Exploring the Use of Spectral Seriation to Uncover Dynamics in Embryonic Development: A Geometric and Probabilistic Approach

by

Roomina Zendehboodi

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

MASTER OF SCIENCE

in

The Faculty of Graduate and Postdoctoral Studies

(Mathematics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

August 2023

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Exploring the Use of Spectral Seriation to Uncover Dynamics in Embryonic Development: A Geometric and Probabilistic Approach

submitted by **Roomina Zendehboodi** in partial fulfillment of the requirements for the degree

of **Master of Science** in **Mathematics**.

Examining Committee:

Geoffrey Schiebinger, Assistant Professor, Mathematics, University of British Columbia
Supervisor

Young-Heon Kim, Professor, Mathematics, University of British Columbia
Supervisory Committee Member

Abstract

Understanding the dynamics of embryonic development is crucial to finding treatments for conditions such as aging and cancer. The development of an embryo can be represented as a curve in the Wasserstein space, and to construct this curve, static snapshots of gene expression profiles are obtained at n selected time points. Since the measurement techniques for obtaining these snapshots are destructive, we have to infer the developmental trajectory using a series of static snapshots of gene expression profiles taken at different time points $t_1, t_2, \ldots, t_n$. To obtain these snapshots, multiple embryos are allowed to develop until each of the desired time points is reached, and the gene expression profile is then captured. However, to reconstruct the curve we need to know which embryo had reached which developmental stage; this information is lost during the measurements. To overcome this, a pairwise similarity function between profiles can be defined, and the profiles can be arranged so that the more similar they are, the closer they are placed together. This is part of a larger class of problems known as the "seriation" problem. In this thesis, the feasibility of using the "spectral seriation" method proposed by Atkins et al. [2] is investigated to recover the order of the profiles based on their similarity, which enables the construction of the curve.

The gene expression profile of an embryo can be seen as a probability measure on a compact subset of $\mathbb{R}^d$. Although the exact measures are unknown, they can be approximated empirically using m samples. In this thesis, we demonstrate that, under reasonable assumptions and with sufficient time points and samples per time point, the spectral seriation method can be effective in sequencing the data. Additionally, we provide tools to determine the number of time points and samples per time point needed to achieve a desired error bound. Furthermore, we investigate how the geometric properties of the curve representing the embryonic development can affect our ability to sequence the data.

Lay Summary

To gain a deeper understanding of embryonic development, it is important to establish a mathematical model that captures the various stages of development. By utilizing mathematical tools, we can better comprehend how an embryo develops. This thesis begins by presenting a mathematical model of embryonic development and investigates the process of reconstructing the curve that represents this development. To achieve this, we must first arrange the data that represent different stages of embryonic development based on their similarities. The thesis provides tools to determine the number of samples and data points required to recover the data's ordering and examines how the properties of the developmental curve may impact our ability to do so.

Preface

This thesis is an original, unpublished, and independent work of the author, Roomina Zendehboodi (RZ), under the supervision of Professor Geoffrey Schiebinger (GS) and in collaboration with Professor Young-Heon Kim (YK).

The presentation of the introductory content in chapter 1 is original and mainly based on [2] and [10]. Chapter 2 provides the mathematical background. Specifically, section 2.1 is based on the book [19] and 2.3 is a summary of the paper [2] and all the material in that section are due to Atkins et al. [2]. 2.2 is based on the book [16].

Chapter 3, 4, and 5 are original, unpublished, and independent work of the author, RZ, under the guidance of GS and YK.

The idea presented at the beginning of chapter 5, connecting the discrete setting to a continuous setting, and the function used in the proof of 5.4.7 are inspired by the paper [12], but the work that follows is original.

Table of Contents

List of Figures

Acknowledgments

I would like to express my sincere gratitude to my supervisor Professor Geoffrey Schiebinger for his unwavering support and patience throughout my Master's program. His optimism and continuous encouragement have kept me inspired to pursue my academic goals with passion and determination, and his expert guidance played a crucial role in keeping me on the right path toward success. I am also extremely grateful to Professor Young-Heon Kim, for being another source of insight and encouragement for me, and whose tireless efforts to guide me in the right direction and continuous support have been instrumental in shaping my academic career. I would also like to thank my family, for always supporting my interests. Especially my sister, for always being there for me and inexplicably putting up with me. In the end, I would like to thank all my friends back home and the friends I made at UBC, for making my academic journey an exciting one, and for keeping life worth living.

Chapter 1

Introduction

The trajectory of embryonic development can be represented as a curve in the Wasserstein space. In order to construct this curve, n time points are selected and the gene expression profiles at those time points are obtained, and each profile is approximated empirically with m samples. However, arranging these profiles in the correct order is crucial to construct the curve, and this information may not always be available. Nonetheless, we can define a pairwise similarity function between the profiles and try to arrange them so that the more similar two profiles are, the closer they are put together. Finding such arrangement belongs to a bigger class of problems, called the *seriation problem*. This thesis investigates the feasibility of using the *spectral seriation* method proposed by Atkins et al. [2] to recover the order of the profiles based on their similarity, thus enabling the construction of the curve. Specifically, we provide mathematical tools enabling us to determine how many time points and samples per time point are needed to achieve a particular error bound in this reconstruction. While in the past there have been attempts for solving the seriation problem for biological developments using heuristics, for example, [3] and [11], they have not provided any theoretical guarantees. Furthermore, to the best of the author's knowledge, there hasn't been any research done on analyzing the feasibility and consistency of the spectral seriation method proposed by 1998 with the added assumption that the data lie on a curve. In chapter 5, we explore how some of the geometric properties of the curve affect the feasibility of this method.

In the first section of this chapter, we outline the biological background and demonstrate how it can be translated into a mathematical problem. The second section discusses the seriation problem in general and how it relates to the problem we are addressing. In the third section, we provide a brief history of seriation and the original contributions of this thesis. In the final section, we provide an overview of this thesis.

1.1 Translating the Biological Question into a Mathematical Problem

Gaining a deeper understanding of embryonic development holds immense potential for advancing our knowledge of biology and medicine. The latest breakthroughs in precision measurement techniques, such as single-cell RNA sequencing (scRNA-seq) and single-cell ATAC-seq, are now empowering us to explore the complexities of biological systems like never before. By using these tools, we may be able to develop more effective treatments and cures for a wide range of

conditions such as cancer and aging.

Understanding the dynamic changes in gene expression that occur during development could immensely improve our understanding of biological systems. Unfortunately, current measurement technologies are destructive and cannot directly observe changes in expression over time. As a result, there has been a surge of research into developing computational methods for inferring trajectories from static snapshots of gene expression profiles. In order to use these methods for understanding disease and developing new therapies, it is important to establish clear guidelines for evaluating their accuracy and reliability.

This thesis explores the feasibility of using the spectral seriation method introduced in [2] to sequence embryos over time in order to gain insights into their development. To construct the developmental trajectory of an embryo, it would be ideal to take many snapshots of its gene expression profile over time. However, current measurement techniques are destructive, making it impossible to take continuous measurements from a single embryo. Instead, we can infer the developmental trajectory using a series of static snapshots of gene expression profiles taken at different time points $t_1, t_2, \ldots, t_n$. To obtain these snapshots, multiple embryos are allowed to develop until each of the desired time points is reached, and the gene expression profile is then captured.

The challenge arises at the end of this process, where it becomes difficult to establish a clear timeline of embryonic development based on the obtained gene expression profiles, as it is not possible to determine which embryo reached which developmental stage. However, we can still determine some information about the relative developmental stages of the embryos by comparing their gene expression profiles. Specifically, we can infer that two embryos must have reached developmental stages that are relatively close in time if their profiles are similar. In other words, we can use similarities in gene expression profiles to establish approximate developmental timelines.

A gene expression profile can be represented as a probability measure ν on a compact subset of $\mathbb{R}^d$, where d is approximately 20000. Although it is impossible to determine ν precisely, we can estimate it using m samples. Specifically, we approximate ν with $\hat{\nu}_m = \frac{1}{m}\sum_{k=1}^{m} \delta_{X_k}$, where X_k are random variables with distribution ν.

With the gene expression profiles now represented mathematically, the task of determining the similarity between them is simplified. Since each profile can be seen as a point in the Wasserstein space, we can utilize the Wasserstein-2 distance to establish the similarity of two gene expression profiles represented by ν and μ, which in this thesis is defined to be $e^{\frac{-W_2^2(\nu,\mu)}{\sigma^2}}$, where $W_2^2(\nu,\mu)$ is the Wasserstein-2 distance between μ and ν, and $\sigma > 0$ is arbitrary.

We have transitioned from the biological inquiry of "how can we establish the developmental trajectory of an embryo" to a purely mathematical problem of "how can we determine the order of n gene expression profiles using a pairwise similarity function". This mathematical inquiry falls into a broader category of problems known as the "seriation problem", which we will explore in the following section.

From now on, we refer to the gene expression profiles obtained at different time points as our data points, and their approximation by m samples as the empirical data points.

1.2 The seriation problem and its relation to the biological problem

The *seriation problem* is a class of problems that involves ordering objects based on a pairwise similarity function. This is a fundamental problem in various areas of science such as archaeology, biology, and computer science, where objects need to be ordered based on some similarity measure. In the context of this thesis, the seriation problem refers to the task of ordering gene expression profiles, representing different developmental stages of an embryo, based on their pairwise similarity.

To put this in mathematical terms, suppose that we are given a set of n elements we need to sequence. Also, suppose that we are given a similarity function $f(i, j)$, which represents the desire of elements i and j to be close. We need to find a permutation π that is consistent with this correlation in the sense that if $\pi(i) < \pi(j) < \pi(k)$ then $f(i, j) \geq f(i, k)$ and $f(j, k) \geq f(i, k)$. There may be an exponential number of such orderings, or there may be none.

Finding a permutation consistent with the correlation function is called the *seriation problem*. If such permutation exists, the problem is called well-posed. We will explore the seriation problem in much more detail in section 2.3, where we present a summary of the paper [2], "A Spectral Algorithm For Seriation And The Consecutive Ones Problem" by Atkins et al. [2].

1.3 The history of seriation and original contributions of this thesis

In this section, we explore the history of seriation and then compare this thesis to related works and explain the original contributions.

1.3.1 A brief history of seriation

The history of seriation can be traced back to the year 1899 when Flinders Petrie wrote the paper "Sequences in Prehistoric Remains" [14]. Many believe this was the first attempt to think about archeological sequencing mathematically, rather than relying on human judgment, which even if professional, can be flawed. In this paper, Petrie argues that even though absolute dating might not be possible, to quote, "the main value of dates is to show the sequence of events." [14]

By this way of thinking, Petrie introduced what he called "sequence dates": His approach to this matter was to compare the number of artifacts found in graves as a way of sequencing them, reasoning that graves with a close number of artifacts should be close in time. Interestingly,

even though he never used the term matrix, he used the notion of matrices by dedicating a slip of card to each grave, dividing each one to nine columns, and dedicating each column to a specific type of pottery. He then wrote the number of each type of pottery found in each grave, and swapped the cards until he came to an arrangement where graves with a close number of pottery types were close.

While seriation has its roots in archeology, many other disciplines noticed the importance and the need for solving the seriation problem as well. Some of these disciplines include cartography, sociology, psychology, ecology, and biology. Many scientists in these fields applied this idea to tackle questions in their respective fields and even created algorithms to find a suitable sequence. Perhaps each field using its own terminology and having different goals led to challenges in communicating approaches that could have been used in all fields.

In 1951, Robinson wrote the paper "A method for chronologically ordering archaeological deposits" [17], coining the phrase *Robinsonian matrix* or *R-matrix*, describing a matrix whose values decrease when moving away from the diagonal. R-matrices are now considered to be the ideal result of reordering a matrix whose entries indicate how similar objects we need to sequence are. The paper also included a manual procedure for rearranging the matrix to get an R-matrix.

About a decade later, Hole and Shaw wrote the paper "Computer analysis of chronological seriation" [5], which studied the best automatic seriation methods available at the time. Following this, many new methods were proposed, such as [1] and [9], which could now be applied in different fields [10].

1.3.2 Original contributions of this thesis and related works

The problem of seriation is well-studied in several disciplines, and many approaches have been taken to address the challenges arising from the data being noisy or approximated. Many of these are based on a paper by Atkins et al. [2], which proposes an elegant solution for the seriation problem assuming the matrix of similarities has a special property, namely if it is an *R-matrix*. We will see a summary of this paper including the definition of an R-matrix in section 2.3.

Since we already have a solution for the seriation problem in case of an R-matrix, there are two main challenges that we should address when dealing with the seriation problem in most settings and applications: The first one is to check when and if our data forms an R-matrix, and the second is to check the consistency of this solution when facing perturbation.

While the seriation problem has been studied in some biological settings, for example [15] and [7], to the best of the author's knowledge, it hasn't been studied for the setting explained in section 1.1. And while this problem has been studied for some specific perturbations, for example in [12] where the noise is Bernoulli, to the best of the author's knowledge it hasn't been studied for the specific type of perturbation, i.e. approximation by m samples, described in section 1.1.

Moreover, to the best of the author's knowledge, the seriation problem hasn't been studied with the added assumption that the data lie on a curve, which is one of the things we will explore in this thesis.

We will also explore how the geometric properties of this curve will affect the consistency; specifically, how the curvature affects the matrix of similarities being close to an R-matrix. In [4] and [8], the authors introduce parameters measuring how close a matrix is to being an R-matrix, and show that under certain assumptions, an arbitrarily close R-matrix can be found. But the normalized nature of these parameters makes them unsuitable for our purpose.

1.4 The structure of the thesis

The structure of this thesis is as follows: in chapter 2, we introduce the mathematical background needed. In chapter 3, we explain the setting of the problem and introduce the notation which will be used. In chapter 4, we explain how the geometric properties of the developmental curve are related to the consistency of the spectral seriation method.

The main contributions and results are presented in chapter 5. In sections 5.1 and 5.2, we explain how we can translate the discrete nature of our problem to a continuous setting. In section 5.3, the first main result of this thesis is presented, which is that the Laplacian operator associated with the empirical data approaches to that of the ground truth data as the number of samples per time point and the number of time points increase. This is shown in Theorem 5.3.8. We also provide the rate of this convergence, which is useful for experiments that need to know how many samples and time points are needed to reach a specific error bound. In section 5.4, we present the second main theorem which provides sufficient conditions for the curve so that the spectral seriation method is consistent. This is presented in Theorem 5.4.7.

Chapter 2

Mathematical Background

In this chapter, we will introduce a few mathematical concepts that we will later use. In section 2.1 we introduce Optimal Transport and the Wasserstein distance, which will be used for defining the similarity function of our data. In section 2.3 we present a summary of the paper [2], in which the spectral seriation method is introduced. In section 2.2 we define the Laplacian operator and the concept of graphons, which are important in the implementation of the spectral seriation method to the biological problem we are addressing.

2.1 Optimal Transport and Wasserstein Distance

In this section, we introduce the metric we use as a way to quantify the similarity (correlation) between the data points: the Wasserstein distance. To do so, we begin with the concept of optimal transport and build toward our metric of choice.

Imagine we have n balls on $\mathbb{R}$ located on points $x_1, x_2, \ldots, x_n$. Suppose we have to change their locations to points $y_1, y_2, \ldots, y_n$; The further we have to move a ball, the harder the move is. What is the easiest way we can move all the balls to the desired locations? There are $n!$ ways to do so, so the problem of determining the easiest way to move the balls isn't trivial.

The above problem has a discrete setting. Now, let's consider a continuous setting: Imagine we have a pile of dirt somewhere and we aim to fill a hole of the same volume while trying to minimize the distance we move the particles of dirt. How can we find the optimal way to do this?

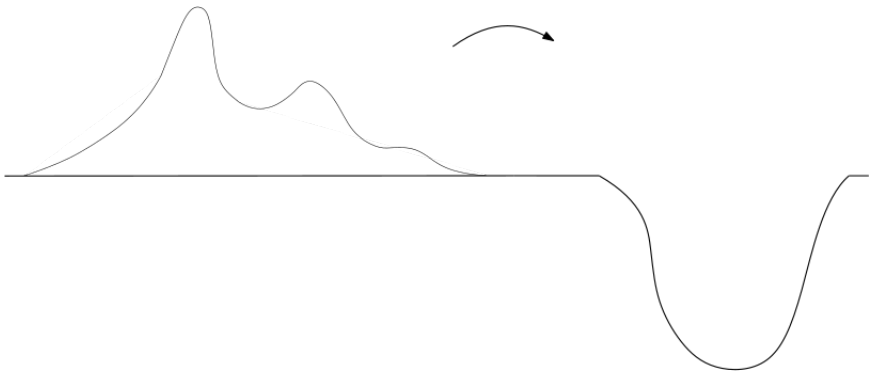

Figure 2.1: Moving a pile to a hole

One way to view the initial pile of dirt is as a probability measure μ, defining $\mu(A)$ for

$A \subset \mathbb{R}^d$ as the normalized amount of dirt above A:

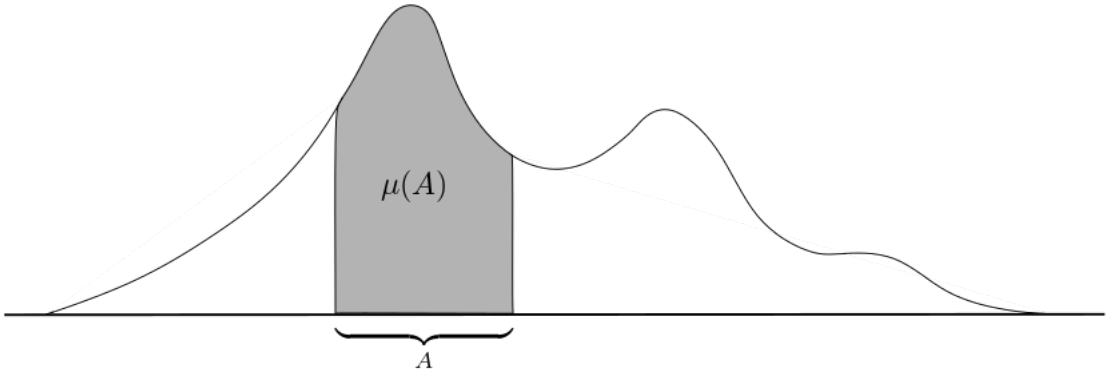

Figure 2.2: Defining measure for a pile of mass

We can think of the hole as a probability measure ν in the same way, where $\nu(A)$ for $A \subset \mathbb{R}^d$ is the normalized volume underneath A in the same fashion as before.

So the problem can be translated as the following:

What is the infimum value of

$$\int_{\mathbb{R}^d} |T(x) - x| \mu(dx)$$

Among all T that "pushforward" μ to ν, meaning that $T : \mathbb{R}^d \to \mathbb{R}^d$ is a measurable function such that

$$\mu(T^{-1}(A)) = \nu(A), \quad \textit{for all Borel measurable subsets } A \textit{ of } \mathbb{R}^d$$

This is precisely what Gaspard Monge, a French mathematician, asked in 1781. The above problem is known as the Monge Problem.

Note that the existence of such T is not guaranteed in all cases. For instance, this is the case when $\mu = \delta_0$ and $\nu = \frac{1}{2}\delta_1 + \frac{1}{2}\delta_2$. In order to overcome this issue, we will introduce a generalization of the problem. To this end, we first introduce the following definition:

Definition 2.1.1. *A coupling π of (μ, ν) is such that if $(X, Y) \in \mathbb{R}^d \times \mathbb{R}^d$ has distribution π, then $X \sim \mu$ and $Y \sim \nu$. The set of all such couplings is represented by $\Pi(\mu, \nu)$. This set is always nonempty since it includes $\mu \otimes \nu$.*

With this definition in mind, we can now introduce the generalized version of the Monge problem known as the Kantorovich problem and also the optimal transport cost:

$$\inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) \pi(dx, dy)$$

Where c can be any cost function, representing how hard it is to move a particle located at x to location y. The Kantrovich problem generalizes the Monge Problem since any pushforward function T defines a coupling: If $X \sim \mu$, then $(X, T(X))$ is a coupling of (μ, ν).

In our problem of interest, we are looking to find a notion of distance between two data points, which are represented by probability measures, namely, a metric structure on the space of probability measures. For a general cost function c, there's no guarantee that the above quantity actually defines a metric, but if the cost function is defined in terms of the distance, it can be proved to satisfy the axioms of a distance function.

Definition 2.1.2. *Let $(\mathcal{X}, d)$ be a complete separable metric space, and let $p \in [1, \infty)$. For any two probability measures μ and ν on $\mathcal{X}$, the Wasserstein distance of order p is defined to be*

$$W_p(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathcal{X} \times \mathcal{X}} d(x, y)^p \pi(dx, dy) \right)^{\frac{1}{p}}$$

The Wasserstein space of order p is defined as

$$P_p(\mathcal{X}) := \left\{ \mu \in P(\mathcal{X}) : \int_{\mathcal{X}} d(x_0, x)^p \mu(dx) < \infty \right\}$$

Here, $x_0 \in \mathcal{X}$ is arbitrary, and the definition doesn't depend on the choice of x_0. It can be proved that W_p defines a metric on $P_p(\mathcal{X})$.

In our case, $\mathcal{X}$ is assumed to be a compact subspace of $\mathbb{R}^d$ where $d \geq 20000$.

For the purpose of this thesis, by the Wasserstein space, we mean $P_p(\mathcal{X})$, where $\mathcal{X}$ is as above.

The elements we wish to sequence are gene expression profiles of embryos represented by points in the Wasserstein space. Using the Wasserstein distance, we can define a similarity function between two gene expression profiles represented by μ and ν. We define this similarity (correlation) function to be $w(\mu, \nu) = e^{\frac{-W_2^2(\mu,\nu)}{\sigma^2}}$ for an arbitrary $\sigma \in \mathbb{R}$.

So far, we have established the mathematical representation of our data, and also have defined a similarity function between them. As such, we have a seriation problem with the additional data that our elements lie on a curve (the trajectory of the embryonic development.) In the last section of this chapter, we will present a summary of the paper [2] where a spectral method is introduced for solving well-posed seriation problems. In the next chapter, we introduce the mathematical background needed for the paper [2].

2.2 The Laplacian and Graphons

In this section, we introduce a concept that will be used abundantly throughout this text, called the *Laplacian*. For this purpose, we refer to the book [16]. We will also introduce a special class of functions, graphons, for which we refer to [4].

Firstly, let's introduce the Laplacian of a symmetric matrix.

Definition 2.2.1. *The Laplacian of a symmetric matrix A is defined to be $L_A = D_A - A$, where D_A is a diagonal matrix with $d_{i,i} = \sum_{j=1}^{n} a_{i,j}$.*

An $n \times n$ real symmetric matrix has real eigenvalues and n eigenvectors that can be constructed to be pairwise orthogonal. The Laplacian of a symmetric matrix is also symmetric, and from the definition, we can see that the vector e, the vector of all ones, is an eigenvector corresponding to the eigenvalue 0. The second smallest eigenvalue (in absolute norm) of the Laplacian is going to be especially important to us, so much that it has its own name:

Definition 2.2.2. *Let A be a symmetric matrix and L_A be its Laplacian. The minimum eigenvalue of L_A with an eigenvector orthogonal to e (the vector of all ones) is called the Fiedler value, and a corresponding eigenvector is called a Fiedler vector. Equivalently, the Fiedler value is given by*

$$\min_{x^T e=0, x^T x=1} x^T L_A x$$

In a small abuse of notation, in this text, we will sometimes refer to the Fiedler value and vector of A when we really mean those of L_A.

It is also possible to define a Laplacian for real valued functions. In this text, we will do so for a special class of functions called *graphons*.

Definition 2.2.3. *A graphon is a measurable function $w : [0,1]^2 \to [0,1]$ which is symmetric in the sense that $w(x,y) = w(y,x)$*

We will be especially interested in R-graphons, coined by Robinson [17], defined as below:

Definition 2.2.4. *A graphon w is an R-graphon (or Robinsonian graphon) if for all $x, y, z \in [0,1]$ we have*

$$y < z < x \Rightarrow w(x,y) \leq w(x,z)$$
$$x < y < z \Rightarrow w(x,y) \geq w(x,z)$$

Now, let's get back to the Laplacians. To define the Laplacian of a graphon and discuss its properties, we first need to be familiar with some concepts from functional analysis.

Definition 2.2.5. *$L^2([0,1])$ is the Hilbert function space with inner product*

$$\langle f, g \rangle := \int_0^1 f(x)g(x)dx$$

and the norm in this space is defined to be $\|f\| := \sqrt{\langle f, f \rangle}$.

Note that the elements of $L^2([0,1])$ are not actually functions, but the equivalence classes of Lebesgue square-integrable functions $f : [0,1] \to \mathbb{R}$. The equivalency classes are defined so that $f \equiv g \Leftrightarrow \|f - g\| = 0$).

We can view vectors as functions whose values are only specified at certain points. As such, we can define the "function version" of the vector e to be $e \in L^2([0,1])$ where $e(x) = 1$ for all $x \in [0,1]$.

We will also be working with *linear operators*, which act on the elements of $L^2([0,1])$. To distinguish them from functions, we will use blackboard letters such as $\mathbb{T}$ to represent them.

The operator norm of a linear operator $\mathbb{T}$ is defined as

$$\|\mathbb{T}\| := \sup_{f \in L^2([0,1]) \text{ s.t. } \|f\|=1} \|\mathbb{T}f\|$$

Now, we are ready to define the Laplacian as a linear operator.

Definition 2.2.6. *The Laplacian $\mathbb{L}$ of a graphon w is a linear operator acting on an element $f \in L^2([0,1])$ in the following way:*

$$\begin{aligned}(\mathbb{L}f)(x) &:= \int_0^1 w(x,y)(f(x) - f(y))dy \\ &= f(x)d(x) - \int_0^1 w(x,y)f(y)dy\end{aligned}$$

where

$$d(x) = \int_0^1 w(x,y)dy$$

Similar to the Laplacian matrix, we can talk about the Eigenvalue and Eigen function of the linear operator Laplacian which are defined in a similar fashion. Specifically, note that

$$\langle \mathbb{L}f, f \rangle = \frac{1}{2} \int_0^1 \int_0^1 w(x,y)(f(x) - f(y))^2 dxdy \geq 0 \tag{2.1}$$

So e is an Eigen function of L with Eigenvalue 0, and if $w(x,y) > 0$ almost everywhere, then 0 is a simple Eigenvalue. Similar to the Laplacian matrix, we can define the Fiedler value as the smallest non-zero Eigenvalue, and the Fiedler function is any corresponding function.

2.3 Summary of the paper "A Spectral Algorithm For Seriation And The Consecutive Ones Problem" [2]

A common problem in many diverse applications, ranging from DNA sequencing to archaeology, is to sequence elements so that similar elements are close to each other. In [2], a spectral algorithm is presented as a solution to this class of problems. Previous to this paper, the traditional methods for solving such problems were rather combinatory whereas this paper uses matrix theory.

In mathematical terms, the *seriation problem* is as follows: Given a set of n elements to sequence and a symmetric real-valued *correlation function* $f(i,j)$ representing the similarity

between elements i and j, we would like to find a permutation π *consistent* with the correlation function, in the sense that if $\pi(i) < \pi(j) < \pi(k)$, then $f(i,j) \geq f(i,k)$ and $f(j,k) \geq f(i,k)$. If such a permutation exists, the problem is *well-posed*. In such cases, there may be an exponential number of consistent permutations; however, it is possible to describe all of them in a compact data structure known as a PQ-tree, which will be defined in 2.3.1.

In this section, we present a summary of the paper "A Spectral Algorithm For Seriation And The Consecutive Ones Problem" by Atkins et al. [2] organized as follows: first, the mathematical notation and some needed results from matrix theory are introduced. Then, a spectral heuristic is described, motivating the approach of their paper. In subsection 2.3.3, the key theorem of the paper [2] is presented and the next subsection is dedicated to addressing some restrictions regarding the key theorem. The last subsection provides a concise algorithm for solving a well-posed seriation problem. All the material in this section is due to Atkins et al. [2], and the proofs presented are as stated in the same paper.

2.3.1 Mathematical Notation

The correlation function for a set of n elements can be viewed as a real symmetric $n \times n$ matrix whose $(i,j)^{th}$ entry is $f(i,j)$. This way, matrix theory can be used to discuss the seriation problem. In doing so, a permutation of elements corresponds to a symmetric permutation of the matrix, i.e. permuting the rows and columns in the same way. Then, whether or not the ordering problem is well-posed becomes a property of the matrix:

Consider a well-posed ordering problem and the corresponding matrix that has been permuted to reflect a consistent solution to the ordering problem. Then, the off-diagonal entries must increase as we move toward the diagonal since moving toward the diagonal corresponds to the correlation of elements more similar to each other. In general, if a matrix has this property, it is called an *R-matrix*. Mathematically:

Definition 2.3.1. *A matrix $A = [a_{ij}]$ is an R-matrix if A is symmetric and*

$$\begin{aligned} a_{i,j} &\leq a_{i,k} \quad \text{for } j < k < i \\ a_{i,j} &\geq a_{i,k} \quad \text{for } i < j < k \end{aligned} \tag{2.2}$$

The diagonal entries of an R-matrix do not need to be specified. A permutation of elements would correspond to a permutation of the corresponding matrix. As such, we have the following definition:

Definition 2.3.2. *For a permutation π, and matrix $A = [a_{ij}]$, A^π is the matrix given by the permutation of rows and columns of A according to π, i.e. $a^\pi_{i,j} = a_{\pi(i),\pi(j)}$. Similarly, for a vector $x = [x_1, \ldots, x_n]^T$, x^π is such that $x^\pi_i = x_{\pi(i)}$.*

For a well-posed problem, permuting the matrix with respect to a consistent permutation should result in a matrix whose entries increase as we move toward the diagonal, i.e. an R-matrix. As such, it is important to classify such matrices:

Definition 2.3.3. *A matrix A is called pre-R if there exists permutation π for which A^π is an R-matrix.*

Note that pre-R matrices correspond precisely to well-posed ordering problems. In addition, the R-matrix property is preserved if we add a constant to all (the off-diagonal) entries of the matrix. As a result, we may assume all the off-diagonal entries are nonnegative.

If a large portion of elements have little to no correlation with another large portion of the elements, it would be reasonable to try and sequence each portion independently. The matrix corresponding to such elements would consist of many zeros, and if permuted with respect to a consistent permutation, it will have blocks of zeros.

Definition 2.3.4. *A symmetric matrix A is reducible if there exists a permutation π such that*

$$A^\pi = \begin{pmatrix} B & 0 \\ 0 & C \end{pmatrix}$$

Where B and C are square matrices. If such permutation doesn't exist, A is called irreducible. If B and C are themselves irreducible, then we refer to them as the irreducible blocks of A.

If $Ax = \lambda x$ for some vector $x \neq 0$, then λ is an eigenvalue of A, and a corresponding vector x is an eigenvector. The algebraic multiplicity of an eigenvalue λ is the number of times the term $z - \lambda$ occurs in the characteristic polynomial $p(z) = \det(A - zI)$. If this only happens once, the eigenvalue is called simple; the eigenvector of a simple eigenvalue is unique up to normalization.

An $n \times n$ real, symmetric matrix has n pairwise orthogonal eigenvectors, and all of its eigenvalues are real. In this paper, it is assumed that the eigenvalues are sorted by increasing value: $\lambda_1 \leq \lambda_2 \leq \ldots \lambda_n$. If $A = [a_{ij}]$ is such that all $a_{ij} \geq 0$, then $A \geq 0$ and A is called nonnegative. A real vector x is monotone if $x_i \leq x_{i+1}$ for all $1 \leq i < n$ or if $x_i \geq x_{i+1}$ for all $1 \leq i < n$.

Next, we remind the definition of the most important concept used in [2], the Laplacian of a matrix.

Definition 2.3.5. *The Laplacian of a symmetric matrix A is defined to be $L_A = D_A - A$, where D_A is a diagonal matrix with $d_{i,i} = \sum_{j=1}^{n} a_{i,j}$. The smallest eigenvalue whose eigenvector is orthogonal to e (the vector of all ones) is called the Fiedler value, and a corresponding eigenvector is called a Fiedler vector. Equivalently, the Fiedler value is given by*

$$\min_{x^T e = 0, x^T x = 1} x^T L_A x$$

A Fiedler vector is any vector x that achieves this minimum while satisfying these constraints. When $A \geq 0$ and irreducible, the Fiedler value is the smallest nonzero eigenvalue and a Fiedler vector is any corresponding eigenvector. To simplify the notations, In this text, the Fiedler value and vector of A are actually meant to be those of L_A.

The next definition is about a structure that will be used later to demonstrate the set of all consistent permutations for a well-posed problem, the PQ-tree:

Definition 2.3.6. *A PQ-tree is a data structure to compactly encode a set of related permutations. A PQ-tree over a set $U = \{u_1, u_2, \ldots, u_n\}$ is a rooted, ordered tree whose leaves are elements of U and whose internal nodes are distinguished as either P-nodes or Q-nodes. A PQ-tree is proper when the following three conditions hold.*

1. *Every element $u_i \in U$ appears precisely once as a leaf.*
2. *Every P-node has at least two children.*
3. *Every Q-node has at least three children.*

Two PQ-trees are said to be equivalent if one can be transformed into the other by applying a sequence of the following two equivalence transformations.

1. *Arbitrarily permute the children of a P-node.*
2. *Reverse the children of a Q-node.*

As it turns out, the equivalence class represented by a PQ-tree corresponds precisely to the set of permutations consistent with correlations of a seriation problem.

2.3.2 A Spectral Heuristic

One way to approach the seriation problem is to minimize the function

$$g(\pi) = \Sigma_{i,j} f(i,j)(\pi_i - \pi_j)^2$$

where π is a permutation. This function can be viewed as a penalty function, as its value gets larger if closely correlated elements are not close to each other. Even if the problem is not well-posed, minimizing g would keep the closely correlated elements close to each other. However, due to the discrete nature of the problem, minimizing g is an NP-hard problem. In order to address this issue, the function h of continuous variables $x = (x_1, \ldots, x_n)$ is considered where

$$h(x) = \Sigma_{i,j} f(i,j)(x_i - x_j)^2$$

In a sense, h is the continuous relaxation of g, and as it turns out, finding its minimizers is easier. The function h doesn't have a unique minimizer, since its value doesn't change if a constant is added to each x component. To avoid this, the constraint $\Sigma_i x_i = 0$ is added to the minimization problem. And to avoid the trivial solution $x = 0$, the constraint $\Sigma_i x_i^2 = 1$ is

added. The resulting minimization problem is now well-defined:

$$\begin{aligned} &\text{minimize } h(x) = \Sigma_{i,j} f(i,j)(x_i - x_j)^2 \\ &\text{subject to } \Sigma_i x_i = 0, \text{ and } \Sigma_i x_i^2 = 1 \end{aligned} \tag{2.3}$$

The solution to this continuous problem can be used as a heuristic for the seriation problem: If the elements x_i of the solution vector x are sorted and the elements are sequenced based upon their sorted order, the closely correlated data are kept close to each other.

This heuristic is specially useful because 2.3 can be rewritten as

$$\min_{x^T e=0, x^T x=1} x^T L_F x$$

Where $F = [f(i,j)]$. As discussed before, this is equivalent to finding the Fiedler vector. Considering all these, it is reasonable to propose that sorting the entries of a Fiedler vector and sequencing based on their sorted order should solve a well-posed ordering problem.

At this point, one might ask the following questions:

- For a well-posed problem, do all the Fiedler vectors lead to the same ordering? As discussed in the next subsection, it turns out that the Fiedler vector of a pre-R matrix, under some reasonable conditions, is actually simple.
- Sequencing based on the order of the Fiedler vector's entries would be consistent with correlations if and only if the Fiedler vector of an R-matrix is monotone. Is that actually the case? The main theorem of this paper [2], stated and proved in the next subsection, answers this question.

2.3.3 The Key Theorem of [2]

The main result of [2] is that after some modification, the simple heuristic presented in section 2.3.2 provides an algorithm for well-posed seriation problems. Of course, to thoroughly prove this one should address the ambiguities of sequencing based on the order of Fiedler vector's entries: What if there is more than one Fiedler vector? and what if there are ties within the Fiedler vector? These issues are addressed in the next subsection.

Before stating the main theorem, a theorem due to Perron and Frobenius is stated, which establishes the core of analysis leading to the main theorem of [2]:

Theorem 2.3.7 (Perron-Frobenius)**.** *Let M be a real, nonnegative matrix. If we define $\rho(M) = \max_i |\lambda_i(M)|$, then there is a vector $x \geq 0$ such that $Mx = \rho(M)x$.*

We can now state the main theorem:

Theorem 2.3.8 (Main Theorem of [2])**.** *If A is an R-matrix then it has a monotone Fiedler vector.*

Proof (as stated in [2]). One way to analyse whether or not a vector $(x_1, x_2, \ldots, x_n)^T$ is monotone is to check if $(x_2 - x_1, \ldots, x_n - x_{n-1})^T$ is nonnegative (or nonpositive). This is where the nonnegative vector in theorem 2.3.7 comes in handy.

To create the vector consisting of differences between neighboring entries in the Fiedler vector, first define the matrix $S \in \mathbb{R}^{(n-1)\times n}$ as

$$S = \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 \\ 0 & -1 & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & -1 & 1 \end{bmatrix}$$

Note that for any vector x, $Sx = (x_2 - x_1, \ldots, x_n - x_{n-1})^T$. Define $T \in \mathbb{R}^{n\times(n-1)}$ by

$$T = \begin{bmatrix} 0 & 0 & \cdots & 0 \\ 1 & 0 & & 0 \\ 1 & 1 & \ddots & \vdots \\ \vdots & \vdots & \ddots & 0 \\ 1 & 1 & \cdots & 1 \end{bmatrix}$$

It can easily be checked that $ST = I_{n-1}$, and that $TS = I_n - ee_1^T$. Now, define $M_A = SL_AT = \{m_{i,j}\}$ and let $L_A = \{l_{i,j}\}$. The following calculation proves that Sx is an eigenvector of M_A if and only if x is an eigenvector of L_A and $x \neq \alpha e$.

$$\begin{aligned} L_Ax = \lambda x, \quad x \neq \alpha e \iff \\ SL_Ax = \lambda Sx, x \neq \alpha e \iff \\ SL_A\left(I - ee_1^T\right)x = \lambda Sx, x \neq \alpha e \iff \\ SL_ATSx = \lambda Sx, x \neq \alpha e \iff \\ M_Ay = \lambda y, \text{ where } y = Sx \neq 0. \end{aligned}$$

The third equivalency is because $L_Ae = 0$. By the above equivalencies, λ is an eigenvalue for both L_A and M_A for eigenvectors of L_A other than e. Hence the eigenvalues of M_A are the same as the eigenvalues of L_A with the zero eigenvalue removed, and the eigenvectors of M_A are differences between neighboring entries of the corresponding eigenvectors of L_A. It is easily seen that $(SL_A)_{i,k} = -l_{i,k} + l_{i+1,k}$ for all i, k, so

$$m_{i,j} = \sum_{k=1}^{n} (SL_A)_{i,k} T_{k,j} = \sum_{k=j+1}^{n} (-l_{i,k} + l_{i+1,k}) = \sum_{k=j+1}^{n} (a_{i,k} - a_{i+1,k}).$$

Since A is an R-matrix, $a_{i,k} \leq a_{i+1,k}$ for $i < k + 1$, and therefore $m_{i,j} \leq 0$ for $i < j$. For $i > j$,

since $\sum_{k=1}^{n} l_{i,k} = 0$ we have

$$m_{i,j} = \sum_{k=j+1}^{n} (-l_{i,k} + l_{i+1,k}) = \sum_{k=1}^{j} (l_{i,k} - l_{i+1,k}) = \sum_{k=1}^{j} (-a_{i,k} + a_{i+1,k}) .$$

Again, since A is an R-matrix, $m_{ij} \leq 0$ for $i > j$. As a result, all the off-diagonal elements in M_A are nonpositive.

Now let β be a value greater than $\max_i \{\lambda_i, m_{ii}\}$, where λ_i are the eigenvalues of M_A. Then $\tilde{M}_A = \beta I - M_A$ is nonnegative with eigenvalues $\tilde{\lambda}_i = \beta - \lambda_i$. Also, $\tilde{M}_A$ and M_A share the same set of eigenvectors. By theorem 2.3.7, there exists a nonnegative eigenvector y of $\tilde{M}_A$ corresponding to the largest eigenvalue of $\tilde{M}_A$. Note that y is also an eigenvector of M_A corresponding to M_A 's smallest eigenvalue which equals Sx, where x is a Fiedler vector of L_A. Since $y = Sx$ is nonnegative, the corresponding Fiedler vector of L_A is nondecreasing and the theorem is proved.

□

Theorem 2.3.9 (Theorem 3.3 of [2]). *Let A be a pre-R-matrix with a simple Fiedler value and a Fiedler vector with no repeated values. Let π_1 (respectively, π_2) be the permutation induced by sorting the values in the Fiedler vector in increasing (decreasing) order. Then A^{π_1} and A^{π_2} are R-matrices, and no other permutations of A produce R-matrices.*

Proof. First, note that if x is the Fiedler vector of A, then x^{π} is the Fiedler vector of A^{π}, meaning that applying a permutation to the matrix only changes the order of entries of the Fiedler vector. Now, let π_* be a permutation such that A^{π_*} is an R matrix. By theorem 2.3.8, x^{π_*} is monotone, since the Fiedler value is simple. But x has no repeated values, so π_* is either π_1 or π_2 □

Theorem 2.3.9 is essential to creating an algorithm for the seriation problem since it shows that under some conditions, the permutation is unique. However, it is too restrictive: the Fiedler value has to be simple and contain no repeated values. In the next subsection, these restrictions are addressed.

2.3.4 Developing the Algorithm and Addressing the Restrictions

In this subsection, the restrictions of Theorem 2.3.9 are addressed while developing the spectral seriation algorithm.

First, observe that adding a constant to all the entries of the correlation matrix wouldn't change whether or not the matrix is pre-R. In fact, it does not change the set of eigenvectors either. This fact can be algebraically stated as the following lemma.

Lemma 2.3.10 (Lemma 4.1 of [2]). *Let A be a symmetric matrix and let $\bar{A} = A - \alpha e e^T$ for some real α. A vector x is a Fiedler vector of A if and only if x is a Fiedler vector of $\bar{A}$. So without loss of generality, we can assume that the smallest off-diagonal entry of A is zero.*

Proof. By the definition of a Laplacian, $L_{\bar{A}} = L_A + \alpha ee^T - \alpha nI$, where n is the dimension of A. Then $L_{\bar{A}}e = 0$, as we would expect, and as for the other eigenvector x of L_A, $L_{\bar{A}}x = L_Ax + 0 - \alpha nx$ since x is orthogonal to e. That is, the eigenvalues are shifted down by αn while the eigenvectors are the same as for L_A. □

To simplify the analysis, the first step of the algorithm is set to be subtracting the smallest correlation from others, which can be done by the above lemma. As such, it can be assumed without loss of generality that a pre-R matrix has zero as the smallest off-diagonal entry.

Next, notice that if A is reducible, a consistent permutation for each of its irreducible blocks can be found and then they can be concatenated in an arbitrary order. Mathematically, this follows from the next lemma.

Lemma 2.3.11 (Lemma 4.2 of [2])**.** *Let $A_i, i = 1, \ldots, k$, be the irreducible blocks of a pre-R-matrix A, and let π_i be a permutation of block A_i such that the submatrix $A_i^{\pi_i}$ is an R-matrix. Then any permutation formed by concatenating the π_i 's will make A become an R matrix. In terms of a PQ-tree, the π_i permutations are children of a single P-node.*

Proof. By the previous lemma, assume all entries in the irreducible blocks are nonnegative. Since the permutations are done symmetrically, any permutation making A_i an R-matrix does not affect the other irreducible blocks. As long as each block is ordered to be an R-matrix, any ordering of blocks will make A an R-matrix. □

With these lemmas in mind, it can now be assumed that the smallest off-diagonal entry is zero and that the matrix is irreducible. As it turns out, The following three lemmas and theorem prove that these conditions are sufficient for the Fiedler value to be simple.

Lemma 2.3.12 (Lemma 4.3 of [2])**.** *Let A be an $n \times n$-matrix with a monotone Fiedler vector x. If $\mathcal{J} = [r, s]$ is a maximal interval such that $x_r = x_s$, then for any $k \notin \mathcal{J}, a_{r,k} = a_{r+1,k} = \cdots = a_{s,k}$*

Proof. Without loss of generality, assume x is nondecreasing since $-x$ is also a Fiedler vector. It suffices to show that $a_{r,k} = a_{s,k}$ for all $k \notin \mathcal{J}$. This is because A is an R-matrix and so all elements between $a_{r,k}$ and $a_{s,k}$ must also be equal. Now, consider rows r and s in the equation $L_Ax = \lambda x$:

$$\sum_{k=1}^{n} (l_{s,k} - l_{r,k})\, x_k = \lambda\,(x_s - x_r) = 0.$$

Since L_A is a Laplacian, $\sum_{k=1}^{n} l_{i,k} = 0$ for all i. Hence,

$$\begin{aligned} 0 &= \sum_{k=1}^{n} (l_{sk} - l_{rk})\,(x_r - x_k) \\ &= \sum_{k=1}^{r-1} \underbrace{(l_{s,k} - l_{r,k})}_{\geq 0}\underbrace{(x_r - x_k)}_{>0} + \sum_{k=s+1}^{n} \underbrace{(l_{s,k} - l_{r,k})}_{\leq 0}\underbrace{(x_r - x_k)}_{<0} \end{aligned}$$

where the fact that x is nondecreasing is used. Because all terms in the sum are nonnegative, all terms must be equal to zero. By assumption, $x_k \neq x_r$ for $k \notin \mathcal{J}$ and as a result, $l_{r,k} = l_{s,k}$ for $k \notin \mathcal{J}$ and the lemma is proved. □

As it turns out, the reverse of this theorem is also true, stated as the next lemma. The proof is not included since it contains detailed algebraic calculations that are not essential to the core content of this text.

Lemma 2.3.13 (Lemma 4.4 of [2]). *Let A be an irreducible $n \times n$ R-matrix with $a_{n,1} = 0$. If $\mathcal{J} = [r, s] \neq [1, n]$ is an interval such that $a_{r,k} = a_{s,k}$ for all $k \notin \mathcal{J}$, then $x_r = x_{r+1} = \cdots = x_s$ for any Fiedler vector x.*

The next lemma states that for an R-matrix even if the Fiedler value is not simple, the index set of repetitive entries of different Fiedler vectors are the same.

Lemma 2.3.14 (Lemma 4.5 of [2]). *Let A be an irreducible R-matrix with $a_{n,1} = 0$, and x, a monotone Fiedler vector of A. If $\mathcal{J} = [r, s]$ is an interval such that $x_r = x_{r+1} = \cdots = x_s$, then for any Fiedler vector $y, y_r = y_{r+1} = \cdots = y_s$*

Proof. By lemma 2.3.12, for any $k \notin \mathcal{J}, a_{r,k} = a_{r+1,k} = \cdots = a_{s,k}$. Since $x^T e = 0$, it follows that $\mathcal{J} \neq [1, n]$. Now, by Lemma 2.3.13 the result follows. □

Now, it can be proved that an irreducible R-matrix has a simple eigenvalue.

Theorem 2.3.15 (Theorem 4.6 of [2]). *If A is an irreducible R-matrix with $a_{n,1} = 0$, then the Fiedler value λ_2 is a simple eigenvalue.*

Proof. (As stated in [12]) Assume that λ_2 is not simple. Then there exist two linearly independent Fiedler vectors x and y, with x non-decreasing by lemma 2.3.8. Define

$$k := \arg\max \left\{ \frac{y_{j+1} - y_j}{x_{j+1} - x_j} \text{ s.t. } j \in \{1, 2, \ldots, n\} \text{ and } x_{j+1} \neq x_j \right\}$$

such k should exist since $x^T e = 0$. Now define $z = (y_{k+1} - y_k)\, x - (x_{k+1} - x_k)\, y$. z is also a Fiedler vector because x and y are linearly independent.

For $j \in \{1, 2, \ldots, n\}$, if $x_{j+1} = x_j$ then by Lemma 2.3.14 $y_{j+1} = y_j$ thus $z_{j+1} = z_j$. If $x_{j+1} \neq x_j$ then

$$z_{j+1} - z_j = (x_{j+1} - x_j)\,(x_{k+1} - x_k) \left(\frac{y_{k+1} - y_k}{x_{k+1} - x_k} - \frac{y_{j+1} - y_j}{x_{j+1} - x_j} \right).$$

then $z_{j+1} - z_j \geq 0$ by definition of k, meaning that z is monotone. On the other hand $z_{k+1} = z_k$, and by Lemma 2.3.14 the indices of any repeated values in z must be repeated in x that is $x_{k+1} = x_k$ which contradicts the definition of k. We conclude that λ_2 should be simple. □

Now, it is time to address the case of repetitive entries of the Fiedler vector.

Theorem 2.3.16 (Theorem 4.7 of [2]). *Let A be a pre-R-matrix with a simple Fiedler value and Fiedler vector x. Suppose there is some repeated value β in x and define $\mathcal{I}$, $\mathcal{J}$, and $\mathcal{K}$ to be the indices for which*

1. *$x_i < \beta$ for all $i \in \mathcal{I}$,*
2. *$x_i = \beta$ for all $i \in \mathcal{J}$,*
3. *$x_i > \beta$ for all $i \in \mathcal{K}$.*

Then π is an R-matrix ordering for A if and only if π or its reversal can be expressed as (π_i, π_j, π_k), where π_j is an R-matrix ordering for the submatrix $A(\mathcal{J}, \mathcal{J})$ of A induced by $\mathcal{J}$, and π_i and π_k are the restrictions of some R-matrix ordering for A to $\mathcal{I}$ and $\mathcal{K}$, respectively.

Proof. By Theorem 2.3.8 for any R-matrix ordering A^π, x^π is monotone, so entries in $\mathcal{I}$ should be located before (or after, depending on whether or not x is non-increasing) entries from $\mathcal{J}$ and elements from $\mathcal{K}$ must appear after (before) elements from $\mathcal{J}$. By Lemma 2.3.12, $a_{ik} = a_{jk}$ for all $i, j \in \mathcal{J}$ and $k \notin \mathcal{J}$. As a result, the ordering of elements inside $\mathcal{J}$ will not affect the ordering outside of $\mathcal{J}$ and vice versa. Hence, the R-matrix ordering of elements in $\mathcal{J}$ depends only on $A(\mathcal{J}, \mathcal{J})$. □

This theorem provides a tool for dealing with the ties in the Fiedler vector: By recursing on the submatrix $A(\mathcal{J}, \mathcal{J})$ where $\mathcal{J}$ corresponds to the set of repeated values, the data can be sequences corresponding to the $\mathcal{J}$ index set. While encoding the consistent permutations in a PQ-tree, the distinct values are combined through a Q-node, and the components of the Q-node must then be expanded recursively.

2.3.5 A Spectral Algorithm for the Seriation Problem

All these results form the main algorithm presented by [2], returning a PQ-tree encoding all the permutations consistent with the correlations of a seriation problem:

```
     Input:    A, an n × n pre-R-matrix
               U, a set of indices for the rows/columns of A
     Output:   T, a PQ-tree that encodes the set of all permutations π
               such that A^π is an R-matrix
     begin
(1)      α := min_{i≠j} a_{i,j}
(1)      A := A − αee^T
(2)      {A_1, ..., A_k} := the irreducible blocks of A
(2)      {U_1, ..., U_k} := the corresponding index sets
(2)      if k > 1
(2)          for j := 1 : k
(2)              T_j := Spectral-Sort(A_j, U_j)
(2)          end
(2)          T := P-node(T_1, T_2, ..., T_k)
         else
(3)          if (n = 1)
(3)              T := u_1
(3)          else if (n = 2)
(3)              T := P-node(u_1, u_2)
             else
(4)              x := Fiedler vector for L_A
(4)              Sort x
(5)              t := number of distinct values in x
(5)              for j := 1 : t
(5)                  V_j := indices of elements in x with jth value
(5)                  T_j := Spectral-Sort(A(V_j, V_j), V_j)
(5)              end
                 T := Q-node(T_1, ..., T_T)
             end
         end
     end
```

Figure 2.3: Spectral-sort algorithm (Fig 1. of [2])

Chapter 3

Geometric Background for the Seriation Problem

The goal of this chapter is to build some initial intuition on the seriation problem with the additional assumption that the data lies on a smooth curve. In this chapter we will discuss how we can approach the problem and why we approach it the way we do without getting into the technicalities, which will be discussed in detail in the next chapters.

3.1 Assumptions

As we discussed before, we aim to recover the ordering of a finite number of data points in the Wasserstein space based on their correlations. Practical limitations prevent us from obtaining an arbitrarily large data set, but for a moment, let's imagine that we could obtain n data points, for n arbitrarily large. What happens when n approaches infinity? The image of the points in the Wasserstein space would then approach a curve, which can be reasonably assumed to be continuous. If the curve doesn't self intersect, a parameterization of this curve would result in two unique orderings for the data (where one is the reverse of the other). Our goal is to recover this ordering using only finitely many points on the curve. In summary, we have the following information and assumption:

Assumption: There is a fixed curve on which all the data points lie. This curve is continuous and does not self-intersect.

Information: We have only finitely many points on this fixed curve.

Even in this early stage of stating the problem, it is clear that we should address two questions:

1. How many data points are "enough" for recovering the ordering?

2. What geometric properties of the curve affect our ability to recover the ordering?

These questions will be addressed in chapter 5. For now, we will begin with a simple analysis to develop a deeper intuition toward the problem at hand and determine an appropriate approach to tackle it.

3.2 Do all curves result in an R-matrix?

As we saw in the summary of [2] presented in section 2.3, if the matrix of correlations corresponding to the data set is a pre-R matrix, then we can find an ordering consistent with the correlations using the spectral seriation algorithm. With our new assumption that the data lie on a curve inducing its own ordering, it is reasonable to ask the following question:

- What properties should the curve have for the matrix of correlations corresponding to a finite set of points on the curve to be an R matrix?

To expand on the question, let's first define the following:

Definition 3.2.1. *Let (X, d) be a metric space. A curve $c : [0, 1] \to X$ is an R-curve if for any $0 \leq t_1 < t_2 < t_3 \leq 1$, we have that $d(c(t_1), c(t_2)) \leq d(c(t_1), c(t_3))$ and $d(c(t_2), c(t_3)) \leq d(c(t_1), c(t_3))$.*

So, the above question is simply asking: Are all curves R-curves?

First, let's look at the problem in a simpler space instead of the Wasserstein space, namely $\mathbb{R}^2$ with the Euclidean distance d. Then we can see that there are curves that are not R-curves:

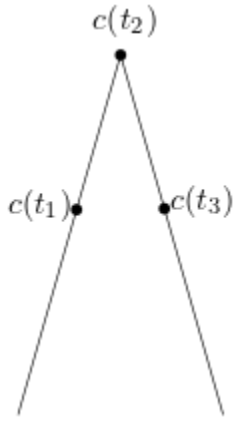

Figure 3.1: An example of a curve that is not an R-curve

Here, even though $c(t_1)$ and $c(t_2)$ are closer in the ordering induced by the curve than $c(t_1)$ and $c(t_3)$, $c(t_1)$ and $c(t_3)$ are closer in Euclidean distance.

Now, imagine a straight line in $\mathbb{R}^2$ as a curve $c : [0, 1] \to \mathbb{R}^2$. We can see that it is an R-curve, as $d(c(t_1), c(t_3)) = d(c(t_1), c(t_2)) + d(c(t_2), c(t_3))$ for any for any $0 \leq t_1 < t_2 < t_3 \leq 1$:

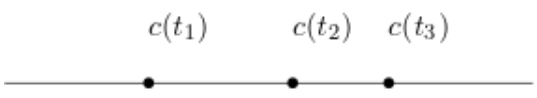

Figure 3.2: A straight line is an R-curve

Now let's bend it very slightly. It remains an R-curve, as distance is a continuous function:

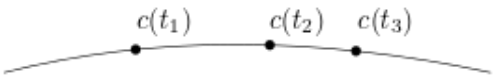

Figure 3.3: A slightly bent line remains an R-curve

however, if we bend it much more, as in the following figure, then it loses this property:

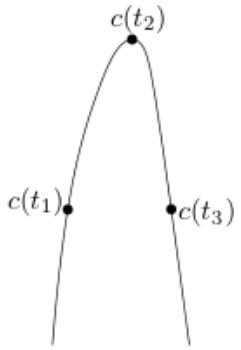

Figure 3.4: A highly bent curve is no longer an R-curve

The question that arises is:

- How much can we "add" to the curvature of a straight line without making it stop being an R-curve?

Note that in a sense, every smooth curve is "locally" an R-curve, since it locally resembles a straight line if we zoom in close enough. How much we should zoom depends on the curvature, since that is what's deviating the curve from being straight and hence, being an R-curve.

Now, we can see that a reasonable guess about the answer of the question we asked earlier, "What geometric properties of the curve affect our ability to recover the ordering?" is curvature.

While we did use $\mathbb{R}^2$ to obtain an intuition about the role curvature plays, in chapter 5, we find a quantitative relation between a generalized notion of curvature and the consistency of the spectral seriation method.

3.3 What should we do if we don't have an R-curve?

As we discussed earlier, since a smooth curve locally resembles a straight line, we can, in theory, solve the ordering problem "locally" and then "patch" the orderings together. This is possible since by theorem 2.3.9, the ordering on each part is unique. In practice, this isn't a pragmatic way, since we wouldn't know how "local" we should get for each point on the curve, or how to order the set of orderings with respect to each other. Even so, knowing that it is in theory possible to do so, hints us toward a practical method which we will discuss in the next chapter.

Chapter 4

Geometric Analysis and it's connection to the spectral gap

While this thesis addresses a question involving the Wasserstein space, the work in this chapter can be applied to any metric space. Throughout this chapter, (M, d) is the metric space our data set belongs to, and $c : [0, 1] \to M$ is a continuous curve.

4.1 Approaching the problem locally

Consider the curve $c : [0, 1] \to M$, Where (M, D) is the metric space.

Definition 4.1.1. *For each $0 \leq t_1 \leq t_2 \leq t_3 \leq 1$, we call (t_1, t_2, t_3) a good triad if:*

$$D(c(t_1), c(t_3)) \geq D(c(t_1), c(t_2))$$
$$D(c(t_1), c(t_3)) \geq D(c(t_2), c(t_3))$$

Definition 4.1.2. *At each point $x \in [0, 1]$, define*

$$d_x := \sup_r \{r | \forall \quad o \leq t_1 \leq t_2 \leq t_3 \leq 1 \quad with \quad c(t_1), c(t_2), c(t_3) \in B_r(x) : (t_1, t_2, t_3) \quad is\ a\ good\ triad\}$$

Where $B_r(x)$ is the ball of radius r around x. Also, define

$$d_c = \min\{d_x : x \in [0, 1]\}.$$

Since the function $x \mapsto d_x$ is continuous, d_c is well-defined.

Consider n data points $\{x_1, \ldots x_n\}$ on the curve, and a similarity function $f(i, j) = F(d(x_i, x_j)$, where F is smooth and $F' < 0$. Now, Let A be the matrix of correlations. Geometrically, d_c measures the size of the local region on which the matrix of correlations of the data, with the ordering induced from the curve, is an R-matrix.

If $d_x > 0$ for all x, then $d_c > 0$. This means that we will have an R-matrix A' around the diagonal of A, the size of which depends on d_c:

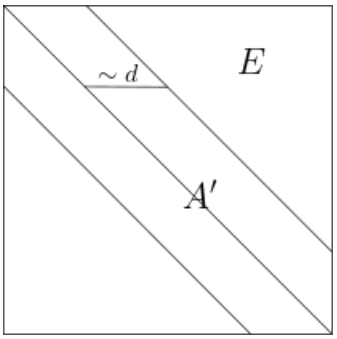

Figure 4.1: $A = A' + E$, A' an R-matrix

Here, $E = A - A'$, and the larger d is, the smaller E is. This means that the matrix of correlations is close to an R-matrix, which we know how to find an ordering for. The question that arises here is:

- What properties should the curve c have for d_c to be large?
- Does d_c being large and consequently the matrix of correlations being "close" to an R-matrix mean that the spectral seriation method will work, at least to some extent, for the data set?

To think about the first question, let's look at the following examples in $\mathbb{R}^2$:

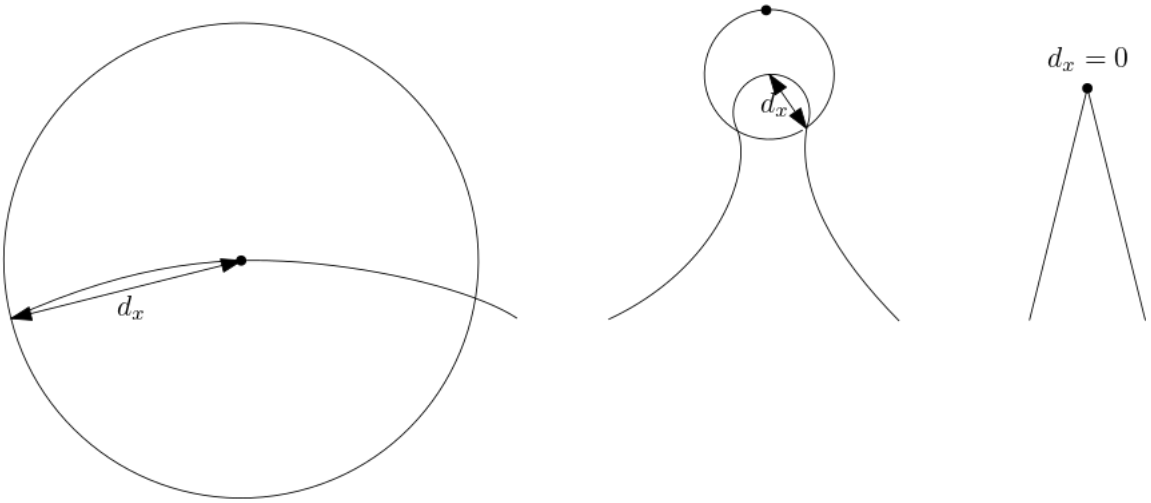

Figure 4.2: The effect of curvature on d_x

We can see that at points where the curvature is small, d_x is larger and when the curvature is large, d_x is small. When we have sharp corners, d_x can be zero. It's reasonable to guess that the smaller the curvature is, the bigger d_x is. This is consistent with our previous observation that while straight lines are R-curves and a slight curvature doesn't change that, larger curvature can mean that the curve is not an R-curve. In the following two sections, we will discuss these questions in more detail.

4.2 Use of Davis-Kahan

So far, we've talked about the intuition behind why for a curve c a small curvature and consequently a large d_c should lead to a correlation matrix that is close to an R-matrix. But how does that guarantee that the spectral seriation method will work for our data set? We know that the seriation method does work for the R-matrix. Also, we know that the ordering gained from the R-matrix alone should coincide with the ground truth ordering, since by only considering the R-matrix, we are finding the ordering locally and then patching them together. For a given data set, we might not be able to know what that inner R-matrix is, so it is worth investigating to see if being close to an R-matrix will guarantee a close Fiedler vector. In this section, we prove that being close enough to an R-matrix does indeed guarantee a close Fiedler vector and a simple Fiedler value. Secondly, we need the Fiedler vector to be monotone, which as we will show can be achieved by being close to an R-matrix. While the calculations in this section result in a rather non-quantitative result, i.e. we can't tell how close the matrix should be to an R-matrix for the method to work, we will have more quantitative results in chapter 5.

Let A be the matrix of correlations for the data set $\{x_1, \dots x_n\}$ on the curve c. As we saw in the previous section, if $d_c > 0$, then we can write

$$A = R + E$$

where R is an R-matrix around the diagonal of A.

A well-known theorem by Davis and Kahan [6] proves that there is an upper bound for the angle between two eigenvectors of two symmetric matrices. A useful variant of this theorem from [21], which we will use, is as follows:

Theorem 4.2.1 (corollary 3 of [21]). *Let L and L' be two $n \times n$ symmetric matrices with eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ and $\lambda'_1 \leq \lambda'_2 \leq \cdots \leq \lambda'_n$, respectively. Fix $j \in \{1, 2, 3.., n\}$, and assume that $\min\{\lambda_j - \lambda_{j-1}, \lambda_{j+1} - \lambda_j\} > 0$, where $\lambda_0 = -\infty$ and $\lambda_{n+1} = \infty$. If $v, v' \in \mathbb{R}^n$ such that $Lv = \lambda_j v$ and $L'v' = \lambda'_j v'$, $v'v^T > 0$, and $||v|| = ||v'|| = 1$, then:*

$$||v - v'|| \leq \frac{2^{1.5}||L - L'||_{op}}{\min\{\lambda_j - \lambda_{j-1}, \lambda_{j+1} - \lambda_j\}}$$

Since we know the Fiedler value of R is isolated, it fits in the above theorem as L. Lets choose $v, v' \in \mathbb{R}^n$ such that $Rv = \lambda_2 v$ and $Av' = \lambda'_2 v'$. Then, theorem 4.2.1 assures that:

$$||v - v'|| \leq \frac{2^{1.5}||L_R - L_A||_{op}}{\min\{\lambda_3 - \lambda_2, \lambda_2\}}$$

where $0 = \lambda_1 < \lambda_2 < \lambda_3$ are the first eigenvalues of L_R and λ'_2 is the Fiedler value of A.
Since R and A are close, so are their Laplacians, and by the above inequality, their corresponding Fiedler vectors.

The closeness of v and v' implies the closeness of λ_2 and λ_2', since:

$$L_R v = \lambda_2 v \implies v^T L_R v = \lambda_2$$
$$L_A v' = \lambda_2' v \implies v'^T L_A v' = \lambda_2'$$

Which implies

$$\begin{aligned}
|\lambda_2 - \lambda_2'| &= |v^T L_R v - v'^T L_A v'| \\
&= |v^T L_R v - v'^T (L_R - (L_R - L_A)v'| \\
&= |v^T L_R v - v'^T L_R v' + v'^T (L_R - L_A)v'| \\
&= |v^T L_R v + v^T L_R v' - v^T L_R v' - v'^T L_R v' + + v'^T (L_R - L_A)v'| \\
&= |v^T L_R (v - v') + (v^T - v'^T) L_R v' + v'^T (L_R - L_A)v'| \\
&\leq |v^T L_R (v - v')| + |(v^T - v'^T) L_R v'| + |v'^T (L_R - L_A)v'|
\end{aligned}$$

Which is small because $v - v'$ and $L_R - L_A$ are small.

With the same argument, we can show that λ_3 and λ_3' are also close. If these are *close enough*, this proves that λ_2' is simple.

Note that by Theorem 2.3.15, the Fiedler vector of R is monotone. This means that if $E = A - R$ is small enough, A's Fiedler vector is also monotone: Let the Fiedler vector of R and A be v and v' respectively. Then since the map $f((x_1, \ldots, x_n)) = (x_2 - x_1, \ldots, x_n - x_{n-1}$ is continuous and $f(v)$ has all the entries positive or negative, then given v and v' are close enough, $f(v')$ should be the same, meaning that v' is also monotone.

In conclusion, if R and A are close enough, the Fiedler vector of A is simple and monotone. This means that we can use the order induced by sorting A's Fiedler vector to effectively sequence the data.

In the following chapter, we will show that by increasing the number of data points and samples taken at each time point, we can get arbitrarily close to an R-matrix that closely approximates the matrix of correlations. This means that, given enough data points and samples, the Fiedler vector of the empirical data can be both simple and monotonic.

Chapter 5

Addressing the limitations of finite data: How many data points and samples are enough?

In this chapter, we investigate how the fact that we have a limited number of data points and samples per data point may affect our results. Specifically, in the first main theorem of this chapter, theorem 5.3.8, we prove that the Laplacian of the empirical data points converges to the Laplacian of the ground truth curve as the number of samples per data point and the number of data points approach infinity. We also quantify the *rate* of this convergence, which is useful in experiments that require knowing how many samples per data point and how many data points are needed to stay within a specific error bound. In the next main theorem of this chapter, Theorem 5.4.7, we prove that having a specific bound on a generalized notion of curvature (the Haantjes curvature [18]) would lead to the Fiedler function associated with the graphon of the curve being simple.

The structure of this chapter is as follows: In the first section, we state some remarks about the setting of the problem. In section two, we explain how we can connect the discrete nature of our limited data points to the continuous setting of the curve they lie on. In the third section, we set the ground for, and state and prove Theorem 5.3.6 and consequently 5.3.8, namely the convergence of the Laplacian in operator norm. In the last section, we state and prove a condition on the curvature for the Fiedler function to be simple in Theorem 5.4.7.

5.1 Overview of the problem in a discrete setting

For the purpose of this chapter, the setting of our problem is as follows:

- There is a curve in the Wasserstein space $\mu : [0,1] \to (\mathcal{W}, W_2)$, representing the development of the embryos
- We have no information about this curve, other than n data points on the curve $\{x_1, x_2 \ldots, x_n\}$
- We may assume that $x_i = \mu_{t_i} := \mu(t_i)$, $1 \leq i \leq n$, and $0 \leq t_1 < t_2 < \ldots t_n \leq 1$
- Each x_i represents an embryo, and in practice, it is approximated by taking m samples. We assume that each $x_i = \mu_{t_i}$ is approximated with $\hat{\mu}_{t_i} = \frac{1}{m}\sum_{k=1}^{m} \delta_{X_k}$, where X_k is a

random variable with distribution μ_{t_i}.

While we do use the notation $x_i = \mu(t_i)$, suggesting that for a particular data point x_i we know the curve μ and the time t_i, in practice we know neither: We don't have an explicit description of μ, and hence, we don't know what order it induces. What we do have is an unordered set of size n, and the fact that they lie on *some* curve μ. However, using this notation helps us analyze the problem. As in the paper [2], we can use the spectral seriation method to try and find an ordering for the data. The question that arises here is, will that ordering be consistent with the ordering induced by the curve, given that we have "enough" samples and data points? The spectral seriation method uses the Fiedler vector of the Laplacian of the correlation matrix. To answer the above question, we should check whether or not the Laplacians for n data points "converge" in some sense, and if they converge to the Laplacian associated with the curve μ. We will restate this question more rigorously in the following section.

5.2 From the discrete setting to the continuous setting

In this section, we pick up the notation and the concept of connecting the discrete setting of data points to the continuous setting of the curve from [13].

As before, $(\mathcal{W}, W_2)$ is the Wasserstein space with the Wasserstein-2 distance. $\mu : [0,1] \to \mathcal{W}$ is a continuous curve, on which we have n data points $\mu_{t_i} := \mu(t_i)$, $1 \leq i \leq n$, and $0 \leq t_i \leq 1$.

To use the spectral seriation method, We first need to define a correlation between the points on μ using the Wasserstein distance as follows:

$$\omega : [0,1]^2 \to [0,1]$$
$$\omega(s,t) = e^{\frac{-W_2^2(\mu(s),\mu(t))}{\sigma^2}}$$

In practice, the only data points we do have some information about are $\mu(t_i)$'s; We can approximate ω using these data points. Let $(I_i)_{i=1}^n$ be a partition of $[0,1]$ such that $t_i \in I_i$. Here we assume that the data set is large enough and distributed in a way that $|I_i| \leq \frac{2}{n}$ for each $1 \leq i \leq n$. Now, let's define:

$$w_n : [0,1]^2 \to [0,1]$$
$$w_n(s,t) = \sum_{i=1}^{n} \sum_{j=1}^{n} \omega(t_i, t_j) \mathbb{1}_{I_j} \mathbb{1}_{I_i}$$

Now, recall that as we discussed in the previous section, μ_{t_i} are approximated by $\hat{\mu}_{t_i}$, and those are the information that we have in practice. Let's assume there is another curve $\hat{\mu} : [0,1] \to X$, on which n points $\hat{\mu}_{t_i} := \hat{\mu}(t_i)$, $1 \leq i \leq n$ lie, Where $\hat{\mu}_{t_i} = \frac{1}{m} \sum_{k=1}^{m} \delta_{X_k}$, X_k's

being random variables with distribution μ_{t_k}. Now, define:

$$\hat{\omega} : [0,1]^2 \to [0,1]$$
$$\hat{\omega}(s,t) = e^{\frac{-W_2^2(\hat{\mu}(s),\hat{\mu}(t))}{\sigma^2}}$$

As before, we can approximate $\hat{\omega}$ with the data we have. Define:

$$\hat{\omega}_n : [0,1]^2 \to [0,1]$$
$$\hat{\omega}_n(s,t) = \sum_{i=1}^{n}\sum_{j=1}^{n} \hat{\omega}(t_i,t_j)\mathbb{1}_{I_j}\mathbb{1}_{I_i}$$

Also, define p^n and $\hat{p}^n$ to be $n \times n$ matrices such that:

$$p_{ij}^n = \omega(t_i,t_j) \tag{5.1}$$
$$\hat{p}_{ij}^n = \hat{\omega}(t_i,t_j) \tag{5.2}$$
$$\tag{5.3}$$

Let L_n and $\hat{L}_n$ be their Laplacian matrices, respectively. Finally, remember that each $f \in L^2([0,1])$, the Laplacian of $a : [0,1]^2 \to [0,1]$ is:

$$\mathbb{L}_a(f) = \int_0^1 a(x,y)(f(x) - f(y))dy$$

Let $\mathbb{L}_n$ and $\hat{\mathbb{L}}_n$ be the Laplacians of ω_n and $\hat{\omega}_n$ respectively, and let $\mathbb{L}$ be that of w. Since in practice, we are only able to calculate $\hat{\mathbb{L}}_n$, it is worth investigating to see if $\mathbb{L}_n$ and $\hat{\mathbb{L}_n}$ converge to each other in (operator) norm as m converges to infinity for a fixed n. Also, we will check that they approach in norm to $\mathbb{L}$.

For the rest of this thesis, let μ be the ground truth curve defined as in this section, $\hat{\mu}$ the curve where the empirical data lie on, and t_i the time points we have information about. Also let n be the number of time points and m the number of samples per time point. We also make the following assumption:

Assumption 1. *The exists a partition $\{I_i\}_{i=1}^n$ of $[0,1]$ such that $t_i \in I_i$ and length of each I_i is less that $\frac{2}{n}$.*

5.3 Approximating the Laplacian using finitely many points: how well can we do that?

In this section, we shall first prove that for a fixed n, $\hat{\mathbb{L}}_n$ and $\mathbb{L}_n$ converge in operator norm with a probability 1 as m approaches infinity. Remember that $\hat{\mathbb{L}}_n$ is a random operator, so we can only talk about the probability of the convergence. Secondly, we will prove that $\mathbb{L}_n$

converges to $\mathbb{L}$ as n approaches infinity, provided the curve is so that the associated graphon is Lipschitz. As a result, provided that the graphon is Lipschitz, the Laplacian associated with our empirical data approaches that of the ground truth curve as m and n approach infinity with probability 1.

Remember that the operator norm of a linear operator T on $L^2([0,1])$ is defined as follows:

$$\|\mathbb{T}\| := \sup_{f \in L^2([0,1]) \text{ s.t. } \|f\|=1} \|\mathbb{T}f\|$$

From now on, by norm, we are referring to the operator norm. To prove that $\hat{\mathbb{L}}_n$ and $\mathbb{L}_n$ converge in norm, we need to be able to talk about the probability of them being close. Theorems 5.3.1 and 5.3.4 below, from the paper [20], provide us with tools to do so.

As we will soon see, to talk about the norm of the difference of $\hat{\mathbb{L}}_n$ and $\mathbb{L}_n$, we need to talk about $W_2(\nu, \hat{\nu}_m)$, where ν is a (Borel) measurable measure in the Wasserstein space and $\hat{\nu}_m$ is its empirical approximation by m samples. Since our Wasserstein space is set on a compact subset of $\mathbb{R}^d$, then it is a known fact ([20]) that

$$W_p(\nu, \hat{\nu}_m) \to 0 \quad \nu\text{-a.s.}$$

Where W_p is the Wasserstein-p distance. In the following text, we will state some results from [20] involving the Wasserstein-p distance, and at the end, we will come back to $p = 2$.

While the above equations guarantee that ν and $\hat{\nu}_m$ can get arbitrarily close as $m \to \infty$, we would like to *quantify* the rate of this convergence. Quantifying the rate of convergence is useful in practice since it gives us a tool to know how many samples are needed to get a desired approximation. Since $\hat{\nu}_m$ is a random measure, it's easier to first talk about the probability of $W_p^p(\nu, \hat{\nu}_m)$ being close to its expectation. The following theorem form [20] is useful for this purpose:

Theorem 5.3.1 (Proposition 20 of [20])**.** *Let ν be a probability measure and $\hat{\nu}_m$ its empirical approximation by m samples. Then For all $0 \leq p < \infty$,*

$$\mathbb{P}\left[W_p^p(\nu, \hat{\nu}_m) \geq \mathbb{E}W_p^p(\nu, \hat{\nu}_m) + t\right] \leq \exp\left(-2mt^2\right).$$

Now that we have an upper bound for the probability of $W_p^p(\nu, \hat{\nu}_m)$ being close to its expectation, we would like to talk about the expectation. Theorem 5.3.4 below, from [20], provides a bound for this expectation. Before stating this theorem, we need to introduce a few definitions and theorems from [20].

Definition 5.3.2 (Definition 2 of [20])**.** *Given a set $S \subseteq X$, the ε-covering number of S, denoted $\mathcal{N}_\varepsilon(S)$, is the minimum m such that there exist m closed balls $B_1, \ldots, B_m$ of diameter ε such that $S \subseteq \bigcup_{1 \leq i \leq m} B_i$.*

In some sense, the covering number $\mathcal{N}_\varepsilon(S)$ measures how many points on a space S are

needed to capture (and possibly reconstruct) the essential essence of the space, as long as we don't distinguish points within ε of each other. Sometimes, we need to only concern ourselves with a part of the set S, in which case the following definition comes in handy:

Definition 5.3.3 (Definition 3 of [20]). *Given a measure ν on X, the (ε, τ)-covering number is*

$$\mathcal{N}_\varepsilon(\nu, \tau) := \inf \{\mathcal{N}_\varepsilon(S) : \nu(S) \geq 1 - \tau\}$$

and the (ε, τ)-dimension is

$$d_\varepsilon(\nu, \tau) := \frac{\log \mathcal{N}_\varepsilon(\nu, \tau)}{-\log \varepsilon}$$

And now we can use these definitions to state the following theorem, bounding the expectation of $W_p^p(\nu, \hat{\nu}_m)$.

Theorem 5.3.4 (Proposition 5 of [20]). *Let $p \in [1, \infty)$. Suppose there exists an $\varepsilon' \leq 1$ and $s > 2p$ such that*

$$d_\varepsilon\left(\nu, \varepsilon^{\frac{sp}{s-2p}}\right) \leq s$$

for all $\varepsilon \leq \varepsilon'$. Then

$$\mathbb{E}\left[W_p^p(\nu, \hat{\nu}_m)\right] \leq C_1 m^{-p/s} + C_2 m^{-1/2},$$

where

$$C_1 = 3^{\frac{3sp}{s-2p}+1}\left(\frac{1}{3^{\frac{8}{2}-p} - 1} + 3\right), \text{ and } C_2 = \left(27/\varepsilon'\right)^{\frac{s}{2}}.$$

Now we have the means to state and prove Lemma 5.3.5. But before that, we need to make the following assumption:

Assumption 2. *There exists an $\varepsilon' \leq 1$ and $s > 2p$ such that*

$$d_\varepsilon\left(\mu(t), \varepsilon^{\frac{sp}{s-2p}}\right) \leq s$$

for all $\varepsilon \leq \varepsilon'$ and $t \in [0, 1]$.

In what follows, let $|||.|||$ be the operator norm for the Laplacians of the continuous graphons, and let $||.||$ be the operator norm for matrices.

We are finally ready to state a lemma that will be useful in proving the first main theorem of this chapter, that the Laplacian of the empirical data converges to that of the data points in operator norm, when the number of samples per time point approaches infinity.

Lemma 5.3.5. *Assume that Assumption 2 holds, and let $\hat{\mathbb{L}}_n$ and $\mathbb{L}_n$ be defined as before. Then the following inequality*

$$|||\hat{\mathbb{L}}_n - \mathbb{L}_n||| \leq 8D(\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t})$$

holds with probability $1-exp(-2mt^2)$*, where* s*,* D*,* C_1 *and* C_2 *are positive constants independent of* m *and* n*.*

Proof. First of all, since the operators have finite rank, note that:

$$\begin{aligned}|||\hat{\mathbb{L}}_n - \mathbb{L}_n||| &\leq \frac{2}{n}||\hat{L}_n - L_n|| \\ &\leq \frac{2}{n}(||\hat{D}^n - D^n|| + ||\hat{P}^n - P^n||)\end{aligned}$$

Where $\hat{D}^n$ and D^n are the degree matrices. From now on, we drop n from the matrices for convenience, but keep in mind that in the following text the matrices have size $n \times n$.

Using the above inequality, we can see that it is enough to prove that $\hat{D}$ and D converge in norm normalized by n, and so do $\hat{P}$ and P. To do so, we find upper bounds for the norms of their differences, and argue that the upper bound approaches 0 with a high probability as m approaches infinity.

Note that by triangle inequality, we have:

$$\begin{aligned}W_2(\hat{\mu}_{t_i}, \hat{\mu}_{t_j}) - W_2(\mu_{t_i}, \mu_{t_j}) &\leq W_2(\hat{\mu}_{t_i}, \mu_{t_i}) + W_2(\hat{\mu}_{t_j}, \mu_{t_j}) \\ W_2(\mu_{t_i}, \mu_{t_j}) - W_2(\hat{\mu}_{t_i}, \hat{\mu}_{t_j}) &\leq W_2(\mu_{t_i}, \hat{\mu}_{t_i}) + W_2(\mu_{t_j}, \hat{\mu}_{t_j})\end{aligned}$$

So,

$$|W_2(\hat{\mu}_{t_i}, \hat{\mu}_{t_j}) - W_2(\mu_{t_i}, \mu_{t_j})| \leq W_2(\mu_{t_i}, \hat{\mu}_{t_i}) + W_2(\mu_{t_j}, \hat{\mu}_{t_j})$$

Define:

$$D = \max_{(a,b)\in A} \Big\{ \frac{e^{\frac{-a^2}{\sigma^2}} - e^{\frac{-b^2}{\sigma^2}}}{a-b} \Big\}$$

Where

$$A = \Big\{ \big(W_2(\mu_s, \mu_t), W_2(\hat{\mu}_s, \hat{\mu}_t)\big) : s, t \in [0,1] \Big\}$$

Since A is compact, D exists and then by the above calculations we have:

$$|(\hat{p} - p)_{ij}| \leq D(W_2(\mu_{t_i}, \hat{\mu}_{t_i}) + W_2(\mu_{t_j}, \hat{\mu}_{t_j})) \tag{5.4}$$

Now, let s be as in assumption 2. Then theorem 5.3.4 holds for each $\mu(t_i)$, and so the

following holds with probability $1 - exp(-2mt^2)$:

$$\begin{aligned} W_2^2(\mu_{t_i}, \hat{\mu}_{t_i}) &< E[W_2^2(\mu_{t_i}, \hat{\mu}_{t_i})] + t \\ &\leq C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t \end{aligned}$$

So,

$$W_2(\mu_{t_i}, \hat{\mu}_{t_i}) < \sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t}$$

which together with 5.4 means:

$$\begin{aligned} |(\hat{p} - p)_{ij}| \leq 2D\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t} \implies \\ ||\hat{p} - p|| \leq 2nD\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t} \end{aligned}$$

Also,

$$\begin{aligned} |(\hat{D} - D)_{ii}| &= |\sum_{j=1}^{n} (\hat{p} - p)_{ij}| \\ &\leq \sum_{j=1}^{n} |(\hat{p} - p)_{ij}| \\ &\leq 2nD\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t} \end{aligned}$$

Since $\hat{D} - D$ is a diagonal matrix, this means that

$$||\hat{D} - D|| \leq 2nD\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t}$$

So, overall we have

$$\begin{aligned} |||\hat{\mathbb{L}}_n - \mathbb{L}_n||| &\leq \frac{2}{n} ||\hat{L}_n - L_n|| \\ &\leq \frac{2}{n} (||\hat{D}^n - D^n|| + ||\hat{P}^n - P^n|| \\ &\leq 8D\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t} \end{aligned}$$

with probability $1 - exp(-2mt^2)$. □

Since t is arbitrary, we can choose t so that the probability in Lemma 5.3.5 is as large as we want. specifically, we have the following theorem:

Theorem 5.3.6. *For a fixed n, $\hat{\mathbb{L}}_n$ and $\mathbb{L}_n$ converge in norm with probability 1 as m approaches infinity.*

Proof. In the previous theorem, Letting $t = m^{\frac{-1}{4}}$, we get that

$$|||\hat{\mathbb{L}}_n - \mathbb{L}_n||| \leq 8D\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + m^{\frac{-1}{4}}}$$

with probability $1 - e^{-2\sqrt{m}}$. As a result as m approaches infinity, $|||\hat{\mathbb{L}}_n - \mathbb{L}_n|||$ approaches 0 with probability 1.

□

The following theorem from [13] shows that if the graphon is Lipschitz, $\mathbb{L}_n$ approaches $\mathbb{L}$ in norm as n goes to infinity.

Theorem 5.3.7 (Theorem 4.2 of [13], modified)**.** *Suppose that the graphon w is Lipschitz, in the sense that there exists a constant $K > 0$ such that*

$$\left|w(x,y) - w\left(x',y'\right)\right| \leq K\left(\left|x - x'\right| + \left|y - y'\right|\right)$$

for all pairs $(x,y) \in [0,1]^2, (x',y') \in [0,1]^2$. Then,

$$|||\mathbb{L}_n - \mathbb{L}||| \leq \frac{64K}{n}$$

Proof. Let $f \in L^2([0,1])$ such that $\|f\| = 1$,

$$\begin{aligned}
\|\mathbb{L}_n f - \mathbb{L}f\|^2 &= \int_0^1 \left(\mathbb{L}_n f - \mathbb{L}f\right)^2(x)dx \\
&= \int_0^1 \left(\int_0^1 \left(f(x) - f(y)\right)\left(w_n(x,y) - w(x,y)\right)dy\right)^2 dx \\
&\leq \int_0^1 \left(\int_0^1 (f(x) - f(y))^2 dy\right)\left(\int_0^1 \left(w_n(x,y) - w(x,y)\right)^2 dy\right)dx \\
&= \sum_{i=1}^n \int_{I_i} \left(\int_0^1 (f(x) - f(y))^2 dy\right)\left(\int_0^1 \left(w_n(x,y) - w(x,y)\right)^2 dy\right)dx
\end{aligned}$$

Whereas in assumption 1, $\{I_i\}_{i=1}^n$ is a partition of $[0,1]$ such that $t_i \in I_i$ and length of each I_i is less than $\frac{2}{n}$. Now, let $x \in I_i$. then,

$$\begin{aligned}
\int_0^1 \left(w_n(x,y) - w(x,y)\right)^2 dy &= \sum_{j=1}^n \int_{I_j} \left(w\left(t_i,t_j\right) - w(x,y)\right)^2 dy \\
&\leq \sum_{j=1}^n \int_{I_j} K^2\left(|t_i - x| + |t_j - y|\right)^2 dy \\
&\leq \sum_{j=1}^n \int_{I_j} \frac{16K^2}{n^2} dy \\
&= \frac{16K^2}{n^2}
\end{aligned}$$

Also,

$$\begin{aligned}\int_0^1 (f(x)-f(y))^2 dy &\leq \int_0^1 2\left(f(x)^2+f(y)^2\right) dy \\ &= 2f(x)^2+2\end{aligned}$$

hence using the above inequality will yield

$$\begin{aligned}\|\mathbb{L}_n f-\mathbb{L}f\|^2 &\leq \frac{32K^2}{N^2}\sum_{i=1}^{n}\int_{I_i}\left(f(x)^2+1\right)dx \\ &= \frac{64K^2}{n^2}\end{aligned}$$

As a result,

$$\|\mathbb{L}_n f-\mathbb{L}f\| \leq \frac{64K}{n}$$

Since this inequality holds for all functions $f \in L^2([0,1])$ with unit norm, the theorem is proved. □

Now, we can combine lemma 5.3.5 and theorem 5.3.7 to conclude that the Laplacian of the empirical data converges to that of the ground truth curve.

Theorem 5.3.8 (Main Theorem)**.** *Suppose $\mu : [0,1] \to \mathcal{W}$ is a continuous curve and the graphon $w(x,y) = e^{\frac{-W_2^2(\mu(x),\mu(y))}{\sigma^2}}$ is Lipschitz. Also, let*

$$D = \max_A \left\{\frac{e^{-a^2}-e^{-b^2}}{a-b} : (a,b) \in A\right\}$$

Where

$$A = \left\{(W_2(\mu_s,\mu_t), W_2(\hat{\mu}_s,\hat{\mu}_t)) : s,t \in [0,1]\right\}$$

Then,

$$|||\hat{\mathbb{L}}_n - \mathbb{L}||| \leq \frac{64K}{n} + 8D\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + t}$$

With probability $1-exp(-2mt^2)$. Also,

$$\lim_{n,m\to\infty} |||\hat{\mathbb{L}}_n - \mathbb{L}||| = 0$$

with probability 1

Proof. The first result follows from Lemma 5.3.5 and 5.3.7 and triangle inequality. By letting $t = m^{\frac{-1}{4}}$, we get that

$$|||\hat{\mathbb{L}}_n - \mathbb{L}||| \leq \frac{64K}{n} + 8D\sqrt{C_1 m^{-\frac{2}{s}} + C_2 m^{-\frac{1}{2}} + m^{\frac{-1}{4}}}$$

with probability $1 - exp(-2\sqrt{m})$.

Now, letting m and n approach infinity would yield the result. □

5.4 Simplicity of the Fiedler Function and the Curvature

As we saw in the last section, if w is Lipschitz, the Laplacian of $\hat{w}_n$, derived from the approximation of data by m samples, approaches the Laplacian of w which is associated with the curve. This is useful since $\hat{w}_n$ is the information we have in practice. Now, if the curve is continuous and an R-curve, so that the matrix of correlations is irreducible and a pre-R matrix, then the Fiedler value is simple by Theorem 2.3.15, and if it doesn't have repeated values, then by Theorem 2.3.9, the permutation induced by sorting the values in the Fiedler vector can be used to sequence the data. In that case, the Laplacian of $\hat{w}_n$ approaching that of the curve would mean that the spectral seriation method would also work with the empirical data.

For a general continuous curve, which we might not know whether or not is an R-curve, it is still worth investigating whether or not the Fiedler value is simple. This is because *if* it was an R-curve, then by theorem 2.3.15 its Fiedler value would be simple. Moreover, if the curve is in some sense *close* to an R-curve, we would expect the permutation induced by the Fiedler vector to still be somewhat consistent with the correlations. Since the Fiedler value of an R-curve is simple, the Fiedler value of the unknown curve not being simple is a sign that it's not close enough to being an R-curve, and hence the spectral seriation method wouldn't be suitable. More importantly, the Fiedler vector being simple means that there is a positive Eigen gap, and as a result, small perturbations will keep the Fiedler vector simple.

In this chapter, we find an explicit condition regarding the curvature under which the Fiedler function (that is, the continuous version of the Fiedler vector) of the curve is simple. We first refer to [12] to investigate a certain class of curves, including geodesics, for which the Fiedler function is simple. Then, we will explore how deviation from geodesics may affect the Fiedler function being simple.

5.4.1 Investigating a Class of Graphons with simple Fiedler Functions

This section is a summary of chapter 5.2 of [12], where it is proved that under certain assumptions, the Fiedler function of a graphon is simple. Consider the following assumptions on a graphon w:

Assumption 3. *$w(x, y) > 0$ for almost all $(x, y) \in [0, 1]^2$.*

Assumption 4. *Assume that for all $x \in [0, 1], w(x, \cdot)$ is continuous almost everywhere. Moreover, assume that the partial derivative $\partial w(x, \cdot)/\partial x$ exists and is non zero and bounded almost everywhere.*

Assumption 5. *Let d be the degree function associated with w defines as $d(x) = \int_0^1 w(x, y)dy$. Assume that the set $N := \{x \in [0, 1] : d'(x) = 0\}$ is countable.*

Assumption 6. *Let*

$$\mathcal{H} = \left\{ f \in L^2([0,1]) : \int_0^1 f(t)dt = 0 \text{ and } \int_0^1 f^2(t)dt = 1 \right\}$$

and $\mathbb{L}$ *be the graphon-Laplacian operator associated with* w*. Assume that*

$$\inf_{f \in \mathcal{H}} \langle \mathbb{L}f, f \rangle < \min_{x \in [0,1]} d(x).$$

The following two theorems in [12] (whose proofs are also due to [12]) show that provided the above assumptions are satisfied, the Fiedler function is simple and monotone.

Theorem 5.4.1 (Theorems 1 and 2 of [12])**.** *Let* w *be an R-graphon satisfying assumptions 3, 4, 5 and 6. Then there exists a Fiedler function* $\phi \in \mathcal{C}^1([0,1])$ *such that* $\phi'(x) > 0$ *for all* $x \in [0,1]$*, and the Fiedler value* λ_2 *is simple.*

Hence, provided the graphon associated with our curve satisfies the above assumptions, the Fiedler function is simple and monotone. But what curves produce such graphons? The following theorem in [12] shows that a certain class of curves, including geodesics, satisfy these conditions.

Theorem 5.4.2 (Theorem 5.2.7 of [12])**.** *Let* $R : [0,1] \to [0,1]$ *be a* $\mathcal{C}^1$ *function such that* $R'(x) < 0$ *for all* $x \in (0,1)$*. Define the graphon* $w(x,y) = R(|x-y|)$ *for all* $x, y \in [0,1]$*, then* w *is a* $\mathcal{R}$*-graphon satisfying Assumptions 3, 4, 5, 6.*

In other words, if the graphon is a smooth function of only the distance (and not the positions), and it has a reverse relation with the distance in the sense that it's larger when the distance is smaller, then the graphon satisfies all the above conditions and hence has a simple and monotone Fiedler function. Note that w, stated as above, is actually an R-graphon:

Let $x, y, z \in [0,1]$; if $y < z < x$ then $|x-y| > |x-z|$ and since R has negative derivative everywhere, $R(|x-y|) \leq R(|x-z|)$. As a result, $w(x,y) \leq w(x,z)$. Using the same argument, if $x < y < z$, then $w(x,y) \geq w(x,z)$.

Proof of Theorem 5.4.2, due to [12]. Assumptions 3 and 4 hold since R is differentiable and satisfies $R'(x) < 0$ for almost all $x \in [0,1]$.

To show that assumption 5 also holds, note that

$$\begin{aligned} d(x) &= \int_0^1 R(|x-y|)dy \\ &= \int_0^x R(x-y)dy + \int_x^1 R(y-x)dy \\ &= \int_0^x R(z)dz + \int_0^{1-x} R(z)dz \end{aligned}$$

As a result, $d'(x) = R(x) - R(1-x)$, meaning that $d'(x) = 0$ if and only if $x = 1/2$ (since $R'(x) < 0$ for all x).

To show that assumption 6 also holds, first note that $\min_{x\in[0,1]} d(x) = d(0) = \int_0^1 R(x)dx$ and so, assumption 4 is equivalent to

$$\inf_{f\in\mathcal{H}} \left(\int_0^1 \int_0^1 R(|x-y|)(f(x)-f(y))^2 dxdy \right) < 2\int_0^1 R(x)dx$$

Let $f \in \mathcal{H}$, note that

$$\int_0^1 \int_0^1 R(|x-y|)(f(x)-f(y))^2 dxdy - 2\int_0^1 R(x)dx = 2D_f(R)$$

where

$$D_f(R) := \int_0^1 \int_x^1 R(y-x)(f(x)-f(y))^2 dydx - \int_0^1 R(z)dz.$$

So, it would suffice to find a function $f \in \mathcal{H}$ such that $D_f(R) < 0$.

To find such f, consider the following family of functions. Let $t \in (0, 1/2)$, and for all $x \in [0,1]$ let

$$f(x) := f_t(x) := \frac{1}{\sqrt{2t}}(I\{x < t\} - I\{x > 1-t\})$$

where I is the indicator. It can be easily checked that $f \in \mathcal{H}$. In addition,

$$2t(f(x)-f(y))^2 = I\{x<t\}I\{y\geq t\} + I\{x\leq 1-t\}I\{y>1-t\} + 2I\{x<t\}I\{y\geq 1-t\}$$

if $0 < x < y < 1$, letting $t \downarrow 0$, we have

$$\int_0^1 \int_x^1 R(y-x)(f(x)-f(y))^2 dydx = \frac{J_1 + J_2 + 2J_3}{2t},$$

where in J_1, J_2 and J_3 the integral is computed on the three above intervals:

$$\begin{aligned}
J_1 &:= \int_0^t \int_t^1 R(y-x)dydx \\
&= \int_0^t \int_{t-x}^{1-x} R(u)dudx \\
&= t\int_0^1 R(z)dz - \int_0^t \left(\int_0^{t-x} R(z)dz + \int_{1-x}^1 R(z)dz \right) dx \\
&= t\int_0^1 R(z)dz - \int_0^t \left(\int_0^{t-x} [R(0)+o_t(1)] + \int_{1-x}^1 [R(1)+o_t(1)] \right) dx \\
&= t\int_0^1 R(z)dz - \frac{t^2}{2}[R(0)+R(1)+o_t(1)]
\end{aligned}$$

here $o_t(1)$ is a quantity that goes to 0 as t goes to 0 . Similarly,

$$\begin{aligned} J_2 &:= \int_{1-t}^{1}\int_{0}^{1-t} R(y-x)dxdy \\ &= t\int_0^1 R(z)dz - \frac{t^2}{2}\left[R(0)+R(1)+o_t(1)\right]; \end{aligned}$$

and

$$\begin{aligned} J_3 &:= \int_0^t \int_{1-t}^{1} R(y-x)dydx \\ &= t^2\left[R(1)+o_t(1)\right]. \end{aligned}$$

Collecting all the pieces, we have

$$D_f(R) = \frac{J_1+J_2+2J_3}{2t} - \int_0^1 R(z)dz = \left[R(1)-R(0)+o_t(1)\right]t/2 < 0$$

for small enough $t > 0$ (because $R(1) < R(0)$). So, assumption 6 is also satisfied.

□

Note that we already knew that w as defined in 5.4.2 is an R-graphon, and so, taking a finite number of data points as in section 5.3 to form P^n would create an R-matrix, for which the Fiedler value is simple. This means that the Fiedler function for w_n, being a finite rank graphone produced by P^n, is also simple and since the Laplacian L_n, the Laplacian of w_n approaches L in operator norm, the Laplacian of w, we could easily use the Davis-Kahan theorem to prove the Fiedler vector associated with w is also simple. However, the above approach gives us an idea of how to deal with non-R graphons, as we will see in the next subsection.

5.4.2 Deviation from Geodesics: The effect of curvature on Fiedler Function being simple

As before, the graphon we work with is $e^{\frac{-W_2^2(\mu_x,\mu_y)}{\sigma^2}}$. As we saw in the previous chapter, if the curve is such that $W_2(\mu_x,\mu_y)$ doesn't depend on the positions of x and y and is solely a function of their distance, then the Fiedler function is simple and monotone. One of the ways this can happen is for the curve to be such that $W_2(\mu_x,\mu_y) = c|x-y|$ for some constant c, i.e., the curve is a geodesic. In this chapter, we explore what happens if the curve deviates from being a geodesic. Specifically, we find a sufficient condition on the curvature guaranteeing the Fidler function being simple.

Definition 5.4.3. *Let $\mu : [0,1] \to \mathcal{W}$ be a continuous curve, parameterized so that the arc length in $\mathcal{W}$ between the points $\mu(x)$ and $\mu(y)$ is $c|x-y|$. Then we say the function $h(x,y) = c|x-y| - W_2(\mu(x),\mu(y))$ is the geodesic deviation of μ.*

Geodesics in a general metric space are equivalent to straight lines in the Euclidean space. Deviation from a straight line in the Euclidean space amounts to positive curvature, and as

it turns out, the geodesic deviation for a curve in a metric space is also related to (one of) the generalized notions of curvature in a metric space. In this section, we use the following definition of metric curvature from [18].

Definition 5.4.4 (Definition 3.4 of [18])**.** *Let (M, d) be a metric space and let $\mu : I = [0, 1] \to M$ be a homeomorphism (so no self-intersections), and let $p, q, r \in \mu(I), q, r \neq p$. Denote by $\widehat{qr}$ the arc of $\mu(I)$ between q and r, and by qr line segment from q to r.*

Then μ has Haantjes (or Finsler-Haantjes) curvature $\kappa_H(p)$ at the point p iff:

$$\kappa_H^2(p) = 24 \lim_{q,r \to p} \frac{l(\widehat{qr}) - d(q, r)}{(l(\widehat{qr}))^3}$$

where "$l(\widehat{qr})$" denotes the arc length - with respect to the intrinsic metric induced by d - of $\widehat{qr}$.

From now on in this text, by curvature k_x, we mean Haantjes curvature $\kappa_H(x)$ at point x.

In the rest of this text, let $\mu : [0, 1] \to \mathcal{W}$ be a smooth curve such that $W_2(\mu(x), \mu(y)) = c|x - y| - h(x, y)$, where h is its geodesic deviation. Also, we assume that the curve is parameterized so that the arc length between points $\mu(x)$ and $\mu(y)$ is $c|x - y|$. Then, the curvature of the curve, using the above definition is

$$k_x = \lim_{x \to y} \sqrt{\frac{24h(x, y)}{c^3|x - y|^3}}$$

As we saw before, in the case of a geodesic where h is zero everywhere and hence so is the curvature, the associated Fiedler function is simple. To explore the cases where the curvature is non-zero, it makes sense to guess that perhaps a small upper bound on curvature would lead to the Fiedler function being simple, i.e. not being so much deviated from a geodesic, should preserve the properties of a geodesic, including the Fiedler function being simple.

Before we state the main theorem, note that we can write the graphon as two parts: one that only depends on the distance, which has already been proven to have a simple Fiedler function, and one that originates from the geodesic deviation being non-zero. The calculations go as follows:

$$\begin{aligned}
W_2^2(\mu_x, \mu_y) &= c^2(x - y)^2 + h^2(x, y) - 2c|x - y|h(x, y) && \implies \\
w = e^{\frac{-W_2^2(\mu_x, \mu_y)}{\sigma^2}} &= e^{\frac{-c^2(x-y)^2 - h^2(x,y) + 2c|x-y|h(x,y)}{\sigma^2}} && \implies \\
e^{\frac{-W_2^2(\mu_x, \mu_y)}{\sigma^2}} &= e^{\frac{-c^2(x-y)^2}{\sigma^2}} + e^{\frac{-c^2(x-y)^2}{\sigma^2}}(g(x, y) - 1) \\
&= R(|x - y|) + E(x, y)
\end{aligned}$$

where

$$g(x,y) = e^{\frac{h(x,y)(2c|x-y|-h(x,y))}{\sigma^2}} \tag{5.5}$$

$$R(z) = e^{\frac{-c^2z^2}{\sigma^2}} \tag{5.6}$$

$$E(x,y) = e^{\frac{-c^2(x-y)^2}{\sigma^2}}(e^{\frac{h(x,y)(2c|x-y|-h(x,y))}{\sigma^2}} - 1) \tag{5.7}$$

As we have discussed before, the Laplacian of R has a simple Fiedler value. We would like to find a condition under which the same would hold for w. The difference of w and R is E, so if E is small enough, w should also have a simple Fiedler value. E is non-zero in the presence of a positive geodesic deviation h, and h is only non-zero when the curvature is. Hence, it is reasonable to think that a small upper bound on the curvature would make E small enough so that w has a simple Fiedler value. This will be proved in the main theorem of this subsection. Before stating the main theorem, we will state the following lemma:

Lemma 5.4.5. *For a curve $\mu : [0,1] \to \mathcal{W}$, let k_x be the (Haantjes) curvature of point $x \in [0,1]$. Let $w = e^{\frac{-W_2^2(\mu_x,\mu_y)}{\sigma^2}}$ and E defined as in 5.5. For each $\tilde{M} > 0$, If $k_x < \tilde{M}$ for every $x \in [0,1]$, then there exists $\epsilon \geq 0$ such that $E(x,y) \leq \frac{2cM\epsilon^4}{\sigma^2}$ when $|x-y| \leq \epsilon$ and $M = c^3\frac{\tilde{M}^2}{24}$.*

Proof. First, note that

$$E(x,y) = e^{\frac{-(c|x-y|-h(x,y))^2}{\sigma^2}} - e^{\frac{-c^2(x-y)^2}{\sigma^2}}$$

Let $f(z) = e^{\frac{-z^2}{\sigma^2}}$. Then $E(x,y) = f(c|x-y|-h(x,y)) - f(c|x-y|)$. By the mean value theorem, there exists $s \in (c|x-y| - h(x,y), c|x-y|)$ such that

$$\begin{aligned} E(x,y) &= f(c|x-y| - h(x,y)) - f(c|x-y|) \\ &= f'(s)(-h(x,y)) \\ &= (\frac{-2s}{\sigma^2})e^{\frac{-s^2}{\sigma^2}}(-h(x,y)) \\ &\leq \frac{2s}{\sigma^2}h(x,y) \end{aligned}$$

meaning that,

$$E(x,y) \leq \frac{2s}{\sigma^2}h(x,y) \tag{5.8}$$

On the other hand, from the bound on curvature we have

$$\begin{aligned} k_x = \lim_{x\to y}\sqrt{\frac{24h(x,y)}{c^3|x-y|^3}} \leq \tilde{M} \quad &\Longrightarrow \\ \forall x\in[0,1] \quad \exists \epsilon_x \quad s.t. \quad \frac{h(x,y)}{|x-y|^3} \leq M \quad if \quad |x-y| \leq \epsilon_x \quad &\Longrightarrow \\ \frac{h(x,y)}{|x-y|^3} \leq M \quad if \quad |x-y| \leq \epsilon := \min_{x\in[0,1]} \epsilon_x \quad &\Longrightarrow \\ h(x,y) \leq M|x-y|^3 \quad if \quad |x-y| \leq \epsilon \end{aligned}$$

And $s < c|x-y|$, so 5.8 becomes

$$E(x,y) \leq \frac{2cM|x-y|^4}{\sigma^2}$$

if $|x-y| \leq \epsilon$, and the result follows. □

To state the main theorem of this subsection, we first need to state a definition and two assumptions:

Definition 5.4.6. *Let the curve μ be defined as before, and assume that it's (Haantjes) curvature is bounded by $\tilde{M} > 0$. Then $\epsilon(\mu) = \epsilon$ is defined as below:*

$$\begin{aligned} k_x = \lim_{x\to y}\sqrt{\frac{24h(x,y)}{c^3|x-y|^3}} \leq \tilde{M} \quad &\Longrightarrow \\ \forall x\in[0,1] \quad \exists \epsilon_x \quad s.t. \quad \frac{h(x,y)}{|x-y|^3} \leq M = c^3\frac{\tilde{M}^2}{24} \quad if \quad |x-y| \leq \epsilon_x \quad &\Longrightarrow \\ \frac{h(x,y)}{|x-y|^3} \leq M \quad if \quad |x-y| \leq \epsilon := \min_{x\in[0,1]} \epsilon_x \quad &\Longrightarrow \\ h(x,y) \leq M|x-y|^3 \quad if \quad |x-y| \leq \epsilon \end{aligned}$$

The following assumptions are used in the next theorem:

Assumption 7. $\min_{x\in[0,1]} \int_0^1 E(x,y)dy = \int_0^1 E(0,y)dy = \int_0^1 E(1,y)dy$

Assumption 8. *$h(x,y) < h(x,y+z)$ for positive z and all $x, y \in [0,1]$.*

Now, we are ready to state the main theorem:

Theorem 5.4.7. *Let μ be such that $w(x,y) = e^{\frac{-W_2^2(\mu_x,\mu_y)}{\sigma^2}}$ is smooth, and that $\epsilon = \epsilon(\mu)$, as defined in definition 5.4.6, is positive. Assume that the (Haantjes) curvature is bounded so that $k_x \leq \frac{\sqrt{1.5}}{8\sigma}e^{-c^2/2\sigma^2}$ for all $x \in [0,1]$. Then, given that assumptions 5, 7, and 8 are satisfied, the Fiedler function of w is simple and monotone.*

Proof. By theorem 5.4.1, it suffices to show that assumptions 3, 4, 5 and 6 are satisfied.

Assumption 3 is satisfied since $w(x,y)$ is positive unless $x = y$, and the diagonal $\{(x,y) : x = y \in [0,1]\}$ has measure zero. Assumption 4 is satisfied since w is assumed to be smooth.

Now, we show that assumption 6 also holds. Let $\tilde{M} = \frac{\sqrt{1.5}}{8\sigma} e^{-c^2/2\sigma^2}$. Then for each $x \in [0,1]$ we have

$$k_x = \lim_{x \to y} \sqrt{\frac{24h(x,y)}{c^3|x-y|^3}} < \tilde{M} \quad \Longrightarrow$$
$$\frac{h(x,y)}{|x-y|^3} < M := c^3 \frac{\tilde{M}^2}{24} \quad if \quad |x-y| \leq \epsilon$$

By Lemma 5.4.5, we know that $E(x,y) \leq \frac{2cM\epsilon^4}{\sigma^2}$ when $|x-y| \leq \epsilon$. Define $E_1 = \frac{2cM\epsilon^4}{\sigma^2}$.

As before, we have

$$R(z) = e^{\frac{-c^2 z^2}{\sigma^2}}$$
$$E(x,y) = e^{\frac{-c^2(x-y)^2}{\sigma^2}} (e^{\frac{h(x,y)(2c|x-y|-h(x,y))}{\sigma^2}} - 1)$$

Let L be the Laplacian of $w(x,y) = R(|x-y|) + E(x,y)$. We can write $L = L_R + L_E$, where L_R and L_E are Laplacians of R and E respectively. Also, define d_R and d_E to be their degree functions.

We wish to show that

$$\inf_{f \in \mathcal{H}} \langle Lf, f \rangle < \min_{x \in [0,1]} d(x) \tag{5.9}$$

As we have already seen, $\inf_{f \in \mathcal{H}} \langle L_R f, f \rangle < \min_{x \in [0,1]} d_R(x)$ holds using the function $f(x) := f_t(x) := \frac{1}{\sqrt{2t}}(I\{x < t\} - I\{x > 1-t\})$ when $t \leq \frac{1}{2}$ is small enough. Now, we show that when the conditions in theorem 5.4.7 hold, the same function works in showing $\inf_{f \in \mathcal{H}} \langle Lf, f \rangle < \min_{x \in [0,1]} d(x)$.

Now, note that

$$\langle L_E f, f \rangle = I = \frac{1}{2} \int_0^1 \int_0^1 E(x,y)(f(x) - f(y))^2 dy dx = I_1 + I_2 + I_3$$

Where,

$$2tI_1 = \int_0^t \int_t^{1-t} E(x,y) dy dx$$

$$2tI_2 = \int_{1-t}^1 \int_t^{1-t} E(x,y) dx dy$$

$$\frac{t}{2}I_3 = \int_0^t \int_{1-t}^1 E(x,y)dydx$$

And by assumption 7, $\min_{x\in[0,1]} \int_0^1 E(x,y)dy = \int_0^1 E(0,y)dy = \int_0^1 E(1,y)dy =: J$.

Remember from the last section that $\min_{x\in[0,1]} d_R(x) = \int_0^1 R(z)dz$.

To show that $\inf_{f\in\mathcal{H}}\langle Lf, f\rangle < \min_{x\in[0,1]d(x)}$, we can show that

$$\langle L_R f, f\rangle - \int_0^1 R(z)dz + \langle L_E f, f\rangle - J < 0$$

for some $f \in \mathcal{H}$. In the previous subsection, it was shown that

$$\langle L_R f, f\rangle - \int_0^1 R(z)dz = [R(1) - R(0) + o_t(1)]t/2$$

Where $o_t(1)$ is a quantity that goes to 0 when t does, and

$$f(x) := f_t(x) := \frac{1}{\sqrt{2t}}(I\{x < t\} - I\{x > 1 - t\})$$

As a result, when t is small enough, $\langle L_R f, f\rangle - \int_0^1 R(z)dz < 0$. So, it is enough to prove that

$$\langle L_E f, f\rangle - J < 0 \tag{5.10}$$

When t is small enough. We do so by bounding I_1, I_2 and I_3.

Bounding I_1

$$\begin{aligned} 2tI_1 &= \int_0^t \int_t^{1-t} E(x,y)dydx \\ &= I_1^1 + I_1^2 \end{aligned}$$

Where

$$\begin{aligned} I_1^1 &= \int_0^t \int_t^{x+\epsilon} E(x,y)dydx \\ &\leq E_1 \int_0^t \int_t^{x+\epsilon} dydx \\ &= E_1(\epsilon t - \frac{t^2}{2}) \end{aligned} \tag{5.11}$$

And,

$$\begin{aligned} I_1^2 &= \int_0^t \int_{x+\epsilon}^{1-t} E(x,y)dydx \\ &= \int_0^t \int_0^1 E(x,y)dydx - \int_0^t \int_{1-t}^1 E(x,y)dydx - \int_0^t \int_0^{x+\epsilon} E(x,y)dydx \\ &= \int_0^t \int_0^1 (E(0,y)+O_t)dydx - \int_0^t \int_{1-t}^1 (E(0,1)+O_t)dydx - \int_0^t \int_0^{x+\epsilon} E(x,y)dydx \\ &= t(J+O_t) - t^2((E(0,1)+O_t)) - \int_0^t \int_0^{x+\epsilon} E(x,y)dydx \end{aligned} \tag{5.12}$$

Where the O_t's are quantities approaching 0 when t does.

So, by 5.11 and 5.12 we have:

$$2tI_1 \leq E_1(\epsilon t - \frac{t^2}{2}) + t(J+O_t) - t^2((E(0,1)+O_t)) - \int_0^t \int_0^{x+\epsilon} E(x,y)dydx \tag{5.13}$$

Bounding I_2

$$\begin{aligned} 2tI_2 &= \int_{1-t}^1 \int_t^{1-t} E(x,y)dxdy \\ &= I_2^1 + I_2^2 \end{aligned}$$

Where

$$\begin{aligned} I_2^1 &= \int_{1-t}^1 \int_{y-\epsilon}^{1-t} E(x,y)dxdy \\ &\leq E_1 \int_{1-t}^1 \int_{y-\epsilon}^{1-t} dxdy \\ &= E_1(\epsilon t - \frac{t^2}{2}) \end{aligned} \tag{5.14}$$

$$\begin{aligned} I_2^2 &= \int_{1-t}^1 \int_t^{y-\epsilon} E(x,y)dxdy \\ &= \int_{1-t}^1 \int_0^1 E(x,y)dxdy - \int_{1-t}^1 \int_0^t E(x,y)dxdy - \int_{1-t}^1 \int_{y-\epsilon}^1 E(x,y)dxdy \\ &= \int_{1-t}^1 \int_0^1 (E(x,1)+O_t)dxdy - \int_{1-t}^1 \int_0^t (E(1,0)+O_t)dxdy - \int_{1-t}^1 \int_{y-\epsilon}^1 E(x,y)dxdy \\ &= t(J+O_t) - t^2(E(0,1)+O_t) - \int_{1-t}^1 \int_{y-\epsilon}^1 E(x,y)dxdy \end{aligned} \tag{5.15}$$

So, by 5.14 and 5.15 we have:

$$2tI_2 \leq E_1(\epsilon t - \frac{t^2}{2}) + t(J + O_t) - t^2((E(0,1) + O_t)) - \int_{1-t}^{1}\int_{y-\epsilon}^{1} E(x,y)dxdy \tag{5.16}$$

Bounding I_3

$$\begin{aligned} I_3 &= \frac{2}{t}\int_0^t \int_{1-t}^{1} E(x,y)dydx \\ &= 2t(E(0,1) + O_t) \end{aligned} \tag{5.17}$$

So, overall, by 5.13, 5.16 and 5.17 we have:

$$\begin{aligned} &I = I_1 + I_2 + I_3 \\ &\leq E_1(\epsilon - \frac{t}{2}) + J + O_t + t(E(0,1) + O_t) - \frac{1}{2t}\int_{1-t}^{1}\int_{y-\epsilon}^{1} E(x,y)dxdy - \frac{1}{2t}\int_0^t \int_0^{x+\epsilon} E(x,y)dydx \end{aligned}$$

As a result, for the inequality 5.10 to hold, it's enough to show that

$$E_1(\epsilon - \frac{t}{2}) + O_t + t(E(0,1) + O_t) - \frac{1}{2t}\int_0^t \int_0^{x+\epsilon} E(x,y)dydx \leq 0 \tag{5.18}$$

If t is small enough.

If we prove that 5.18 holds as $t \to 0$, then it should also hold for some small t. Keeping that in mind, we now bound $\frac{1}{2t}\int_0^t \int_0^{x+\epsilon} E(x,y)dydx$. To do so, first define

$$\begin{aligned} E(x,y) = l^2(x,y) &= \frac{e^{h(x,y)(2c|x-y| - h(x,y))/\sigma^2} - 1}{e^{c^2|x-y|^2/\sigma^2}} \\ g^2(x,y) &= e^{c^2|x-y|^2/\sigma^2} \end{aligned}$$

By Cauchy - Schwartz:

$$\left(\int lg\right)^2 \leq \int l^2 \int g^2$$

which can be rewritten as

$$\int l^2 \geq \frac{\left(\int lg\right)^2}{\int g^2}$$

which means

$$\int_0^t \int_0^{x+\epsilon} e^{-c^2|x-y|^2/\sigma^2} \left(e^{h(x,y)(2|x-y|-h(x,y))/\sigma^2} - 1\right) \geq \frac{\left(\int_0^t \int_0^{x+\epsilon} \sqrt{e^{h(x,y)(2c|x-y|-h(x,y))/\sigma^2} - 1}\right)^2}{\int_0^t \int_0^{x+\epsilon} e^{c^2|x-y|^2/\sigma^2}}$$

$$\geq \frac{\left(\int_0^t \int_0^{x+\epsilon} \sqrt{e^{h^2/\sigma^2} - 1}\right)^2}{\int_0^t \int_0^{x+\epsilon} e^{c^2/\sigma^2}}$$

$$\geq \frac{\left(\int_0^t \int_0^{x+\epsilon} e^{h^2/2\sigma^2} - 1\right)^2}{\left(\frac{t^2}{2} + t\epsilon\right) e^{c^2/\sigma^2}}$$

$$\geq \frac{\left(\int_0^t \int_{\frac{\epsilon}{2}}^{\epsilon} h^2/2\sigma^2\right)^2}{\left(\frac{t^2}{2} + \epsilon t\right) e^{c^2/\sigma^2}}$$

$$\geq \frac{t\frac{\epsilon^2}{16} h^4\left(x^*, y^*\right)}{\left(\frac{t}{2} + \epsilon\right) e^{c^2/\sigma^2}\sigma^4}$$

Where (x^*, y^*) in the last line comes from using the mean value theorem for integrals, and $0 \leq x^* \leq t$ and $\frac{\epsilon}{2} \leq y^* \leq \epsilon$. So, at the end we have

$$\frac{1}{2t}\int_0^t \int_0^{x+\epsilon} e^{-c^2|x-y|^2/\sigma^2} \left(e^{h(x,y)(2c|x-y|-h(x,y))/\sigma^2} - 1\right) \geq \frac{\frac{\epsilon^2}{32} h^4\left(x^*, y^*\right)}{\left(\frac{t}{2} + \epsilon\right)\left(e^{c^2/\sigma^2}\sigma^4\right)} \tag{5.19}$$

Now, remember that we want to show 5.18, which is

$$E_1(\epsilon - \frac{t}{2}) + O_t + t(E(0,1) + O_t) - \frac{1}{2t}\int_0^t \int_0^{x+\epsilon} E(x,y)dydx\frac{t}{2} \leq 0$$

and for this to happen, by 5.19 it is sufficient to have

$$E_1(\epsilon - \frac{t}{2}) + O_t + t(E(0,1) + O_t) - \frac{\frac{\epsilon^2}{32} h^4\left(x^*, y^*\right)}{\left(\frac{t}{2} + \epsilon\right)\left(e^{c^2/\sigma^2}\sigma^4\right)} \leq 0 \tag{5.20}$$

If we prove that the above holds when $t \to 0$, then for some small t, inequality 5.20 holds. So, it is sufficient to have

$$E_1\epsilon \leq \frac{\frac{\epsilon^2}{32} h^4\left(x^*, y^*\right)}{\epsilon\left(e^{c^2/\sigma^2}\sigma^4\right)} \tag{5.21}$$

When $t \to 0$, $x^* \to 0$ since $0 \leq x^* \leq t$, and we have:

$$\frac{2cM\epsilon^4}{\sigma^2} \leq \frac{h^4\left(0, y^*\right)}{\left(32\sigma^4\right)\left(e^{c^2/\sigma^2}\right)}$$

Using the assumption on h, i.e. 8, and since $\frac{\epsilon}{2} \leq y^* \leq \epsilon$ it suffices to show that

$$\frac{2cM\epsilon^4}{\sigma^2} \leq \frac{h^4(0, \frac{\epsilon}{2})}{\left(32\sigma^4\right)\left(e^{c^2/\sigma^2}\right)}$$

Or, equivalently,

$$\frac{\epsilon^4}{h^4(0, \frac{\epsilon}{2})} \leq \frac{1}{64c\sigma^2 M e^{\frac{c^2}{\sigma^2}}} \tag{5.22}$$

Now, remember that ϵ was defined so that

$$h(x, y) \leq M|x - y|^3 \quad when \quad |x - y| \leq \epsilon$$

The existence of ϵ is justified using the limit definition of the curvature. As such, if we find $\epsilon_0 < \epsilon$ for which 5.22 holds, our desired outcome, 5.10 will still be achieved.

Note that $\lim_{\epsilon \to 0} \frac{\epsilon^4}{h^4(0, \frac{\epsilon}{2})} = \frac{16}{c^4}$, and if $\frac{16}{c^4} \leq \frac{1}{64c\sigma^2 M e^{\frac{c^2}{\sigma^2}}}$, then for some ϵ near zero we have $\frac{\epsilon^4}{h^4(0, \frac{\epsilon}{2})} \leq \frac{\sigma^2}{64cM e^{\frac{c^2}{\sigma^2}}}$.

Hence, if $\tilde{M} \leq \frac{\sqrt{1.5}}{8\sigma} e^{-c^2/2\sigma^2}$ or equivalently, $16 \leq \frac{c^3}{64\sigma^2 M e^{\frac{c^2}{\sigma^2}}}$, then for some $\epsilon > 0$ we have

$$\frac{\epsilon^4}{h^4(0, \frac{\epsilon}{2})} \leq \frac{1}{64c\sigma^2 M e^{\frac{c^2}{\sigma^2}}}$$

which by 5.22 is a sufficient condition for having the inequality 5.9 to hold, and the theorem is proved.

□

Bibliography

[1] Ascher, M. and R. Ascher (1963). Chronological ordering by computer. *American Anthropologist 65*(5), 1045–1052.

[2] Atkins, J. E., E. G. Boman, and B. Hendrickson (1998). A spectral algorithm for seriation and the consecutive ones problem. *SIAM Journal on Computing 28*(1), 297–310.

[3] Calderon, D., R. Blecher-Gonen, X. Huang, S. Secchia, J. Kentro, R. M. Daza, B. Martin, A. Dulja, C. Schaub, C. Trapnell, et al. (2022). The continuum of drosophila embryonic development at single-cell resolution. *Science 377*(6606), eabn5800.

[4] Chuangpishit, H., M. Ghandehari, M. Hurshman, J. Janssen, and N. Kalyaniwalla (2015). Linear embeddings of graphs and graph limits. *Journal of Combinatorial Theory, Series B 113*, 162–184.

[5] Cowgill, G., F. Hole, and M. Shaw (1968, 10). Computer analysis of chronological seriation. *American Antiquity 33*, 517.

[6] Davis, C. and W. M. Kahan (1970). The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis 7*(1), 1–46.

[7] Feng, B., L. Zhou, and J. Tang (2017). Ancestral genome reconstruction on whole genome level. *Current Genomics 18*(4), 306–315.

[8] Ghandehari, M. and J. Janssen (2019). An optimization parameter for seriation of noisy data. *SIAM Journal on Discrete Mathematics 33*(2), 712–730.

[9] Kuzara, R. S., G. R. Mead, and K. A. Dixon (1966). Seriation of anthropological data: A computer program for matrix-ordering. *American Anthropologist 68*(6), 1442–1455.

[10] Liiv, I. (2010). Seriation and matrix reordering methods: An historical overview. *Statistical Analysis and Data Mining: The ASA Data Science Journal 3*(2), 70–91.

[11] Mittnenzweig, M., Y. Mayshar, S. Cheng, R. Ben-Yair, R. Hadas, Y. Rais, E. Chomsky, N. Reines, A. Uzonyi, L. Lumerman, et al. (2021). A single-embryo, single-cell time-resolved model for mouse gastrulation. *Cell 184*(11), 2825–2842.

[12] Natik, A. (2019). *Consistency of the spectral seriation algorithm.* Ph. D. thesis, Université d'Ottawa/University of Ottawa.

[13] Natik, A. and A. Smith (2021). Consistency of spectral seriation. *arXiv preprint arXiv:2112.04408*.

[14] Petrie, W. F. (1899). Sequences in prehistoric remains. *Journal of the Anthropological Institute of Great Britain and Ireland*, 295–301.

[15] Recanati, A., N. Servant, J.-P. Vert, and A. d'Aspremont (2018). Robust seriation and applications to cancer genomics. *arXiv preprint arXiv:1806.00664*.

[16] Reed, M. (2012). *Methods of modern mathematical physics: Functional analysis*. Elsevier.

[17] Robinson, W. S. (1951). A method for chronologically ordering archaeological deposits. *American antiquity 16*(4), 293–301.

[18] Saucan, E. (2015). Metric curvatures and their applications i. *Geometry, Imaging and Computing 2*(4), 257–334.

[19] Villani, C. et al. (2009). *Optimal transport: old and new*, Volume 338. Springer.

[20] Weed, J. and F. Bach (2017, 06). Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli 25*.

[21] Yu, Y., T. Wang, and R. J. Samworth (2015). A useful variant of the davis–kahan theorem for statisticians. *Biometrika 102*(2), 315–323.