{k=1}^{d} g_{ik} x_k^{''} + \sum_{j=1}^{d}\sum_{k=1}^{d} \frac{1}{2}[g_{ik,j} + g_{ij,k} - g_{jk,i}] x_j^{'} x_k^{'}$$

for all $i = 1(1)d$. In compact form, these system of d equations can be re-written as

$$\beta[(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}^{'}) G\boldsymbol{x}^{'} - (\boldsymbol{x}^{'T} G\boldsymbol{x}^{'})\frac{\partial h}{\partial \boldsymbol{x}}] = G\boldsymbol{x}^{''} + \boldsymbol{c}$$

where $\frac{\partial h}{\partial \boldsymbol{x}} = \begin{pmatrix} \frac{\partial h}{\partial x_1} \\ . \\ . \\ \frac{\partial h}{\partial x_d} \end{pmatrix}$, $\boldsymbol{x}^{'} = \begin{pmatrix} x_1^{'} \\ . \\ . \\ x_d^{'} \end{pmatrix}$, $\boldsymbol{c} = \begin{pmatrix} \sum_{j=1}^{d}\sum_{k=1}^{d} \frac{1}{2}[g_{1k,j} + g_{1j,k} - g_{jk,1}] x_j^{'} x_k^{'} \\ . \\ . \\ \sum_{j=1}^{d}\sum_{k=1}^{d} \frac{1}{2}[g_{dk,j} + g_{dj,k} - g_{jk,d}] x_j^{'} x_k^{'} \end{pmatrix}$.

Multiplying by G^{-1} yields,

$$\beta[(\frac{\partial h}{\partial \boldsymbol{x}^{'}}^T \boldsymbol{x}^{'}) \boldsymbol{x}^{'} - (\boldsymbol{x}^{'T} G\boldsymbol{x}^{'}) G^{-1}\frac{\partial h}{\partial \boldsymbol{x}}] = \boldsymbol{x}^{''} + \Gamma$$

where

$$\Gamma = G^{-1}\boldsymbol{c} = \begin{pmatrix} \sum_{j=1}^{d}\sum_{k=1}^{d}\sum_{i=1}^{d} \frac{1}{2}[g_{ik,j} + g_{ij,k} - g_{jk,i}] g^{i1} x_j^{'} x_k^{'} \\ . \\ . \\ \sum_{j=1}^{d}\sum_{k=1}^{d}\sum_{i=1}^{d} \frac{1}{2}[g_{ik,j} + g_{ij,k} - g_{jk,i}] g^{id} x_j^{'} x_k^{'} \end{pmatrix}$$

$$= \begin{pmatrix} \sum_{j=1}^{d}\sum_{k=1}^{d} \Gamma_{jk}^{1} x_j^{'} x_k^{'} \\ . \\ . \\ \sum_{j=1}^{d}\sum_{k=1}^{d} \Gamma_{jk}^{d} x_j^{'} x_k^{'} \end{pmatrix}.$$

□

4.4.2 Interpretation, role of β and $\gamma_\infty(t)$

Proposition 4.8

$(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}')\boldsymbol{x}' - (\boldsymbol{x}'^T G \boldsymbol{x}')G^{-1}\frac{\partial h}{\partial \boldsymbol{x}} = \mathbf{0}$ if and only if $\frac{\partial h}{\partial \boldsymbol{x}} \propto G\boldsymbol{x}'$. □

Proof 4.12

Applying Cauchy-Schwarz inequality on $G^{-\frac{1}{2}}\frac{\partial h}{\partial \boldsymbol{x}}$ and $G^{\frac{1}{2}}\boldsymbol{x}'$ yields

$$(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}')^2 \leq (\boldsymbol{x}'^T G \boldsymbol{x}')(\frac{\partial h}{\partial \boldsymbol{x}}^T G^{-1}\frac{\partial h}{\partial \boldsymbol{x}})$$

Thus, $(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}')\boldsymbol{x}' - (\boldsymbol{x}'^T G \boldsymbol{x}')G^{-1}\frac{\partial h}{\partial \boldsymbol{x}} = \mathbf{0}$ if and only if equality in Cauchy-Schwarz inequality holds, i.e. $\frac{\partial h}{\partial \boldsymbol{x}} \propto G\boldsymbol{x}'$. □

Remark 4.2

Note that

$$\frac{\partial h}{\partial \boldsymbol{x}} \propto G\boldsymbol{x}' = \frac{\partial |\gamma_t'|}{\partial \boldsymbol{x}'}$$

generalizes the notion of flow-lines induced by the gradient of h in a Riemannian manifold. Just like in $M = \mathbb{R}^d$ case, we see that β controls the amount of influence of the default geodesic (corresponding to $\beta = 0$) given by the Einstein's equation $\boldsymbol{x}'' + \Gamma = \mathbf{0}$ (Herman, 2008) and union of flow-lines (corresponding to $\beta = \infty$) $\frac{\partial h}{\partial \boldsymbol{x}} \propto G\boldsymbol{x}'$ in $\gamma_\beta(t)$. Also, which union of flow lines does $\gamma_\infty(t)$ correspond to is very similar to the $\mathbb{R}^d$ scenario as discussed in the previous section and is omitted here for the sake of brevity. □

4.4.3 Some examples

Example 4.4

$\mathbb{R}^d$:

Clearly for $\mathbb{R}^d$ with Euclidean metric, $G \equiv I$ and all Christoffel symbols are 0,

thus in case of $\mathbb{R}^d$ we have

$$\beta[(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}^{'})\boldsymbol{x}^{'} - (\boldsymbol{x}^{'T}\boldsymbol{x}^{'})\frac{\partial h}{\partial \boldsymbol{x}}] = \boldsymbol{x}^{''}$$

Re-parametrizing $x_2, .., x_n$ w.r.t. x_1, denoting $\boldsymbol{q}^T = (x_2, .., x_n)$, $\boldsymbol{q}^{'T} = (x_2^{'}, .., x_n^{'})$, letting $\gamma^2 = 1 + \boldsymbol{q}^{'T}\boldsymbol{q}^{'}$ and noting $x_i^{'} = x_1^{'}q_i^{'}$, $x_i^{''} = x_1^{'2}q_i^{''}$ for $i = 2, 3, ..d$ we obtain,

$$\gamma^2\beta[\frac{1}{\gamma^2}(\frac{\partial h}{\partial x_1}\gamma^2)\boldsymbol{q}^{'} - \frac{\partial h}{\partial \boldsymbol{q}}] = \boldsymbol{q}^{''}$$

$$\Rightarrow \boldsymbol{q}^{''} = \gamma^2\beta[\frac{\partial h}{\partial x_1}\boldsymbol{q}^{'} - \frac{\partial h}{\partial \boldsymbol{q}}]$$

Thus we arrive at theorem 4.2. □

Example 4.5

2 dimensional sphere embedded in $\mathbb{R}^3$: Consider a 2 dimensional sphere $M = \mathbb{S}_2$ embedded in $\mathbb{R}^3$ with radius $a > 0$ parametrized by θ and s as follows:

$$\mathbb{S}_2 = \{(a\,cos\theta\,cos\frac{s}{a}, a\,sin\theta\,cos\frac{s}{a}, asin\frac{s}{a}) : 0 \leq \theta < 2\pi, -\frac{a\pi}{2} \leq s \leq \frac{a\pi}{2}\}$$

and a density function f on it with $f(x) > 0$ for any $x \in \mathbb{S}_2$. Here the tangent bundle T_pM has 2 basis vectors:

$$\frac{\partial}{\partial \theta} = (-a\,sin\theta\,cos\frac{s}{a}, a\,cos\theta\,cos\frac{s}{a}, 0)$$

and

$$\frac{\partial}{\partial s} = (-cos\theta sin\frac{s}{a}, -sin\theta sin\frac{s}{a}, cos\frac{s}{a})$$

It is not hard to verify that under this parametrization,

$$G = \begin{pmatrix} a^2\,cos^2\frac{s}{a} & 0 \\ 0 & 1 \end{pmatrix}$$

and $\Gamma^s_{\theta,\theta} = a\,cos\frac{s}{a}sin\frac{s}{a}$, $\Gamma^\theta_{\theta,s} = \Gamma^\theta_{s,\theta} = -\frac{1}{a}\,tan\frac{s}{a}$ and the remaining Christoffel symbols are 0. Thus, 2 equations of geodesics $\gamma_\beta(t)$ are

$$\beta[(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}^{'})\boldsymbol{x}^{'} - (\boldsymbol{x}^{'T}G\boldsymbol{x}^{'})G^{-1}\frac{\partial h}{\partial \boldsymbol{x}}] = \boldsymbol{x}^{''} + \Gamma$$

where $h = logf$, $\boldsymbol{x}^{'} = \begin{pmatrix} \theta^{'} \\ s^{'} \end{pmatrix}$, $\boldsymbol{x}^{''} = \begin{pmatrix} \theta^{''} \\ s^{''} \end{pmatrix}$, $\Gamma = \begin{pmatrix} -\frac{2}{a} tan\frac{s}{a} s^{'}\theta^{'} \\ acos\frac{s}{a} sin\frac{s}{a} \theta^{'2} \end{pmatrix}$, $G = \begin{pmatrix} a^2 \cos^2\frac{s}{a} & 0 \\ 0 & 1 \end{pmatrix}$ and $G^{-1} = \begin{pmatrix} \frac{1}{a^2 cos^2\frac{s}{a}} & 0 \\ 0 & 1 \end{pmatrix}$ with the boundary condition $\gamma_\beta(0) = p$ and $\gamma_\beta(1) = q$.

After some algebraic manipulation, the 2 equations reduce to

$$\beta[\frac{\partial h}{\partial s} s^{'}\theta^{'} - \frac{\partial h}{\partial \theta} \frac{s^{'2}}{a^2 cos^2\frac{s}{a}})] = \theta^{''} - \frac{2}{a} tan\frac{s}{a} s^{'}\theta^{'}$$

and

$$\beta[\frac{\partial h}{\partial \theta} s^{'}\theta^{'} - \frac{\partial h}{\partial s} a^2 cos^2\frac{s}{a}\theta^{'2}] = s^{''} + a\, cos\frac{s}{a} sin\frac{s}{a}\theta^{'2}$$

Note that the default geodesic (corresponding to $\beta = 0$) on a sphere joining p and q is the smaller section of the unique "great circle" connecting x and y. □

Example 4.6

2 dimensional torus embedded in $\mathbb{R}^3$: Consider a 2 dimensional torus $M = \mathbb{T}_2$ embedded in $\mathbb{R}^3$ parametrized by θ and s for $a > b$ as follows :

$$\mathbb{T}_2 = \{(a + b\, coss)cos\theta, (a + b\, coss)sin\theta, bsins) : 0 \leq \theta < 2\pi, -0 \leq s < 2\pi\}$$

and a density function f on it with $f(x) > 0$ for any $x \in \mathbb{T}_2$. Here the tangent bundle T_pM has 2 basis vectors:

$$\frac{\partial}{\partial \theta} = (-(a + b\, coss)\, sin\theta, (a + b\, coss)\, cos\theta, 0)$$

and

$$\frac{\partial}{\partial s} = (-b\, sins\, cos\theta, -b\, sins\, sin\theta, bcoss)$$

It is not hard to verify that under this parametrization,

$$G = \begin{pmatrix} (a + b\, coss)^2 & 0 \\ 0 & b^2 \end{pmatrix}$$

and $\Gamma^{\theta}_{\theta,s} = \Gamma^{\theta}_{s,\theta} = -\frac{b\,sins}{a+b\,coss}$, $\Gamma^{s}_{\theta,\theta} = \frac{1}{b}sins(a+b\,coss)$ and the remaining Christoffel symbols are 0. Thus, 2 equations of geodesics $\gamma_\beta(t)$ from p to q are

$$\beta[(\frac{\partial h}{\partial \boldsymbol{x}}^T \boldsymbol{x}^{'})\boldsymbol{x}^{'} - (\boldsymbol{x}^{'T} G \boldsymbol{x}^{'})G^{-1}\frac{\partial h}{\partial \boldsymbol{x}}] = \boldsymbol{x}^{''} + \Gamma$$

where $h = logf$, $\boldsymbol{x}^{'} = \begin{pmatrix}\theta^{'}\\ s^{'}\end{pmatrix}$, $\boldsymbol{x}^{''} = \begin{pmatrix}\theta^{''}\\ s^{''}\end{pmatrix}$, $\Gamma = \begin{pmatrix}-\frac{2bsins}{a+bcoss}s^{'}\theta^{'}\\ \frac{1}{b}sins(a+bcoss)\theta^{'2}\end{pmatrix}$, $G = \begin{pmatrix}(a+bcoss)^2 & 0\\ 0 & b^2\end{pmatrix}$ and $G^{-1} = \begin{pmatrix}\frac{1}{(a+bcoss)^2} & 0\\ 0 & \frac{1}{b^2}\end{pmatrix}$ with the boundary condition $\gamma_\beta(0) = p$ and $\gamma_\beta(1) = q$.

After some algebraic manipulation, the 2 equations reduce to

$$\beta[\frac{\partial h}{\partial s}s^{'}\theta^{'} - \frac{\partial h}{\partial \theta}\frac{b^2}{(a+bcoss)^2}s^{'2})] = \theta^{''} - \frac{2bsins}{a+bcoss}s^{'}\theta^{'}$$

and

$$\beta[\frac{\partial h}{\partial \theta}s^{'}\theta^{'} - \frac{\partial h}{\partial s}\frac{(a+bcoss)^2}{b^2}\theta^{'2}] = s^{''} + \frac{1}{b}sins(a+bcoss)\theta^{'2}$$

Note that finding the default geodesic (corresponding to $\beta = 0$) on a torus joining p and q is by itself a non-trivial problem and is analyzed in detail in Irons (2005); Jantzen (2012). □

Remark 4.3

Uniqueness of $\{\gamma_\beta(x,y)\}$: Since γ_β, $\beta > 0$ is an interpolation between γ_0 (default geodesic) and γ_∞, γ_β is unique if γ_0 and γ_∞ are both unique. In-general, neither of them have to be unique. For example γ_0 is locally unique in a Riemannian manifold because of the Picard–Lindelöf theorem for a system of first order ODEs. Recall that the default geodesic equation $x^{''} + \Gamma = 0$ can be recast as a system of first order ODEs. But in-general, γ_0 is not globally unique. Take, for example, north pole and south pole in a sphere, any meridian is a γ_0 geodesic.

On the other hand, any 2 connected basin of attractions can have more than one conductance point and there can be more than 1 path in G as introduced in proposition 4.6 that attains bottleneck value making γ_∞ not unique either. For example, consider $f \sim \frac{1}{4}\sum_{i=1}^{4} N(\mu_i, I)$ with $\mu_1 = (0,0)$, $\mu_2 = (a,a)$, $\mu_3 = (a,-a)$ and $\mu_4 = (2a,0)$ for sufficiently large a so that there are 4 basin of attractions A_1, A_2, A_3, A_4 induced by 4 maximas of f m_1, m_2, m_3, m_4 which are close to 4 means respectively and $x = (-1,0)$ $y = (2a+1,0)$. Clearly by symmetry, $\gamma_\infty(x,y)$ is not unique since $x \to m_1 \to C(A_1,A_2) \to m_2 \to C(A_2,A_4) \to m_4 \to y$ and $x \to m_1 \to C(A_1,A_3) \to m_3 \to C(A_3,A_4) \to m_4 \to y$ are both candidate $\gamma_\infty(x,y)$.

However, as shown in proposition 4.4, if x and y are in the same basin of attraction then $\gamma_\infty(x,y)$ is unique. Thus, the conclusion is $\gamma_\beta(x,y)$ is in-general not unique but it is locally unique if the entire domain is in a common basin of attraction. □

Chapter 5

Spectral Clustering using g-distance and LLPD

5.1 Introduction, our contribution and organization

In this chapter we derive performance guarantee conditions of several spectral clustering (SC) algorithms using density sensitive distances like normalized g distance (L_g^N) and LLPD. We formulate *continent problem* first and show that SC with eigen-gap heuristics using suitably scaled version of these distances can efficiently solve it. Next we derive the exact and asymptotic performance guarantee criterion for SC with eigen-gap heuristic using L_g^N and LLPD under different data generating models and discuss the trade offs of different distances. The data generating models we have considered in this chapter are

- Data generated from a density f in a D dimensional Riemannian manifold.
- Data generated from multiple disjoint lower dimensional Riemannian manifold with or without D-dimensional ambient noise, known as low dimensional large noise (LDLN) model in literature (Little et al., 2017).

- The previous settings with the additional requirement of being able to detect intersections between manifolds, known as multi-manifold clustering (MMC) (Wang et al., 2010) in literature.

Summary of the results in this chapter are as follows:

1. In context of continent problem as well as LDLN model LLPD SC is equivalent to prominence/density based clustering, is robust to cluster imbalances both in size and shape, do not require clusters to be separated by a distance, and can detect lower dimensional clusters from high dimensional background noise while having significantly higher variance (than other distances) in clustering accuracy and is very susceptible to lower dimensional outliers and any small perturbation in data generating model.

2. On the other hand, SC with normalized g distance with small b_1 (recall that b_1 is the smallest power of g) and Euclidean distance (which is equivalent to K-means algorithm) find it difficult to detect lower dimensional structure amongst higher dimensional noise, require clusters to be separated by relatively a large distance compared to their sizes, clusters to be fairly balanced both in size and shape. But SC using these distances have lower variance in clustering accuracy than LLPD SC and are more robust towards small changes in data generating model and lower-dimensional outliers.

Remark 5.1

Although some of these results in this chapter are already known via simulation experiments (Mckenzie and Damelin, 2019; Little et al., 2020) or have been qualitatively justified using heuristics, to our knowledge no rigorous theoretical justification of these results or exact cost of the trade-offs between different distances have yet been provided. These exact cost of trade-offs are very important since it can effectively determine optimal balance between geometry and

density given a particular data generating model which will maximize overall expected clustering accuracy. □

3. We conclude by showing that normalized g-distance, if chosen properly (for example exponential distance is a good candidate), do not need large sample to cluster efficiently unlike LLPD or power-weighted distances and can outperform aforementioned distances in small/medium sample scenario.

The organization of this chapter is as follows: In section 2 we provide a brief introduction to g distance and LLPD and state some results related to the computation efficiency of them in SC algorithms, which has been used later. In section 3 we provide a brief overview of the SC algorithm. Section 4 and 5 are dedicated to providing the background necessary for the remainder of the chapter. In particular in section 4 we provide a brief introduction to different data generative models considered in this chapter and in section 5 we provide a brief overview on some concepts such as ridges, level set clustering, mode based clustering and prominence based clustering. In section 6 continent problem is introduced. In section 7 and 8 we derive performance guarantees of continent problem for LLPD SC with L_{sym} and L respectively. In section 9 and 10 we compare performance of SC with normalized power-weighted distance SC in context of solving continent problem. In section 11, 12 and 13 we apply g-distance and LLPD in other data generating models such as LDLN and MMC as introduced in section 4.

5.2 Normalized g-distance and LLPD

5.2.1 Brief introduction to g-distance and LLPD

Recall that given dataset χ_n, LLPD is defined by

$$LLPD(x, y, \chi_n) \coloneqq \min_p \max_i dist_1(x_i, x_{i+1})$$

where p is a path from $x = x_0 \to x_1 \to x_2 .. \to x_r \to x_{r+1} = y$ such that $x_i \in \chi_n$ for all $i = 1(1)r$, r is free to vary. Under same settings, normalized g-distance L_g^N is defined as

$$L_g^N(x, y, \chi_n) \coloneqq \min_p g^{-1}\big(\sum_{i=0}^{r} g(dist_1(x_i, x_{i+1}))\big)$$

for a continuous, super-additive function $g : [0, \infty) \to [0, \infty)$ with $g(0) = 0$ as a generalization of power-weighted and other shortest path distances proposed before.

In chapter 2 we have shown the following:

If $g(t) = O(t^k)$ as $t \to 0^+$, then

$$\lim_{n\to\infty} n^{\frac{k-1}{d\,k}}\, L_g^N\big(x, y, \chi_n\big) \overset{a.s.}{\to} \begin{cases} C(d, g)\, dist_k(x, y) & \text{if } k > 1 \\ g^{-1}(dist_1(x, y)) & \text{if } k = 1. \end{cases}$$

where $C(d, g) > 0$ is a constant depending on d, g and $dist_k(x, y)$ is a population level density sensitive distance defined by

$$dist_k(x, y) \coloneqq \big(\inf_\gamma \int_0^1 f(\gamma_t)^{\frac{1-k}{d}} |\gamma_t'| dt\big)^{\frac{1}{k}}$$

where γ_t is a piece-wise continuous curve from x to y.

In chapter 3 we have shown that,

$$\lim_{n\to\infty} \big(\frac{n}{log\, n}\big)^{\frac{1}{d}}\, LLPD\big(x, y, \chi_n\big) \overset{a.s.}{\to} V_d^{-\frac{1}{d}}\, dist_\infty(x, y)$$

where V_d is the volume of d-dimensional unit ball and $dist_\infty(x,y) = \lim_{k\to\infty} dist_k(x,y)$ is defined as

$$dist_\infty(x,y) \coloneqq \left(\max_\gamma \min_t f(\gamma_t)\right)^{-\frac{1}{d}}$$

where γ_t is any piece wise continuous curve from x to y.

Remark 5.2

Broadly speaking, both of these convergence limit results (for g^N and for LLPD respectively) can be used in two ways for clustering applications.

1. Firstly, we note that un-scaled g^N or LLPD converge to 0 uniformly which is very useful for LDLN type models. Here we exploit the fact that in these models within cluster distances approach 0 while between cluster distances remain bounded away from 0 as $n \to \infty$.

2. Suitably scaled LLPD and g^N distances on the other hand accurately capture the density-geometry trade off inherently present in the data generating model. This approach of clustering lead to the formulation and solution of the continent problem as presented later this chapter.

5.2.2 g-distance and LLPD restricted to k-NN graph

Given data χ_n, k Nearest Neighbor (k-NN) graph C_k is $(V = \chi_n, E)$ where x_i and x_j are joined ($i \neq j$) if at least one of them is within k nearest neighbor $dist_1$ distance of another.

The main result in this sub-section is the following: As $n \to \infty$, both L_p^N and LLPD on complete graph C are very well-approximated by their restriction to k-nearest neighbor (k-NN) graphs C_k for $k \geq c\,log\,n$ where $c > 0$ depends on dimension d, type of distance, f and manifold M.

These results will be used in the SC algorithms later and, in particular, lead to significant reduction in computational complexity. In particular, cost of computing g or LLPD distance matrix is $O(n^2 logn\, d)$ for d dimensional data using $O(logn)$-NN graph, which is only marginally expensive than Euclidean distance computational cost $O(n^2\, d)$.

We provide the following result for restriction of LLPD to k-NN graph.

Proposition 5.1

If $k > \frac{1}{2|log(1-\frac{\rho_d}{2})|}\, logn$ where $\rho_d := \frac{\mu_d(B(0,1)\cap B(e_1,1))}{\mu_d(B(0,1)\cup B(e_1,1))}$, μ_d denoting d-dimensional Lebesgue measure, $e_1 = (1, 0, ..0)^T$ then w.p. $\to 1$, $LLPD|C = LLPD|C_k$. □

Proof 5.1

From proposition 2.2 in González-Barrios and Quiroz (2003) we see that w.p. $\to 1$ minimal spanning tree is contained in k-NN graph for $k > \frac{1}{2|log(1-\frac{\rho_d}{2})|}\, logn$. The result follows from this. □

The following 2 results for L_p^N are provided in this section. We note that similar result holds for g-distance as well since in large sample size ($n \to \infty$) regime normalized g-distance is equivalent to $L_{b_1}^N$ where b_1 is the smallest power of g.

Groisman et al. (2018) proposition 2.11: if $k > c_1\, log(\frac{n}{\epsilon}) + c_2$ for some $c_1, c_2 > 0$ then $\mathbb{P}(L_p^N|C = L_p^N|C_k) > 1 - \epsilon$.

This result was further improved by **Little et al. (2020), Theorem 3.9.** which states the following: If $k > 4k_0^2[\frac{f^0}{f_0}][\frac{4}{4^{1-\frac{1}{p}}-1}]^{\frac{d}{2}} logn$ then w.p. at least $1 - \frac{1}{n}$, $L_p^N|C = L_p^N|C_k$. where $k_0 > 0$ is a constant, f^0 and f_0 are maxima and minima of f in M respectively.

5.3 Spectral clustering: A brief overview

In this section we provide a brief introduction on SC.

Given the dataset $\chi_n \coloneqq \{x_1, x_2, .., x_n\}$ and a symmetric similarity measure $s(.,.)$ one first constructs an adjacency matrix $((W))_{ij} \coloneqq s(x_i, x_j)$. For example, a popular choice is

$$s(x_i, x_j) = exp(-\frac{|x_i - x_j|^2}{\sigma^2})$$

for some specified $\sigma > 0$. Let D be the degree matrix defined by $((D))_{ii} = \sum_j W_{ij}$ for $i = 1(1)n$ and 0's in the off diagonal positions.

The next step to construct a graph Laplacian from W and D. In literature there are 3 distinct graph Laplacians available: (Von Luxburg, 2007)

$$L \coloneqq D - W,$$

$$L_{sym} \coloneqq D^{-\frac{1}{2}} L D^{-\frac{1}{2}} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}},$$

and

$$L_{rw} \coloneqq D^{-1} L = I - D^{-1} W.$$

L is called the un-normalized graph Laplacian whereas L_{sym} and L_{rw} are called normalized and random-walk graph Laplacians respectively.

The next step is to compute the k eigenvectors $u_1 = \frac{1}{\sqrt{n}} 1, u_2, .., u_k$ corresponding to k smallest eigenvalues $\lambda_1 = 0 \leq \lambda_2 \leq .. \leq \lambda_k$, of a graph Laplacian for some pre-specified $k < n$ and form the matrix $U \in \mathbb{R}^{n \times k}$ with $u_1, u_2, .., u_k$ as columns. Let $y_i \in \mathbb{R}^k$ be the i-th row of U for $i = 1(1)n$. y_i is called the *spectral embedding* of x_i.

The final step is to compute cluster labels by applying the K-means on the data $\{y_i\}_{i=1}^n$ if using L or $\{\frac{y_i}{||y_i||_2}\}_{i=1}^n$ as proposed by Ng et al. (2002) if using L_{sym} or L_{rw}.

Note that k, the number of clusters, need to be pre-specified which is usually not known apriori. One well-used heuristic for choosing k is the eigen-gap statistic which suggest $\hat{k} = \underset{i=1(1)n-1}{arg\ max}\, \lambda_{i+1} - \lambda_i$ where λ_i is the i-th smallest eigenvalue of graph Laplacian. While it is unclear why k well-separated clusters should imply $\hat{k} =$

Algorithm 1 SC with L.

Input: $W \in \mathbb{R}^{n \times n}$, the adjacency matrix of the data $\chi_n = \{x_1, x_2, .., x_n\}$.
Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data.
1: Compute the un-normalized Laplacian $L = D - W$
2: Compute the eigen decomposition $\{(\phi_k, \lambda_k)\}_{k=1}^{n}$ of L, sorted so that $0 = \lambda_1 \leq \lambda_2 \leq .. \leq \lambda_n$.
3: Estimate the number of clusters k as $\hat{k} = \underset{i=1(1)n-1}{arg\ max}\, \lambda_{i+1} - \lambda_i$. Let $U_{n\times k} = (\phi_1 : \phi_2 : .. : \phi_k)$.
4: For $1 \leq i \leq n$, let $y_i = U_{i,.}$ define the spectral embedding.
5: Compute the label Z by running K-means on $\{y_i\}_{i=1}^{n}$ using k as the number of clusters and $Z_{i,.}$ is label of y_i as obtained by the K-means.

$\underset{i=1(1)n-1}{arg\ max}\, \lambda_{i+1} - \lambda_i$, one can argue that if there are truly k well-separated clusters, λ_i's for $i \leq k$ should all be very small while λ_{k+1} is relatively large (Von Luxburg, 2007). i.e. one should expect $\hat{k} = \underset{i=1(1)\hat{k}}{arg\ max}\, \lambda_{i+1} - \lambda_i$. Thus $\lambda_{k+1} - \lambda_k$ is the first prominent eigen-gap under the k clusters assumption. However, that alone cannot be used to determine correct k and in practice we at-least need $\hat{k} = \underset{i=1(1)\hat{k}+m}{arg\ max}\, \lambda_{i+1} - \lambda_i$ for some m independent of $\hat{k}$ to correctly determine k.

We provide the formal spectral clustering (SC) algorithm using L and L_{sym} (SC with L_{rw} is similar to SC with L_{sym}, thus we shall only discuss SC with L and L_{sym} here) with the eigen-gap heuristic below (algorithm 1 and 2 respectively). Both algorithms have their own graph cut (for example, SC with L and L_{sym} are relaxation of the NP-hard problems of minimizing ratio cut and normalized cut of a graph respectively), random walk, perturbation theory point of view. We recommend readers Von Luxburg (2007) for a comprehensive review on these perspectives. General consensus is that L_{sym} is better candidate than L for SC purpose (Von Luxburg et al., 2008).

Algorithm 2 SC with L_{sym}. L_{rw} is similar

Input: $W \in \mathbb{R}^{n\times n}$, the adjacency matrix of the data $\chi_n = \{x_1, x_2, .., x_n\}$.
Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data.
1: Compute the normalized Laplacian $L_{sym} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$
2: Compute the eigen decomposition $\{(\phi_k, \lambda_k)\}_{k=1}^{n}$, sorted so that $0 = \lambda_1 \leq \lambda_2 \leq .. \leq \lambda_n$.
3: Estimate the number of clusters k as $\hat{k} = \underset{i=1(1)n-1}{arg\ max}\, \lambda_{i+1} - \lambda_i$. Let $U_{n\times k} = (\phi_1 : \phi_2 : .. : \phi_k)$.
4: For $1 \leq i \leq n$, let $y_i = \frac{U_{i,.}}{||U_{i,.}||_2}$ define the normalized spectral embedding.
5: Compute the label Z by running K-means on $\{y_i\}_{i=1}^{n}$ using k as the number of clusters and $Z_{i,.}$ is label of y_i as obtained by the K-means.

5.4 Different data generating models

Objective of this section is to introduce data generating models (in order of increasing complexity) we shall consider throughout this chapter and the potential challenges of clustering data χ_n generated from them. We shall primarily focus on density based clustering methods i.e. clustering methods which, partially or entirely, consider the underlying density of the data while trying to cluster it and clusters are loosely based on the criterion of being in "high density region" of the data. Throughout this section we shall use $\mu_D(.)$ to denote D dimensional Lebesgue measure and reserve D and d for ambient and intrinsic dimension respectively ($d < D$).

5.4.1 A density f supported in $M \subseteq \mathbb{R}^D$ where $\mu_D(M) > 0$

This is the simplest and most common data generating model for density based clustering algorithms. Some assumptions about M and f mostly seen in the literature are

A1: M is connected, compact and smooth i.e. having a $D-1$ dimensional boundary.

We assume f to be *nicely behaved* i.e.

A2: $\int_L f = 0$ for any $L \subset M$ so that $\mu_D(L) = 0$. This in particular ensures that f has no lower dimensional structure.

A3: f is Morse-Smale. This allows us to split M (except for a lower dimensional set) into finite number of "basin of attractions" $A_1, A_2, ..A_K$ which are candidate clusters for mode based clustering algorithms such as mean-shift, quick shift and variants of it as discussed in Chen et al. (2016); Cheng (1995) etc.

In this scenario, clusters are loosely defined as connected components of a suitably chosen λ super-level set of f defined as

$$M_\lambda := \{x \in M : f(x) \geq \lambda\}, \tag{5.1}$$

or basin of attractions as introduced later.

This approach of clustering is used in DBSCAN (Schubert et al., 2017), cluster tree (Chaudhuri and Dasgupta, 2010)), most mode based clustering algorithms such as mean shift, quick shift or persistence based clustering (Chazal et al., 2013; Genovese et al., 2016)). The primary challenges (depending on the algorithm) include choosing a suitable scaling parameter (bandwidth in case of kernel based estimator, k in case of k-NN based estimator or ϵ in DBSCAN) to estimate f which preserves it's cluster structure, choosing appropriate level λ of the super level set, curse of dimensionality and computational complexity (of estimating f, ∇f or "prominence" of basin of attractions).

5.4.2 Manifold hypothesis

We are increasingly confronted with very high dimensional data in speech signals, images, gene expression, and other sources. Most of those data are supported on a d dimensional Riemannian manifold in $\mathbb{R}^D$ with $d << D$. In literature this is known as manifold hypothesis (MH) (Narayanan and Mitter, 2010). We assume

A1: $M \subset \mathbb{R}^D$ is connected, compact and smooth i.e. having a $d-1$ dimensional boundary with $\mu_d(M) > 0$.

A2: $\int_L f = 0$ for any $L \subset M$ so that $\mu_d(L) = 0$. This in particular ensures that f has no lower dimensional structure.

A3: f is Morse-Smale. This allows us to split M (except for a lower dimensional set) into finite number of "basin of attractions" $A_1, A_2, ..A_K$.

In addition to the challenges mentioned in the previous sub-section we additionally have the challenge to estimate the geometry information present in M from the data χ_n in order to cluster meaningfully. Thus problem of clustering and manifold learning becomes intertwined. Typically clustering is done by performing standard non-linear dimension reduction techniques such as ISOMAP (Tenenbaum et al., 2000), Laplacian eigenmaps (LE) (Belkin and Niyogi, 2003), locally linear embedding (LLE) (Saul and Roweis, 2003) Hessian LLE (HLLE) (Donoho and Grimes, 2003) on the original data followed by doing SC or density based clustering on the transformed data (Goh and Vidal, 2008).

5.4.3 MH with noise

The data generating model is

$$f(x) = \alpha\, 1(x \in M_1) f_1(x) + (1-\alpha)\, 1(x \in M \setminus M_1) f_N(x)$$

for $0 < \alpha \leq 1$, f_1 and f_N being densities in M_1 and $M \setminus M_1$ respectively.

We generalize the model in the previous sub-section by adding a D dimensional "noise cloud" around M_1, a d dimensional Riemannian manifold with the usual assumptions about M_1 and f_1 as introduced in the previous sub-section. The added challenge is to de-noise the data first (effectively labeling all data in $M \setminus M_1$ as "noise") appropriately followed by clustering procedure. The de-noising procedure is typically carried out by labeling those data-points are "noise" whose distance from

it's k-nearest neighbor is above certain threshold for suitably chosen k. The challenges associated with the de-noising procedure is to correctly choose k and de-noise the points in $M \setminus M_1$ accurately.

5.4.4 MH with multiple manifolds + noise: LDLN model

LDLN (low dimensional large noise) model (Arias-Castro, 2011; Little et al., 2017) further generalizes the model in previous sub-section by considering multiple disjoint manifolds with different dimensions. Formally the data generating model is

$$f(x) = \sum_{i=1}^{k} \alpha_i \, 1(x \in M_i) f_i(x) + (1 - \sum_{i=1}^{k} \alpha_i) \, 1(x \in M \setminus (\cup_{i=1}^{k} M_i)) f_N(x) \qquad (5.2)$$

such that $\alpha_i > 0$, $\sum_{i=1}^{k} \alpha_i \leq 1$, f_i and f_N are densities supported on M_i and $M \setminus (\cup_{i=1}^{k} M_i)$ respectively $\forall i = 1(1)k$. We further assume $\dim(M_i) = d_i < D$ and $\underset{i \neq j}{min} \, dist_1(M_i, M_j) = \delta > 0$. The clustering procedure at stage 1 should be able to de-noise the data first and identify $M_1, .., M_k$. At stage 2, we can further refine clustering procedure by clustering within individual M_i's based on geometry of M_i and density f_i as discussed in the previous sub-section.

5.4.5 Multi-Manifold Clustering

Multi-manifold clustering (MMC) (Wang et al., 2010; Arias-Castro et al., 2017; Trillos et al., 2021) further generalizes LDLN by removing the assumption that manifolds need to be disjoint or not self-intersecting. Thus 2 manifolds can intersect transversely and when they do we need to recognize it. This is typically carried out by local PCA Arias-Castro et al. (2017) or by estimating local tangent space for all data points/ data points in a small neighborhood Trillos et al. (2021); Wang et al. (2010), essentially detecting sharp change of curvature as a hint for intersection of manifolds.

5.5 Some prerequisites

In this section we briefly introduce some important concepts such as the notions of ridge, level-set clustering, mode based clustering, Morse-Smale function and 0-dimensional persistence homology (prominence based clustering) associated with it. These concepts will be used in future sections.

5.5.1 Ridge

Ridge generalize the notion of a local maxima (which are 0-dimensional ridges in our formulation) of f. Formally, following Genovese et al. (2014); Wasserman (2018) a d-dimensional ***rigde*** in a D dimensional ambient space is a distribution W_σ defined by

$$W_\sigma := W * \Phi_\sigma \tag{5.3}$$

where W is a distribution supported on $M_1 \subset \mathbb{R}^D$, a d-dimensional Riemannian manifold and Φ_σ is a noise-distribution (i.e. $\mathbb{E}[\Phi_\sigma] = 0$) on $\mathbb{R}^D$ parametrized by $\sigma > 0$ and $*$ denote the convolution operator. A typical choice for Φ_σ is $N_D(0, \sigma^2 I)$.

In the context of manifold hypothesis, thus ridges can be viewed either as lower dimensional clusters or lower dimensional systematic noise depending on context.

5.5.2 Level set clustering

Following Hartigan (1981), given data χ_n in level set clustering paradigm the clusters are defined as the data residing in the connected components of super level-set M_λ for some fixed $\lambda \in (f_0, f^0)$ and data in $M \setminus M_\lambda$ is defined as "noise". Once a suitable λ is chosen, clusters and noise are estimated via a consistent density estimator such

as kernel density estimator or k-NN density estimator (for example, it is sufficient to choose $k = O(logn)$). See Maier et al. (2009); Chaudhuri et al. (2014); Zhao and Lai (2020) for consistency results on k-NN density estimator, for instance, and algorithms for level set clustering.

The main challenge in level set clustering is to choose an appropriate value of λ. Two methods for choosing λ is popular. One possibility is to choose λ to cover some prescribed fraction $1-\beta$ of the total mass; thus we choose λ to satisfy $\int_{M_\lambda} f = 1-\beta$ (a typical choice would be $\beta = .05$). Another idea is to consider connected components of super-level set at all levels $\lambda \in (f_0, f^0)$. This leads us to the idea of *cluster trees* as developed in Chaudhuri and Dasgupta (2010).

5.5.3 Mode based clustering: Morse Smale functions

An alternative but related idea is the mode based clustering and to view clusters as basin of attractions.

For any $x \in \mathbb{R}^D$ we let $g(x) \coloneqq \nabla f(x)$ and $H(x)$ denote the gradient vector and Hessian matrix of f at x respectively. Following Chen et al. (2016) we define f to be ***Morse*** if for any of it's finitely many critical points i.e. any x^c satisfying $g(x^c) = 0$ of f, we have $H(x^c)$ non-singular. In this case, x^c is a local maxima if and only if $g(x^c) = 0$ and $\lambda_1(x^c) < 0$ where $\lambda_1(x)$ denote the largest eigenvalue of $H(x)$.

An ***integral line/flow-line*** $\lambda : \mathbb{R} \to \mathbb{R}^D$ is a curve with $\frac{d\lambda}{ds}(s) = \nabla f(\lambda(s))$. We define $s(\lambda) = \lim_{s\to-\infty} \lambda(s)$ and $d(\lambda) = \lim_{s\to\infty} \lambda(s)$ as ***source*** and ***destination*** of the integral line respectively. Source and destinations are both, by definition, critical points of f. Define the ***ascending manifold*** $A(x)$ and ***descending manifold*** $D(x)$ corresponding to a critical point x of f as

$$A(x) \coloneqq \{\lambda : s(\lambda) = x\} \tag{5.4}$$

$$D(x) \coloneqq \{\lambda : d(\lambda) = x\} \tag{5.5}$$

A Morse function f is ***Morse-Smale*** if the ascending and descending manifolds intersect transversely only.

In most mode-seeking algorithms such as mean shift and variants of it, clusters are defined by set of all D dimensional descending manifolds, which are also known as ***basin of attractions*** in literature. I.e. we have the following decomposition

$$\mathbb{R}^D = \cup_{i=1}^{K} A_i \cup L$$

where $A_1, .., A_K$ are basin of attractions (candidate clusters) corresponding to K maximas of f, say $m_1, .., m_K$, and L consists of union of lower dimensional descending manifolds and forms the boundary of $A_1, .., A_K$.

However, it is well known that basin of attractions can be arbitrarily small in Lebesgue measure, very complicated in shape and notoriously unstable under small perturbation of f. See for example Chazal et al. (2013). All of these can cause a basin of attraction to remain undetected in finite sample.

Therefore to cluster meaningfully, a notion of cluster salience need to be additionally developed in order to distinguish a stable basin of attraction from an unstable basin of attraction which is discussed in the next sub-section.

Another issue with mode based clustering is that it is impossible to distinguish between background noise and periphery of basin of attractions since both are characterized by arbitrarily low density regions. Thus, in most mode seeking algorithms one only considers the basin of attractions restricted to some super-level set as clusters.

5.5.4 Persistent homology: prominence based clustering

Persistent homology is one of the key concepts in topological data analysis (TDA) literature. Following Chazal et al. (2013); Genovese et al. (2016), in context of mode based clustering, we start with the super level set M_λ as described in equation (5.1) and we let λ decrease from f^0 to f_0. We can think of λ as "time" but, in this case, time runs backwards since it starts at f^0 and decrease to f_0. Formally $\{M_\lambda\}_{\lambda \in (f_0, f^0)}$ is called ***super-level sets filtration*** of f. For any $\lambda \in (f_0, f^0)$ and $x \in M$, let $C(x, \lambda) \subseteq M_\lambda$ denote the connected component of M_λ that contains x. Morse theory tells us that when a maxima m_p of f enters the super level-sets filtration, at time $\lambda = b(m_p)$, a new path-connected component $C(m_p, \lambda)$ appears in the superlevel-set M_λ.

Formally, in homological terms, the peak m_p is called the generator of the component born at time $b(m_p)$. $C(m_p, \lambda)$ ceases to be an independent connected component in M_λ when it gets connected to another component generated by a higher peak m_q. At that particular time $\lambda = d(m_p)$, we say that the component $C(m_p, \lambda)$ gets merged into $C(m_q, \lambda)$. While m_q remains the generator of the component $C(m_q, \lambda)$, m_p ceases to be a generator, and we call m_q its parent, noted $m_q = r(m_p)$. The lifespan of m_p as a generator is encoded in a persistence diagram by the point p of coordinates $p_x = b(m_p)$ and $p_y = d(m_p) \leq p_x$. The difference $d_p = p_x - p_y \geq 0$ between birth and death times is called the ***prominence*** of the peak m_p. Thus given a thresholding parameter $\tau \geq 0$, we restrict our focus only to the peaks m_p of f of prominence at least τ and form clusters by merging basins of attractions of peaks having prominence $< \tau$ into the basin of attractions of their roots which have prominence $\geq \tau$. For more details on theory as well as implementation on real data-sets, interested readers are referred to Chazal et al. (2013); Gerber et al. (2013); Genovese et al. (2016).

The following corollary relates prominence of local maximas with the bottleneck values between maximas as introduced in chapter 3 and 4.

Corollary 5.1

Let f be Morse-Smale with K local maximas $m_1 > m_2 > .. > m_K$ (we can always perturb f by an arbitrarily small amount to make sure the strict inequality holds). Then prominence of m_i is given by

$$\text{Prominence}(m_i) = \begin{cases} \infty & i = 1 \\ f(m_i) - \underset{j>i}{max}\ \text{Bottleneck}\,(m_i, m_j) & i > 1 \end{cases}$$

□

Proof 5.2

Follows immediately from the definition of bottleneck value and prominence. □

5.6 Continent Problem

Throughout this section we shall assume the data generating model same as discussed in sub-section 5.4.1, D known and use the similarity measure $s(.,.)$ in the SC algorithm as

$$s(x,y) \coloneqq exp\left(-\frac{(\frac{n}{log\, n})^{\frac{2}{D}}\, LLPD(x,y,\chi_n)^2}{\sigma^2}\right) \tag{5.6}$$

for some $\sigma = O(1) > 0$.

We additionally assume that

A4: Level sets of f, $S_\lambda = \{x \in M : f(x) = \lambda\}$ are finite union of $D-1$ dimensional Riemannian manifolds for any $\lambda \in (f_0, f^0)$.

Recall that we made the same assumption while proving LLPD's convergence limit in chapter 3.

5.6.1 ϵ-stable clusters

As discussed in the level set clustering sub-section of the previous section, given data χ_n while it is possible to estimate the connected components of super-level set M_λ for a fixed level λ or the entire cluster tree consistently (see algorithm 1, 2 and theorem 3.3, 3.4 in Chaudhuri et al. (2014)) without calculating LLPDs, the main challenge is to choose an appropriate value or the set of appropriate values of $\lambda \in (f_0, f^0)$. While cluster tree bypasses this problem by considering all values of λ, and thus producing a "hierarchical cluster structure", the problem of choosing appropriate values of λ still remains since not all values of λ in this hierarchical cluster structure will result in *stable* (i.e. which can be reliably estimated from a finite sample) clusters.

As discussed in persistent homology sub-section of the previous section, prominence of the maximas of the connected components of super-level set at some specified level provides us a way to tell if the resulting clusters are stable. In particular, if we decide to cluster at a fixed level b for some $b \in (f_0, f^0)$ and if $\{C(m_i, b)\}_{i=1(1)k}$ ($k := k(b)$, the number of connected components, depends on b) denote the connected components of M_b corresponding to maximas $\{m_i\}_{i=1(1)k}$ of these components then we need to make sure that none of these components merge sufficiently close to b. I.e. formally we need to have

$$\underset{i \neq j,\, i,j \in \{1,2,..,k\}}{max} \text{Bottleneck}\,(m_i, m_j) < a \tag{5.7}$$

for some $a < b$ with $b - a = \epsilon > 0$ sufficiently large since this will ensure that prominence of $m_i \geq \epsilon$ for $i = 1(1)k$. Alternatively, in terms of persistence homology we require that topology of M_λ should not change in the range $[a, b]$.

When equation (5.7) is satisfied, we denote the clusters (connected components of M_b) corresponding to level b being *"ϵ-stable"* for $\epsilon = b - a$.

Note that this notion of stability of level-set clusters is not new in literature. As discussed here, this idea is intimately related to the persistent homology of level set filtration and is already introduced in definition 3.2 of Chaudhuri et al. (2014) as "(σ, ϵ)-separability", albeit in slightly different notation. In fact as discussed in theorem 3.3 in Chaudhuri et al. (2014), ϵ-stability (for some $\epsilon > 0$) is a fundamental requirement for consistent estimation of level set clusters. However, one important distinction from a persistence homology standpoint is ϵ also need to be sufficiently big in order for clusters to be *topologically stable*.

We provide 2 examples involving Gaussian mixtures in 1 dimension to illustrate the notion of ϵ-stability further.

Example 5.1

Consider

$$f(x) \sim .5N_1(0, 1) + .5N_1(2.1, 1)$$

It is well known that $f(x) \sim .5N_1(0, 1) + .5N_1(a, 1)$ for $a > 0$ is bi-modal iff $a > 2$. Thus, in this case we have 2 modes m_1 and m_2 with bottleneck point 1.05. However, in this case $f(m_1)/f(m_2)$ is extremely close to $f(1.05)$. Thus, only M_b for $b < f(1.05)$ satisfy the stability criterion as discussed above resulting in 1 cluster.

On the other hand, for

$$f(x) \sim .5N_1(0, 1) + .5N_1(3, 1)$$

$f(m_1)$ or $f(m_2)$ is sufficiently bigger than the bottleneck value between m_1 and m_2, namely $f(1.5)$. Thus in this case M_b corresponding to $b \in (f(1.5) + \epsilon, f(m_1))$ for some $\epsilon > 0$ (corresponds to 2 clusters) is a valid choice for clustering as well along with $b < f(1.5)$. □

Example 5.2

$f(x) \sim .45N_1(0,1) + .1N_1(3.25,1) + .45N_1(6.5,1)$

This gaussian mixture has 3 maximas m_1, m_2, m_3 near $0, 3.25, 6.5$ respectively with $f(m_1) = f(m_3) \approx .18$ and $f(m_2) \approx .03$. Moreover, it turns out that Bottleneck (m_1, m_2)= Bottleneck $(m_2, m_3) \coloneqq a$ is very close to $f(m_2)$, i.e. $f(m_2) - a > 0$ is very small. Thus, only M_b corresponding to $b \in (f(m_2) + \epsilon, f(m_1))$ for some $\epsilon > 0$ (this lead to to 2 clusters corresponding to m_1 and m_3) and $b < a$ (this corresponds to 1 cluster) results in stable clusters. □

Given this notion of stability of clusters in level-set clustering or cluster-tree literature, we need to be able to estimate bottleneck values between maxima's of f from data χ_n consistently. In traditional prominence based clustering methods such as algorithm 1 of Chazal et al. (2013), this is done by estimating f from data χ_n resulting in huge computational complexity. This is exactly where LLPD comes into picture.

5.6.2 Estimating ϵ-stable clusters by LLPD

Recall, from LLPD's convergence limit and definition of bottleneck value as introduced in , given b (and under the additional assumption **A4**) for any $x_k \in C(m_i, b) \cap \chi_n$ and $x_l \in C(m_j, b) \cap \chi_n$, $i \neq j$ we have

$$\big(\frac{n}{logn}\big)^{\frac{1}{D}} LLPD(x_k, x_l, \chi_n) \to \big(\frac{V_D D}{2(D-1)}\big)^{-\frac{1}{D}} \text{Bottleneck}(m_i, m_j)^{-\frac{1}{D}} \text{ a.s.}$$

$$\Rightarrow \big(\frac{n}{logn}\frac{V_D D}{2(D-1)}\big)^{-1} LLPD(x_i, x_j, \chi_n)^{-D} \to \text{Bottleneck}(m_i, m_j) \text{ a.s.}$$

Thus, a single LLPD computation between *any* 2 data-points belonging to different connected components of M_b (clusters) provide a consistent estimate of bottleneck value between those 2 clusters. This, along with proposition 5.1 (LLPDs can be estimated from $O(logn)$-NN graphs) results in huge computational efficiency as compared

Algorithm 3 ϵ-stable b level-set clusters of f

Input: Data χ_n, $b \in (f_0, f^0)$,$\epsilon > 0$.
Output: Indicator of whether ϵ-stable cluster exists at level b, label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data or indicator of noise.
1: Compute r-NN distance $\sigma_{i,r}$ from each point x_i. Label a point x_i as "noise" if $\sigma_{i,r} > (\frac{r}{bnV_D})^{\frac{1}{D}}$. Let remaining data points be χ_N with $|\chi_N| = N$.
2: Form the r-NN graph C_r of χ_N. I.e. for $i \neq j$, join x_i, x_j in χ_N if $|x_i - x_j| \leq max\{\sigma_{i,r}, \sigma_{j,r}\}$.
3: Obtain the connected components of C_r. Let the connected components be $\{A_1, .., A_{k(b)}\}$.
4: Pick $y_i \in A_i$ for $i = 1(1)k(b)$. Calculate $LLPD(y_i, y_j, C_r)$ for $i \neq j$.
5: Check if $\underset{i\neq j}{max}\ (\frac{nV_D D}{logn2(D-1)})^{-1} LLPD(y_i, y_j, C_r)^{-D} < b - \epsilon = a$.
6: If no, output: "No ϵ-stable cluster at level b". If yes, output: Label $Z_{N\times 1}$with the label of $x_i \in \chi_N$ as the label of the connected component it is in C_r.

to estimating f to obtain the bottleneck values. In particular, in order to verify (5.7), we only need $\frac{k(k-1)}{2}$ LLPD computations from $O(logn)$-NN graphs.

Below we provide an algorithm (algorithm 3) which produces level set clusters corresponding to level $b \in (f_0, f^0)$ and then checks if the resulting clusters are ϵ-stable. This algorithm can be easily generalized (by repeating this algorithm at different levels) to produces ϵ-stable cluster trees of f.

Now we provide the result related to consistency of algorithm 3.

Lemma 5.1

For any $\delta > 0$, $\exists C(\delta) > 0$ depending on δ ($C(\delta) \to \infty$ as $\delta \to 0$) so that as $n \to \infty$ if $r = C(\delta)\, logn$, w.p. at least $1 - \delta$ connected components of C_r in algorithm 3 correspond to the connected components of $M_{b-a(\delta)}$ where $a(\delta) \to 0$ as $\delta \to 0$. Moreover, if step 5 of algorithm 3 is satisfied, then resulting clusters are ϵ-stable. □

Proof 5.3

Follows from definition of k-NN density estimator ($f_{kNN}(x_i) \coloneqq \frac{k}{nV_D \sigma_{i,k}^D}$), theorem 3.4 in Chaudhuri et al. (2014)/ theorem 5 in Zhao and Lai (2020) (in particular, theorem

Continent Problem description: Density f on a D dimensional Riemannian manifold M.
Input: *Sea-level* $b > 0$ for $b \in (f_0, f^0)$ and *islands* i.e. k connected components $C(m_1, b), .., C(m_k, b)$ of super-level set $M_b := \{x \in M : f(x) \geq b\}$, m_i being the maxima of $C(m_i, b)$, $f(m_i)$'s are in $\downarrow$ order.
Objective: Obtain $\leq k$ *continents* from the islands.

Sample level problem description: Input: Data $\chi_n = \{x_1, x_2, .., x_n\}$.
Output: Label $Z_{n\times 1}$ with i-th entry being the *continent label* of i-th data or indicator of *noise* i.e. points which are below b-level.

5 in Zhao and Lai (2020) states that with probability at least $1 - \delta$, $\sup_x |f(x) - f_{kNN}(x)| \leq O(k^{-\frac{1}{2}}\sqrt{ln\frac{n}{\delta}})$. I.e. in particular for $\delta > 0$, we can choose $C(\delta) = O(log^2\frac{1}{\delta})$, $a(\delta) = O(\frac{1}{log\frac{1}{\delta}})$.) along with discussion in previous sub-subsection (how LLPDs and bottleneck values are related) and proposition 5.1. □

Alternatively, in case we don't find ϵ-stable cluster at level b, we can merge A_i and A_j for any $i \neq j$ as one cluster if $\frac{nV_D D}{logn2(D-1)})^{-1} LLPD(y_i, y_j, C_r)^{-D} > a$ and consider clusters as the all merged connected components of C_r. This produces clusters which are analogous to persistence based clusters as discussed in Chazal et al. (2013). However, choice of ϵ (equivalently a) is subjective and moreover algorithm 3 do not suggest how to merge the clusters when they are not ϵ-stable. Therefore, we want an algorithm which will reliably produce stable level set clusters or cluster tree without any additional suggestion of ϵ. As we shall see in this chapter, LLPD (or density sensitive distance) SC with eigen-gap heuristic achieves that.

5.6.3 Continent problem: formulation

The problem of finding stable level set clusters corresponding to a particular level is analogous to determining *continents* from all landmasses/islands on a planet. Thus we refer to this as the ***continent problem***.

We discussed about topographic prominence of an island in the previous subsection. However, depending on application we might consider other factors as well. 3 factors historically in considerations have been

1. **Relative size** of $C(m_i, b)$, given by

$$p_i \coloneqq \frac{\int_{C(m_i,b)} f}{\int_{M_b} f}.$$

   Small sized islands should be part of a bigger continent.

2. **Topographic isolation** between $C(m_i, b)$ and $C(m_j, b)$, given by

$$dist_1(C(m_i, b), C(m_j, b)) \coloneqq \min_{x \in C(m_i,b),\, y \in C(m_j,b)} dist_1(x, y).$$

   2 islands with small topographic isolation should be in the same continent.

3. **Topographic prominence** of $C(m_i, b)$, given by

$$\text{Prominence}\big(C(m_i, b)\big) \coloneqq \begin{cases} \infty & i = 1 \\ f(m_i) - \max\limits_{j>i} f_{m_i,m_j} & i = 2, 3, .., k. \end{cases}$$

   where $f_{x,y} \coloneqq \max\limits_{\gamma:\gamma_0=x,\gamma_1=y} \min\limits_{t} f(\gamma_t)$ is the bottleneck value of x and y.

   Islands with small prominence should be part of a bigger continent.

Continent problem naturally arises in many geostatistics, genetics, ecology applications.

5.7 LLPD SC with L_{sym}

In this section we derive the theoretical performance guarantee of LLPD SC with L_{sym} using the similarity measure as proposed in equation (5.6) in relation to the

continent problem.

5.7.1 Assumptions

We list our assumptions here again for the sake of completeness.

A1: $M \subseteq \mathbb{R}^D$ is connected, compact and smooth i.e. having a $D-1$ dimensional boundary. We shall work with $M = \mathbb{R}^D$ for sake of simplicity, although our result applies in the general case.

A2: $\int_L f = 0$ for any $L \subset M$ so that $\mu_D(L) = 0$. This in particular ensures that f has no ridges or lower dimensional structure.

A3: f is Morse-Smale. This allows us to split M (except for a lower dimensional set) into finite number of "basin of attractions" $A_1, A_2, ..A_K$.

A4: Level sets of f, $S_\lambda = \{x \in M : f(x) = \lambda\}$ are finite union of $D-1$ dimensional Riemannian manifolds for any $\lambda \in (f_0, f^0)$.

As discussed in last section, in order to make the notion of well-separated clusters precise we assume that $\exists\ b \in (f_0, f^0)$ and $k \leq K$ *"significant"* local maximas of f, namely $m_1, m_2, ..m_k$, so that m_i is the maxima of $C(m_i, b)$ where $C(m_i, b)$ denote the connected component of M_b containing m_i as defined in the previous section. I.e. we assume

$$\textbf{A5: } f(m_i) > b\, \forall i = 1, 2, .., k; M_b = \cup_{i=1}^{k} C(m_i, b). \tag{5.8}$$

We further assume that all $\frac{k(k-1)}{2}$ bottleneck values between these k significant local maximas are between d and a for some $f_0 \leq d \leq a < b$. I.e.,

$$\textbf{A6:} d = \min_{i \neq j,\, i,j \in \{1,2,..,k\}} \text{Bottleneck}(m_i, m_j) \leq \max_{i \neq j,\, i,j \in \{1,2,..,k\}} \text{Bottleneck}(m_i, m_j) = a < b. \tag{5.9}$$

Rationale behind **A5** and **A6** has already been discussed in the previous section. In particular, from **A5** and **A6** and corollary 5.1 it follows that

$$\text{Prominence}(m_i) > b - a \tag{5.10}$$

for all $i = 1(1)k$.

We order $m_1, m_2, .., m_k$ in terms of the cluster volume. I.e.

$$\int_{C(m_1,b)} f \geq \int_{C(m_2,b)} f \geq .. \geq \int_{C(m_k,b)} f$$

.

Next we define some parameters in this context which will be useful in our analysis

$$a^* := e^{-\frac{\frac{2(D-1)}{D}^{\frac{2}{D}}}{\sigma^2}}.$$

$$b^* := e^{-\frac{\frac{2(D-1)}{Db\theta_D}^{\frac{2}{D}}}{\sigma^2}}.$$

$$d^* := e^{-\frac{\frac{2(D-1)}{Dd\theta_D}^{\frac{2}{D}}}{\sigma^2}}.$$

$$\gamma := \max_{i=1(1)k} \frac{N}{|A_i|} = \frac{N}{|A_k|} = \frac{\int_{M_b} f}{\int_{C(m_k,b)} f} \geq k.$$

$$\beta := \min_{i=1(1)k} \frac{N}{|A_i|} = \frac{N}{|A_1|} = \frac{\int_{M_b} f}{\int_{C(m_1,b)} f} \leq k.$$

$$R := \frac{\beta}{\gamma} = \frac{\min\limits_{i=1(1)k} |A_i|}{\max\limits_{i=1(1)k} |A_i|} \to \frac{\int_{C(m_k,b)} f}{\int_{C(m_1,b)} f} \leq 1.$$

$$Q := \frac{b^* - d^*}{1 - a^*}.$$

We formally present the algorithm for LLPD SC with L_{sym} below (algorithm 4).

Algorithm 4 Continent problem with LLPD SC using L_{sym}

Input: Data $\chi_n = \{x_1, x_2, .., x_n\}$, $b \in (f_0, f^0)$.
Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data or indicator of noise.
1: Compute $r = C\,logn$-NN graph C_r on χ_n.
2a: Compute r-NN distance $\sigma_{i,r}$ from each point x_i. Label a point x_i as noise if $\sigma_{i,r} > (\frac{r}{bn\theta_D})^{\frac{1}{D}}$. Let the remaining data points be χ_N with $|\chi_N| = N$.
2b: Form $W_{N\times N}$ with $((W))_{ij} = e^{-\frac{((\frac{n}{logn})^{\frac{1}{D}}\, LLPD(x_i,x_j,C_r))^2}{\sigma^2}}$ for some $\sigma > 0$.
3: Form the normalized Laplacian $L_{sym} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$
4: Compute the eigen decomposition $\{(\phi_k, \lambda_k)\}_{k=1}^N$, sorted so that $0 = \lambda_1 \leq \lambda_2 \leq .. \leq \lambda_N$.
5: Estimate the number of clusters k_c by *eigen-gap statistic* $k_c := \underset{i}{argmax}\, \lambda_{i+1} - \lambda_i$. Let $U_{N\times k} = (\phi_1 : \phi_2 : .. : \phi_{k_c})$.
6: Compute the label $Z_{N\times 1}$ by running K-means on $\{U_{i,.}\}_{i=1}^N$ using k_c as the number of clusters and $Z_{i,.}$ is label of $U_{i,.}$ as obtained by the K-means.

Note that, we are not performing the normalization of spectral embedding step (step 4 of algorithm 2) since it is not required to obtain a better performance of LLPD SC (as noted in Little et al. (2017) as well).

Also, this algorithm (algorithm 4) can be run at different levels to obtain stable cluster trees of f.

5.7.2 Main result: equivalence with prominence based clustering

Our main result on LLPD SC with L_{sym} in context of continent problem is loosely the following:

Theorem 5.1

If $b - a$ is sufficiently large (i.e. all islands are prominent) and $a - d$ is sufficiently small (i.e. bottleneck values between islands are of the same order), then $\lambda_{k+1} - \lambda_k$ is indeed the largest eigen-gap of L_{sym} (i.e. $k_c = k$) and algorithm 4 achieves near

100% accuracy in clustering $\{C(m_i, b)\}_{i=1}^{k}$ when we choose $r \geq C\, logn$ for some sufficiently large $C > 0$. In this case, islands becomes continents and no merging occurs. Otherwise $k_c < k$ i.e. merging of some islands occurs. Almost surely, $k_c > k$ never occurs. □

Thus, LLPD SC with L_{sym} is very similar to prominence based clustering without it's computational burden.

Remark 5.3

Let $\{A_1, .., A_{k_r}\}$ denote the connected components of χ_N in C_r. From theorem 5 in Zhao and Lai (2020) (or theorem 3.4 in Chaudhuri et al. (2014)) it follows that w.p. at least $1-\delta$, as $n \to \infty$ these k_r connected components correspond to the connected components of $M_{b-O(\frac{1}{log\frac{1}{\delta}})}$ if $r = O(log^2\frac{1}{\delta})\, logn$. Therefore from **A6** it follows that with very high probability, $k_r = k$ for sufficiently large C and n. We assume this set-up for our analysis on the performance of SC. Thus in this case we have (after a possible re-labeling of A_i's), $N \approx n\int_{M_b} f$ and $|A_i| \approx n\int_{C(m_i,b)} f$ for $i = 1(1)k$. □

5.7.3 Derivation of main result

Recall from LLPD's convergence limit result in chapter 3 that under the assumption **A4**

$$LLPD(x, y, \chi_n) \to (\frac{logn}{n\, V_D}\frac{2(D-1)}{D\, f_{x,y}})^{\frac{1}{D}} \text{a.s.} \tag{5.11}$$

where $f_{x,y} = \underset{\gamma}{max}\underset{t}{min} f(\gamma_t)$, γ_t is any piece-wise continuous curve with $\gamma(0) = x$ and $\gamma(1) = y$.

As a consequence of this and our assumptions we have

Proposition 5.2

As $n \to \infty$, $C \to \infty$, for $x_i, x_j \in \chi_N$ we have a.s.

$$(\frac{n\,V_D}{log\,n})^{\frac{1}{D}}\,LLPD(x_i,x_j,C_r) \in \begin{cases} (0,\frac{2(D-1)}{Db}^{\frac{1}{D}}] & \text{if } x_i \text{ and } x_j \text{ are from same c.c.} \\ (\frac{2(D-1)}{Da}^{\frac{1}{D}},\frac{2(D-1)}{Dd}^{\frac{1}{D}}] & \text{if } x_i \text{ and } x_j \text{ are from diff. c.c.} \end{cases}$$

where c.c. denotes connected components of C_r. □

Proof 5.4

If x_i and x_j are in the same connected component of C_r (i.e. in the same connected component of M_b for sufficiently large C and n) then for any piecewise continuous path γ_t joining x_i x_j and staying in the common connected component that both x_i and x_j are a part of, $f(\gamma_t) \geq b\ \forall t \in [0,1]$ by **A5**. Thus for sufficiently large n, $LLPD(x_i,x_j,C_r) \leq (\frac{logn}{n\,V_D}\frac{2(D-1)}{D})^{\frac{1}{D}}b^{-\frac{1}{D}}$. For the lower bound, if x_i and x_j are nearest neighbor (then of course direct edge joining x_i and x_j will be the LLPD between them) then $|x_i - x_j| = O_p(n^{-\frac{1}{D}})$, thus in that case $(\frac{n\,V_D}{logn})^{\frac{1}{D}}LLPD(x_i,x_j,C_r) = O_p(logn^{-\frac{1}{D}}) \to 0$ as $n \to \infty$.

On the other hand, if x_i and x_j are in different connected components of C_r (i.e. in different connected components of M_b say $x_i \in C(m_l,b)$ and $x_j \in C(m_k,b)$), then any LLPD path joining x_i and x_j have to pass through a bottleneck point having the bottleneck value Bottleneck (m_k,m_l). Thus, by **A6** we have $(\frac{logn}{n\,V_D}\frac{2(D-1)}{D})^{\frac{1}{D}}a^{-\frac{1}{D}} \leq LLPD(x_i,x_j,\chi_n) \leq (\frac{logn}{n\,V_D}\frac{2(D-1)}{D})^{\frac{1}{D}}d^{-\frac{1}{D}}$ when C and n are sufficiently large. □

Therefore, from our choice of similarity measure as in equation (5.6) we have as $n \to \infty$, a.s.

$$((W))_{ij} \in \begin{cases} [b^*,1) & \text{If } x_i \text{ and } x_j \text{ are from same connected component of } C_r \\ [d^*,a^*) & \text{If } x_i \text{ and } x_j \text{ are from different connected component of } C_r \end{cases} \tag{5.12}$$

Note that we have the following series of inequalities

$$0 < d^* \leq a^* < b^* < 1. \tag{5.13}$$

Note that, (5.12) is similar to the structure of W which has been analyzed in Little et al. (2017)'s section 5.1 'Ultrametric Cluster Model'. Our analysis of SC, in this section as well as in this chapter in-general, will follow some of the techniques developed by them in order to derive theoretical performance guarantee.

Here we state our main result for LLPD SC with L_{sym} formally

Theorem 5.2

As $n \to \infty$, $C \to \infty$, w.p. $\geq 1 - O(logn^{-B}) \to 1$ for any $B < \frac{e(D-1)}{2}$,

i)$\lambda_{k+1} - \lambda_k$ is the largest gap (i.e. $k_c = k$) in the eigenvalues of L_{sym} provided

$$\frac{(p-q)}{\gamma} > max\{\frac{3P_2+1}{2}, 4P_2\} \tag{5.14}$$

where $p \coloneqq \frac{b^*}{\frac{1}{\beta}+(1-\frac{1}{\beta})a^*}$ and $q \coloneqq \frac{a^*}{\frac{1}{\gamma}b^*+(1-\frac{1}{\gamma})d^*}$ and P_2 is defined in equation (5.17). Moreover if

ii)

$$\frac{C}{k^3\gamma^3} \geq P_2 \tag{5.15}$$

where C is an absolute constant, then algorithm 4 perfectly recovers cluster labels on χ_N as the connected components of M_b. □

Proof 5.5

In the remaining of this section, we shall ignore the r-NN graph C_r construction and work with $LLPD(x, y, \chi_n)$ instead in order to derive the theoretical performance guarantee. Recall that from proposition 5.1 we know w.p. $\to 1$ $LLPD(x, y, C_r) =$

$LLPD(x, y, \chi_n)$ for $r = C\,logn$ for sufficiently large $C > 0$, thus this won't affect our result in anyway.

We start by decomposing $W_{N\times N} = U + V$ where $U \coloneqq \mathbb{E}[W]$ and $V \coloneqq W - U$ and focus on V first.

Lemma 5.2

As $n \to \infty$ we have $\mathbb{P}(|\underset{i,j}{max}((V))_{ij}| > \epsilon) \leq A^{'} n^{-B\epsilon}$ for some $A > 0$ and any $B < \frac{e(D-1)}{2}$. In particular, w.p. $\geq 1 - O(logn^{-B})$, $|\underset{i,j}{max}((V))_{ij}| < \frac{loglogn}{logn} \to 0$. □

Proof 5.6

W.l.g. we can focus on only those data points x_i, x_j in χ_N so that $|x_i - x_j| = O(1)$ since contribution of other data points are negligible (i.e. $\frac{|\{(i,j):|x_i - x_j| < O(1)\}|}{N} \to 0$) in W as $n \to \infty$.

We know from LLPD's convergence result in chapter 3 that when $|x - y| = O(1)$,

$$\mathbb{P}(\underset{x,y}{sup}|\frac{LLPD(x, y, \chi_n)}{(\frac{logn}{n\,V_D}\frac{2(D-1)}{D\,f_{x,y}})^{\frac{1}{D}}} - 1| > \epsilon) \leq An^{-(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}} \tag{5.16}$$

for some $A > 0$ and any $\eta < 1$.

Let $g(z) = e^{-\frac{z^2}{\sigma^2}}$ for $z \in \mathbb{R}$. Then from (5.16) it follows,

$$\mathbb{P}\big(\underset{x,y}{sup}|(\frac{n\,V_D}{logn})^{\frac{1}{D}} LLPD(x, y, \chi_n) - (\frac{2(D-1)}{D\,f_{x,y}})^{\frac{1}{D}}| > \epsilon\big)$$

$$\leq An^{-(\frac{2(D-1)}{D\,f_{x,y}})^{-\frac{1}{D}}(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}}.$$

$$\Rightarrow \mathbb{P}\big(\underset{x,y}{sup}|g((\frac{n\,V_D}{logn})^{\frac{1}{D}} LLPD(x, y, \chi_n)) - g([\frac{2(D-1)}{D\,f_{x,y}}]^{\frac{1}{D}})| > \epsilon\,|g^{'}([\frac{2(D-1)}{D\,f_{x,y}}]^{\frac{1}{D}})|\big)$$

$$\leq A^{'} n^{-(\frac{2(D-1)}{D\,f_{x,y}})^{-\frac{1}{D}}(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}}$$

where $|g^{'}(z)| = \frac{2z}{\sigma^2} e^{-\frac{z^2}{\sigma^2}}$.

$$\Rightarrow \mathbb{P}\big(\underset{x,y}{sup}|g((\frac{n\,V_D}{logn})^{\frac{1}{D}} LLPD(x, y, \chi_n)) - g([\frac{2(D-1)}{D\,f_{x_i,x_j}}]^{\frac{1}{D}})| > \epsilon\big) \leq$$

$$A^{'} n^{-(\frac{2(D-1)}{D f_{x,y}})^{-\frac{1}{D}}(D-1)\eta\epsilon \frac{\sigma^2}{2[\frac{2(D-1)}{D f_{x,y}}]^{\frac{1}{D}}} e^{\frac{[\frac{2(D-1)}{D f_{x,y}}]^{\frac{2}{D}}}{\sigma^2}}} log\, n^{\frac{D-1}{D}}$$

$$= A^{'} n^{-(\frac{2(D-1)}{D f_{x,y}})^{-\frac{2}{D}} \frac{\sigma^2}{2} e^{\frac{[\frac{2(D-1)}{D f_{x,y}}]^{\frac{2}{D}}}{\sigma^2}}(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}}$$

Now since $h(z) = ze^{\frac{b}{z}} 1(z > 0)$ has a unique minima at $z = b$ with $z = (\frac{2(D-1)}{D f_{x,y}})^{-\frac{2}{D}}$ and $b = \frac{1}{\sigma^2}$, we conclude that $h(z) \geq h(\frac{1}{\sigma^2})$. Thus conclusion follows from observing

$$A^{'} n^{-(\frac{2(D-1)}{D f_{x,y}})^{-\frac{2}{D}} \frac{\sigma^2}{2} e^{\frac{[\frac{2(D-1)}{D f_{x,y}}]^{\frac{2}{D}}}{\sigma^2}}(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}} \leq A^{'} n^{-\frac{e}{2}(D-1)\eta\epsilon} log\, n^{\frac{D-1}{D}} \leq A^{'} n^{-B\epsilon}$$

for any $B < \frac{e}{2}(D-1)$. □

Proposition 5.3

As $n \to \infty$, largest eigen-gap of L_{sym} is determined by largest eigen-gap of $I - D^{-\frac{1}{2}} U D^{-\frac{1}{2}}$. □

Proof 5.7

Given our definition of U and V we have

$$W = U + V$$

Thus,

$$L_{sym} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}} = I - D^{-\frac{1}{2}} U D^{-\frac{1}{2}} - D^{-\frac{1}{2}} V D^{-\frac{1}{2}}$$

We are interested in the spectral perturbation P_0 in the eigenvalues of $I - D^{-\frac{1}{2}} U D^{-\frac{1}{2}}$ due to the matrix $D^{-\frac{1}{2}} V D^{-\frac{1}{2}}$. Clearly P_0 is bounded by

$$||D^{-\frac{1}{2}} V D^{-\frac{1}{2}}||_2 \leq N||D^{-\frac{1}{2}} V D^{-\frac{1}{2}}||_{max} \leq N\frac{||V||_{max}}{||D||_{min}}$$

$$\leq N\frac{||V||_{max}}{N(b^*\frac{1}{\gamma} + d^*(1-\frac{1}{\gamma}))} \leq \frac{||V||_{max}}{d^*}.$$

Thus, from lemma 5.2 it follows that w.p. $\geq 1 - O((logn)^{-B}) \to 1$, P_0 is bounded by $O(\frac{log\, logn}{logn}) \to 0$ as $n \to \infty$. Thus, largest eigen-gap of L_{sym} is determined by the eigen-gaps of $I - D^{-\frac{1}{2}}UD^{-\frac{1}{2}}$ and the perturbation due to $D^{-\frac{1}{2}}VD^{-\frac{1}{2}}$ goes to 0 as $n \to \infty$. □

Hence, for the largest eigen-gap analysis, it is sufficient to work with $I - D^{-\frac{1}{2}}UD^{-\frac{1}{2}}$, or equivalently with $D^{-\frac{1}{2}}UD^{-\frac{1}{2}}$.

W.l.g. we assume that χ_N is permuted so that $\chi_N = \{A_1, .., A_k\}$ with $|A_1| \geq |A_2|.. \geq |A_k|$. Then U have the following structure as specified by (5.12)

$$U = \begin{bmatrix} U_{A_1,A_1} & U_{A_1,A_2} & . & . & U_{A_1,A_k} \\ U_{A_2,A_1} & U_{A_2,A_2} & . & . & U_{A_2,A_k} \\ . & . & . & . & . \\ . & . & . & . & . \\ U_{A_k,A_1} & U_{A_k,A_2} & . & . & U_{A_k,A_k} \end{bmatrix} = \begin{bmatrix} U_{A_1,A_1} & u^{1,2}J & . & . & u^{1,k}J \\ u^{2,1}J & U_{A_2,A_2} & . & . & u^{2,k}J \\ . & . & . & . & . \\ . & . & . & . & . \\ u^{k,1}J & u^{k,2}J & . & . & U_{A_k,A_k} \end{bmatrix},$$

where U_{A_i,A_j} is a $|A_i| \times |A_j|$ matrix for $i, j = 1(1)k$ with

$$u^{i,j} = e^{-\frac{\frac{2(D-1)}{D\, Bottleneck(m_i,m_j)\, V_D}\frac{2}{D}}{\sigma^2}}$$

and J is matrix of 1's. Note that by **A6** we have $0 < d^* \leq u^{i,j} \leq a^*$ for $i \neq j$,

Note that $D^{-\frac{1}{2}}UD^{-\frac{1}{2}}$ can be written as

$$D^{-\frac{1}{2}}UD^{-\frac{1}{2}} := \begin{bmatrix} D_{A_1}^{-\frac{1}{2}}U_{A_1,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_1}^{-\frac{1}{2}}U_{A_1,A_2}D_{A_2}^{-\frac{1}{2}} & . & . & D_{A_1}^{-\frac{1}{2}}U_{A_1,A_k}D_{A_k}^{-\frac{1}{2}} \\ D_{A_2}^{-\frac{1}{2}}U_{A_2,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_2}^{-\frac{1}{2}}U_{A_2,A_2}D_{B_2}^{-\frac{1}{2}} & . & . & D_{A_2}^{-\frac{1}{2}}U_{A_2,A_2}D_{A_k}^{-\frac{1}{2}} \\ . & . & . & . & . \\ . & . & . & . & . \\ D_{A_k}^{-\frac{1}{2}}U_{A_k,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_k}^{-\frac{1}{2}}U_{A_k,A_2}D_{A_2}^{-\frac{1}{2}} & . & . & D_{A_k}^{-\frac{1}{2}}U_{A_k,A_k}D_{A_k}^{-\frac{1}{2}} \end{bmatrix}$$

Proposition 5.4

In this notation, we have

$$((D^{-\frac{1}{2}}UD^{-\frac{1}{2}}))_{mn} \in \begin{cases} \left[\frac{b^*}{\frac{N}{\beta}+(N-\frac{N}{\beta})a^*}, \frac{1}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}\right) & \text{if } x_m \text{ and } x_n \text{ are from same c.c.} \\ \left[\frac{d^*}{\frac{N}{\beta}+(N-\frac{N}{\beta})a^*}, \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}\right) & \text{if } x_m \text{ and } x_n \text{ are from diff. c.c.} \end{cases}$$

where c.c. denotes connected components of C_r. □

Proof 5.8

Recall that

$$\gamma \coloneqq \underset{i=1(1)k}{max} \frac{N}{|A_i|} \geq k$$

and

$$\beta \coloneqq \underset{i=1(1)k}{min} \frac{N}{|A_i|} \leq k.$$

From equation (5.12), it follows that for any data point in A_i it's degree is within $[|A_i|b^* + (N - |A_i|)d^*, |A_i| + (N - |A_i|)a^*]$. Thus, for any $x_m \in A_i$, $x_n \in A_i$ and $i = 1(1)k$ we have

$$\frac{b^*}{|A_i| + (N - |A_i|)a^*} \leq ((D^{-\frac{1}{2}}UD^{-\frac{1}{2}}))_{mn}$$

$$= ((D_{A_i}^{-\frac{1}{2}}U_{A_i,A_i}D_{A_i}^{-\frac{1}{2}}))_{mn} < \frac{1}{|A_i|b^* + (N - |A_i|)d^*}.$$

Furthermore, for all $i = 1(1)k$ we have

$$\frac{b^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*} \leq \frac{b^*}{|A_i| + (N - |A_i|)a^*}$$

and

$$\frac{1}{|A_i|b^* + (N - |A_i|)d^*} \leq \frac{1}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*}.$$

Thus, for any $i = 1(1)k$ and $m \neq n$

$$\frac{b^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*} \leq ((D^{-\frac{1}{2}}UD^{-\frac{1}{2}}))_{mn} < \frac{1}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*}.$$

Similarly, for any $x_m \in A_i$ and $x_n \in A_j$, $i \neq j$ we have

$$\frac{d^*}{\sqrt{(|A_i| + (N - |A_i|)a^*)\,(|A_j| + (N - |A_j|)a^*)}} \leq ((D^{-\frac{1}{2}}UD^{-\frac{1}{2}}))_{mn}$$

$$< \frac{a^*}{\sqrt{(|A_i|b^* + (N - |A_i|)d^*)(|A_j|b^* + (N - |A_j|)d^*)}}$$

and for any $i \neq j$ we have

$$\frac{d^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*} \leq \frac{d^*}{\sqrt{(|A_i| + (N - |A_i|)a^*)\,(|A_j| + (N - |A_j|)a^*)}}$$

and

$$\frac{a^*}{\sqrt{(|A_i|b^* + (N - |A_i|)d^*)(|A_j|b^* + (N - |A_j|)d^*)}} \leq \frac{a^*}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*}.$$

Thus, for any $x_m \in A_i$ and $x_n \in A_j$, $i \neq j$ we have

$$\frac{d^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*} \leq ((D^{-\frac{1}{2}}UD^{-\frac{1}{2}}))_{mn} < \frac{a^*}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*}.$$ □

Now we consider the following decomposition of $D^{-\frac{1}{2}}UD^{-\frac{1}{2}}$.

$$D^{-\frac{1}{2}}UD^{-\frac{1}{2}} = \begin{bmatrix} D_{A_1}^{-\frac{1}{2}}U_{A_1,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_1}^{-\frac{1}{2}}U_{A_1,A_2}D_{A_2}^{-\frac{1}{2}} & . & . & D_{A_1}^{-\frac{1}{2}}U_{A_1,A_k}D_{A_k}^{-\frac{1}{2}} \\ D_{A_2}^{-\frac{1}{2}}U_{A_2,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_2}^{-\frac{1}{2}}U_{A_2,A_2}D_{B_2}^{-\frac{1}{2}} & . & . & D_{A_2}^{-\frac{1}{2}}U_{A_2,A_2}D_{A_k}^{-\frac{1}{2}} \\ . & . & . & . & . \\ . & . & . & . & . \\ D_{A_k}^{-\frac{1}{2}}U_{A_k,A_1}D_{A_1}^{-\frac{1}{2}} & D_{A_k}^{-\frac{1}{2}}U_{A_k,A_2}D_{A_2}^{-\frac{1}{2}} & . & . & D_{A_k}^{-\frac{1}{2}}U_{A_k,A_k}D_{A_k}^{-\frac{1}{2}} \end{bmatrix}$$

$$\xrightarrow{P_1} \begin{bmatrix} \frac{b^*}{\frac{N}{\beta}+(N-\frac{N}{\beta})a^*}J_{|A_1|\times|A_1|} & \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_1|\times|A_2|} & . & . & \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_1|\times|A_k|} \\ \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_2|\times|A_1|} & \frac{b^*}{\frac{N}{\beta}+(N-\frac{N}{\beta})a^*}J_{|A_2|\times|A_2|} & . & . & \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_2|\times|A_k|} \\ . & . & . & . & . \\ . & . & . & . & . \\ \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_k|\times|A_1|} & \frac{a^*}{\frac{N}{\gamma}b^*+(N-\frac{N}{\gamma})d^*}J_{|A_k|\times|A_2|} & . & . & \frac{b^*}{\frac{N}{\beta}+(N-\frac{N}{\beta})a^*}J_{|A_k|\times|A_k|} \end{bmatrix}$$

$$\coloneqq B_{N\times N}$$

and consider the spectral perturbation of $I - B$ due to P_1 and eigen-gaps of $I - B$, which is our *idealized* L_{sym}.

Recall the definition of cluster balance parameter,

$$R \coloneqq \frac{\beta}{\gamma} = \frac{\underset{i=1(1)k}{min}|A_i|}{\underset{i=1(1)k}{max}|A_i|} \to \frac{\int_{C(m_k,b)} f}{\int_{C(m_1,b)} f}.$$

Clearly $0 \leq R \leq 1$. $R \approx 1$ when all k clusters (connected components of M_b) have similar D dimensional Lebesgue measure.

and cluster similarity parameter

$$Q \coloneqq \frac{b^* - d^*}{1 - a^*}.$$

Clearly $0 \leq Q \leq \frac{1}{1-a^*}$. $Q \approx 1$ if $1 - b^*$ and $a^* - d^*$ are small relative to $1 - a^*$.

Then we have

Proposition 5.5

$$P_1 \leq \frac{(a^* - d^*b^*) + \frac{1}{\beta}(1 - a^*)\, max\{a^* - d^*QR, 1 - b^*QR\}}{(\frac{RQ}{\beta}(1 - a^*) + d^*)(\frac{1}{\beta}(1 - a^*) + a^*)} \coloneqq P_2 \tag{5.17}$$

In particular, P_1 is small when $Q \approx 1$ and $R \approx 1$, i.e. when all k clusters (connected components of M_b) have similar D dimensional Lebesgue measure and $1 - b^*$ and $a^* - d^*$ are small relative to $1 - a^*$. □

Proof 5.9

From proposition 5.4, P_1 is bounded by

$$N\, max\{\frac{1}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*} - \frac{b^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*}, \frac{a^*}{\frac{N}{\gamma}b^* + (N - \frac{N}{\gamma})d^*} - \frac{d^*}{\frac{N}{\beta} + (N - \frac{N}{\beta})a^*}\}$$

Now,

$$\frac{1}{\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*} - \frac{b^*}{\frac{1}{\beta} + (1-\frac{1}{\beta})a^*}$$

$$= \frac{\frac{1}{\beta} + (1-\frac{1}{\beta})a^* - b^*(\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*)}{(\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*)(\frac{1}{\beta} + (1-\frac{1}{\beta})a^*)}$$

$$= \frac{(a^* - b^*d^*) + \frac{1}{\beta}(1-a^*) - \frac{RQ}{\beta}b^*(1-a^*)}{(\frac{RQ}{\beta}(1-a^*) + d^*)(\frac{1}{\beta}(1-a^*) + a^*)}$$

$$= \frac{(a^* - b^*d^*) + \frac{1}{\beta}(1-a^*)(1-b^*QR)}{(\frac{RQ}{\beta}(1-a^*) + d^*)(\frac{1}{\beta}(1-a^*) + a^*)}$$

Similarly,

$$\frac{a^*}{\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*} - \frac{d^*}{\frac{1}{\beta} + (1-\frac{1}{\beta})a^*}$$

$$= \frac{a^*(\frac{1}{\beta} + (1-\frac{1}{\beta})a^*) - d^*(\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*)}{(\frac{RQ}{\beta}(1-a^*) + d^*)(\frac{1}{\beta}(1-a^*) + a^*)}$$

$$\leq \frac{(a^{*2} - d^{*2}) + \frac{1}{\beta}(1-a^*)(a^* - d^*QR)}{(\frac{RQ}{\beta}(1-a^*) + d^*)(\frac{1}{\beta}(1-a^*) + a^*)}$$

Conclusion follows by noting

$$a^{*2} - d^{*2} = a^* - b^*d^* - (1-a^*)(a^* - d^*Q) \leq a^* - b^*d^*$$

In particular, it is easy to see that P_1 can be made arbitrarily small by choosing $a^* - d^*$, $1 - b^*$and $1 - R$ arbitrarily small. □

Now we derive the eigenvalues and eigenvectors for our *idealized* Laplacian $I - B$.

Lemma 5.3

Let

$$C_{N\times N} = \begin{bmatrix} pJ_{o_1\times o_1} & qJ_{o_1\times o_2} & . & . & qJ_{o_1\times o_k} \\ qJ_{o_2\times o_1} & pJ_{o_2\times o_2} & . & . & qJ_{o_2\times o_k} \\ . & . & . & . & . \\ . & . & . & . & . \\ qJ_{o_k\times o_1} & qJ_{o_k\times o_2} & . & . & pJ_{o_k\times o_k} \end{bmatrix}$$

$$= qJ_{N\times N} + (p-q) \text{ Blockdiag } (J_{o_1\times o_1}, J_{o_2\times o_2}, .., J_{o_k\times o_k}) = B_1 + B_2$$

for $p > q$ and $N = \sum_{i=1}^{k} o_i$. Also $o_1 \geq o_2 \geq .. \geq o_k$.

Suppose eigenvalues of $I - C$ are $\lambda_1 \leq \lambda_2 .. \leq \lambda_N$. Then we have

$$\lambda_i(I-C) = \begin{cases} 1-(Nq + o_1(p-q)) & i = 1 \\ 1 - o_i(p-q) & i = 2(1)k \\ 1 & i = k+1(1)N \end{cases}$$

□

Proof 5.10

B_1 has leading eigenvalue Nq with corresponding eigenvector $\propto \frac{1}{\sqrt{N}} 1_{N\times 1}$ and remaining $N-1$ eigenvalues are 0 with corresponding eigenvectors $\perp 1_{N\times 1}$. Similarly, B_2 has top k eigenvalues as $o_i(p-q)$ with the requirements for the corresponding eigenvectors to be $\propto 1$ on the i-th block respectively, and $o_i - 1$ eigenvalues as 0 with the requirements for the corresponding eigenvectors to be $\perp 1$ on the i-th block for $i = 1(1)k$.

Let eigenvalues of C be $\kappa_1 \geq \kappa_2 .. \geq \kappa_N$ and corresponding eigenvectors are $\phi_1 = \frac{1}{\sqrt{N}} 1_{N\times 1}$, $\phi_2, .., \phi_N$.

We claim that

$$\kappa_i = \begin{cases} Nq + o_1(p-q) & i = 1 \\ o_i(p-q) & i = 2(1)k \\ 0 & i = k+1(1)N \end{cases}$$

This means eigenvalues of C are just sum of eigenvalues of B_1 and B_2. But for that we need to be able to choose ϕ_i to be $\perp 1_{N\times 1}$ as well as (w.l.g.) constant on i-th block for $i = 2(1)k$ and ϕ_i to be $\perp 1_{N\times 1}$ and also $\perp$ on 2nd block for $i = k+1(1)k+o_1-1$, $\perp$ on the 3rd block for $i = k+o_1(1)k+o_1+o_2-2$ and so on. Thus in-particular ϕ_i's can be chosen like this for $i = 1(1)N$:

$$\phi_1 = \frac{1}{\sqrt{N}} 1_{N\times 1}$$

$$\phi_2 = \frac{1}{\sqrt{o_1 \frac{\sum_{j=1}^k o_j}{\sum_{j=2}^k o_j}}} (1_{1\times o_1} : -\frac{o_1}{\sum_{j=2}^k o_j} 1_{1\times o_2} : -\frac{o_1}{\sum_{j=2}^k o_j} 1_{1\times o_3} : .. : -\frac{o_1}{\sum_{j=2}^k o_j} 1_{1\times o_k})^T$$

$$\phi_3 = \frac{1}{\sqrt{o_2 \frac{\sum_{j=2}^k o_j}{\sum_{j=3}^k o_j}}} (0_{1\times o_1} : 1_{1\times o_2} : -\frac{o_2}{\sum_{j=3}^k o_j} 1_{1\times o_3} : -\frac{o_2}{\sum_{j=3}^k o_j} 1_{1\times o_4} : .. : -\frac{o_2}{\sum_{j=3}^k o_j} 1_{1\times o_k})^T$$

$$\phi_4 = \frac{1}{\sqrt{o_3 \frac{\sum_{j=3}^k o_j}{\sum_{j=4}^k o_j}}} (0_{1\times o_1} : 0_{1\times o_2} : 1_{1\times o_3} : -\frac{o_3}{\sum_{j=4}^k o_j} 1_{1\times o_4} : .. : -\frac{o_3}{\sum_{j=4}^k o_j} 1_{1\times o_k})^T$$

.

.

$$\phi_k = \frac{1}{\sqrt{o_{k-1} \frac{o_{k-1}+o_k}{o_k}}} (0_{1\times o_1} : 0_{1\times o_2} : 0_{1\times o_3} : .. : 1_{1\times o_{k-1}} : -\frac{o_{k-1}}{o_k} 1_{1\times o_k})^T$$

and

$$\phi_i = (c_{i1\times o_1} : 0_{1\times o_2} : 0_{1\times o_3} : .. : 0_{1\times o_k})^T \text{ for } i = k+1(1)k+o_1-1$$

where c_i's are $o_1 - 1$ orthonormal contrasts in o_1 dimension.

$$\phi_i = (0_{1\times o_1} : c_{i1\times o_2} : c_{i\ 1\times o_3} : .. : 0_{1\times o_k})^T \text{ for } i = k + o_1(1)k + o_1 + o_2 - 2$$

where c_i's are $o_2 - 1$ orthonormal contrasts in o_2 dimension and so on..

$$\phi_i = (0_{1\times o_1} : 0_{1\times o_2} : 0_{1\times o_3} : .. : c_{i1\times o_k})^T \text{ for } i = \sum_{j=1}^{k-1} o_j + 1(1)N$$

where c_i's are $o_k - 1$ orthonormal contrasts in o_k dimension. It is easy to verify that ϕ_i's are all orthonormal vectors for $i = 1(1)N$.

Conclusion follows after noticing if eigenvalues of $I - C$ are $\lambda_1 \leq \lambda_2 .. \leq \lambda_N$, then we have $\lambda_i = 1 - \kappa_i$ for $i = 1(1)N$ and eigenvector corresponding to λ_i is ϕ_i. □

Corollary 5.2

Suppose eigenvalues of the idealized Laplacian $I - B$ are $\lambda_1 \leq \lambda_2 .. \leq \lambda_N$ with corresponding eigenvectors ϕ_i for $i = 1(1)N$. Then we have

$$\lambda_i(I - B) = \begin{cases} 1 - (q + p_1(p - q)) & i = 1 \\ 1 - p_i(p - q) & i = 2(1)k \\ 1 & i = k + 1(1)N \end{cases}$$

with $p = \frac{b^*}{\frac{1}{\beta}+(1-\frac{1}{\beta})a^*}$ and $q = \frac{a^*}{\frac{1}{\gamma}b^*+(1-\frac{1}{\gamma})d^*}$, $\frac{|A_i|}{N} = \frac{\int_{C(m_i,b)} f}{\int_{M_b} f} := p_i$ for $i = 1(1)k$, with the understanding that $p > q$ (this happens when $b^* - a^*$ is sufficiently large, $a^* - d^*$ is sufficiently small and R is not too small from 1). Note that we have $p_1 \geq p_2 .. \geq p_k$, $p_1 = \frac{1}{\beta}$ and $p_k = \frac{1}{\gamma}$, $\sum_{i=1}^{k} p_i = 1$. Additionally, ϕ_i's are same as given by lemma 5.3 with $o_i = |A_i|$. □

Proof 5.11

Follows immediately from lemma 5.3. □

Lemma 5.4

If $\frac{(p-q)}{\gamma} > max\{\frac{3P_2+1}{2}, 4P_2\}$, then indeed $\triangle_k$ is the largest eigen-gap of L_{sym}. □

Proof 5.12

Since $|\lambda_i(L_{sym}) - \lambda_i(I-B)| \leq ||B - D^{-\frac{1}{2}}UD^{-\frac{1}{2}}||_2 \leq P_1$ and L_{sym} is positive semi-definite, by Hoffman-Wielandt theorem,

$$\lambda_i(L_{sym}) \in \begin{cases} 0 & i = 1 \\ [0 \vee 1 - p_i(p-q) - P_1, 1 - p_i(p-q) + P_1] & i = 2(1)k \\ [1 - P_1, 1 + P_1] & i = k+1(1)N \end{cases}$$

Thus, for $i = 1(1)k-1$

$$\triangle_i(L_{sym}) \leq 1 - p_{i+1}(p-q) + P_1 \leq 1 - p_k(p-q) + P_1$$

for $i = k$,

$$\triangle_i(L_{sym}) \geq 1 - P_1 - (1 - p_k(p-q) + P_1)$$

$$= p_k(p-q) - 2P_1$$

and for $i > k$,

$$\triangle_i(L_{sym}) \leq 2P_1$$

Thus $\triangle_k$ is the largest eigen-gap of L_{sym} if

$$1 - p_k(p-q) + P_1 < p_k(p-q) - 2P_1$$

$$\Rightarrow p_k(p-q) > \frac{3P_1 + 1}{2}$$

and

$$2P_1 < p_k(p-q) - 2P_1$$

i.e.

$$\Rightarrow p_k(p-q) > 4P_1$$

Hence, the condition for $\triangle_k$ being the largest eigen-gap of L_{sym} reduces to

$$\frac{(p-q)}{\gamma} > max\{\frac{3P_1+1}{2}, 4P_1\}$$

□

Proposition 5.6

If $1-R>0$, $1-b^*>0$ and $a^*-d^*>0$ are all sufficiently small with $1-a^*$ being sufficiently away from 0, then our eigen-gap condition holds. □

Proof 5.13

Clearly, from equation (5.17) it follows that $P_1 \to 0$ as $max\{1-R, 1-b^*, a^*-d^*\} \to 0$.

Then additionally suppose $a^* < \frac{1}{2}$. Then by recalling $p = \frac{b^*}{\frac{1}{\beta}+(1-\frac{1}{\beta})a^*}$ and $q = \frac{a^*}{\frac{1}{\gamma}b^*+(1-\frac{1}{\gamma})d^*}$ we have,

$$\frac{p-q}{\gamma} \approx \frac{1-a^*}{\frac{1}{k}+(1-\frac{1}{k})a^*} > \frac{1-a^*}{(1+\frac{1}{k})(1-a^*)} = \frac{k}{k+1} \geq \frac{1}{2}$$

Thus, indeed $\exists R, a^*, b^*, d^*$ so that

$$\frac{(p-q)}{\gamma} > max\{\frac{3P_1+1}{2}, 4P_1\}$$

□

This concludes the proof of the part i of theorem 5.2.

For part ii, we proceed similarly as in Little et al. (2017).

Let columns of $\phi_{N\times k} = (\phi_1, .., \phi_k)$ are first k eigen-vectors of B and columns of $\tilde{\phi}_{N\times k}$ are first k eigen-vectors of $D^{-\frac{1}{2}}WD^{-\frac{1}{2}} = I - L_{sym}$.

We define coherence of ϕ by

$$coh(\phi) \coloneqq \frac{N}{k} \underset{i=1(1)N}{max} \sum_{j=1}^{k} \phi_{ij}^2$$

Proposition 5.7

$coh(\phi) \leq \gamma$. □

Proof 5.14

Recall that $p_i = \frac{|A_i|}{N}$ with $\sum_{i=1}^{k} p_i = 1$ and $\frac{1}{p_1} = \beta$ and $\frac{1}{p_k} = \gamma$.

Note that if i is in 1st cluster, then

$$\frac{N}{k}\sum_{j=1}^{k}\phi_{ij}^2 = \frac{N}{k}(\frac{1}{N} + \frac{1}{Np_1}\sum_{i=2}^{k} p_i) = \frac{1}{k}(1 + \frac{\sum_{i=2}^{k} p_i}{p_1}) \leq \frac{2}{k}\gamma \leq \gamma$$

If i is in 2nd cluster, then

$$\frac{N}{k}\sum_{j=1}^{k}\phi_{ij}^2 = \frac{N}{k}(\frac{1}{N} + \frac{p_1}{N\sum_{i=2}^{k} p_i} + \frac{\sum_{i=3}^{k} p_i}{Np_2\sum_{i=2}^{k} p_i})$$

$$= \frac{1}{k}(1 + \frac{p_1}{\sum_{i=2}^{k} p_i} + \frac{\sum_{i=3}^{k} p_i}{p_2\sum_{i=2}^{k} p_i}) \leq \frac{3}{k}\gamma \leq \gamma$$

If i is in 3rd cluster, then

$$\frac{N}{k}\sum_{j=1}^{k}\phi_{ij}^2 = \frac{N}{k}(\frac{1}{N} + \frac{p_1}{N(\sum_{i=2}^{k} p_i)} + \frac{p_2}{N(\sum_{i=2}^{k} p_i)(\sum_{i=3}^{k} p_i)} + \frac{\sum_{i=4}^{k} p_i}{Np_3\sum_{i=3}^{k} p_i})$$

$$= \frac{1}{k}(1 + \frac{p_1}{(\sum_{i=2}^{k} p_i)} + \frac{p_2}{(\sum_{i=2}^{k} p_i)(\sum_{i=3}^{k} p_i)} + \frac{\sum_{i=4}^{k} p_i}{p_3\sum_{i=3}^{k} p_i}) \leq \frac{4}{k}\gamma \leq \gamma$$

and so on..

Finally if i is in k-th cluster, then

$$\frac{N}{k}\sum_{j=1}^{k}\phi_{ij}^2 = \frac{N}{k}(\frac{1}{N} + \frac{p_1}{N(\sum_{i=2}^{k} p_i)} + \frac{p_2}{N(\sum_{i=2}^{k} p_i)(\sum_{i=3}^{k} p_i)} + .. + \frac{p_{k-1}}{N(\sum_{i=k-1}^{k} p_i)\,p_k})$$

$$\leq \frac{1}{k}(\gamma + ..k \text{ times} + \gamma) = \gamma$$

Since $\frac{p_j}{\sum_{j}^{k} p_i} \leq 1$ and $\frac{1}{\sum_{j+1}^{k} k} \leq \frac{1}{p_k} = \gamma$.

In general, when i is in j th cluster, inside parentheses there will be $max\{j+1, k\}$ terms all of which are bounded by γ. So

$$coh(\phi) = \frac{N}{k} \underset{i=1(1)N}{max} \sum_{j=1}^{k} \phi_{ij}^2 \leq \gamma.$$

□

Proposition 5.8

When equation (5.15) holds, then indeed algorithm 4 perfectly recovers cluster labels on χ_N as the connected components of M_b. □

Proof 5.15

We apply Theorem 2 from Fan et al. (2018) to conclude that $\exists$ a rotation R such that

$$||\tilde{\phi}R - \phi||_{max} = O(\frac{k^{\frac{5}{2}} coh^2(\phi)\, ||(I-B) - L_{sym}||_\infty}{\lambda_k(B)\sqrt{N}}).$$

Since $||V||_{max} \to 0$ a.s., we have

$$||(I-B) - L_{sym}||_\infty \leq N||B - D^{-\frac{1}{2}} U D^{-\frac{1}{2}}||_{max} \leq P_2.$$

Thus $\exists\ C' > 0$ so that,

$$||\tilde{\phi}R - \phi||_{max} \leq \frac{C' k^{\frac{5}{2}} \gamma^3 P_2}{\sqrt{N}} := P_3$$

Thus if, $\{r_i\}_{i=1}^N$ are rows of ϕ and $\{\tilde{r}_i\}_{i=1}^N$ are rows of $\tilde{\phi}R$ we have $||r_i - \tilde{r}_i||_2 \leq \sqrt{k} P_3$ for all i.

Note that $|r_i - r_j| \geq O(\frac{1}{\sqrt{N}})$ for i and j belonging to different clusters and thus $O(\frac{1}{\sqrt{N}}) > 10\sqrt{k} P_3$ is sufficient to ensure that clusters are well separated to be perfectly classified by K means step, which translate into $\frac{C}{k^3\gamma^3} \geq P_2$ for some $C > 0$. □

This concludes proof of part ii of theorem 5.2 □

5.7.4 Eigen-gap statistic k_c

In this sub-section we investigate the conditions under which eigen-gap statistic k_c produces different values in $1(1)k$ and role of σ.

Suppose P_1 is relatively small i.e. eigen-gaps of L_{sym} are roughly determined by eigen-gaps of $I - B$.

We have

$$\triangle_i(I-B) = \begin{cases} q + (p_1 - p_2)(p-q) & i = 1 \\ (p_i - p_{i+1})(p-q) & i = 2(1)k-1 \\ p_k(p-q) = \frac{p-q}{\gamma} & i = k \\ 0 & i = k+1(1)N-1 \end{cases}$$

We define $p_{k+1} = 0$.

Then $k_c = 1$ if

$$q + (p_2 - p_1)(p-q) > \underset{i=2(1)k}{max}(p_i - p_{i+1})(p-q)$$

This happens when $q = \frac{a^*}{\frac{1}{\gamma}b^* + (1-\frac{1}{\gamma})d^*}$ is much bigger than $p - q = \frac{b^*}{\frac{1}{\beta}+(1-\frac{1}{\beta})a^*} - \frac{a^*}{\frac{1}{\gamma}b^*+(1-\frac{1}{\gamma})d^*}$ and/or $(p_2 - p_1)$ is much bigger than $\underset{i=2(1)k}{max}(p_i - p_{i+1})$. First condition holds when $b \approx a \approx d$ i.e. when none of $m_1, .., m_k$ are prominent. Second condition holds if $p_1 >> p_2 \approx p_3 .. \approx p_k$ i.e. when there is only 1 dominant cluster in terms of cluster size.

We have $1 < k_c = k_0 < k$ if

$$q + (p_2 - p_1)(p-q) < (p_{k_0} - p_{k_0+1})(p-q)$$

as well as

$$\underset{i=2(1)k, i\neq k_0}{max}(p_i - p_{i+1})(p-q) < (p_{k_0} - p_{k_0+1})(p-q)$$

occurs.

These 2 conditions hold if and only if $p_1 \approx p_2 \approx .. \approx p_{k_0} >> p_{k_0+1} \approx .. \approx p_k \approx 0$ (i.e. there are k_0 dominant clusters in terms of cluster size) as well as q is much smaller than $p - q$ i.e. when $b >> a \approx d$ (i.e. these k_0 dominant clusters are all prominent).

Similarly, $k_c = k$ if

$$q + (p_2 - p_1)(p - q) < p_k(p - q)$$

as well as

$$\underset{i=2(1)k-1}{max}(p_i - p_{i+1})(p - q) < p_k(p - q)$$

These 2 conditions hold if and only if $p_1 \approx p_2 \approx .. \approx p_{k_0} \approx p_k$ and $p_k >> 0$ (i.e. there are k dominant clusters in terms of cluster size) as well as q is much smaller than $p - q$ i.e. when $b >> a \approx d$ (i.e. these k dominant clusters are all prominent).

We also see that $k_c > k$ never occurs.

Thus, LLPD SC using L_{sym} is able to successfully merge k-connected components of M_b depending on the prominences of corresponding maximas as well cluster size of those components as mentioned in theorem 5.1.

Remark 5.4

Choice of σ: As long as $b - a$ is large, (all k clusters have high prominence) $a - d$ is small and clusters are fairly balanced i.e. $R \approx 1$, a wide range of σ should be able to satisfy eigen-gap condition i.e. equation (5.14). In practice, eigen-gap heuristics works well even when eigen-gap condition is not satisfied since our eigen-gap condition is very conservative in practice.

Unlike other clustering methods such as DBSCAN, kernel based or mode based clustering where σ (smoothing parameter) is used to estimate the density of f, and choosing correct σ is very important and difficult, performance of LLPD SC do not depend on σ for a large range of σ since LLPD itself is being used to estimate the useful information (bottleneck values) of f. □

5.8 LLPD SC with L

In this section we consider LLPD SC with L in the context of continent problem. To be precise, we consider algorithm 4 again but instead of calculating normalized Laplacian $L_{sym} \coloneqq I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$ from W in step 3, we shall work with the raw Laplacian $L \coloneqq D - W$.

Theorem 5.3

As $n \to \infty$ $C \to \infty$, with probability $\geq 1 - O(logn^{-B}) \to 1$ for any $B < \frac{e}{2}(D-1)$,

i) if

$$\begin{aligned} 6\, max\big\{a^* - d^*, 1 - b^*\big\} + max\big\{a^*, (b^* - a^*) \max_{i=1(1)k-1} (p_i - p_{i+1}) \\ +2\, max\{a^* - d^*, 1 - b^*\}\big\} < (b^* - a^*) p_k, \end{aligned} \tag{5.18}$$

then $k = \underset{i=1(1)N}{arg\ max}\, \lambda_{i+1} - \lambda_i$.

ii) Moreover, if

$$\frac{1}{k^3 \gamma^3} \geq C\, max\{a^* - d^*, 1 - b^*\} \tag{5.19}$$

for an absolute constant $C > 0$ then LLPD SC with L perfectly recovers all cluster labels as the connected components of M_b. □

Proof 5.16

In this case we work with the following decomposition of W

$$W \overset{P_0'}{\to} U = \mathbb{E}[W] \overset{P_1'}{\to} B' \coloneqq \begin{bmatrix} b^* J_{|A_1|\times|A_1|} & a^* J_{|A_1|\times|A_2|} & . & . & a^* J_{|A_1|\times|A_k|} \\ a^* J_{|A_2|\times|A_1|} & b^* J_{|A_2|\times|A_2|} & . & . & a^* J_{|A_2|\times|A_2|} \\ . & . & . & . & . \\ . & . & . & . & . \\ a^* J_{|A_k|\times|A_1|} & a^* J_{|A_k|\times|A_2|} & . & . & b^* J_{|A_k|\times|A_k|} \end{bmatrix}$$

From lemma 5.2 it immediately follows that w.p. $\geq 1 - O(logn^{-B}) \rightarrow 1$, $P_0^{'} = O(N^{\frac{loglogN}{logN}})$.

In this case we have $L = D - W$ as the graph laplacian where D is the degree matrix corresponding to W. Let $D^{'}$ be the degree matrix correspond to the *ideal* similarity matrix $B^{'}$ and $L^{'} = D^{'} - B^{'}$ be the corresponding *ideal* graph Laplacian. Our first objective will be to find the eigenvalues and eigenvalues of $L^{'}$.

Proposition 5.9

Let eigenvalues of $L^{'}$, the ideal Laplacian be $\lambda_1 = 0 \leq \lambda_2 \leq .. \leq \lambda_N$ and corresponding eigenvectors are $\phi_1 = \frac{1}{\sqrt{N}} 1_{N \times 1}, \phi_2, .., \phi_N$. Then we have,

$$\lambda_i(L^{'}) = \begin{cases} 0 & \text{for } i = 1 \\ Na^* & \text{for } i = 2(1)k \\ Na^* + |A_k|(b^* - a^*) & \text{for } i = k+1(1)k + |A_k| - 1 \\ Na^* + |A_{k-1}|(b^* - a^*) & \text{for } i = k + |A_k|(1)k + |A_k| + |A_{k-1}| - 2 \\ . & . \\ Na^* + |A_1|(b^* - a^*) & \text{for } i = N - |A_1| + 2(1)N \end{cases}$$

and ϕ_i's are same as in lemma 5.3 with o_i's replaced by A_i's. □

Proof 5.17

We can write

$$B^{'} = (a^*)J_{N \times N} + (b^* - a^*) \text{ Blockdiag } (J_{|A_1| \times |A_1|}, J_{|A_2| \times |A_2|}, .., J_{|A_k| \times |A_k|}) = B^1 + B^2$$

, Thus, $L^{'} = L^1 + L^2$ where $L^1 = D^1 - B^1$ denote the graph Laplacian corresponding to B^1 and $L^2 = D^2 - B^2$ denote the graph Laplacian corresponding to B^2. Note that

$$L^1 + L^2 = (D^1 + D^2) - (B^1 + B^2) = D^{'} - B^{'} = L^{'}$$

indeed holds.

Note that L^1 has 1 eigenvalue 0 with corresponding eigenvector $\frac{1}{\sqrt{N}}1_{N\times 1}$ and $N-1$ eigenvalues Na^* with corresponding eigenvectors which are $\perp 1_{N\times 1}$. Similarly L^2 has k eigenvalues 0 with the requirements for the corresponding eigenvectors to be $\propto 1$ on the k blocks respectively, and $|A_i|-1$ eigenvalues as $|A_i|(b^*-a^*)$ with the requirements for the corresponding eigenvectors to be $\perp 1$ on the i-th block for $i = 1(1)k$.

It is easy to check that the eigenvectors will be same as lemma 5.3 since they need to satisfy the same constrains as lemma 5.3 and eigenvalues of $L^{'}$ will then be given by the sum of eigenvalues from 2 graph Laplacians L^1 and L^2. □

Now we have already shown that, w.p. $\rightarrow 1$

$$P_0^{'} \leq N\, O(\frac{log\, logN}{logN})$$

Thus eigenvalues in $L = D - W$ are bounded by eigenvalues of $L^{'} = D^{'} - B^{'}$ plus the spectral perturbations due to $B^{'} - W = P_1^{'}$ and $D - D^{'} = P_2^{'}$.

and it is not hard to see that,

$$P_1^{'} \leq N\, max\{1 - b^*, a^* - d^*\}$$

and

$$P_2^{'} \leq N\, max\{1 - b^*, a^* - d^*\}$$

and contribution from $P_0^{'}$ is negligible compared to $P_1^{'}$ and $P_2^{'}$ as $N \rightarrow \infty$.

Thus combined spectral perturbation of $B^{'} - W$ and $D - D^{'}$ is bounded by

$$P^{'} = P_1^{'} + P_2^{'} \leq 2N\, max\{a^* - d^*, 1 - b^*\}.$$

Proposition 5.10

When the condition in equation (5.18) holds, $k_c = k$. □

Proof 5.18

By Hoffman-Wielandt theorem, we have $\lambda_i(L) =$

$$\begin{cases} 0 & \text{for } i = 1 \\ [0 \vee Na^* - P^{'}, Na^* + P^{'}] & \text{for } i = 2(1)k \\ [0 \vee Na^* + (b^* - a^*)|A_k| - P^{'}, Na^* + (b^* - a^*)|A_k| + P^{'}] & \text{for } i = k+1(1)k + |A_k| - 1 \\ . & . \\ . & . \\ [0 \vee Na^* + (b^* - a^*)|A_1| - P^{'}, Na^* + (b^* - a^*)|A_1| + P^{'}] & \text{for } i = N - |A_1| + 2(1)N \end{cases}$$

Thus for $i = 1(1)k - 1$

$$\triangle_i(L) \leq Na^* + P^{'}.$$

while

$$\triangle_k(L) \geq (b^* - a^*)|A_k| - 2P^{'}.$$

assuming

$$Na^* + (b^* - a^*)|A_k| - 2N\, max\{a^* - d^*, 1 - b^*\} > 0$$

$$\Rightarrow 2\, max\{a^* - d^*, 1 - b^*\} - a^* < (b^* - a^*)p_k.$$

Thus, for $i = k + 1(1)N$

$$\triangle_i(L) \leq \underset{i=1(1)k-1}{max} (b^* - a^*)(|A_i| - |A_{i+1}|) + 2P^{'}$$

Thus, if $k = \underset{i=1(1)N}{arg\, max}\, \lambda_{i+1} - \lambda_i$, we need to have

$$max\{Na^* + P^{'}, \underset{i=1(1)k-1}{max} (b^* - a^*)(|A_i| - |A_{i+1}|) + 2P^{'}\} < (b^* - a^*)|A_k| - 2P^{'}.$$

Which reduces to

$$a^* + 6max\{a^* - d^*, 1 - b^*\} < (b^* - a^*)p_k.$$

as well as

$$(b^* - a^*) \underset{i=1(1)k-1}{max} (p_i - p_{i+1}) + 8\, max\{a^* - d^*, 1 - b^*\} < (b^* - a^*)p_k.$$

i.e. overall we need to have (5.18). □

This condition is very similar to the eigen-gap condition as derived in equation (5.14) for L_{sym}. In order for this condition to hold, cluster sizes need to be balanced as well as $b^* - a^*$ need to be large compared to $max\{a^* - d^*, 1 - b^*\}$.
Proof of part ii of theorem 5.3 follows similarly as the proof of part ii of theorem 5.2 and is omitted. □

5.8.1 Eigen-gap statistic k_c

The main difference between LLPD SC with L and L_{sym} is that while eigen-gap statistic $k_c = k$ implies k dominant clusters in both the cases, when there are fewer than k dominant clusters eigen-gap statistic for LLPD SC with L_{sym} can correctly predict number of clusters while eigen-gap statistic for LLPD SC with L cannot. Moreover in this case eigen-gap statistic can be more than k, the number of islands.

As usual we shall assume that P' i.e., $a^* - d^*$ and $1 - b^*$ are relatively small and eigen-gaps of L are determined by eigen-gaps of L'.

Clearly, from proposition 5.9 we have

$$\frac{\triangle_i(L')}{N} = \begin{cases} a^* & \text{for } i = 1 \\ 0 & \text{for } i = 2(1)k-1 \\ p_k(b^*-a^*) & \text{for } i = k \\ 0 & \text{for } i = k+1(1)k+|A_k|-2 \\ (p_{k-1}-p_k)(b^*-a^*) & \text{for } i = k+|A_k|-1 \\ 0 & \text{for } i = k+|A_k|(1)k+|A_k|+|A_{k-1}|-3 \\ . & . \\ . & . \\ (p_2-p_1)(b^*-a^*) & \text{for } i = N-|A_1|+1 \\ 0 & \text{for } i = N-|A_1|+2(1)N-1 \end{cases}$$

Eigen-gap statistic correctly predicts 1 cluster when a^* is significantly bigger than b^*-a^* i.e. when none of the k clusters have significant prominence. When $p_k >> \underset{i=2(1)k}{max}\, p_i - p_{i+1}$ and (b^*-a^*) is significantly bigger than a^* (i.e. all k clusters have comparable size and significant prominence), then eigen-gap statistic correctly predicts k clusters as well.

However in the intermediate regime, i.e. when there are only k_0 dominant clusters in terms of size for some $1 \leq k_0 < k$, i.e. when we have $p_1 \approx p_2 \approx .. \approx p_{k_0} >> p_{k_0+1} \approx .. \approx p_k \approx 0$ and all of them are prominent, $k_c = k_0 + \sum_{i=k_0+1}^{k} |A_i|$ instead of k_0 which depends on $|A_i|$'s (cluster degrees) as well and therefore is unhelpful.

5.9 Advantages and disadvantages of LLPD SC

In this section we summarize the advantages and disadvantages of LLPD SC.

LLPD SC allows us to perform robust (based on prominence as well as size of clusters) density based clustering in the context of continent problem. Note that although we are doing SC here, we are essentially estimating bottleneck values between basin of attractions and level sets which almost all density or mode based clustering procedures such as DBSCAN (Schubert et al., 2017), mean-shift, quick shift (Vedaldi and Soatto, 2008; Tran et al., 2006; Cheng, 1995), kernel based density clustering (Tran et al., 2006), persistence based clustering (Chazal et al., 2013) try to do. Unlike those methods we do not need to perform computation intensive density or gradient estimation procedure or ponder over smoothing parameter choices. This is a huge advantage in terms of computational efficiency specially since LLPD can be accurately computed from $O(logn)$-NN graph.

However, two major disadvantages of using LLPD SC (whether using L or L_{sym}) in the context of continent problem are as follows:

1. **Slow convergence rate and high variance:** LLPD's very slow $O(\frac{loglogn}{logn})$ convergence rate, that too with probability $1-O(logn^{-B})$ for as shown in theorem 5.2 and 5.3 is problematic in real data application. This cannot be improved because of the underlying distributional property of LLPD as discussed in chapter 3. This issue has been noted in the simulation studies performed in Little et al. (2017); Mckenzie and Damelin (2019) using LLPD SC as well albeit in different settings.

2. **No consideration of geometry** The continents formed by LLPD SC are entirely topological prominence or density based which, depending on f, can lead to very complicated partition of M_b. For example, if we consider a small ball around $x \in M_b$ for any x, it can contain as many as K clusters, where K is the number of maximas of f. For many practical purposes we may additionally want that if 2 islands are sufficiently near in terms of Euclidean distance or default

geodesic distance between them, then with high probability they should belong in the same continent. I.e. we might want to consider size and topographical proximity as well along with topographical prominence.

Both of these problems can be solved by considering L_p^N or suitable g^N-distances instead of LLPD which performing SC.

5.10 SC using L_p^N, $p \geq 1$

In this section we shall derive performance guarantee conditions for SC with L_p^N distance in the context of continent problem and compare the conditions with the conditions derived in the previous section for LLPD.

Recall that given the data $\chi_n = \{x_1, .., x_n\}$ for $p \geq 1$,

$$L_p^N(x, y, \chi_n) \coloneqq \big(\min_r \sum_{i=0}^{m} |x_{i+1} - x_i|^p\big)^{\frac{1}{p}}$$

where r is any piecewise continuous path $x_0 = x \to x_1 .. \to x_m \to x_{m+1} = y$, $x_i \in \chi_n$ for $i = 1(1)m$ and m is free to vary.

We know that

$$L_p^N(x, y, \chi_n) \to C(D, p) n^{\frac{1-p}{pD}} \big(\inf_\gamma \int_0^1 f^{\frac{1-p}{D}}(\gamma_t) |\gamma_t^{'}| dt\big)^{\frac{1}{p}} \text{ a.s.} \tag{5.20}$$

where $C(D, p) > 0$ is a constant depending on D and p and γ is any piecewise continuous path within M with $\gamma_0 = x$ and $\gamma_1 = y$.

The similarity measure we'll be using in this section is

$$s(x, y) \coloneqq exp\big(-\frac{(n^{\frac{p-1}{pD}} L_p^N(x, y, \chi_n))^2}{\sigma^2}\big) \tag{5.21}$$

Here we formally provide algorithm for SC with L_p^N distance using L_{sym} (algorithm 5). The algorithm using L is similar.

Algorithm 5 L_p^N $(p \geq 1)$ SC using L_{sym} at level b

Input: Data $\chi_n = \{x_1, x_2, .., x_n\}$, $b \in (f_0, f^0)$.
Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data or indicator of noise.
1: Compute $r = C\,logn$-NN graph C_r on χ_n.
2a: Compute r-NN distance $\sigma_{i,r}$ from each point x_i. Label a point x_i as noise if $\sigma_{i,r} > (\frac{r}{bn\theta_D})^{\frac{1}{D}}$. Let the remaining data points be χ_N with $|\chi_N| = N$.
2b: Form $W_{N\times N}$ with $((W))_{ij} = exp\,(-\frac{(n^{\frac{p-1}{pD}}\, L_p^N(x,y,C_r))^2}{\sigma^2})$ for some $\sigma > 0$.
3: Form the normalized Laplacian $L_{sym} = I - D^{-\frac{1}{2}} W D^{-\frac{1}{2}}$
4: Compute the eigen decomposition $\{(\phi_k, \lambda_k)\}_{k=1}^N$, sorted so that $0 = \lambda_1 \leq \lambda_2 \leq .. \leq \lambda_N$.
5: Estimate the number of clusters k_c by *eigen-gap statistic* $k_c := \underset{i}{argmax}\, \lambda_{i+1} - \lambda_i$. Let $U_{N\times k} = (\phi_1 : \phi_2 : .. : \phi_{k_c})$.
6: Compute the label $Z_{N\times 1}$ by running K-means on $\{U_{i,.}\}_{i=1}^N$ using k_c as the number of clusters and $Z_{i,.}$ is label of y_i as obtained by the K-means.

Summary of the performance of this algorithm in context of continent problem is as follows:

Theorem 5.4

When all k islands are balanced in shape and size, sufficiently prominent and sufficiently distant from one another; algorithm 5 produces $k_c = k$ i.e. no merging occurs and continents are islands w.p. 1 as $n \to \infty$. Otherwise $k_c < k$ i.e. merging occurs. Almost surely $k_c > k$ do not occur.

For small values of p, merging is more likely to occur when some islands are not balanced in shape and size or not sufficiently distant from one another while for higher values of p, merging is more likely to occur when some islands have small prominence. □

5.10.1 Assumptions

We list our assumptions in this section below:

A1: $M \subseteq \mathbb{R}^D$ is connected, compact and smooth i.e. having a $D-1$ dimensional boundary. We shall work with $M = \mathbb{R}^D$ for sake of simplicity, although our result still applies in the general case.

A2: $\int_L f = 0$ for any $L \subset M$ so that $\mu_D(L) = 0$. This in particular ensures that f has no lower dimensional structure.

A3: f is Morse-Smale. This allows us to split M (except for a lower dimensional set) into finite number of "basin of attractions" $A_1, A_2, ..A_K$.

A4:

$$f(m_i) > b \, \forall i = 1, 2, .., k; M_b = \cup_{i=1}^{k} C(m_i, b).$$

A5:

$$\min_{i \neq j} \min_{x \in C(m_i,b), y \in C(m_j,b)} dist_1(x, y) := \delta(b) > 0.$$

Notice that we have removed **A4** (dominance of $D-1$ dimensional boundary effect in level sets) and **A6** (assumptions on bottleneck values) from the previous section since unlike LLPD, performance of SC with L_p^N do not depend on those. Instead we have added the separation requirement of the connected components of level sets in **A5**. Although not strictly needed for our analysis we still recall from previous section that

$$\max_{i \neq j, \, i,j \in \{1,2,..,k\}} \text{Bottleneck}(m_i, m_j) = a < b$$

5.10.2 Theoretical guarantee for L_{sym} and L

We formally state the theoretical guarantee results for eigen-gap statistic to return k as optimal number of clusters and perfect recovery of cluster lables as the connected compoments of M_b.

We define

$$d^* := e^{-\frac{C^2(D,p) f_0^{\frac{2(1-p)}{pD}} (\gamma_\beta(M_b))^{\frac{2}{p}}}{\sigma^2}}.$$

$$a^* := e^{-\frac{C^2(D,p)y_\beta(b)^{\frac{2(1-p)}{pD}}(\delta(b))^{\frac{2}{p}}}{\sigma^2}}.$$

and

$$b^* := e^{-\frac{C^2(D,p)x_\beta(b)^{\frac{2(1-p)}{pD}}(\gamma_0(b)I_\beta(b))^{\frac{2}{p}}}{\sigma^2}}.$$

where $C(D,p) > 0$ is the same constant as in equation (5.20), $\delta(b)$ is the same constant as in **A5** and $\gamma_\beta(M_b)$, $I_\beta(b)$, $x_\beta(b)$ and $y_\beta(b)$ are quantities related to within and between cluaster distances and densities and are provided by equations (5.26), (5.27), (5.28), (5.29) respectively in this section.

Theorem 5.5

As $n \to \infty$ $C \to \infty$, with probability w.p. $\geq 1 - O(exp(-n^r) \to 1$ for any $r < \frac{1}{D+2p}$,

i)$\lambda_{k+1} - \lambda_k$ is the largest gap (i.e. $k_c = k$) in the eigenvalues of L_{sym} provided

$$\frac{(p-q)}{\gamma} > max\{\frac{3P_2+1}{2}, 4P_2\} \tag{5.22}$$

Moreover if

ii)

$$\frac{C}{k^3\gamma^3} \geq P_2 \tag{5.23}$$

□

where C is an absolute constant, then algorithm 5 perfectly recovers cluster labels on χ_N as the connected components of M_b.

Theorem 5.6

As $n \to \infty$ $C \to \infty$, with probability w.p. $\geq 1 - O(exp(-n^r) \to 1$ for any $r < \frac{1}{D+2p}$,

i) if

$$6\,max\big\{a^* - d^*, 1 - b^*\big\} + max\big\{a^*, (b^* - a^*) \underset{i=1(1)k-1}{max}(p_i - p_{i+1}) \\ +2\,max\{a^* - d^*, 1 - b^*\}\big\} < (b^* - a^*)p_k, \tag{5.24}$$

then $k = \underset{i=1(1)N}{arg\ max}\, \lambda_{i+1} - \lambda_i$.

ii) Moreover, if

$$\frac{1}{k^3\gamma^3} \geq C\,max\{a^* - d^*, 1 - b^*\} \tag{5.25}$$

for an absolute constant $C > 0$ then L_p^N SC with L perfectly recovers all cluster labels as the connected components of M_b. □

Note that except for the convergence rate and value of a^*, b^*, d^* the statements of theorem 5.5 and 5.6 are exactly same to the statements of theorem 5.2 and 5.3 respectively and proof technique is also exactly the same. Thus it is sufficient to derive the values of a^*, b^* and d^* and provide the proof of the convergence rate.

Proof 5.19

We derive the convergence rate of L_p^N SC first. The value of a^*, b^* and d^* will be derived subsequently.

From Hwang et al. (2016), we know

$$\mathbb{P}\big(\underset{x,y}{sup}\,|\frac{L_p^N(x, y, \chi_n)}{n^{\frac{1-p}{pD}}\,dist_p(x, y)^{\frac{1}{p}}} - C(D, p)| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\alpha})$$

for small $\epsilon > 0$, $\theta_0^{'} > 0$ and $\alpha = \frac{1}{D+2p}$.

It follows that

$$\mathbb{P}\big(\underset{x,y}{sup}\,|exp\,\big(-\frac{(n^{\frac{p-1}{pD}}\,L_p^N(x, y, \chi_n))^2}{\sigma^2}\big) - exp\,\big(-\frac{C(D, p)^2 dist_p(x, y)^{\frac{2}{p}}}{\sigma^2}\big)| > \epsilon\big)$$

$$\leq exp(-\theta_1^{'}\epsilon^2 n^{\alpha})$$

for small $\epsilon > 0$, $\theta_1^{'} > 0$ and $\alpha = \frac{1}{D+2p}$.

In particular,

$$exp\left(-\frac{(n^{\frac{p-1}{pD}}\ L_p^N(x,y,\chi_n))^2}{\sigma^2}\right) \to exp\left(-\frac{C(D,p)^2 dist_p(x,y)^{\frac{2}{p}}}{\sigma^2}\right) \text{ a.s.}$$

Let $W = U + V$ as usual where

$$((V))_{ij} = exp\,(-\frac{(n^{\frac{p-1}{pD}}\ L_p^N(x_i,x_j,\chi_n))^2}{\sigma^2}) - exp\,(-\frac{C(D,p)^2 dist_p^{\frac{2}{p}}}{\sigma^2}) = ((W - \mathbb{E}[W]))_{ij}$$

and

$$((U))_{ij} = \mathbb{E}[W]_{ij} = exp\,(-\frac{(n^{\frac{p-1}{pD}}\ L_p^N(x_i,x_j,\chi_n))^2}{\sigma^2})$$

.

Thus as $n \to \infty$ if $|x_i - x_j| = O(1)$, a.s. we have

$$\mathbb{P}(|\max_{i,j}((V))_{ij}| > \epsilon) \leq exp(-\theta_1^{'}\epsilon^2 n^{\alpha}).$$

Also w.l.g. we can ignore $|x_i - x_j| < O(1)$ case. So, in particular if we let $\epsilon = O(n^{-\vartheta})$, then we can say that w.p. $1 - O(exp(-n^{\alpha-2\vartheta}))$, spectral perturbation due to $D^{-\frac{1}{2}}VD^{-\frac{1}{2}}$ is bounded above by $O(n^{-\vartheta})$. Thus we can say w.p.$\to 1$, spectral perturbation due to $D^{-\frac{1}{2}}VD^{-\frac{1}{2}}$ is bounded above by $O(n^{-\vartheta})$ for any $0 < \vartheta < \frac{1}{2(D+2p)}$. The proof for the convergence rate of SC using L is similar and is omitted. □

5.10.3 Interpolation between K-means and prominence based clustering

In this sub-section, we derive the values of a^*, b^*, d^* and argue that L_p^N SC is an interpolation between K-means and prominence based clustering in the context of continent problem.

Proof 5.20

Recall that

$$n^{\frac{p-1}{pD}} L_p^N(x, y, \chi_n) \to C(D,p)(\inf_{\gamma} \int_0^1 f^{\frac{1-p}{D}}(\gamma_t)|\gamma_t^{'}|dt)^{\frac{1}{p}} \text{ a.s.}$$

$$= C(D,p)(\int_0^1 f^{-\beta}(\gamma_\beta(t))|\gamma_\beta^{'}(t)|dt)^{\frac{1}{p}} \text{ a.s.}$$

Where following the notations introduced in chapter 4, $\gamma_\beta(t)$ denote any piecewise continuous path minimizing $\int_0^1 f^{-\beta}(\gamma_t)|\gamma_t^{'}|dt$, $C(D,p) > 0$ and $\beta \coloneqq \frac{p-1}{D}$.

Our analysis will be focused on $dist_p(.,.)$ which is defined as

$$dist_p(\gamma_\beta(0), \gamma_\beta(1)) \coloneqq \int_0^1 f^{-\beta}(\gamma_\beta(t))|\gamma_\beta^{'}(t)|dt$$

Now we introduce some measures related to clusters $C(m_i, b)$ $i = 1(1)k$ which will be useful for our analysis.

For $i = 1(1)k$ and $\beta \in [0, \infty]$ we define

$$\gamma_\beta(C(m_i, b)) \coloneqq \max_{\gamma_\beta(0), \gamma_\beta(1) \in C(m_i,b)} |\gamma_\beta(t)|$$

From chapter 4, we know that $\lim_{\beta \to \infty} \gamma_\beta$ exists and we shall call it γ_∞ (this correspond to $\beta = \infty$).

Also, in this notation it is clear that $\gamma_0(C(m_i, b))$ is just the diameter of $C(m_i, b)$ and from lemma 4.1 in chapter 4 it follows that

$$\gamma_0(C(m_i, b)) \leq \gamma_\beta(C(m_i, b)) \leq \gamma_\infty(C(m_i, b))$$

Define

$$\gamma_\beta(b) \coloneqq \max_{i=1(1)k} \gamma_\beta(C(m_i, b)) \tag{5.26}$$

as the maximum "within cluster" arc-length of γ_β and

$$I_\beta(b) \coloneqq \underset{i=1(1)k}{max} \frac{\gamma_\beta(C(m_i,b))}{\gamma_0(C(m_i,b))} \geq 1 \tag{5.27}$$

which is a measure of within cluster density-geometry interaction.

Remark 5.5

Note that we are using $\gamma_\beta(0)$ and $\gamma_\beta(1)$ to denote the beginning and endpoint of the $\gamma_\beta(t)$ curve. This is not related in any way to $\gamma_\beta(b)$. □

Clearly for any $0 < \beta < \infty$ we have

$$\gamma_0(b) \leq \gamma_\beta(b) \leq \gamma_\infty(b).$$

Now we define $\hat{\gamma}_{i\beta}$ as any arc that maximizes within $C(m_i,b)$ $dist_p(.,.)$. I.e.

$$\hat{\gamma}_{i\beta} \coloneqq \underset{\gamma_\beta:\gamma_\beta(0),\gamma_\beta(1)\in C(m_i,b)}{arg\ max} dist_p(\gamma_\beta(0),\gamma_\beta(1)).$$

Note that we do not require $\hat{\gamma}_{i\beta}$ to be entirely within $C(m_i,b)$. The only requirement we have is $\gamma_\beta(0) \in C(m_i,b)$ and $\gamma_\beta(1) \in C(m_i,b)$.

We now define

$$x_\beta(b) \coloneqq \underset{i=1(1)k}{min} \Big(\frac{\int_0^1 f^{-\beta}(\hat{\gamma}_{i\beta}(t))|\hat{\gamma}_{i\beta}^{'}(t)|dt}{|\hat{\gamma}_{i\beta}(t)|}\Big)^{-\frac{1}{\beta}} \tag{5.28}$$

as the " minimum average within cluster density" w.r.t. L_p^N metric and

$$y_\beta(b) \coloneqq \underset{\gamma_\beta(0)\in\partial C(m_i,b),\gamma_\beta(1)\in\partial C(m_j,b),i\neq j}{sup} \Big(\frac{\int_0^1 f^{-\beta}(\gamma_\beta(t))|\gamma_\beta^{'}(t)|dt}{|\gamma_\beta(t)|}\Big)^{-\frac{1}{\beta}} \tag{5.29}$$

as the "maximum average between cluster density" w.r.t. L_p^N metric.

Remark 5.6

Note that for most types of $C(m_i, b)$'s and/or for sufficient large β, $x_\beta(b)$ will be $\geq b$ while $y_\beta(b)$ will be always be $< b$, although it can be arbitrarily close to b depending on f. Indeed, any γ_β path joining 2 points within the same cluster should mostly stay within that cluster (i.e. will stay within $\geq b$ density region), specially when β is large, unless those 2 points happen to be very close in Euclidean distance and the straight line between them travels through a very low density region i.e. $x_\beta(b) < b$ occurs only when β is not too large and/or at least one of $C(m_i, b)$'s are highly non-convex. □

Remark 5.7

Also, $x_\infty(b) = b$ since γ_∞ path within a cluster always stays within that cluster whereas $y_\infty(b) = \underset{i \neq j}{max}\, \text{Bottleneck}(m_i, m_j) = a$ as introduced in the previous section.□

Our goal will be to find an upper bound of $n^{\frac{p-1}{pD}}\, L_p^N(x, y, \chi_n)$ when x and y are from same cluster and a lower bound of $n^{\frac{p-1}{pD}}\, L_p^N(x, y, \chi_n)$ when x and y are from different clusters.

Clearly, as $n \to \infty$ we have

$$\underset{i=1(1)k}{max} \underset{\gamma_\beta(0), \gamma_\beta(1) \in C(m_i, b)}{sup} \int_0^1 f^{\frac{1-p}{D}}(\gamma_\beta(t)) |\gamma_\beta'(t)| dt \leq x_\beta(b)^{-\beta} \gamma_\beta(b) \leq x_\beta(b)^{-\beta} \gamma_0(b) I_\beta(b)$$

On the other hand

$$\underset{\gamma_\beta(0) \in C(m_i, b), \gamma_\beta(1) \in C(m_j, b), i \neq j}{inf} \int_0^1 f^{\frac{1-p}{D}}(\gamma_\beta(t)) |\gamma_\beta'(t)| dt \geq y_\beta(b)^{-\beta} \delta(b)$$

If we denote

$$\gamma_\beta(M_b) := \underset{\gamma_0, \gamma_1 \in M_b}{max} |\gamma_\beta(t)| \tag{5.30}$$

as $n \to \infty$, for $x_i, x_j \in \chi_N$ we have a.s.

$$e^{-\frac{\left(n^{\frac{p-1}{pD}}\, L_p^N(x_i, x_j, \chi_n)\right)^2}{\sigma^2}} \in \begin{cases} [b^*, 1) & \text{If both } x_i \text{ and } x_j \text{ are in same cluster} \\ (d^*, a^*] & \text{If } x_i \text{ and } x_j \text{ are from different clusters} \end{cases}$$

Where

$$d^* := e^{-\frac{C^2(D,p) f_0^{\frac{2(1-p)}{pD}} (\gamma_\beta(M_b))^{\frac{2}{p}}}{\sigma^2}}.$$

$$a^* := e^{-\frac{C^2(D,p) y_\beta(b)^{\frac{2(1-p)}{pD}} (\delta(b))^{\frac{2}{p}}}{\sigma^2}}.$$

and

$$b^* := e^{-\frac{C^2(D,p) x_\beta(b)^{\frac{2(1-p)}{pD}} (\gamma_0(b) I_\beta(b))^{\frac{2}{p}}}{\sigma^2}}.$$

Thus we have arrived at the values of a^*, b^*, d^* as mentioned in the beginning of the previous sub-section. □

In the rest of this sub-section, we focus on the role of p and how clusters (connected components of b-level sets) are merged based on prominence, size as well as their geometric characteristics depending on it.

Note that for $k_c = k$ (i.e. for eigen-gap heuristics to retain all the islands as continents), we need to have maximal within cluster L_p^N distance much smaller than minimum between cluster L_p^N distance. I.e.

$$x_\beta(b)^{-\beta} \gamma_0(b) I_\beta(b) << y_\beta(b)^{-\beta} \delta(b)$$

$$\Rightarrow \left(\frac{x_{\frac{p-1}{D}}(b)}{y_{\frac{p-1}{D}}(b)}\right)^{\frac{1}{D}} << \left(\frac{\delta(b)}{\gamma_0(b)} \frac{1}{I_{\frac{p-1}{D}}(b)}\right)^{\frac{1}{p-1}} \tag{5.31}$$

Remark 5.8

This equation demonstates the interplay between distance and density in L_p^N SC. Notice that when $p \to \infty$ RHS $\to 1$, therefore the requirement for $k_c = k$ reduce to the density requirement of

$$a = \lim_{p \to \infty} x_{\frac{p-1}{D}}(b) << y_{\frac{p-1}{D}}(b) = b.$$

i.e. all k clusters are significantly prominent- which is the requirement we had in LLPD SC. This is not surprising since $\lim_{p\to\infty} L_p^N = LLPD$. □

Remark 5.9

On the other hand, when p is close to 1, requirement for $k_c = k$ is dominated by the geometry requirement $\delta(b) > \gamma_0(b)$ so that RHS $\to \infty$ i.e. the minimum distance between different clusters should be bigger than maximum within cluster diameter which is a sufficient requirement for K-means clustering algorithm to return perfect clusters as shown in proposition 5.11. □

Proposition 5.11

When $\delta(b) > \gamma_0(b)$ then K-means algorithm with number of clusters=k perfectly clusters dataset. □

Proof 5.21

For dataset $\chi_n = \{x_1, .., x_n\}$ sampled from f in M, let $C(m_i, b) \cap \chi_n = A_i$ as usual with $\sum_{i=1}^k |A_i| = N \leq n$.

K-means objective function in this case is given by

$$J(k) \coloneqq \sum_{j=1}^{k} \sum_{x \in S_i} ||x - \mu_i||^2$$

where $\{S_1, .., S_k\}$ is a partition of $\{\cup_{i=1}^k C(m_i, b)\} \cap \chi_n$ and μ_i is the mean of points in S_i. We want to show that (w.l.g.) $S_i = C(m_i, b)$ lead to minimization of $J(k)$.

Let the partition P_0 correspond to the partition $S_i = C(m_i, b)$ for $i = 1(1)k$. W.l.g. suppose $x_1 \in C(m_1, b)$ and let P_1 denote the *switched* partition $S_1 = C(m_1, b) \setminus \{x_1\}$, $S_2 = C(m_2, b) \cup \{x_1\}$,..,$S_k = C(m_k, b)$. W.l.g. we can also assume n is large so that $\mu_i|_{P_0} \approx \mu_i|P_1$ for $i = 1, 2$.

Now

$$J(k)|_{P_1}$$

$$= J(k)|_{P_0} + ||x_1 - \mu_2||^2 - ||x_1 - \mu_1||^2$$

$$\geq J(k)|_{P_0} + \delta(b)^2 - \gamma_0(b)^2 > J(k)|_{P_0}$$

since $||x_1 - \mu_2|| \geq \delta(b)$ and $||x_1 - \mu_1|| \leq \gamma_0(b)$ from their definition and $\delta(b) \geq \gamma_0(b)$ from our assumption.

The proof follows by noting that any generic permutation P can be reached from P_0 after finitely many *switches* like $P_0 \to P_1$. □

Remark 5.10

Now we focus on L_p^N SC for a specific $1 < p < \infty$. From RHS of equation (5.31), it is clear that $\gamma_0(b)$ (the maximum within cluster diameter) need to be minimized. It is not difficult to show that for a fixed cluster volume $\int_{M_b} f$ and number of cluster k, $\gamma_0(b)$- the maximum within cluster diameter is minimized when $C(m_i, b)$'s are of same volume and spherical in shape. □

Remark 5.11

Moreover when p is close to 1, more emphasis is put into "cluster balance" characteristics in terms of geometry (size and shape) and minimum cluster separation $\delta(b)$- which are all geometrical property of $C(m_i, b)$'s. Any elongated or large cluster is discouraged and any 2 clusters which are close (in terms of Euclidean distance) will get merged. On the other hand as p becomes large, all geometric characteristics of $C(m_i, b)$'s becomes gradually irrelevant (i.e. more elongated and larger clusters can be tolerated) and only the prominences of these clusters become important. This interplay between geometry and density (or equivalently between prominence based and K means clustering) is the key theme in L_p^N SC. □

Remark 5.12

Choice of σ: Similar to LLPD SC, σ do not play an important role and in principle should not change the underlying clustering structure. It is a good practice to choose

$\sigma = A\,C(D,p)$ for some $A > 0$ to compensate for the constant $C(D,p)$ in equation (5.20). □

5.11 LLPD/L_p^N SC in LDLN model

In many important application (speech, images, gnome, cosmology data etc.) clusters are defined as lower dimensional (nonlinear) manifolds sitting inside higher dimensional ambient space with or without noise around them. In literature, this is known as manifold hypothesis (MH). In this section we shall focus on how to cluster (i.e. identify lower dimensional structure) from data generating models as discussed in sub-sections 5.4.2, 5.4.3 and 5.4.4.

The key difference between the family of SC algorithms in this section and the previous section(s) is that we shall be working with the un-scaled version of g^N and LLPD and mostly will be using the fact that these distances approach uniformly to 0 as $n \to \infty$.

5.11.1 LDLN model

Our ambient space A is a D dimensional subset of $\mathbb{R}^D$ with $\mu_D(A) < \infty$. Following Arias-Castro (2011) we let $M^{'} \subset A$ an element of $V(d,\kappa,\epsilon_0)$ for $d < D$, $\kappa \geq 1$ and $\epsilon_0 > 0$ if it is connected, $\mu_d(M^{'}) < \infty$ and for all $x \in M^{'}$, $\epsilon \in (0,\epsilon_0)$ we have

$$\kappa^{-1}\epsilon^d \leq \frac{\mu_d(M^{'} \cap B_D(x,\epsilon))}{\mu_d(B_d(0,1))} \leq \kappa\epsilon^d \tag{5.32}$$

where $B_D(x,\epsilon)$ denotes D dimensional ball of radius ϵ around x and $\mu_d(.)$ denote d dimensional Lebesgue measure. Or equivalently

$$g_1(x,y) \leq |x-y|(1+\kappa|x-y|^2) \text{ for any } |x-y| \leq \epsilon_0 \tag{5.33}$$

where $g_1(x, y)$ denote the default geodesic distance between x and y.

Also, we define

$$B(M_i, \epsilon) \coloneqq \{x \in \mathbb{R}^D : dist_1(x, M_i) \leq \epsilon).$$

Note that $V(d, \kappa, \epsilon_0)$ includes d-dimensional smooth compact Riemannian manifolds with positive reach (i.e. this condition in equation (5.32) prevents $M^{'}$ from being too narrow or self-intersecting in some places.) (Federer, 1959) and in most practical situations the lower dimensional manifolds we encounter do belong in $V(d, \kappa, \epsilon_0)$ class.

In this context, the LDLN data model as introduced in sub-section 5.4.4 is as follows

$$X_i \overset{iid}{\sim} (1 - \sum_{i=1}^{k} \alpha_i)Y + \sum_{i=1}^{k} \alpha_i (W_i * \phi_i(\sigma_i)) \text{ for } i = 1(1)n \tag{5.34}$$

where $\alpha_i > 0$ for $i = 1(1)k$ such that $\sum_{i=1}^{k} \alpha_i \leq 1$, Y is the "background noise" (uniform or close to uniform) distribution with the density f_Y supported on the ambient space A in $\mathbb{R}^D$, W_i is a distribution with density f_{W_i} supported on M_i. We assume that

$$\textbf{A1: } M_i \in V(d_i, \kappa, \epsilon_0) \, \forall i = 1(1)k$$

, $d_i < D$ and $\phi_i(\sigma_i)$ is "surrounding noise" distribution for M_i with $\mathbb{E}[\phi_i(\sigma_i)] = 0$, density f_{ϕ_i} supported on $\mathbb{R}^D$ or a D dimensional subset of it, having maxima at 0 and σ_i denoting it's dispersion parameters for $i = 1(1)k$ and $*$ denoting the convolution operator. To be more precise, density function of $W_i * \phi_i(\sigma_i)$, $f_{W_i * \phi_i}$ at a point $x \in A$ is given by

$$f_{W_i * \phi_i}(x) = \int_{M_i} f_{\phi_i}(x - z) f_{W_i}(z) dz \tag{5.35}$$

i.e. M_i's are d_i dimensional "ridges" in D dimensional background as introduced

in sub-section 5.1.

We additionally need to have M_i's well-separated and $\phi_i(\sigma_i)$'s to be concentrated around 0 (i.e. dispersion parameter σ_i's to be small) to make the LDLN model identifiable. Thus we also assume

$$\textbf{A2: } \min_{i \neq j} d_1(M_i, M_j) = \delta > 0 \tag{5.36}$$

where $d_1(X, Y)$ denotes the minimum $dist_1$ distance between points in X and Y and σ_i's to be small.

Some choices for f_{ϕ_i} include $f_{\phi_i}(x) = \delta_0(x)$ where δ_0 is the Dirac delta function at 0. This choice corresponds to no surrounding noise/smoothing for the density f_{W_i} on the lower dimensional manifold M_i. Other choices, which provide smoothing around f_{W_i} can be $\phi_i \sim N_D(0, \sigma_i^2 I)$ or $\phi_i \sim \sigma_i Z_i$ where Z_i is a density supported on $B_D(0, 1)$, $\sigma_i > 0$.

Our clustering procedure for a data χ_n generated from LDLN data model will have 2 stages. Our goal in stage 1 will be to identify lower dimensional manifolds M_i's for $i = 1(1)k$. In stage 2, we further cluster within each M_i's based on it's density f_{W_i}. Note that in most real data application the ambient dimension D is known, but $d_i < D$'s are unknown. We shall assume the same and w.l.g. shall work with the "no surrounding noise" version i.e. $\sigma_i = 0$ for all $i = 1(1)k$ in equation (5.34).

5.11.2 LLPD/L_p^N SC algorithm for LDLN data and comparison

Little et al. (2017) advocate for using LLPD SC with eigen-gap heuristics in order to identify lower dimensional clusters (i.e. in stage 1) M_i's for $i = 1(1)k$ following a denoising procedure. However, strictly speaking we don't need to perform LLPD SC, in particular calculate LLPDs, in order to identify lower dimensional clusters.

Algorithm 6 LLPD/L_p^N SC with L_{sym} for LDLN data

Input: Data $\chi_n = \{x_1, x_2, .., x_n\}$, $r \in \mathbb{N}$ sufficiently large, D- ambient dimension.
Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data or indicator of noise.
1: **Denoising step** Compute r-th nearest neighbor distance $\sigma_{i,r}$ from each point x_i. Label a point x_i as noise if $\sigma_{i,r} > \epsilon_D := n^{-\frac{1}{(D-\frac{1}{2})}}$. Let the remaining data points be χ_N with $|\chi_N| = N$.
2: **Level 1 clustering step** Obtain ϵ_D-NN graph C_{ϵ_D} of χ_N. I.e. join x_i and x_j in χ_N if and only if $|x_i - x_j| \leq \epsilon_D$. Obtain the connected components of C_{ϵ_D}. Let the connected components be $\{A_1, .., A_k\}$ with $\chi_N = \cup_{i=1}^k A_i$. Declare $A_1, .., A_k$ as stage 1 clusters.
3: **Dimension determination step** Compute ϵ_j-NN graph C_{ϵ_j} on each of$A_1, .., A_k$ for $j = 2(1)D-1$. Set $\dim(A_i) := \{ \min_{j\in 2(1)D} : A_i \text{ is connected in } C_{\epsilon_j}\} - 1$.
4: **Level 2 clustering step** Perform step 3-6 of algorithm 4/5 on A_i with $dim(A_i)$ and chosen level b_i for $i = 1(1)k$.

Instead in order to save computational resources, we can opt for computing the connected components of ϵ_D-NN graph of the denoised data instead for some ϵ_D to be discussed later to identify lower dimensional clusters and then perform appropriately scaled LLPD/L_p^N SC with eigen-gap heuristics on individual lower dimensional manifolds essentially solving continent problem on them. Below we provide an outline for LLPD/L_p^N SC algorithm (algorithm 6) on LDLN data.

Theorem 5.7

As $n \to \infty$ algorithm 6, w.p. $\to 1$ recovers all lower dimensional manifolds and correctly determines their dimension. □

Proof 5.22

Note that when local dimension is D, connectivity distance as well as largest k-nearest neighbor distance are $O(\frac{logn}{n}^{\frac{1}{D}})$ Penrose (1999); Penrose et al. (1999) for $k = O(1)$. Since $O(\frac{logn}{n}^{\frac{1}{D-1}}) < \epsilon_D = n^{-\frac{1}{(D-\frac{1}{2})}} < O(\frac{logn}{n}^{\frac{1}{D}})$ and $\epsilon_D < \delta = \min_{i\neq j} d_1(M_i, M_j)$ for all sufficiently large n, it follows that as $n \to \infty$ the denoising procedure removes all

D-dimensional points while leaving behind all lower (upto $D-1$) dimensional point and connected components of ϵ_D-NN graphs are indeed data points residing in lower dimensional manifolds.

Similarly, if A_i has dimension $d_i < D$, then since connectivity distance is $O(\frac{logn}{n}^{\frac{1}{d_i}})$ it follows that $d_i + 1$ is the smallest interger so that $\epsilon_{d_i+1} = \frac{1}{n}^{\frac{1}{d_i+\frac{1}{2}}} > O(\frac{logn}{n}^{\frac{1}{d_i}})$. Thus indeed dimensions of all lower dimensional manifolds are correctly determined. □

Remark 5.13

Little et al. (2017) advocate for LLPD SC with similarity measure

$$w(x_i, x_j) = e^{-\frac{LLPD^2(x_i,x_j,\chi_N)}{\sigma^2}}$$

and

$$\underset{i=1(1)k}{max}\ O\big(\frac{logn}{n}^{\frac{1}{d_i}}\big) < \sigma < O(1)$$

(remark 3.14 in their paper). It is not difficult to see that this choice of σ lead to the *perfect* $\mathbb{E}[W]$ with $a^* = d^* \to 0$ and $b^* \to 1$ as $n \to \infty$ (since "within cluster LLPD" is bounded above by $\underset{i=1(1)k}{max}\ O(\frac{logn}{n}^{\frac{1}{d_i}})$ and "between cluster LLPD" is bounded below by δ.) thus leading to a perfect clustering. Thus both Little et al. (2017) and our approaches are able to find lower dimensional clusters accurately as $n \to \infty$. □

Remark 5.14

Also, since $\frac{\underset{i=1(1)k}{max}\ O(\frac{logn}{n}^{\frac{1}{d_i}})}{\frac{logn}{n}^{\frac{1}{D}}} \to 0$, the de-noising procedure can separate lower dimensional clusters from arbitrary high amount (i.e. $\sum_{i=1}^{k} \alpha_i$ can be arbitrarily close to 0) of background noise in LDLN model. Alternatively in terms of sample size n, number of background noise points exponentially $\frac{D}{\underset{i=1(1)k}{max}\ d_i}$ times larger than the number of cluster points can be tolerated when LLPD SC is used in the step 1 of clustering procedure in LDLN data. □

Remark 5.15

On the flip side, this also reveals that LLPD metric is very vulnerable to the presence of any arbitrarily small amount of lower dimensional "systematic noise" inside ambient high dimensional clusters and thus provides a justification of no ridge assumption (assumption **A2**) in section 5.7. □

Remark 5.16

It is not hard to show that if we instead use L_p^N SC in the step 1 of clustering procedure, then number of background noise points exponentially $\frac{(p-1)D}{p \underset{i=1(1)k}{max} d_i}$ times larger than the number of cluster points can be tolerated. This make L_p^N less suitable than LLPD for LDLN model but more robust towards lower dimensional "systematic noise". □

5.12 Multimanifold clustering (MMC)

For our 2 stage procedure of MMC clustering as proposed in sub-section 5.4.5 where we not only wish to detect (possibly intersecting) lower dimensional manifolds but also wish to detect cluster structures in them, we propose the following *hybrid* algorithm (algorithm 7) below. Note that this particular algorithm combines the idea developed in Arias-Castro et al. (2017) with the ideas developed in this chapter so far. The similarity measure proposed in equations (5.37) and (5.38) detect any sudden change of curvature of the manifolds as well as geometry and density information in it.

Remark 5.17

We look at the similarity measure in (5.38)

$$W_{ij} = exp\big(-\frac{[n^{\frac{p-1}{p\,dim(A_i)}} L_p^N(x_i, x_j, C_s)]^2}{\sigma^2}\big)\, exp\big(\frac{||Q_i - Q_j||^2}{\eta^2}\big).$$

Algorithm 7 LLPD/L_p^N SC with L_{sym} for MMC

Input: Data $\chi_n = \{x_1, x_2, .., x_n\}$, $r \in \mathbb{N}$ sufficiently large, $s = O(logn)$, D ambient dimension, $\sigma > 0$ spatial scale, $\eta > 0$ projection scale.

Output: Label $Z_{n\times 1}$ with i-th entry being the cluster label of i-th data or indicator of noise.

1: **Denoising step** Compute r-th nearest neighbor distance $\sigma_{i,r}$ from each point x_i. Label a point x_i as noise if $\sigma_{i,r} > \epsilon_D \coloneqq n^{-\frac{1}{(D-\frac{1}{2})}}$. Let the remaining data points be χ_N with $|\chi_N| = N$.

2: **Level 1 clustering step** Obtain ϵ_D-NN graph C_{ϵ_D} of χ_N. I.e. join x_i and x_j in χ_N if and only if $|x_i - x_j| \leq \epsilon_D$. Obtain the connected components of C_{ϵ_D}. Let the connected components be $\{A_1, .., A_k\}$ with $\chi_N = \cup_{i=1}^k A_i$. Declare $A_1, .., A_k$ as stage 1 clusters.

3: **Dimension determination step** Compute ϵ_j-NN graph C_{ϵ_j} on each of$A_1, .., A_k$ for $j = 2(1)D - 1$. Set $\dim(A_i) \coloneqq \{ \underset{j\in 2(1)D}{min} : A_i \text{ is connected in } C_{\epsilon_j}\} - 1$.

4: **Level 2 clustering step** Perform step 3-6 of algorithm 4/5 on A_i with $dim(A_i)$ and chosen level b_i for $i = 1(1)k$ with the following affinity measures:

$$W_{ij} = exp\big(-\frac{[(\frac{logn}{n})^{-\frac{1}{dim(A_i)}} LLPD(x_i, x_j, C_s)]^2}{\sigma^2}\big)\, exp\big(\frac{||Q_i - Q_j||^2}{\eta^2}\big) \tag{5.37}$$

or

$$W_{ij} = exp\big(-\frac{[n^{\frac{p-1}{p\, dim(A_i)}} L_p^N(x_i, x_j, C_s)]^2}{\sigma^2}\big)\, exp\big(\frac{||Q_i - Q_j||^2}{\eta^2}\big) \tag{5.38}$$

where $Q(.)$ is the local PCA measure as introduced in Arias-Castro et al. (2017).

This similarity measure is close to 1 if both $n^{\frac{p-1}{p\, dim(A_i)}} L_p^N(x_i, x_j, C_s)$ and $||Q_i - Q_j||$ are close to 0. This can only happen when x_i and x_j are in same high density region, are sufficiently close to one another (small $nn^{\frac{p-1}{p\, dim(A_i)}} L_p^N(x_i, x_j, C_s)$ ensures this) and there is no sudden changes of curvature in the underlying manifold between these 2 data points (small $||Q_i - Q_j||$ ensures this). If one of these criterion fails, the similarity measure in (5.38) will be small. Thus this measure is effective in MMC scenario. Similar justification can be provided for the similarity measure in (5.37).□

5.13 Comparison of LLPD, L_p^N and g^N SC

We conclude this chapter by showing that after the de-noising procedure is applied, for LDLN model g^N distance is able to distinguish clusters w.p. 1 and compare LLPD, L_p^N and L_g^N distances in this scenario.

5.13.1 Noiseless LDLN model

In the noiseless scenario as discussed in Mckenzie and Damelin (2019) we assume clusters $B_1, ..B_K$ are all $\leq D$ dimensional open and bounded subsets embedded in $\mathbb{R}^D$ with $\min\limits_{i\neq j} d_1(B_i, B_j) = \delta$, $\max\limits_{i} \text{diam}(B_i) = \triangle$ where $diam(A) \coloneqq \max\limits_{x,y\in A} dist_1(x, y)$. I.e. in this case we sample χ_n from the model

$$f(x) = c\sum_{i=1}^{k} f_i(x)1_{B_i}$$

for some normalizing constant c and $f_i > 0$ for all i. We apply algorithm 1 and 2 (for L_{sym} and L respectively) with the similarity measure $e^{-\frac{d(x_i,x_j,\chi_n)^2}{\sigma^2}}$, where $d(.,.)$ is $LLPD(x_i, x_j, \chi_n)$, $L_p^N(x_i, x_j, \chi_n)$ or $g^N(x_i, x_j, \chi_n)$ depending on the context.

Theorem 5.8

Under this data generating model, SC with LLPD, L_p^N and g^N with L or L_{sym} and eigen-gap heuristics is able to perfectly identify number of clusters and cluster the labels. □

Proof 5.23

As $n \to \infty$,

1. For LLPD SC, we have within cluster distances $\leq O((\frac{logn}{n})^{\frac{1}{D}})$ and between cluster distances $\geq \delta$ with within cluster distance $\leq O(n^{-r})$ from it's expected value for any $r > \frac{1}{D}$ w.p. 1.

2. For L_p^N as well as g^N with lowest power $1 < p < \infty$, we have within cluster distances $\leq O(n^{\frac{1-p}{pD}})$ and between cluster distances $\geq \delta$ with within cluster distance $\leq O(n^{-r})$ from it's expected value for any $r > \frac{p-1}{pD} + \frac{1}{2(D+2p)}$ w.p. 1.

   For both of these 2 cases, within cluster distance $\to 0$ while between cluster distance remains bounded below

3. For g^N with lowest power 1 we have within cluster distances $\leq g^{-1}(a_1 \triangle)$, which can be made arbitrarily small by suitably choosing g while between cluster distance remains $\geq \delta$ with within cluster distance $\leq O(n^{-r})$ from it's expected value for any $r > \frac{1}{D+b_2+1}$ w.p. 1 where a_1 is the coefficient corresponding to the linear term of g and b_2 is the 2nd lowest power of g. □

Remark 5.18

Proof of these 3 results follows from the convergence rates of LLPD and g^N-distance as derived in chapter 2 and 3. We show the L_p^N/g^N scenario with $p > 1$ here. From the result in chapter 2,

$$\mathbb{P}\big(\sup_{x,y}|\frac{L_p^N(x,y,\chi_n)}{n^{\frac{1-p}{pD}}\, dist_p(x,y)^{\frac{1}{p}}} - C(D,p)| > \epsilon\big) \leq exp(-\theta_0^{'}\epsilon^2 n^{\alpha})$$

for $\alpha = \frac{1}{D+2p}$.

$$\Rightarrow \mathbb{P}\big(\sup_{x,y}|L_p^N(x,y,\chi_n) - C(D,p)\, n^{\frac{1-p}{pD}}\, dist_p(x,y)^{\frac{1}{p}}| > \epsilon\big) \leq exp(-\theta_0^{''}\epsilon^2 n^{\alpha+\frac{2(p-1)}{pD}})$$

Thus, w.p. $\to 1$, $|L_p^N(x,y,\chi_n) - C(D,p)\, n^{\frac{1-p}{pD}}\, dist_p(x,y)^{\frac{1}{p}}| < n^{-r}$ for any $r > \frac{p-1}{pD} + \frac{1}{2(D+2p)}$. □

Remark 5.19

Note that if we let $\sigma = O(1)$ and sufficiently small, then in LLPD as well as in L_p^N scenario, we have within cluster distance $\to 0$ while between cluster distances remaining $\geq \delta$ as $n \to \infty$ resulting in a W with $a^* << 1$ and $b^* = 1$ scenario. Thus,

SC with LLPD, L_p^N and g^N , with L or L_{sym} and eigen-gap heuristics will be able to perfectly identify number of clusters as well as the cluster labels. Also note that if p and D are sufficiently large, $\frac{p-1}{pD} + \frac{1}{2(D+2p)} > \frac{1}{D}$ making SC with L_p^N more robust than LLPD SC. But on the other hand, LLPD SC have within cluster faster convergence rate $O((\frac{logn}{n})^{\frac{1}{D}})$ as opposed to L_p^N SC, which is $O(n^{\frac{1-p}{pD}})$. This in particular provides the theoretical justification to the simulation observations in Mckenzie and Damelin (2019). □

Remark 5.20

Note that for g^N-distance with lowest power = 1, the upper bound of within cluster distance $g^{-1}(a_1\triangle)$ can be made arbitrarily small by choosing g with sufficiently higher powers and/or higher coefficients corresponding to higher powers- resulting in a highly "contraction mapping".

For instance, for $g(t) = e^{at} - 1$ and $g(t) = at + bt^2$, which has been widely used in SC/K-means algorithms; under this settings we have within cluster distances $\leq \frac{1}{a} ln(1+a\triangle)$ and $\leq \frac{\sqrt{a^2+4ab\triangle}-a}{2b}$ respectively, both of which can be made arbitrarily small by choosing a and $\frac{b}{a}$ sufficiently large respectively while the between cluster distances always remains bigger than δ. Thus in this scenario as well we obtain a W with $a^* << 1$ and $b^* \approx 1$. Thus, SC g^N when the lowest power is 1 , with L or L_{sym} and eigen-gap heuristics will be able to perfectly identify number of clusters as well as the cluster labels. □

Chapter 6

High dimensional, robust, unsupervised record linkage

6.1 Introduction

In recent times, owing to rapid advancement of a variety of technological resources and services, and increasingly digitally connected environment, numerous kinds of datasets are available. For example, for a given community of individual's, there may be very high dimensional data available on each individual's (i) online shopping patterns, as well as on their (ii) social media presence and usage. It may of interest to businesses to understand their customers better based on their social and cultural backgrounds, consequently it is of interest to link an online shopper's profile with their social media data. Owing to privacy rights of individuals and confidentiality concerns, identifying information may not be available to the statistician linking the records.

This chapter is primarily on a methodology for linking high dimensional datasets of above type. Many existing approaches for entity resolution and record linkage are applicable only on low dimensional datasets, and where the datasets have shared features or variables. We do not require the two datasets to have a common set of features and in fact, present our discussion for the case where the datasets have no

common variable.

In this context, we also develop a mathematical framework for the topic of record linkage, for better understanding and tractability of the theoretical properties of such linking algorithms. Parts of the existing literature on record linkage and entity resolution are based on *ad hoc* principles, and we hope to address some foundational challenges in this topic.

Additionally, it is routinely observed that high dimensional datasets can contain outliers or aberrant observations. Consequently a major aspect of our proposed methodology is to develop *robust* techniques for data linkage. Our proposal borrows recent studies on high-dimensional principal components and extends them to the case of robust principal components.

Another aspect of this chapter is that we propose a computationally much simpler and easily implementable method for linking records than the available Bayesian approaches such as the ones given in Liseo and Tancredi (2013). Other machine learning approaches involving graphs and networks that are sometimes adopted for entity resolution, also require heavy machinery computing. It is not clear if such extremely computation intensive methodology is either necessary, or whether there is a principled statistical foundation to such methodology. Apropos of this, the computational burden from our proposal is significantly lower.

This chapter advocates a principled approach. Our approach is broadly as follows: we implement a robust, sparse, high-dimensional principal component analysis (PCA often hereafter) on both datasets, and consolidate the information about each observation (that is, each row of both matrices) into a low-dimensional vector (p_0 in the notations of this chapter). Then, we compute correlations between these p_0 dimensional vectors from the two matrices, with the understanding that an existing linkage will show up as a highly correlated entry. The threshold for the correlation is based on the training set in the current chapter, but our principle is workable even

when there is no training set available. Owing to the facts that (a) our proposal requires no common features or variables common to both datasets, and (b) we do not require a training dataset, we call our proposal *unsupervised* record linkage.

In order to ensure clarity of our presentation and to keep the technicalities at a reasonable level, in this chapter we only present results on unsupervised record linkage where all the variables are continuous in nature. In particular, commonly used variables for record linkage, like name, date of birth, address, do not satisfy our technical assumptions. In practice, we may use a traditional method using the nominal and ordinal variables to do a preliminary sub-setting of potential linkages, after which the unsupervised record linkage method may be used on the continuous variables. Also, it is possible in some cases to use continuous variables as underlying latent variables governing the behavior of a categorical random variable. These directions of research will be part of our future work.

The rest of this chapter is organized as follows: In Section 6.2 we present a brief and necessarily incomplete state of the literature on record linkage, in order to clarify how our contribution differs from the advancements on this topic thus far. Then, in Section Section 6.3 we discuss notations, the conceptual framework of the datasets that we propose to link, and the linking model. Following that, in Section 6.4 we present our statistical model and technical arguments. In particular, Section 6.4.1 contains the record linking algorithm, and Section 6.4.2 contains the theoretical framework and justifications for our statistical model and algorithmic steps. Then, in Section 6.5 we present some numeric results based on simulation studies, along with additional comments on practical implementation of our proposed methodology. A final Section 6.6 collects our concluding remarks.

6.2 A broad overview of record linkage

The seminal article by Fellegi and Sunter (1969) presented the first mathematical model for this topic, based on earlier work by Newcombe and Kennedy (1962) for one-to-one entity resolution across two databases in terms of Neyman-Pearson hypothesis testing.

In this brief review, we focus primarily on *bipartite record linkage*, where the key assumption is that each entity is recorded *at-most* once in each files. Most of the literature (including our set-up) on record linkage falls in this scenario. This assumption implies a maximum one-to-one restriction in the linkage, that is, a record from one file can be linked with maximum one record from the other file.

The main principle of bipartite record linkage may be described as follows: Consider two data files $Y_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$ that record information from two overlapping sets of individuals or entities. These data files contain n_ℓ records respectively (without loss of generality we assume $n_1 \leq n_2$) for $\ell = 1, 2$ with n_0 being the number of entities simultaneously recorded in both files, hence $0 \leq n_0 \leq n_1$.

In the bipartite record linkage context, we can think of the records from files Y_1 and Y_2 as two disjoint sets of nodes, where an edge between two records represents them referring to the same entity, which we also call being co-referent or being a match. Formally, this match can be encoded into a matrix $\triangle_{n_1 \times n_2}$as follows:

$$\triangle_{ij} = \begin{cases} 1 & \text{if records } i \in Y_1 \text{ and } j \in Y_2 \text{ represent the same entity} \\ 0 & \text{otherwise} \end{cases}$$

The characteristics of a bipartite matching imply that at-most one entry in each column and each row of $\triangle$ can be equal to 1. The goal of bipartite record linkage is to estimate $\triangle$ using the information contained in Y_1 and Y_2.

The set of ordered record pairs $Y_1 \times Y_2$ can be thought as the union of the set of matches $M = \{(i, j) : i \in Y_1,\ j \in Y_2,\ \triangle_{ij} = 1\}$ and the set of non-matches

$U = \{(i,j) : i \in Y_1,\ j \in Y_2,\ \triangle_{ij} = 0\}$. Thus, the problem of estimating $\triangle$ from Y_1 and Y_2 can be seen as identifying the sets M and U. When record pairs are estimated to be matches they are called links and when estimated to be non-matches they are called non-links.

6.2.1 The Fellegi–Sunter approach of record linkage

The key idea of the Fellegi-Sunter approach is as follows: Comparison vectors $\boldsymbol{\gamma}_{ij}$ are obtained for each record pair (i,j) in $Y_1 \times Y_2$ with the goal of finding evidence of whether they represent matches or not. These vectors can be written as $\boldsymbol{\gamma}_{ij} = (\gamma_{ij}^1, \ldots, \gamma_{ij}^f, \ldots, \gamma_{ij}^F)$, where F denotes the number of criteria used to compare the records. Traditionally, these F criteria correspond to one comparison for each variable that the data files have in common.

Let $S_f(i,j)$ denote a similarity measure computed from field f of records i and j. The range of S_f can be divided into L_f+1 intervals $I_{f0}, I_{f1}, \ldots, I_{fL_f}$, which represent different disagreement levels. In this construction, the interval I_{f0} represents the highest level of agreement, which includes total agreement, and the last interval I_{fL_f} represents the highest level of disagreement.

In their chapter, Fellegi and Sunter (1969), the authors propose to the log-likelihood ratios

$$w_{ij} = \log \frac{\mathbb{P}[\boldsymbol{\gamma}_{ij} | \triangle_{ij} = 1]}{\mathbb{P}[\boldsymbol{\gamma}_{ij} | \triangle_{ij} = 0]}$$

as weights to estimate which record pairs are matches. The expression for w_{ij} assumes that $\boldsymbol{\gamma}_{ij}$ is a realization of a random vector, say, G_{ij} whose distribution depends on the matching status $\triangle_{ij}$ of the record pair. Similar to the Neyman-Pearson hypothesis testing, if this ratio is large we favor the hypothesis of the pair being a match.

When $\mathbb{P}[\boldsymbol{\gamma}_{ij} | \triangle_{ij} = 1]$ and $\mathbb{P}[\boldsymbol{\gamma}_{ij} | \triangle_{ij} = 0]$ are known, the procedure orders the possible values of $\boldsymbol{\gamma}_{ij}$ by their weights w_{ij} in non-increasing order, indexing by the

subscript h, and determines two values, $h^{'}$ and $h^{''}$, such that $\sum_{h\leq h^{'}-1}\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=0] < \mu \leq \sum_{h\leq h^{'}}\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=0]$ and $\sum_{h\geq h^{''}}\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=1] \geq \lambda > \sum_{h\geq h^{''}+1}\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=1]$, where $\mu = \mathbb{P}[\text{assign } (i,j) \text{ as link}|\triangle_{ij}=0]$ and $\lambda = \mathbb{P}[\text{assign } (i,j) \text{ as non-link}|\triangle_{ij}=1)$ are two admissible "type 1" and "type 2" error levels.

Finally, the record pairs are classified into 3 groups:

1. Those with $h \leq h^{'}-1$ are declared links
2. Those with $h \geq h^{''}+1$ are non-links and
3. Those with configurations between $h^{'}$ and $h^{''}$ require clerical review.

Fellegi and Sunter showed that this decision rule is optimal in the sense that for fixed values of μ and λ it minimizes the probability of sending a pair to clerical review.

However, in practice, $\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=1]$ and $\mathbb{P}[\boldsymbol{\gamma}_{ij}|\triangle_{ij}=0]$ are not known, and have to be estimated from Y_1 and Y_2. So, Jaro (1989); Larsen and Rubin (2001) proposed to model the comparison data using mixture models of the type

$$
\begin{aligned}
G_{ij}|\triangle_{ij}=1 &\overset{iid}{\sim} M(\boldsymbol{m}) \\
G_{ij}|\triangle_{ij}=0 &\overset{iid}{\sim} U(\boldsymbol{u}), \\
\triangle_{ij} &\overset{iid}{\sim} \text{Bernoulli}(\theta)
\end{aligned}
$$

for comparison variables G_{ij}, some distributions $M(\boldsymbol{m})$ and $U(\boldsymbol{u})$, and $\theta \in (0,1)$. Estimation of M and U is usually done by EM-type algorithms.

6.2.2 Machine learning/classification approach

In-general, record linkage task becomes quickly infeasible with size (n_ℓ) as well as the dimension (p_ℓ) of data files. A common solution to this problem is to partition the data files into blocks (e.g. geography, or gender and year of birth) of records

determined by information that is thought to be accurately recorded in both data files, and then solve the task only within blocks. See Christen (2011); Steorts et al. (2014) for extensive surveys. Earlier development into blocking is presented in Herzog et al. (2007), who also discuss the use of blocking to identify duplicate list entires and for matching records between two sample surveys.

Recently a common approach of tackling the record linkage problems has been to treat it as a traditional supervised or semi-supervised classification problem: we need to classify record pairs into matches and non-matches. If we have a sample of record pairs for which the true matching statuses are known, we can train a classifier on this sample using comparisons between the pairs of records as our predictors, and then predict the matching statuses of the remaining record pairs. See Martins (2011); Torvik and Smalheiser (2009); Treeratpituk and Giles (2009); Ventura et al. (2015) for some examples.

6.2.3 Bayesian methods

Bayesian methods have a long history of use in record linkage models. A major advantage of Bayesian methods is their natural handling of uncertainty quantification for the resulting estimates. For a review of recent development in Bayesian methods, see Liseo and Tancredi (2013). While some of the Bayesian work incorporates the record data only through pairwise similarity scores Sadinle (2017); Sadinle and Fienberg (2013), other works Steorts et al. (2016) directly model the actual record data which usually requires crafting specific models for each type of field, and therefore mostly deal with categorical information. However, recently Steorts and Others (2015) has generalized Bayesian methods to incorporate string variables such as addresses, phone numbers, or dates.

In addition, Sadinle and Fienberg (2013) has extended the Fellegi-Sunter approach to linking records across more than two databases. Also, Singla and Domingos (2006);

Enamorado et al. (2018) generalized the underlying mixture models (specially the i.i.d. assumptions) in the Fellegi-Sunter approach.

6.3 Notations, the data and linking model

The main focus of this chapter is to *link* observations from two datasets. Both datasets are matrices, with iid rows. However, the same observational units may have been used for both datasets. For example, the 17-th row of the first dataset and 47-th row of the second dataset may belong to the same individual. For a number of cases (n in the notation used below) we know the linkage, and thus can match and pair the information from both datasets. However, such linkage is not known for many other rows of both datasets. The main goal of a ***record linkage*** exercise is to establish such linkages.

It is also commonly understood that most observations from both datasets are linked, and only a handful of observations from either dataset do not have an entry on the other. For this chapter, we assume that the datasets do not have any common variables. We also assume that both datasets considered here are high-dimensional.

6.3.1 Notations

Since the data, various parameters and latent variables will have multiple indices, we establish some notations first. The notation $a_{\square}$ denotes a vector, of dimension that will be determined by the context. All vectors are column vectors, and the notation $a_{\square}^T$ or a^T denotes the transpose, and $|a|$ denotes its Euclidean norm. The $n \times m$ matrix A has column vectors denoted by $A_{\square,1}, \ldots, A_{\square,m} \in \mathbb{R}^n$, and row vectors denoted by

$A_{1,\square}, \ldots, A_{n,\square} \in \mathbb{R}^m$, thus

$$A = (A_{\square,1} : A_{\square,2} : \ldots : A_{\square,m})_{n\times m} = \begin{pmatrix} A_{1,\square} \\ A_{2,\square} \\ . \\ \\ . \\ . \\ A_{n,\square} \end{pmatrix}_{n\times m} .$$

We index the datasets used in this chapter by $\ell = 1, 2$, and the notation Y_ℓ stands for the ℓ-th dataset with dimensions $n_\ell \times p_\ell$, consisting of n_ℓ independent observations of the p_ℓ features that are stacked as row vectors of the Y_ℓ matrix. The k-dimensional multivariate Normal distribution with mean $\mu \in \mathbb{R}^p$ and variance $\Sigma \in \mathbb{R}^{p\times p}$ will be denoted by $N_p(\mu, \Sigma)$. The notation $X_i \overset{i.i.d.}{=} \mathbb{F}$ denotes that the X_i's are independent, identically distributed according to $\mathbb{F}$.

6.3.2 The data and the linking framework

We consider two datasets for linkage, $Y_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$ for $\ell = 1, 2$. Without loss of generality, $n_1 \leq n_2$. In both datasets, each row represents an observation, and each column a feature. We assume, for mathematical simplicity, that there is no duplication of features in the two datasets. Within each matrix, the rows are independent. However, a pair of rows, one from each matrix, may have dependency.

For any positive integer k, let $\mathbb{N}_k = \{1, 2, \ldots, k\}$, the set of positive integers or natural numbers up to and including k. For any finite set $\mathcal{S}$ (for example, $\mathbb{N}_k$), let $\sigma(\mathcal{S})$ be any permutation of the elements of $\mathcal{S}$.

Suppose that $n_0 \leq n_1 (\leq n_2)$ is the unknown number of *linked* observations between the datasets. Define $\tilde{n}_\ell = n_\ell - n_0$ for $\ell = 1, 2$, denoting the number of un-

matched observations from either dataset. Define $N = n_0 + \tilde{n}_1 + \tilde{n}_2$, this is the total number of observational units we consider, and we label the observational units with the index set $\mathbb{N}_N$. The i-th observation of the first dataset, $Y_{1,i,\square} \in \mathbb{R}^{p_1}$, corresponds to the unit whose index matches with the i-th element of $\sigma\big(\{1, \ldots, n_0 + \tilde{n}_1\}\big)$, a random permutation of the index subset $\mathcal{S}_1 = \{1, \ldots, n_0 + \tilde{n}_1\}$. The index subset of the second data is $\mathcal{S}_2 = \{1, \ldots, n_0, n_0 + \tilde{n}_1 + 1, \ldots, N\}$. The i-th observation of the second dataset, $Y_{2,i,\square} \in \mathbb{R}^{p_2}$, corresponds to the unit whose index matches with the i-th element of $\sigma\big(\mathcal{S}_2\}\big)$, a random permutation of the index subset $\mathcal{S}_2$. Note that $\mathcal{S}_1$ and $\mathcal{S}_2$ have exactly n_0 elements in common, reflecting the n_0 matched observational units for the two datasets.

In a few cases, the linkage between some observations units is known, and forms the training set. A training set is not needed for the present chapter, but if indeed we have known and established linkages between, say $n \leq n_0$, observational units, without loss of generality we stack these known linkage cases as the first n rows of Y_ℓ, $\ell = 1, 2$.

We assume that the observations satisfy

$$Y_{\ell,i,\square} \overset{i.i.d.}{=} N_{p_\ell}\Big(0, \Sigma_\ell\Big), i = 1, \ldots, n_\ell, \quad \Sigma_\ell \text{ unknown}, \ \ell = 1, 2.$$

We use the spectral representation

$$\begin{aligned}
&\Sigma_\ell = \Gamma_\ell \Lambda_\ell \Gamma_\ell^T, \text{ where} \\
&\Lambda_\ell = diag\Big(\lambda_{\ell,1}, \ldots, \lambda_{\ell,p_\ell}\Big), \text{ with} \\
&\quad \lambda_{\ell,1} \geq \lambda_{\ell,2} \geq \ldots \geq \lambda_{\ell,p_0} \gg \lambda_{\ell,p_0+1} \geq \ldots \lambda_{\ell,p_\ell}, \\
&\Gamma_\ell = \Big[\gamma_{\ell,\square,1} : \ldots : \gamma_{\ell,\square,p_\ell}\Big] \in \mathbb{R}^{p_\ell \times p_\ell}.
\end{aligned}$$

Thus, Λ_ℓ is the diagonal matrix of the eigenvalues of Σ_ℓ, and the columns of Γ_ℓ contain

the corresponding eigenvectors. It is assumed that the first p_0 eigenvalues are considerably higher than the rest, and contain information relevant for linking records. Also for mathematical simplicity, we assume henceforth that $\lambda_{1,j} = \lambda_{2,j}$ for $j = 1, \ldots, p_0$. That is, the top p_0 eigenvalues are the same. This is not a necessary assumption, but makes the presentation and technicalities of the developments presented below considerably simpler. We assume that the top p_0 eigenvectors of both $\Sigma_\ell, \ell = 1, 2$ are sparse, in that all but κ_ℓ of the entries in these eigenvectors are zero. This is a necessary assumption to obtain statistically consistent and computationally obtainable estimators of the principal components that we use in this chapter, see Wang et al. (2016) for further details.

There are multiple record linkage contexts in which the above framework may be useful. First, traditional linkage techniques that rely on nominal and ordinal variables like names, addresses and so on often result in plausible subsets of observations from one dataset linked to each unit of the other dataset. At that stage, a further analysis based on the continuous variables as described here may be useful. Second, due to confidentiality and privacy considerations, datasets are often anonymized. In such cases, the model presented above may be extremely useful, either directly for modeling the reported continuous variables, or in conjunction with other continuous but latent variables. Third, our framework allows the scope of record linkage to extend beyond the traditional applications of linking sample surveys involving individuals or households, into linking data from multitude of sources, like social media, online shopping platforms, and electronic records of various kinds (Li et al., 2020; Fatemi et al., 2018). In many such contexts, the observed and often suitably anonymized data may be modeled using high-dimensional continuous (observed or latent) variables. With such an extended scope, record linkage may provide an increase in precision and accuracy of recommender systems (Drachsler et al., 2010; Shabtai et al., 2013; Slokom, 2018), for providing online security and privacy (Zhu et al., 2016; Salas, 2019), for

transfer learning (Rong et al., 2012) and for distributed computing and related technical developments.

6.4 The statistical model

Without loss of generality and to considerably simplify the presentation below, we assume that the first n_0 rows of $Y_\ell, \ell = 1, 2$ are linked. To relate the two datasets $Y_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$, $\ell = 1, 2$, we define the following quantities:

$$\begin{pmatrix} Z_{1,i,j} \\ Z_{2,i,j} \end{pmatrix} \overset{i.i.d.}{=} N_2 \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix} \right), \text{ when } i = 1, \ldots, n_0 \text{ and } j = 1, \ldots, p_0,$$
$$Z_{\ell,i,j} \overset{i.i.d.}{=} N(0,1), \text{ for } \ell = 1, 2, \text{ when } i > n_0 \text{ or } j > p_0.$$

Thus, $Z_{\ell,i,j}$ are all standard normal random variables, and for the case $i = 1, \ldots, n_0$ and $j = 1, \ldots, p_0$, the two random variables $Z_{1,i,j}$ and $Z_{2,i,j}$ share a correlation ρ between them. We arrange the $Z_{\ell,i,j}$ into two matrices of dimensions identical to those of our datasets Y_ℓ. Thus, $Z_{\ell,i,j}$ is the (i,j)-th element of the matrix $Z_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$, $\ell = 1, 2$. It can be seen that each matrix Z_ℓ has iid $N(0,1)$ entries, but the top left corners of Z_1 and Z_2 are related.

We model the data as

$$Y_{\ell,i,\square} = \Gamma_\ell \Lambda_\ell^{1/2} Z_{\ell,i,\square}, \text{ where}$$
$$Z_{\ell,i,\square} \overset{i.i.d.}{=} N_{p_\ell}\Big(0, \mathbb{I}_{p_\ell}\Big), i = 1, \ldots, n_\ell$$

Algorithm 8 Record Linking Algorithm

1. Scale each row of the two datasets by their respective norms, to get $X_{\ell,i,\square} = Y_{\ell,i,\square}/|Y_{\ell,i,\square}|$, for $\ell = 1, 2$ and $i = 1, \ldots, n_l$. Collect these in the matrices $X_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$, $\ell = 1, 2$.

2. Run the high dimensional sparse PCA algorithm due to Wang et al. (2016) on $X_\ell, \ell = 1, 2$. This obtains the leading eigenvalue and eigenvector for these matrices. Project the data on the orthogonal space to this estimated eigenvector, and repeat the process to obtain the leading p_0 eigenvectors.

3. Obtain the coefficients $W_{\ell,i,\square} \in \mathbb{R}^{p_0}$ (given in (6.1) below) for each $i = 1, \ldots, n_\ell$, $\ell = 1, 2$ from the projections of the observations on the top p_0 eigenvectors.

4. Obtain the correlations $C(i, \tilde{i})$ of $W_{1,i,\square}$ and $W_{2,\tilde{i},\square}$.

5. Arrange the correlations in descending order. Based on the values corresponding to the training set and a pre-set value for the maximum proportion of false positive matches, select a correlation threshold. If there are multiple matches above this threshold for any $i \in \{1, \ldots, n_1\}$ or $\tilde{i} \in \{1, \ldots, n_2\}$, the match with the higher correlation value is chosen. The proportion of false negatives is estimated from the number of training sample matches below the threshold.

described above. In matrix terms, this then translates to

$$Y_\ell = Z_\ell \Lambda_\ell^{1/2} \Gamma_\ell^T, \text{ where}$$
$$Z_{\ell,i,j} \overset{i.i.d.}{=} N(0, 1).$$

Note, however, that we do not imply with the above that the matrices $Z_\ell \in \mathbb{R}^{n_\ell \times p_\ell}$ are independent of each other, and indeed they are not.

6.4.1 The record linking algorithm

Our proposed algorithm (algorithm 8) is as follows:

We now discuss in details the steps outlined above. First, the scaling $X_{\ell,i,\square} =$

$Y_{\ell,i,\square}/|Y_{\ell,i,\square}|$ ensures that each $X_{\ell,i,\square}$ has unit norm, thus ensuring that outliers do not affect the PCA and subsequent computations. It is well known that PCA is very sensitive to outliers. The second step of the above algorithm is about computation of the high dimensional principal components using an established procedure. The third through last steps of the algorithm are about using the principal component scores to obtain the record linkage. There can be considerable variation in the details in these steps depending on the context, we have presented one simple procedure.

The above algorithm used the training data only for the last step of setting a threshold for the correlations. We can easily formulate a variation, where a training set is not needed. Since we compute $n_1 n_2$ correlations of which only $\min(n_1, n_2)$ can possibly correspond to linked data, a threshold can be determined based a change-point in the correlation values, or on multiple matches. We illustrate this aspect in simulations reported later in this chapter.

6.4.2 Theoretical properties

The justification for using the above algorithm rests on the fact that there is a clear separation of the correlation values between the linked data-pairs, as opposed to the correlation between the not-linked cases. We establish this fact in the following result:

Theorem 6.1

Under the conditions of the model, the population correlation value for each linked pair of observations is ρ, and is zero for two observations that are not linked. $\square$

Thus, there is a clear separation of the correlation values from the linked data-pairs from the rest. There would be sample variations, and the value of ρ is not known. Consequently, either a training set-based threshold or a change detection technique can be used to sort the true linkages from the rest.

Proof 6.1 (Proof of Theorem 6.1)

Let

$$\tilde{\Gamma}_\ell = \left[\gamma_{\ell,\square,1} : \ldots : \gamma_{\ell,\square,p_0}\right] \in \mathbb{R}^{p_\ell \times p_0},$$

the first p_0 columns of Γ_ℓ, for which the eigenvalues are considerably higher than the rest. We project the datapoints $Y_{\ell,i,\square}$ on the column space of $\tilde{\Gamma}_\ell$, for $i = 1, \ldots, n_\ell$. Note that all columns of $\tilde{\Gamma}_\ell$ are orthonormal (by construction, since these are estimators of successive eigenvectors, so they are orthogonal to each other and have unit norm). Given this, it is easy to see that the projection of $Y_{\ell,i,\square}$ is

$$\begin{aligned}\tilde{Y}_{\ell,i,\square} &= \sum_{j=1}^{p_0} \left\langle Y_{\ell,i,\square}, \gamma_{\ell,\square,j} \right\rangle \gamma_{\ell,\square,j} \\ &= \tilde{\Gamma}_\ell \tilde{\Gamma}_\ell^T Y_{\ell,i,\square} \\ &= \tilde{\Gamma}_\ell W_{\ell,i,\square}.\end{aligned}$$

The relevant information about the projections is carried in the low-dimensional ***weights*** $W_{\ell,i,\square} \in \mathbb{R}^{p_0}$, consequently we develop our analysis based on these below. Putting these weights as rows in a matrix, we have

$$\begin{aligned}W_\ell &= Y_\ell \tilde{\Gamma}_\ell \\ &= Z_\ell \Lambda_\ell^{1/2} \Gamma_\ell^T \tilde{\Gamma}_\ell \in \mathbb{R}^{n_\ell \times p_0}\end{aligned} \tag{6.1}$$

This last expression can be simplified further, since

$$\begin{aligned}\Gamma_\ell^T\tilde{\Gamma}_\ell &= \begin{pmatrix} \gamma_{\ell,\square,1}^T \\ \gamma_{\ell,\square,2}^T \\ \cdot \\ \cdot \\ \cdot \\ \gamma_{\ell,\square,p_\ell}^T \end{pmatrix} \left[\gamma_{\ell,\square,1} : \ldots : \gamma_{\ell,\square,p_0}\right] \\ &= \begin{pmatrix} \mathbb{I}_{p_0} \\ \mathbf{0}_{(p_\ell - p_0)\times p_0} \end{pmatrix}.\end{aligned}$$

Thus we have

$$\begin{aligned}&\Lambda_\ell^{1/2}\Gamma_\ell^T\tilde{\Gamma}_\ell \\ &\qquad = \begin{pmatrix} \lambda_{\ell,1}^{1/2} & 0 & 0 & 0 \\ 0 & \lambda_{\ell,2}^{1/2} & 0 & 0 \\ \cdot & & & \\ 0 & 0 & 0 & \lambda_{\ell,p_0}^{1/2} \\ 0 & 0 & 0 & 0 \\ \cdot & & & \\ \cdot & & & \\ 0 & 0 & 0 & 0 \end{pmatrix} \\ &\qquad = \begin{pmatrix} \tilde{\Lambda}_\ell^{1/2} \\ \mathbf{0}_{(p_\ell - p_0)\times p_0} \end{pmatrix} \in \mathbb{R}^{p_\ell \times p_0}.\end{aligned}$$

Define

$$\tilde{Z}_\ell = \left[Z_{\ell,\square,1} : Z_{\ell,\square,2} : \ldots : Z_{\ell,\square,p_0}\right] \in \mathbb{R}^{n_\ell \times p_0}.$$

Consequently, we have

$$\begin{aligned} W_\ell &= Y_\ell \tilde{\Gamma}_\ell \\ &= Z_\ell \Lambda_\ell^{1/2} \Gamma_\ell^T \tilde{\Gamma}_\ell \\ &= \left[\lambda_{\ell,1}^{1/2} Z_{\ell,\square,1} : \lambda_{\ell,2}^{1/2} Z_{\ell,\square,2} : \ldots : \lambda_{\ell,p_0}^{1/2} Z_{\ell,\square,p_0} \right] \\ &= \tilde{Z}_\ell \tilde{\Lambda}_\ell^{1/2} \in \mathbb{R}^{n_\ell \times p_0}. \end{aligned}$$

We thus have, for any $i \in \{1, \ldots, n_0\}$

$$\begin{aligned} \mathbb{E} W_{\ell,i,\square} &= 0, \\ \mathbb{V} W_{\ell,i,\square} &= \tilde{\Lambda}_\ell, \\ \mathbb{E} W_{1,i,\square} W_{2,i,\square}^T &= \tilde{\Lambda}_1 \mathbb{E} \tilde{Z}_{1,i,\square} \tilde{Z}_{2,i,\square}^T \tilde{\Lambda}_2, \\ \mathbb{E} W_{1,i,\square}^T W_{2,i,\square} &= \mathbb{E} \tilde{Z}_{1,i,\square}^T \tilde{\Lambda}_1 \tilde{\Lambda}_2 \tilde{Z}_{2,i,\square}^T = \sum_{j=1}^{p_0} \lambda_{1,j}^{1/2} \lambda_{1,j}^{1/2} \mathbb{E} \tilde{Z}_{1,i,j} \tilde{Z}_{2,i,j}. \end{aligned}$$

Then it follows that

$$Cor\big(W_{1,i,\square}, W_{2,i,\square}\big) = \rho. \qquad \square$$

The algebra for the observations that are not linked is similar and omitted here.

One important consideration for our framework is to ensure that the robustness procedure we implemented in the first step does not alter the eigenvector structure of the original data. That is, we need to ensure that the eigenvectors of Σ_ℓ match those of the variance of X_ℓ, $\ell = 1, 2$. This is ensured in the following result:

Theorem 6.2

Suppose $X \in \mathbb{R}^p$ is a random vector with variance Σ_X, and let the variance of $U = X/|X|$ be denoted by Σ_U.

(i) When X has an elliptically symmetric distribution and zero mean, the eigenvectors of Σ_X and Σ_U are identical.

(ii) If $\mathbb{E}U = 0 \in \mathbb{R}^p$, $\mathbb{E}U|X| = 0 \in \mathbb{R}^p$ and $\mathbb{E}|X|^2UU^T = \mathbb{E}|X|^2\mathbb{E}UU^T \in \mathbb{R}^{p\times p}$, then again the eigenvectors of Σ_X and Σ_U are identical. Moreover, if the eigenvalues of Σ_X are $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_p > 0$, then the eigenvectors of Σ_U are $\frac{\lambda_1}{\sum_{i=1}^p \lambda_i} \geq \ldots \frac{\lambda_p}{\sum_{i=1}^p \lambda_i} > 0$. □

Note that a p-dimensional random vector X is said to elliptically distributed if there exist a vector $\mu \in \mathbb{R}^p$, a positive semi-definite matrix $\Sigma \in \mathbb{R}^{p\times p}$ and a function $\phi : (0, \infty) \to \mathbb{R}$ such that the characteristic function of X is $\exp\{it^T\mu\}\phi(t^T\Sigma t)$ for $t \in \mathbb{R}^p$. See Fang et al. (1990) for several alternative and equivalent definitions of the elliptically contoured family, as well as for additional details. An example of elliptically contoured distribution is the multivariate Gaussian distribution, thus the framework adopted in this chapter satisfies the first condition of Theorem 6.2. The second part of Theorem 6.2 is for general interest, in case an elliptic distributional assumption is not satisfied.

Proof 6.2 (Proof of Theorem 6.2)

First, consider the case where X has an elliptically symmetric distribution with mean zero. In such cases, we may write $X = R\Gamma\Lambda^{1/2}E$, where Γ is a rotation matrix, Λ is a diagonal matrix with positive elements, E is uniformly distributed on the unit sphere and R is a positive random variable that is independent of E. Then, we have $|X|^2 = R^2E^T\Lambda E$. Let $\tilde{E} = \frac{\Lambda^{1/2}E}{|\Lambda^{1/2}E|}$. Consequently, $U = X/|X| = \Gamma\tilde{E}$ is a function of E alone. Note that in the circularly symmetric case where $\Lambda = \lambda\mathbb{I}$, we now have $|X|$ independent of U, and the above conditions are trivially satisfied. For general Λ, note that

$$\mathbb{E}UU^T = \Gamma\mathbb{E}\tilde{E}\tilde{E}^T\Gamma^T,$$

and it can be easily shown that $\mathbb{E}\tilde{E}\tilde{E}^T$ is a diagonal matrix. Thus, $\Sigma_X = \Gamma\Lambda\Gamma^T$ and Σ_U have the same eigenvectors in this case. This proof is reminiscent of the arguments used in Taskinen et al. (2012).

Under the assumptions of the second part, that is, $\mathbb{E}U = 0 \in \mathbb{R}^p$, $\mathbb{E}U|X| = 0 \in \mathbb{R}^p$ and $\mathbb{E}|X|^2UU^T = \mathbb{E}|X|^2\mathbb{E}UU^T \in \mathbb{R}^{p\times p}$, U and $|X|$ are uncorrelated, as is $|X|^2$ and UU^T. This immediately implies that $\mathbb{E}X = \mathbb{E}U|X| = 0$, thus we have $\Sigma_X = \mathbb{E}XX^T$ and $\Sigma_U = UU^T$. We also easily have

$$\begin{aligned}\Sigma_X &= \mathbb{E}XX^T \\ &= \mathbb{E}|X|^2UU^T \\ &= \mathbb{E}|X|^2\mathbb{E}UU^T.\end{aligned}$$

Thus, the eigenvectors of Σ_X and Σ_U are identical, since $\mathbb{E}|X|^2$ is a scalar.

Now, note that

$$\Sigma_U = \frac{\Sigma_X}{\mathbb{E}|X|^2} = \frac{\Sigma_X}{\mathbb{E}[\text{Trace } XX^T]} = \frac{\Sigma_X}{\text{Trace } \Sigma_X} = \frac{\Sigma_X}{\sum_{i=1}^p \lambda_i}.$$ □

The rest of this section is on the estimation of the high dimensional principal components. We present the results only for the first principal component. Our development here closely follows that of Wang et al. (2016), and we essentially make use of their theoretical machinery and algorithm for the rest of this chapter. The results below are primarily designed to show that the technical conditions of Wang et al. (2016) hold for our case, and the eigenvector estimation algorithm they established also works for us. We omit many algebraic details, since they are similar to those of Wang et al. (2016).

Our first result is to show that U has a sub-Gaussian distribution. This is immediate, since U is bounded. We have multiple proofs of this result with sharp bounds on the constant σ^2, but present the simplest one here for clarity. Ensuring that U

has a sub-Gaussian distribution facilitates the use of various known concentration inequality and other probabilistic results.

Lemma 6.1

$U \in$ Sub-Gaussian(2). □

Proof 6.3 (Proof of Lemma 6.1)

We recall the definition of Sub-Gaussian distributions

$$X \in \text{Sub-Gaussian}(\sigma^2) \text{ if } \forall u \in \mathbb{R}^p,\ \mathbb{E}[e^{u^T X}] \leq e^{\frac{\sigma^2 |u|^2}{2}}$$

Since $|U| = 1$, we have for $|u| \geq 1$

$$\begin{aligned} u^T U \leq |u|^2 &\text{(C-S inequality) for any } u \text{ on } \mathbb{R}^p \\ \Rightarrow e^{u^T U} \leq e^{|u|^2} &\text{ for any } u \text{ on } \mathbb{R}^p \\ \Rightarrow U \in &\text{ Sub-Gaussian(2).} \end{aligned}$$

The case for $|u| < 1$ is more delicate, but can be handled with some routine algebra. We omit the details here.

We now recall some definitions from Wang et al. (2016) for developing our next set of results. Since our framework is high dimensional, we need structural assumptions on the nature of the eigenvectors of Σ_U (or Σ_X), and the most common and convenient assumption here is one of sparsity. We define the sparse unit ball in p-dimensions having at most k non-zero entries as follows:

$$B_0(k) = \Big\{ x \in \mathbb{R}^p : |x| = 1, \sum_{j=1}^{p} \mathcal{I}_{\{x_i \neq 0\} \leq k} \Big\}.$$

Based on this and sample size n, for any $j \in \{1, \ldots, p\}$ and $C > 0$, a probability measure P is said the satisfy the *Restricted Covariance Condition* (RCC) with parameters

p, n, j and C, and written as $P \in RCC_p(n, j, C)$ if

$$P\{\sup_{u \in B_0(l)} |\hat{V}(u) - V(u)| \geq C \max(\sqrt{\frac{j \log(p/\delta)}{n}}, \frac{j \log(p/\delta)}{n})\} \leq \delta$$

for all $\delta > 0$, where $V(u) = \mathbb{E}u^T \Sigma u$, $\hat{V}(u) = \mathbb{E}u^T \hat{\Sigma} u$ and $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^{n} U_i U_i^T$ for $U_1, U_2, \ldots, U_n \overset{iid}{\sim} P$. We also define

$$RCC_p(C) = \bigcap_{l=1}^{p} \bigcap_{n=1}^{\infty} RCC_p(n, l, C).$$

Suppose, associated with a generic distribution $\mathbb{P}$ on $\mathbb{R}^p$, is the variance matrix Σ with the j-th eigenvalue and eigenvector respectively being λ_j and $\gamma_{\square,j}$, $j = 1, \ldots, p$. The results for the rest of this section are valid for the following class of probability measures. For $\theta > 0$, define

$$\mathcal{P}_p(n, k, \theta) = \Big\{\mathbb{P} \in RCC_p(n, 2, 1) \cap RCC_p(n, 2k, 1) : \gamma_{\square,1} \in B_0(k), \lambda_1 - \lambda_2 \geq \theta\Big\}.$$

Our next result states that U, after suitable scaling, has a distribution that satisfies the restricted covariance condition with appropriate selection of constants.

Lemma 6.2

For the random variable $U \in \mathbb{R}^p$ with $p \geq 2$, assume that $\gamma_{\square,1} \in B_0(k)$ and $\theta = \big(\lambda_1(\Sigma_X) - \lambda_2(\Sigma_X)\big) > 0$. Then $Z = \frac{U}{22} \in \mathcal{P}_p(n, k, \frac{\theta}{22\text{Trace}(\Sigma_X)})$. □

Proof 6.4 (Proof of Lemma 6.2)

Recall from Proposition 1 of Wang et al. (2016) that for every $\sigma > 0$,

$$\text{Sub-Gaussian}(\sigma^2) \subseteq RCC_p\big(16\sigma^2(1 + \frac{9}{\log(p)})\big).$$

Therefore using Lemma 6.1, we have that

$$\begin{aligned} & U = \frac{X}{|X|} \in RCC_p\big(32\,(1+\frac{9}{\log(p)})\big), \\ \Rightarrow & \frac{U}{\sqrt{32\,(1+\frac{9}{\log(p)})}} \in RCC_p(1), \\ \Rightarrow & \frac{U}{22} \in RCC_p(1), \\ \Rightarrow & \frac{U}{22} \in \mathcal{P}_p(n,k,\frac{\theta}{22\text{Trace}(\Sigma_X)}). \end{aligned}$$

□

For a symmetric matrix $A \in \mathbb{R}^{p\times p}$, we define

$$\hat{\gamma}^k_{\square,max}(A) = \text{sargmax}_{u\in B_0(k)} u^T A u$$

to be the k-sparse maximum eigenvector of A, where *sargmax* denotes the smallest element of the argmax in the lexicographic ordering. We use A as an argument of $\hat{\gamma}^k_{\square,max}(\cdot)$ here to distinguish between the estimated eigenvectors from various matrices. Also, between 2 unit vectors u and v , we define the loss function

$$L(u,v) = \big(1-(u^T v)^2\big)^{1/2}.$$

Note that $\hat{\gamma}^k_{\square,max}(A)$ are all identical for $A = \hat{\Sigma}_U, \hat{\Sigma}_X, \hat{\Sigma}_Z$. Our next result is the main consistency and probabilistic guarantee result on the sample version of the k-sparse maximum eigen-vector. This result ensures in particular that under suitable conditions with $\log(p) = o(n)$, the sample k-sparse maximum eigenvector is consistent for the population maximum eigenvector.

Theorem 6.3

For $2k\,\log(p) \le n$, the k-sparse empirical maximum eigen-vector, $\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_U)$ satisfies

$$\mathbb{E}L\big(\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_U), \gamma_{\square,1}(\Sigma_U)\big) \le 44\sqrt{2}(1+\frac{1}{\log(p)})\sqrt{\frac{k\,\log(p)}{n\theta^2}}\ \text{Trace}(\Sigma_X).$$

□

Proof 6.5 (Proof of Theorem 6.3)

We apply Theorem 2 of Wang et al. (2016) on Z and note that for $2k\,\log(p) \leq n$, the k-sparse empirical maximum eigen-vector, $\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_Z)$ satisfies

$$\mathbb{E}L\big(\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_Z), \gamma_{\square,1}(\Sigma_Z)\big) \leq 2\sqrt{44}(1 + \frac{1}{\log(p)})\sqrt{\frac{k\,\log(p)}{n\theta^2}}\,\text{Trace}(\Sigma_X).$$

Proof follows by noting that

$$\mathbb{E}L\big(\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_Z), \gamma_{\square,1}(\Sigma_Z)\big) = \mathbb{E}L\big(\hat{\gamma}^k_{\square,max}(\hat{\Sigma}_U), \gamma_{\square,1}(\Sigma_U)\big).$$ □

Let $U_i = \frac{X_i}{|X_i|}$ for $i = 1, 2, \ldots, n$ where $X_i \overset{iid}{\sim} \mathbb{P}$, and we denote $\hat{\gamma}^{SDP}_{\square,1}(U)$ for the output of the SDP algorithm of Wang et al. (2016). with input $U = (U_1, \ldots, U_n)^T$, $\lambda = 4\sqrt{\frac{\log(p)}{n}}$ and $\epsilon = \frac{\log(p)}{4n}$.

While Theorem 6.3 provided a general probabilistic guarantee on the error of the sample k-sparse maximum eigenvector, we need a similar result for the sparse maximum eigenvector that is obtained using the SDP algorithm. Note that the SDP algorithm allows for computation of a sparse maximum eigenvector in real time, and is thus both practical and is of theoretical relevance. The following result establishes a probabilistic guarantee and consistency for this version of the sparse empirical maximum eigenvector.

Theorem 6.4

If $4\,\log(p) \leq n \leq k^2p^2\theta^{-2}\log(p)$, then

$$\mathbb{E}L\big(\hat{\gamma}^{SDP}_{\square,1}(U), \gamma_{\square,1}(\Sigma_U)\big) \leq (352\sqrt{2} + 44)\,\text{Trace}\,\Sigma_X\sqrt{\frac{k^2\,\log(p)}{n\theta^2}}.$$ □

Proof 6.6 (Proof of Theorem 6.4)

We apply Theorem 5 of Wang et al. (2016) on Z and note that

$$\mathbb{E}L\big(\hat{\gamma}^{SDP}_{\square,1}(Z), \gamma_{\square,1}(\Sigma_Z)\big) \leq 22(16\sqrt{2} + 2)\sqrt{\frac{k^2\,\log(p)}{n\theta^2}}\text{Trace}(\Sigma_X).$$ □

for $Z = (Z_1, \ldots, Z_n)^T$ where $Z_i = \frac{U_i}{22}$. The result is immediate.

In general, $\hat{\gamma}_{\square,1}^{SDP}(U)$ is not a sparse estimator. However, it turns out that a k-sparse version of $\hat{\gamma}_{\square,1}^{SDP}(U)$, that is, some $\hat{\gamma}_{\square,1}^{SDP,k}(U) \in B_0(k)$, may be obtained by setting all but the top k coordinates of $\hat{\gamma}_{\square,1}^{SDP}(U)$ in absolute value to zero and renormalizing the vector. In particular, $\hat{v}_0^{SDP}$ is computable in polynomial time and under the same condition as in Theorem 6.4.

6.5 Some simulation results

In this section, we present a simulation exercise to illustrate the performance of the proposed record linkage methodology, and also to illustrate some practical implementation steps.

To generate data, we followed the framework laid out at the start of Section Section 6.4. That is, we generated a set of independent bivariate Gaussian random variables with common correlation ρ, and several independent univariate standard Normal random variables, and used these to populate the two data matrices. We tested various choices of sample sizes, dimensions, correlation ρ. For brevity, we report the case where the two matrix datasets that we use are of dimensions $n_1 = 60, p_1 = 100$, and $n_2 = 70, p_2 = 120$ of independent rows each, and $\rho = 0.8$. The first $n_0 = 50$ entries of these two matrices are linked to each other. The rest (10 for the first matrix, 20 for the second matrix) are not linked. We use the first $n = 20$ observations for training in the version of the algorithm where a training set is used, thus leaving the last 30 linked data points for testing. We also demonstrate the performance of our method when no training dataset is available. We fix $p_0 = 10$ for this exercise, and repeated the entire simulation 100 times.

As practical steps, we found that using a sparse version of the estimated eigenvalues, as proposed in Theorem 6.4, considerably improves performance, owing to reduction of the effect of noise terms in the eventual linkage. Also, the estimated

principal components for the two datasets may not have the same orientation and may not appear in the same order. Hence, when needed, a principled permutation and sign reversal of the estimated eigenvectors of the second dataset is done to improve linkage accuracy. While in theory the estimators of the eigenvalues are not required for the linkage steps, using those as weights improves linkage.

Over 100 replications of the simulation experiment, the correct linkage established on the test set by our proposed method was about 43.5% times, with a standard error of about 5.37%. When no training sample is used and instead a threshold for the correlation estimated from the data, the correct linkage percentage increases to about 56%, however, the standard error also increases to about 14.8%. The estimated threshold for correlation was at a lower value than the case with training data: a pattern that we noticed in multiple simulations, which we will study further later. The linkage accuracy may seem low, however, we need to remember that this is a ***unsupervised*** framework, involving high dimensional datasets with no common variables, and minimal or no training data.

6.6 Conclusions and future work

Record linkage is in-general a difficult exercise. Bayesian models are often too complex for practical purposes and some Bayesian formulations fail to accommodate non-categorical variables. Supervised classification methods assume the existence of large, accurate sets of training data, which are often difficult and/or expensive to obtain. Also, in those methodologies, it is difficult to guarantee maximum one-to-one assignment constrain of bipartite record linkage. While some theoretical modifications have been developed to ensure an one-to-one assignment(Sadinle, 2017), typically some subsequent post-processing step is required to solve these inconsistencies.

In view of these difficulties, in this chapter we propose a completely new approach

towards record linkage. We do not require a common set of variables between the two datasets, we do not require a training set, and the dimensionality and sample sizes can both be large. Naturally, our methodology extends to cases where a common set of variables exist, we will elaborate on this in a future work. If a training set exists, we can make use of it, as illustrated in this chapter. We have presented the case for the bipartite record linkage, but our model conceptually extends to other cases as well.

Some of the technical assumptions of this chapter, like the two covariances matrices having the same set of leading eigenvalues, or the number leading eigenvalues p_0 being known, or the latent random variables that link the two datasets having the same correlation ρ, can be addressed with some additional work and methodological developments. The assumption of multivariate normality of the data is not critical: our proposal only depends on robust, high dimensional principal components, and these are available for data from many distributions with both discrete and continuous components. The assumption of sparsity in the leading eigenvectors is owing to the fact that for high dimensional modeling, some structural assumptions are needed since the sample size is not adequate to estimate all relevant unknown parameters. In any case, there are considerable challenges to estimating high dimensional principal components, see Paul (2007).

The robust, high dimensional principal component we use is built on the work by Wang et al. (2016). The credit for both the theoretical framework and the algorithm goes to that work primarily. Our setup differs from Wang et al. (2016) in the detail that for robustness purposes we transform each observation to be on the unit sphere. One future work for us is to establish the theoretical results under weaker assumptions than Wang et al. (2016), or to show better theoretical properties.

We have used a simple method for linking observations in this chapter, using correlations. A correlation-based linkage is not critical to our primary methodological

steps. More complex and realistic measures of linkage will be studied in the future. The case where no training data is present needs further investigation, which will also be part of our future work. In absolute terms, our simulation results are not excellent; however, we do not have a baseline for comparison since most other papers on record linkage do not use as general a framework as ours with (*i*) high dimensional data, (*ii*) no common set of features, and (*iii*) possibly no training set. Our framework may be termed *unsupervised learning of record linkage*, and in the unsupervised learning framework, our numeric results are perhaps acceptable. However, considerable fine tuning and experimentation with the algorithm for record linkage is needed. We have ensured that our high-dimensional, robust and potentially sparse principal component estimator is highly accurate, and some of our studies (not reported here) suggest that using a small number of common features dramatically increases linkage accuracy. A part of our future work is on including nominal and categorical variables for linkages in our framework, which will make our proposed approach more aligned with traditional record linkage techniques. In this context, we will also investigate how much additional gain results from using PCA in addition to available matching fields, compared to the traditional Felligi-Sunter method.

An important topic to consider in future from this chapter is on statistical inference based on linked datasets. This is a non-trivial task, since the datasets are used multiple times in the process of linking, estimation of various quantities of interest, and then inference. The article Han and Lahiri (2019) provides review of the current state of the art in this direction of work. Some alternatives to fully Bayesian methods, for example regression analysis using linked data Lahiri and Larsen (2005); Scheuren and Winkler (1997, 1993), have both computational efficiency and analytical tractability, which may make them attractive practical choices for applications. Comparisons with such alternatives is an additional future work.

Chapter 7

Future Work

We list our future work as follows.

- We would like to work on the "inverse problem" of chapter 2. I.e. given a generic function h as in equation (1.2) $d_h(x,y) \coloneqq \underset{\gamma}{inf} \int_0^1 h(f(\gamma_t))|\gamma_t'|dt$, can we find a sample based distance that converge to this distance as $n \to \infty$? If yes, we would like to derive the convergence rate and properties of this sample based distance.

- Given a data-set, a clustering objective and a particular form of g-distance, we would like to formulate and derive a theoretical framework for optimally choosing the coefficients/powers of the g-distance and perform numerical experiments comparing performance of different g.

- In many applications the underlying network/graph is directed. We would like to generalize the notion of g distance (i.e. capturing geometry and intensity information present in the network simultaneously) for directed data and develop relevant statistical methodologies.

- We would like to study the optimality of SC using density sensitive distances when data is generated from a known mixture distribution, for example from a Gaussian mixture model. In this scenario, distribution based clustering methods

such as EM algorithm should outperform SC. However, it will be interesting to derive the performance guarantee condition for SC and compare it with existing performance guarantee conditions for SC using Euclidean distance and EM algorithm.

- We would like to do extensive numerical work to understand the practical implications of the theoretical results presented in this thesis. In particular, we would like to perform numerical experiments and simulation studies comparing the performance of the algorithms developed in chapter 5 on synthetic as well as real datasets. Also, we would like to do extensive numerical experiments to compare performance of different g distances and LLPD in varying sample size and density regime scenarios in relation to our work in chapter 2.

- We would like to extend the methodology proposed in chapter 6 to multiple datasets. We will also work on real data examples, which has not been possible owing to data access limitations. It will be of interest to compare our unsupervised record linkage approach with more traditional record linking algorithms.

References

Arias-Castro, E. (2011). Clustering based on pairwise distances when the data is of mixed dimensions. *IEEE Transactions on Information Theory*, 57(3):1692–1706.

Arias-Castro, E., Lerman, G., and Zhang, T. (2017). Spectral clustering based on local pca. *The Journal of Machine Learning Research*, 18(1):253–309.

Banyaga, A. and Hurtubise, D. (2013). *Lectures on Morse homology*, volume 29. Springer Science & Business Media.

Belkin, M. and Niyogi, P. (2003). Laplacian eigenmaps for dimensionality reduction and data representation. *Neural computation*, 15(6):1373–1396.

Bijral, A. S., Ratliff, N., and Srebro, N. (2012). Semi-supervised learning with density based distances. *arXiv preprint arXiv:1202.3702*.

Bose, P., Devroye, L., Evans, W., and Kirkpatrick, D. (2002). On the spanning ratio of gabriel graphs and β-skeletons. In *Latin American Symposium on Theoretical Informatics*, pages 479–493. Springer.

Cardinal, J., Collette, S., and Langerman, S. (2009). Empty region graphs. *Computational geometry*, 42(3):183–195.

Chang, H. and Yeung, D.-Y. (2008). Robust path-based spectral clustering. *Pattern Recognition*, 41(1):191–203.

Chaudhuri, K. and Dasgupta, S. (2010). Rates of convergence for the cluster tree. In *Advances in neural information processing systems*, pages 343–351.

Chaudhuri, K., Dasgupta, S., Kpotufe, S., and Von Luxburg, U. (2014). Consistent procedures for cluster tree estimation and pruning. *IEEE Transactions on Information Theory*, 60(12):7900–7912.

Chazal, F., Guibas, L. J., Oudot, S. Y., and Skraba, P. (2013). Persistence-based clustering in riemannian manifolds. *Journal of the ACM (JACM)*, 60(6):41.

Chen, Y.-C., Genovese, C. R., Wasserman, L., et al. (2016). A comprehensive approach to mode clustering. *Electronic Journal of Statistics*, 10(1):210–241.

Chen, Y.-C., Genovese, C. R., Wasserman, L., et al. (2017). Statistical inference using the morse-smale complex. *Electronic Journal of Statistics*, 11(1):1390–1433.

Cheng, Y. (1995). Mean shift, mode seeking, and clustering. *IEEE transactions on pattern analysis and machine intelligence*, 17(8):790–799.

Christen, P. (2011). a survey of indexing techniques for scalable record linkage and deduplication. *Ieee Transactions on Knowledge and Data Engineering*, 24(9):1537–1555.

Correa, C. D. and Lindstrom, P. (2012). Locally-scaled spectral clustering using empty region graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1330–1338. ACM.

Devillers, O. and Duménil, C. (2021). Stochastic analysis of empty-region graphs. In *CCCG 2021-33rd Canadian Conference on Computational Geometry*.

Donoho, D. L. and Grimes, C. (2003). Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596.

Drachsler, H., Bogers, T., Vuorikari, R., Verbert, K., Duval, E., Manouselis, N., Beham, G., Lindstaedt, S., Stern, H., Friedrich, M., and Others (2010). Issues and considerations regarding sharable data sets for recommender systems in technology enhanced learning. *Procedia Computer Science*, 1(2):2849–2858.

Enamorado, T., Fifield, B., and Imai, K. (2018). Using a probabilistic model to assist merging of large-scale administrative records. *Available at Ssrn 3214172*.

Fan, J., Wang, W., and Zhong, Y. (2018). An l$infty$ eigenvector perturbation bound and its application to robust covariance estimation. *Journal of Machine Learning Research*, 18(207):1–42.

Fang, K.-t., Kotz, S., and Ng, K.-w. (1990). *Symmetric Multivariate and Related Distributions*. Crc Press.

Fatemi, B., Kazemi, S. M., and Poole, D. (2018). Record linkage to match customer names: A probabilistic approach. *Arxiv Preprint Arxiv:1806.10928*.

Federer, H. (1959). Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491.

Fefferman, C., Mitter, S., and Narayanan, H. (2016). Testing the manifold hypothesis. *Journal of the American Mathematical Society*, 29(4):983–1049.

Fellegi, I. P. and Sunter, A. B. (1969). a theory for record linkage. *Journal of the American Statistical Association*, 64(328):1183–1210.

Fischer, B., Roth, V., and Buhmann, J. M. (2004). Clustering with the connectivity kernel. In *Advances in neural information processing systems*, pages 89–96.

Genovese, C. R., Perone-Pacifico, M., Verdinelli, I., and Wasserman, L. (2016). Non-parametric inference for density modes. *Journal of the Royal Statistical Society: Series B: Statistical Methodology*, pages 99–126.

Genovese, C. R., Perone-Pacifico, M., Verdinelli, I., Wasserman, L., et al. (2014). Nonparametric ridge estimation. *Annals of Statistics*, 42(4):1511–1545.

Gerber, S., Rübel, O., Bremer, P.-T., Pascucci, V., and Whitaker, R. T. (2013). Morse–smale regression. *Journal of Computational and Graphical Statistics*, 22(1):193–214.

Goh, A. and Vidal, R. (2008). Clustering and dimensionality reduction on riemannian manifolds. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–7. IEEE.

González-Barrios, J. M. and Quiroz, A. J. (2003). A clustering procedure based on the comparison between the k nearest neighbors graph and the minimal spanning tree. *Statistics & Probability Letters*, 62(1):23–34.

Groisman, P., Jonckheere, M., and Sapienza, F. (2018). Nonhomogeneous euclidean first-passage percolation and distance learning. *arXiv preprint arXiv:1810.09398*.

Han, Y. and Lahiri, P. (2019). Statistical analysis with linked data. *International Statistical Review*, 87:S139–s157.

Hartigan, J. A. (1981). Consistency of single linkage for high-density clusters. *Journal of the American Statistical Association*, 76(374):388–394.

Herman, R. L. (2008). Derivation of the geodesic equation and defining the christoffel symbols. *University of North Carolina Wilmington*.

Herzog, T. N., Scheuren, F. J., and Winkler, W. E. (2007). *Data Quality and Record Linkage Techniques*. Springer Science & Business Media.

Howard, C. D. and Newman, C. M. (2001). Geodesics and spanning trees for euclidean first-passage percolation. *Annals of Probability*, pages 577–623.

Hwang, S. J. (2012). Geometric representations of high dimensional random data.

Hwang, S. J., Damelin, S. B., Hero III, A. O., et al. (2016). Shortest path through random points. *The Annals of Applied Probability*, 26(5):2791–2823.

Ichino, M. and Sklansky, J. (1985). The relative neighborhood graph for mixed feature variables. *Pattern recognition*, 18(2):161–167.

Irons, M. L. (2005). The curvature and geodesics of the torus. *Retrieved January*, 20:2008.

Jantzen, R. T. (2012). Geodesics on the torus and other surfaces of revolution clarified using undergraduate physics tricks with bonus: nonrelativistic and relativistic kepler problems. *arXiv preprint arXiv:1212.6206*.

Jaro, M. a. (1989). Advances in record-linkage methodology as applied to matching the 1985 census of tampa, florida. *Journal of the American Statistical Association*, 84(406):414–420.

Katz, S. and Tal, A. (2017). On visibility and empty-region graphs. *Computers & Graphics*, 66:45–52.

Lahiri, P. and Larsen, M. D. (2005). Regression analysis with linked data. *Journal of the American Statistical Association*, 100(469):222–230.

Lan, Y.-D., Deng, H., and Chen, T. (2011). A new method of distance measure for graph-based semi-supervised learning. In *2011 International Conference on Machine Learning and Cybernetics*, volume 4, pages 1444–1448. IEEE.

Larsen, M. D. and Rubin, D. B. (2001). Iterative automated record linkage using mixture models. *Journal of the American Statistical Association*, 96(453):32–41.

Li, J., Dou, Z., Zhu, Y., Zuo, X., and Wen, J.-r. (2020). Deep cross-platform product matching in e-commerce. *Information Retrieval Journal*, 23(2):136–158.

Liseo, B. and Tancredi, A. (2013). Some advances on Bayesian record linkage and inference for linked data. *Url Http://Www. Ine. Es/E/Essnetdi_ws2011/Ppts/Liseo_tancredi. Pdf*.

Little, A., Maggioni, M., and Murphy, J. M. (2017). Path-based spectral clustering: Guarantees, robustness to outliers, and fast algorithms. *arXiv preprint arXiv:1712.06206*.

Little, A., McKenzie, D., and Murphy, J. (2020). Balancing geometry and density: Path distances on high-dimensional data. *arXiv preprint arXiv:2012.09385*.

Maggioni, M. and Murphy, J. M. (2018). Learning by unsupervised nonlinear diffusion. *arXiv preprint arXiv:1810.06702*.

Maier, M., Hein, M., and Von Luxburg, U. (2009). Optimal construction of k-nearest-neighbor graphs for identifying noisy clusters. *Theoretical Computer Science*, 410(19):1749–1764.

Martins, B. (2011). a supervised machine learning approach for duplicate detection over gazetteer records. In *International Conference on Geospatial Sematics*, pages 34–51. Springer.

Mckenzie, D. and Damelin, S. (2019). Power weighted shortest paths for clustering euclidean data. *Foundations of Data Science*, 1(3):307.

Narayanan, H. and Mitter, S. (2010). Sample complexity of testing the manifold hypothesis. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems-Volume 2*, pages 1786–1794.

Newcombe, H. B. and Kennedy, J. M. (1962). Record linkage: Making maximum use of the discriminating power of identifying information. *Communications of the Acm*, 5(11):563–566.

Ng, A. Y., Jordan, M. I., and Weiss, Y. (2002). On spectral clustering: Analysis and an algorithm. In *Advances in neural information processing systems*, pages 849–856.

O'neill, B. (1983). *Semi-Riemannian geometry with applications to relativity*, volume 103. Academic press.

Orlitsky, A. et al. (2005). Estimating and computing density based distance metrics. In *Proceedings of the 22nd international conference on Machine learning*, pages 760–767. ACM.

Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17(4):1617–1642.

Penrose, M. D. (1997). The longest edge of the random minimal spanning tree. *The annals of applied probability*, pages 340–361.

Penrose, M. D. (1998). Extremes for the minimal spanning tree on normally distributed points. *Advances in Applied Probability*, 30(3):628–639.

Penrose, M. D. (1999). A strong law for the largest nearest-neighbour link between random points. *Journal of the london mathematical society*, 60(3):951–960.

Penrose, M. D. et al. (1999). A strong law for the longest edge of the minimal spanning tree. *The Annals of Probability*, 27(1):246–260.

Rong, S., Niu, X., Xiang, E. W., Wang, H., Yang, Q., and Yu, Y. (2012). a machine learning approach for instance matching based on similarity metrics. In *International Semantic Web Conference*, pages 460–475. Springer.

Sadinle, M. (2017). Bayesian estimation of bipartite matchings for record linkage. *Journal of the American Statistical Association*, 112(518):600–612.

Sadinle, M. and Fienberg, S. E. (2013). a generalized fellegi–sunter framework for multiple record linkage with application to homicide record systems. *Journal of the American Statistical Association*, 108(502):385–397.

Salas, J. (2019). Sanitizing and measuring privacy of large sparse datasets for recommender systems. *Journal of Ambient Intelligence and Humanized Computing*, pages 1–12.

Saul, L. K. and Roweis, S. T. (2003). Think globally, fit locally: unsupervised learning of low dimensional manifolds. *Departmental Papers (CIS)*, page 12.

Scheuren, F. and Winkler, W. E. (1993). Regression analysis of data files that are computer matched. *Survey Methodology*, 19:39 – 58.

Scheuren, F. and Winkler, W. E. (1997). Regression analysis of data files that are computer matched-ii. *Survey Methodology*, 23:157 –165.

Schubert, E., Sander, J., Ester, M., Kriegel, H. P., and Xu, X. (2017). Dbscan revisited, revisited: why and how you should (still) use dbscan. *ACM Transactions on Database Systems (TODS)*, 42(3):1–21.

Shabtai, A., Rokach, L., and Elovici, Y. (2013). Occt: A one-class clustering tree for implementing one-to-many data linkage. *Ieee Transactions on Knowledge and Data Engineering*, 26(3):682–697.

Singla, P. and Domingos, P. (2006). Entity resolution with markov logic. In *Sixth International Conference on Data Mining (Icdm'06)*, pages 572–582. Ieee.

Slokom, M. (2018). Comparing recommender systems using synthetic data. In *Proceedings of the 12th Acm Conference on Recommender Systems*, pages 548–552.

Steorts, R. C., Hall, R., and Fienberg, S. E. (2016). a bayesian approach to graphical record linkage and deduplication. *Journal of the American Statistical Association*, 111(516):1660–1672.

Steorts, R. C. and Others (2015). Entity resolution with empirically motivated priors. *Bayesian Analysis*, 10(4):849–875.

Steorts, R. C., Ventura, S. L., Sadinle, M., and Fienberg, S. E. (2014). a comparison of blocking methods for record linkage. In *International Conference on Privacy in Statistical Databases*, pages 253–268. Springer.

Supowit, K. J. (1983). The relative neighborhood graph, with an application to minimum spanning trees. *Journal of the ACM (JACM)*, 30(3):428–448.

Talagrand, M. (1994). Concentration of measure and isoperimetric inequalities in product spaces.

Tao, X., Wang, R., Chang, R., Li, C., Liu, R., and Zou, J. (2019). Spectral clustering algorithm using density-sensitive distance measure with global and local consistencies. *Knowledge-Based Systems*, 170:26–42.

Taskinen, S., Koch, I., and Oja, H. (2012). Robustifying principal component analysis with spatial sign vectors. *Statistics & Probability Letters*, 82(4):765 – 774.

Tenenbaum, J. B., De Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500):2319–2323.

Torvik, V. I. and Smalheiser, N. R. (2009). Author name disambiguation in medline. *Acm Transactions on Knowledge Discovery from Data (Tkdd)*, 3(3):1–29.

Tran, T. N., Wehrens, R., and Buydens, L. M. (2006). Knn-kernel density-based clustering for high-dimensional multivariate data. *Computational Statistics & Data Analysis*, 51(2):513–525.

Treeratpituk, P. and Giles, C. L. (2009). Disambiguating authors in academic publications using random forests. In *Proceedings of the 9th Acm/Ieee-cs Joint Conference on Digital Libraries*, pages 39–48.

Trillos, N. G., He, P., and Li, C. (2021). Large sample spectral analysis of graph-based multi-manifold clustering. *arXiv preprint arXiv:2107.13610*.

Vedaldi, A. and Soatto, S. (2008). Quick shift and kernel methods for mode seeking. In *European conference on computer vision*, pages 705–718. Springer.

Ventura, S. L., Nugent, R., and Fuchs, E. R. (2015). Seeing the non-stars:(some) sources of bias in past disambiguation approaches and a new public tool leveraging labeled records. *Research Policy*, 44(9):1672–1701.

Vincent, P. and Bengio, Y. (2003). Density-sensitive metrics and kernels. In *Snowbird Learning Workshop*.

Von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416.

Von Luxburg, U., Belkin, M., and Bousquet, O. (2008). Consistency of spectral clustering. *The Annals of Statistics*, pages 555–586.

Wang, L., Bo, L., and Jiao, L. (2006). A modified k-means clustering with a density-sensitive distance metric. In Wang, G.-Y., Peters, J. F., Skowron, A., and Yao, Y., editors, *Rough Sets and Knowledge Technology*, pages 544–551, Berlin, Heidelberg. Springer Berlin Heidelberg.

Wang, T., Berthet, Q., and Samworth, R. J. (2016). Statistical and computational trade-offs in estimation of sparse principal components. *the Annals of Statistics*, 44(5):1896–1930.

Wang, W., Li, X.-Y., Moaveninejad, K., Wang, Y., and Song, W.-Z. (2003). The spanning ratio of β-skeletons. In *Proceedings of the Canadian Conference on Computational Geometry (CCCG)*.

Wang, Y., Jiang, Y., Wu, Y., and Zhou, Z.-H. (2010). Multi-manifold clustering. In *Pacific Rim International Conference on Artificial Intelligence*, pages 280–291. Springer.

Wasserman, L. (2018). Topological data analysis. *Annual Review of Statistics and Its Application*, 5:501–532.

Zhang, W. and King, I. (2002). A study of the relationship between support vector machine and gabriel graph. In *Proceedings of the 2002 International Joint Conference on Neural Networks. IJCNN'02 (Cat. No. 02CH37290)*, volume 1, pages 239–244. IEEE.

Zhao, P. and Lai, L. (2020). Analysis of knn density estimation. *arXiv preprint arXiv:2010.00438*.

Zhu, J., Zhang, S., Singh, L., Yang, G. H., and Sherr, M. (2016). Generating risk reduction recommendations to decrease vulnerability of public online profiles. In *2016 Ieee/Acm International Conference on Advances in Social Networks Analysis and Mining (Asonam)*, pages 411–416. Ieee.