Network Representation Matrices and their Eigenproperties: A Comparative Study

a thesis presented for the degree of

Doctor of Philosophy of Imperial College London

and the

Diploma of Imperial College

by

Johannes F. Lutzeyer

Department of Mathematics

Imperial College

180 Queen's Gate, London SW7 2BZ

January 2020

I certify that this thesis, and the research to which it refers, are the product of my own work, and that any ideas or quotations from the work of other people, published or otherwise, are fully acknowledged in accordance with the standard referencing practices of the discipline.

Signed: _____

Copyright

Acknowledgments

I want to express my deepest gratitude to Andrew Walden, not only, for being an excellent supervisor in every metric, but also, a person I could comfortably look up to and learn from in every aspect of life. Your constructive attitude, patience and focus have made this work possible. I thoroughly enjoyed our many, long meetings and conversations.

I would like to thank the EPSRC for funding this project.

I am grateful to Axel Gandy, Edwin Hancock and Ed Cohen for their insightful comments given on numerous occasions including my PhD milestone assessments and viva. I would also like to thank Dean Bodenham for his support and in particular for his help with the LaTeX implementation of a list of symbols.

A special thanks goes to the great community of Statistics and Mathematics PhD students I got to be a part of. Our daily conversations were, both personally and mathematically, incredibly enriching. I am very thankful to have been accompanied by such great friends on this journey. Of course this thanks extends to the wider Statistics section community at Imperial College; I am extremely grateful to have been a part of such a social and productive environment.

I would like to say what a pleasure it was to work with Simon, Francesco, Marie, Victor, Hannah and all four great senior teams. Thank you very much for making Gabor hall my home in these past four years.

I am indebted to my family for their unconditional support throughout these years, for always helping me recharge and get back to work with a fresh perspective.

Finally, I would like to sincerely thank Victoria for her love and support throughout these years.

Johannes Lutzeyer
London, UK
January 2020

Thesis advisor: Professor Andrew T. Walden Johannes F. Lutzeyer

Network Representation Matrices and their Eigenproperties: A Comparative Study

Abstract

Typically, network structures are represented by one of three different representation matrices: the adjacency matrix, the unnormalised and the normalised graph Laplacian matrices.

We review their spectral (eigenvalue) and eigenvector properties and recall several relations to graph theoretic quantities. We compare the representation spectra pairwise by applying an affine transformation to one of them, which enables comparison whilst preserving certain key properties such as normalised eigengaps. Bounds are given on the eigenvalue and normalised eigengap differences thus found, which depend on the minimum and maximum degree of the graph. It is found that if the degree extreme difference is large, different choices of representation matrix may give rise to disparate inference drawn from network analysis; smaller degree extreme differences result in consistent inference, whatever the choice of representation matrix.

In order to compare the representation matrix eigenvectors, we extend the applicability and tighten the Davis–Kahan theorem by making use of polynomial matrix transformations. Our extended version of the Davis–Kahan theorem enables the comparison of the spaces spanned by sets of eigenvectors of any two symmetric matrices, such as the representation matrices. We then discuss how globally optimal bound values of this extended theorem can be computed in practice for affine transformations using fractional programming theory. We make use of our implementation of the extended Davis–Kahan theorem to compare the spaces spanned by the eigenvectors of the different graph representation matrices for a range of different stochastic blockmodel graphs and of covariance matrices in a spiked covariance model.

Contents

List of Figures

List of Tables

List of Frequent Symbols

w_i i^{th} eigenvector of the matrix Ψ 98

U_j matrix with eigenvectors $u_{j+1}, \ldots, u_{j+r}$ as columns 98

W_j matrix with eigenvectors $w_{j+1}, \ldots, w_{j+r}$ as columns 98

$\mathcal{U}_j$ space spanned by the columns of U_j 103

$\mathcal{W}_j$ space spanned by the columns of W_j 103

$O(r)$ group of $r \times r$ orthogonal matrices 101

$\mathbb{V}_{n,r}$ Stiefel manifold 101

$\mathbb{G}_{n,r}$ Grassmann manifold 101

$\rho_1(\mathcal{U}, \mathcal{W}), \rho_2(\mathcal{U}, \mathcal{W})$ distance metrics on the spaces $\mathcal{U}, \mathcal{W}$ 103

$S_1 = [a, b]$ interval used in the context of the Davis–Kahan theorem 106

S_2 interval used in the context of the Davis–Kahan theorem 106

δ separation between intervals S_1 and S_2 106

N sample size 147

Introduction

In this thesis we work on the mathematical representation of graph (or network) structures using matrices. Graphs can be, and often are, represented via the adjacency matrix A, which has as $(i,j)^{\text{th}}$ entry the edge weight between nodes v_i and v_j in the graph. From the adjacency matrix we can define the diagonal degree matrix D where the i^{th} diagonal entry is equal to the sum of all edge weights connecting node v_i to the rest of the graph. Using the degree matrix and the adjacency matrix we define the three graph Laplacians, which are used to represent graphs in this thesis, namely the unnormalised Laplacian, $L = D - A$, the symmetric normalised Laplacian, $L_{sym} = D^{-1/2}LD^{-1/2}$, and the random walk normalised Laplacian, $L_{rw} = D^{-1}L$. We will refer to the adjacency matrix and these three Laplacian matrices as the *representation matrices*. In the signal processing literature the representation matrices are commonly referred to as graph shift operator matrices (Ortega et al., 2018).

The main focus of this thesis is to investigate differences of the representation matrices and provide tools to measure and quantify these differences. As pointed out in Zenil et al. (2016) the information content in the representation matrices is equal, in the sense that they each unambiguously define the graph structure they correspond to. However, in graph analysis, graph learning and graphical inference the eigenvalues and the eigenvectors of the representation matrices are commonly used. In this thesis we refer to the spectrum (eigenvalues) and eigenvectors of a representation matrix, and their significance in representing the graph structure, as the representation matrices' *eigenproperties*. For general graphs, it is well known that the eigenproperties of the representation matrices are different (Butler (2008) stated that for general graphs "it will make a big difference which spectrum is used, and some results which might hold for one spectrum may not hold for another."). Therefore, the representation matrix choice is consequential. In all chapters of this thesis we concern ourselves with the eigenproperties of the representation matrices, in fact the work in Chapter 3 applies to any two symmetric matrices, which

includes the representation matrices.

Which representation matrix to use, and when, and which factors to consider when making this choice are still unknown. Recently several authors have discussed this issue. Zenil et al. (2016) state that it remains to be ascertained which representation matrix spectrum conserves more information. When introducing the representation matrices Butler and Chung (2017, p. 47-1) state: "No one matrix is best because each matrix has its own limitations in that there is some property which the matrix cannot always determine (this shows the importance of understanding all the different matrices)." A similar approach has recently become popular in the graph signal processing literature, where – as mentioned above – the representation matrices are referred to as graph shift operators. Deri and Moura (2017) and Ortega et al. (2018) state that the representation matrix choice involves "different tradeoffs" and leads to different signal models. Therefore, they recommend using whichever representation matrix works best in a particular analysis or learning task. They confirm that making the choice of representation matrix and its normalisation appropriately in a given application remains an open question (Ortega et al., 2018, p. 816). Of course it is application specific to determine the quality of representation of a certain representation matrix. We believe that this fundamental choice needs to be addressed. In this thesis we provide an extensive account of the representation matrix eigenproperties, frameworks under which the representation eigenproperties can be compared and an improved understanding of the magnitude of the differences of the representation matrix eigenproperties. These contributions are of real use in informing the representation matrix choice. The fact that we have investigated the choice of representation matrix has already been recognised as important in Bohannon et al. (2019), Dankulov et al. (2019) and Clark and Macdonald (2018).

In addition to the uncertainty about how to make the representation matrix choice, it is also largely unknown to what extent results developed based on one representation matrix carry over to analyses based on another representation matrix. This thesis addresses this issue. In Chapter 2 we provide bounds on the maximal differences of the eigenvalues and normalised eigengaps. In Chapters 3 and 4 the Davis–Kahan theorem is used in the comparison of the representation matrix eigenvectors. The Davis–Kahan theorem bounds the

distance of the spaces spanned by the eigenvectors of two symmetric matrices if their eigenvalues satisfy a given interval structure. It is a simple analytic bound involving the matrix error term, i.e., the matrix difference, and a 'separation of intervals term', where intervals are most naturally defined via the matrix spectra.

An instrumental tool in all of these comparisons is the use of matrix transforms to enable a sensible comparison of the representation eigenproperties. In Chapter 2 we make use of affine transforms to increase the comparability of the representation spectra and to enable further analysis of their differences. In Chapter 3 we make use of polynomial matrix transformations to extend and tighten the Davis–Kahan theorem to apply to the comparison of the spaces spanned by any two consecutive corresponding sets of eigenvectors of symmetric matrices, such as the representation matrices. In Chapter 4 we consider the special case of affine transformations again and develop an efficient solution approach to calculating our tightened extended Davis–Kahan bounds.

So, in addition to the eigenproperties of the representation matrices and the use of matrix transforms to enable the representation matrix comparison, the chapters in this thesis have another common theme, namely the use of theoretical bounds stemming from perturbation theory to provide maximally possible differences of the representation matrix eigenproperties. These bounds help quantify to what extent results developed for one representation matrix carry over to another.

In a particular analysis, given a choice of one representation matrix, the eigenproperties of the matrix are often utilised as in, for example, the spectral clustering algorithm (von Luxburg, 2007), graph wavelets (Tremblay and Borgnat, 2014) and the much utilised graph Fourier transforms (Shuman et al., 2013; Dong et al., 2019). Tremblay and Borgnat (2014) use their construction of graph wavelets, utilising both the spectrum and the eigenvectors in their design, to obtain clusters on several different scales. The graph Fourier transform finds uses in a variety of applications such as complex brain network analysis in Wang et al. (2018), to understand centrality patterns over several types of networks in Singh et al. (2016) and to define a graph similarity distance from the graph Fourier bases in Lagunas et al. (2018).

It is important to differentiate the use of eigenproperties in graphical inference

or analysis, as just discussed, and graph learning. Dong et al. (2019) and Mateos et al. (2019) provide a good overview of graph learning algorithms from a graph signal processing perspective. When fitting a graph to time series, the criteria for the choice of representation matrix are very different to the ones when analysing a given graph. As described in Dong et al. (2019) and Mateos et al. (2019) criteria for choosing a graph representation matrix when learning a graph include complexity of the fitting algorithms and desired properties of the fitted graphs. The eigenproperties of several representation matrices are starting to be combined and leveraged in graph learning (Kumar et al., 2019). Kumar et al. (2019) make the point that while the requirements on the graph spectra are different in graph learning in comparison to graphical analysis, the eigenproperties of the representation matrices can be made use of in both contexts.

For completeness we list several alternative representation matrices, which are less frequently used in practice, but can be found in the literature.

Among the main three representation matrices, Stanić (2015) includes a review of the eigenproperties of the *signless Laplacian* defined as $D + A$ in his book; this representation matrix can also be found in Butler and Chung (2017) and Brouwer and Haemers (2011). Aouchiche and Hansen (2013) define the *distance Laplacian* and the *distance signless Laplacian* as generalisations of the distance matrix of a graph and discuss their spectral properties. Bühler and Hein (2009) define the *p-Laplacians*, which arise as a generalisation of the quadratic forms of both the unnormalised and normalised Laplacian. In Banks et al. (2014) the *z-Laplacian* is defined as $zA - D$ and its eigenproperties are discussed. The z-Laplacian is said to share important spectral properties with another representation matrix, the *non-backtracking matrix*, the spectral properties of which are discussed in Krzakala et al. (2013) and Bordenave et al. (2015).

Potentially the most prominent alternative representation matrix is the *modularity matrix* (Newman, 2018, p. 505; Fortunato, 2010, p. 103; van Mieghem, 2011, p. 97). The $(i, j)^{\text{th}}$ entry of the modularity matrix is defined as the $(i, j)^{\text{th}}$ adjacency matrix entry minus the number of expected edges between two nodes given their degrees.

It is likely that the methodology developed in this thesis carries over to the

comparison of the eigenproperties of the majority of these less commonly used representation matrices.

Contributions

The contributions made in this thesis can be summarised as follows:

Chapter 1 contributes an extensive review of the different, known eigenproperties of the graph representation matrices and their significance is emphasised by highlighting several of their relations to different graph parameters and operations.

Chapter 2 contributes a framework under which the representation matrix spectra can be compared. This leads to bounds on the individual eigenvalue and eigengap differences, which turn out to be in terms of the degree extreme difference of the corresponding graph.

Chapter 3 contributes the tightening and extension of the Davis–Kahan theorem to apply to the comparison of the spaces spanned by any two consecutive corresponding sets of eigenvectors of two symmetric matrices (excluding degenerate examples).

Chapter 4 contributes a fractional programming implementation according to which optimal affine transformation parameters can be chosen for our extended Davis–Kahan theorem. The developed algorithm enables computation of our extended and sharpened Davis–Kahan bound in practice.

Chapter summaries

Chapter 1 begins by formally introducing this thesis' notation. It then continues by giving an introduction to graph theory, expander graphs, random graph models, especially the prominent stochastic blockmodel, and referencing several real datasets. Then we list the eigenproperties of the adjacency matrix and introduce homo- and heterophilic cluster structure via the adjacency version

of the Cheeger inequality and the Expander Mixing Lemma. This is followed by a list of the eigenproperties of the unnormalised Laplacian matrix together with a discussion of the unnormalised Laplacian Cheeger inequality and its recent extension. Furthermore, the use of L's eigenvectors for graph drawing and their signal processing interpretation as frequency components on the graph is discussed. We begin our review of the normalised Laplacians by relating and then listing their eigenproperties. We briefly introduce the normalised Laplacian version of Cheeger's inequality. In the final three sections of Chapter 1 we introduce the interaction of several graph operations with the representation spectra, introduce twin vertices and how their presence is manifested in the representation spectra and finally review existing work in the literature on relating the representation spectra.

In *Chapter 2* spectral bounds on the individual eigenvalue and normalised eigengap differences of the representation matrices are derived. In these pairwise comparisons one of the two matrices is always transformed using an affine transform to increase their comparability. All bounds derived are in terms of the degree extreme difference of the graphs. Therefore, we classify graphs according to their degree extreme difference, which helps us explain the monotonicity observed in the eigenvalue bounds. We visualise all bounds using several simulated graphs and a real social network. Finally, we visualise the impact of the representation matrix choice by applying the spectral clustering algorithm to a simulated graph and a real social network.

In *Chapter 3* the Davis–Kahan theorem is extended to apply to the comparison of the spaces spanned by the eigenvectors of any two symmetric matrices (excluding degenerate examples). The main tool in the proof of this extended theorem is the use of polynomial matrix transformations. Previously the Davis–Kahan theorem could only be applied to symmetric matrices if their eigenvalues satisfied a specific structure; we were able to move the assumptions from the eigenvalues on to the transformation parameters. This chapter ends with a proof of concept example where the spaces spanned by the eigenvectors of the representation matrices corresponding to a d-regular graph are compared.

In *Chapter 4* we frame an optimisation problem over affine transformation parameters which result in optimal bounds from our extended Davis–Kahan

theorem. We show that this problem falls into the class of fractional programs and discuss the implementation of its solution. Then we use this implementation to provide bound values from our tightened and extended Davis–Kahan theorem in the comparison (i) of the spaces spanned by the eigenvectors of the representation matrices, (ii) of the representation matrices compared to their corresponding generating matrices in a stochastic blockmodel and (iii) of the sample and population covariance matrix in a spiked covariance model.

List of publications

J. F. Lutzeyer & A. T. Walden, "Comparing Spectra of Graph Shift Operator Matrices," In: H. Cherifi, S. Gaito, J. F. Mendes, E. Moro, L. M. Rocha (eds.) *Complex Networks and Their Applications VIII,* Cham, Switzerland: Springer, pp. 191–202, 2020. https://doi.org/10.1007/978-3-030-36683-4_16

J. F. Lutzeyer & A. T. Walden, "Extending the Davis-Kahan theorem for comparing eigenvectors of two symmetric matrices I: Theory," *arXiv:1908.03462 [math.ST]*, 2019.

J. F. Lutzeyer & A. T. Walden, "Extending the Davis-Kahan theorem for comparing eigenvectors of two symmetric matrices II: Computation and applications," *arXiv:1908.03465 [math.ST]*, 2019.

1

Known Eigenproperties of the Graph Representation Matrices

In this chapter we review the known eigenproperties of the graph representation matrices, how they relate to other graph constants and their interactions with operations applied to the underlying graph. Note that in this section we aim to give reference to the publications which introduce concepts most comprehensively and do not always reference the first publication of a concept.

Specifically, in Section 1.1 we formally introduce the notation we use throughout this thesis. In Section 1.2 we review the necessary graph theory background for the work in this thesis and provide an overview of graph theoretic graph families, expander graphs, random graph models, especially the stochastic blockmodel and real world graph datasets. In Section 1.3 we review the eigenproperties of the adjacency matrix and especially highlight their relation to homophilic and heterophilic cluster structures. Then, in Section 1.4 we review the eigenproperties of the unnormalised graph Laplacian and discuss using its eigenvectors as a map to a latent space. In Section 1.5 we discuss the eigenproperties of the normalised Laplacians and how their spectra relate. Following this review of the representation matrix eigenproperties, we discuss the interaction of operations on the graph structure with the repres-

entation spectra in Section 1.6 and review the concept of twin vertices and its implications on the representation eigenproperties in Section 1.7. We end this chapter with a discussion of existing work on relating the representation matrix spectra in Section 1.8.

1.1 Graph representation matrices and their notation

We chose the notation in this thesis as close as possible to the literature standard so readers experienced with graphs and networks should find no surprises here.

In this thesis we consider graphs or networks $G = (V, E)$ with vertex set V and edge set E. Throughout the thesis the terms graph and network will be used interchangeably, referring to the same structure. The number of elements in V, denoted by $|V|$, is equal to n. In an *undirected* graph edges connect vertices without any direction to the connection and in *simple* graphs any two vertices are connected by one or fewer edges and there are no edges connecting a vertex to itself. We will assume our graphs to be undirected and simple. An edge between vertices v_i and v_j is denoted to have edge weight a_{ij}, which is defined to be 0 if there exists no edge between v_i and v_j.

For clarity we repeat here the definitions of the representation matrices that we gave in the introduction. The matrix holding the edge weights is called the *adjacency matrix* $A \in [0, 1]^{n\times n}$ of a graph and is defined to have entries $A_{ij} = a_{ij}$ (the maximal weight of an edge, unity, can be achieved by normalising by the maximal edge weight if necessary). Since we are working with simple graphs we find that the adjacency matrix has a diagonal of zeros. The adjacency matrix is one of the standard graph representation matrices considered here.

The degree d_i of the i^{th} vertex is defined to be the sum of the weights of all edges connecting this vertex with others, i.e., $d_i = \sum_{j=1}^{n} a_{ij}$ and the *degree matrix* D is defined to be the diagonal matrix $D = \text{diag}(d_1, \ldots, d_n)$. The degree sequence consists of the vertex degrees $\{d_1, \ldots, d_n\}$. The minimal and maximal degree in the degree sequence is denoted $d_{\min}$ and $d_{\max}$, respectively. An important special case is that of *d-regular graphs*, in which each vertex is

of degree d, so that $d_{\max} = d_{\min} = d$, and the degree matrix is a scalar multiple of the identity matrix, $D = dI$, where I is the identity matrix.

For a general D matrix, the *unnormalised graph Laplacian*, L, is defined as

$$L = D - A.$$

In the literature there are two *normalised Laplacians* which are considered (von Luxburg, 2007); these follow from the unnormalised graph Laplacian by use of the inverse of the degree matrix as follows,

$$L_{rw} = D^{-1}L; \qquad L_{sym} = D^{-1/2}LD^{-1/2}.$$

Remark 1.1. The graph learning literature (Dong et al., 2019) has recently started talking about the matrices representing a signal on a graph called signal representation matrices, these matrices are different to the representation matrices we consider in this thesis. Signal representation matrices are equal to a generating function acting on a graph and describe a signal arising on a graph. Contrariwise, the representation matrices we work with in this thesis describe the graph structure itself and not data generated on a graph. ◁

We use the common nomenclature that the set of eigenvalues of a matrix is called its *spectrum*. In order to distinguish eigenvalues of A, L and L_{rw} we use μ for an eigenvalue of A, λ for an eigenvalue of L and η for one of L_{rw}. In the literature it is standard procedure to primarily consider the largest eigenvalues of the adjacency matrix (e.g. Fairbanks et al. (2017, p. 553)) and the smallest eigenvalues of the Laplacian matrices (e.g. von Luxburg (2007)). The negative sign in the relation of the adjacency and the unnormalised Laplacian matrix $L = D - A$ causes the adjacency spectrum to be reflected around 0 and therefore, there is a rough correspondence of the larger end of the adjacency matrix spectrum to the smaller end of the Laplacian spectra. To reflect this correspondence, we have chosen to index the spectra in opposing order, the eigenvalues of the adjacency matrix are ordered in *descending* order, i.e., $\mu_1 \geq \ldots \geq \mu_n$, while the Laplacian eigenvalues are indexed in *ascending* order, i.e., $\lambda_1 \leq \ldots \leq \lambda_n$ and $\eta_1 \leq \ldots \leq \eta_n$.

1.2 Graph theory and models for graphs

In this section we give a review of the necessary graph theory to understand this thesis and review several types of graphs and graph models. Readers familiar with graph theory, graphical models and the literature standard graph datasets should find no surprises in this section. We recommend these readers to keep reading from Section 1.3 and refer back to the current section if they encounter graph theoretic concepts or graphical models they are unfamiliar with.

1.2.1 Graph theory background material

We begin by introducing the necessary graph theoretic background as well as a few graph theoretic families of graphs, which will be used throughout the thesis, especially in Chapter 2, to visualise our results. The definitions in this section are taken from the graph theory book by Bondy and Murty (2008). We begin by defining four families of graphs.

Definition 1.2.

1. In a *bipartite graph* the vertices can be divided into two disjoint subsets $\mathcal{S}_1$ and $\mathcal{S}_2$ such that every edge connects a vertex in $\mathcal{S}_1$ to a vertex in $\mathcal{S}_2$.
2. In a *star graph* T_{n-1} all edges are incident to a single vertex and connect this single vertex to all other vertices in the graph.
3. In a *complete graph* K_n all vertices are connected to one another.
4. In a *cycle* C_n the vertices can be arranged in a cyclic sequence such that vertices are connected by edges if and only if they are consecutive in the sequence. ◁

In Figure 1.1 we display examples of the graph families defined in Definition 1.2. Note that this particular example of a bipartite graph on 8 vertices in Figure 1.1(a) is also a cycle and is hence denoted C_8. It can easily be shown that all cycle graphs on an even number of vertices are 2-regular bipartite graphs. In Figure 1.1(b) we observe that all *star* graphs are special cases of

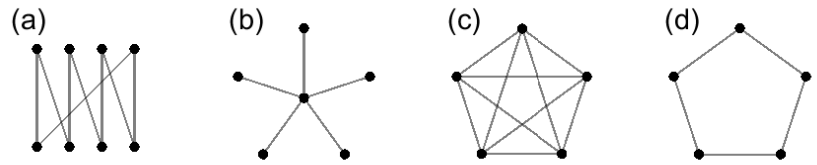

Figure 1.1: Example graphs: (a) A cycle on 8 nodes C_8, which is also a bipartite graph. (b) T_5, the star graph on 6 nodes. (c) K_5, the complete graph on 5 nodes. (d) C_5, the cycle graph on 5 nodes.

bipartite graphs, where one of the partitioning vertex sets contains only the central vertex.

When working with graph spectra from a graph theoretic standpoint it is a common tool to consider the graph's complement, which we define now.

Definition 1.3. Let $G = (V, E)$ be a graph. Then the *complement* $\overline{G}$ of G is the graph whose vertex set is V and whose edges are the pairs of vertices not connected by an edge in G. ◁

We furthermore add the definitions of cliques and stable sets, these will help us interpret the clustering results first mentioned in Section 1.3.

Definition 1.4. A *clique* is a subset of vertices in which any two vertices are connected by an edge. ◁

Definition 1.5. A *stable set* is a subset of vertices in which no two vertices are connected by an edge. Note that stable sets are sometimes referred to as independent sets of a graph. ◁

Definitions 1.4 and 1.5 are complementary in the sense that a set $\mathcal{S}$ of vertices is a clique of a simple graph G if and only if it is a stable set of the complement $\overline{G}$. Due to the complementary structure of cliques and independent sets one might expect a graph to either contain a large clique or a large stable set. This rationale is formally treated by the concept of Ramsey numbers.

Theorem 1.6. (Bondy and Murty, 2008, p. 309) Given any two positive integers k and l, there exists a smallest integer $r(k, l)$ such that every graph on

$r(k, l)$ vertices contains either a clique on k vertices or a stable set of l vertices. The numbers $r(k, l)$ are known as *Ramsey numbers*. ◁

In their book Bondy and Murty (2008) show all the Ramsey numbers known for which $3 \leq k \leq l$. The largest numbers they show are $r(3, 9) = 36$ and $r(4, 5) = 25$. We are able to deduce that if we were to fit a triangle free graphical model to data and we had more than 36 vertices, then, we would have an independent set of at least 9 vertices. Hence any graph of sufficient size must contain some sort of "cluster structure" on a small scale, which is known to exist from Theorem 1.6.

We now define the concept of connected components in a graph, which will be very useful to us throughout this work.

Definition 1.7. A *path* on a graph G is an ordered list of unique vertices such that consecutive vertices are connected by an edge. A vertex set $\mathcal{S}$ is called a *connected component* if there exists a path between any two vertices $v_i, v_j \in \mathcal{S}$ and there exists no path from any $v_i \in \mathcal{S}$ to any $v_j \notin \mathcal{S}$. ◁

Remark 1.8. The adjacency matrix of a graph with several connected components can be arranged to be block-diagonal, e.g., to be of the form $A = (A_1, 0; 0, A_2)$. In contrast, the adjacency matrix of a bipartite graph can be arranged to have 0 matrices on the block diagonal, e.g., to be of the form $A = (0, B; B^T, 0)$ (Brouwer and Haemers, 2011, p. 6). In Section 1.3 we find that these two structures are identified by eigenvectors corresponding to eigenvalues on the opposing extremes of the adjacency spectrum of a graph. ◁

Finally, we define the diameter of a graph. The diameter of a graph turns out to have several interesting relations with the representation spectra which will be discussed in Sections 1.3 and 1.4.

Definition 1.9. The *distance* of two vertices is equal to the number of edges in the shortest path joining them. The *diameter* of a graph, denoted by Δ, is defined to be the equal to the maximal distance between any two vertices in the graph. ◁

1.2.2 Expander graphs

Now we continue with a less standard family of graphs, the expander graphs. The papers by Hoory et al. (2006) and Lubotzky (2012) offer a good introduction to these graphs from several different angles. Lubotzky (2012) defines expander graphs to be *highly connected sparse graphs*, in conjunction these properties are highly desirable in a variety of contexts such as for example efficient communication networks. More concise definitions are given depending on the scientific field in which expander graphs are used. The first formal definition given in Lubotzky (2012), displayed here as Definition 1.11, relies on the definition of the boundary of a vertex set.

Definition 1.10. The *boundary* of a set $\mathcal{S} \subseteq V$ denoted $\partial\mathcal{S}$ is the set of all vertices, which are connected to a vertex in $\mathcal{S}$ by an edge but are not in $\mathcal{S}$. The boundary of a single vertex v_i is called its *neighbourhood* and is denoted $\mathcal{N}(v_i)$. ◁

Definition 1.11. For $0 < \delta \in \mathbb{R}$, G is an δ-*expander graph* if for all $\mathcal{S} \subseteq V$ such that $|\mathcal{S}| \leq \frac{|V|}{2}$ we have $|\partial\mathcal{S}| \leq \delta|\mathcal{S}|$. ◁

This definition gives rise to the definition of the expansion ratio of a graph. Note that the expansion ratio defined here is also commonly referred to as the isoperimetric constant or Cheeger constant (van Mieghem, 2011, p. 95; Butler and Chung, 2017, p. 47-11).

Definition 1.12. The *expansion ratio* $h(G)$ is defined to be the minimal δ such that G is an δ-expander graph or as Hoory et al. (2006, p. 452) put it

$$h(G) = \min_{\{\mathcal{S}:|\mathcal{S}|\leq\frac{n}{2}\}} \frac{|\partial\mathcal{S}|}{|\mathcal{S}|}. \tag{1.1}$$

◁

Intuitively a high expansion ratio means that each "small" set of vertices has a rather large neighbourhood in comparison. This ties in nicely with the qualitative definition as highly connected sparse graphs by Lubotzky (2012) given earlier.

Note that in the denominator of Equation (1.1) the measure for the size of a vertex set is the number of vertices in the set, independent of how well they are

connected, which would be measured by the so called *volume* of a vertex set, defined as $\text{vol}(\mathcal{S}) = \sum_{i:v_i \in \mathcal{S}} d_i$ for $\mathcal{S} \subset V$. The volume of a vertex set is used in Butler and Chung (2017, p. 47-11) to define their own Cheeger constant of a graph $h'(G)$ as

$$h'(G) = \min_{\mathcal{S} \subset V} \left(\frac{|\partial \mathcal{S}|}{\min(\text{vol}(\mathcal{S}), \text{vol}(\bar{\mathcal{S}}))} \right), \tag{1.2}$$

where $\bar{\mathcal{S}} = V \backslash \mathcal{S}$, i.e., $\bar{\mathcal{S}}$ is the complement set of $\mathcal{S}$ in V.

In Sections 1.3.2, 1.4.2 and 1.5.3 we will describe the relation of these two expansion ratios to the representation matrix spectra. These relations will help us motivate the use of the representation matrix spectra to recover cluster structure in networks.

1.2.3 Random graph models

Besides these non-stochastic graph structures there is also a multitude of random graph models, which will be reviewed in this section.

von Luxburg (2007, p. 396) discusses the following three types of similarity graphs, which are based on a set of data points.

- In the *ϵ-neighbourhood graph* data points are joined by an edge if their pairwise distance is smaller than ϵ.
- In the *k-nearest neighbour graphs* each data point or vertex v_i is connected to another vertex v_j if v_j is amongst the k-nearest neighbours of v_i. This leads to a directed graph and hence von Luxburg (2007, p. 396) presents two methods of making these graphs undirected. In what is usually called the *k-nearest neighbour graph* the direction of the edges is just ignored. In the *mutual k-nearest neighbour graphs* an edge is only drawn if there exists a directed edge both ways.
- In the *fully connected graph* all pairs of data points with a positive similarity in a given similarity measure are connected. All edges are weighted by their similarity. von Luxburg (2007, p. 396) states that that these graphs are only sensible to consider if some locality is encoded in the

similarity function as for example in the Gaussian similarity function $s(x_i, x_j) = \exp(-\|x_i - x_j\|^2/(2\sigma^2))$, where the parameter σ controls the width of the neighbourhood.

In the following list we define graph generating models which are not based on an underlying set of data points.

- In *Erdös-Renyi-Gilbert* random graphs edges between any pair of nodes in the vertex set are sampled from independent identically distributed Bernoulli trials with success probability p (Newman et al., 2001; Goldenberg et al., 2010, p. 22).
- In a graph sampled from a *configuration model* we sample n degrees from a given degree distribution. Then each node is assigned a number of "half edges" according to the sampled degrees and half edges are connected at random to form edges. In an optional final step self-loops, multiple edges and a possibly remaining single "half edge" are removed (Molloy and Reed, 1995; Spricer and Britton, 2015, p. 966; Ball and Neal, 2017, p. 1058; Newman, 2018, p. 369).
- In a *preferential attachment model* nodes are added individually to the network. For each added node edges from that node to all existing nodes in the network are sampled randomly with the probability of an edge depending on the degree of the node existing in the network (Albert and Barabási, 2002, p. 71; Goldenberg et al., 2010, p. 42; Newman, 2018, p. 435).

The preferential attachment and configuration model are part of a broader class of models called *scale free networks* (Albert and Barabási, 2002). In scale free networks the degree distribution follows a power law distribution as the number of nodes tends to infinity. The internet graph, where routers are represented by nodes and are connected if there exists a physical connection between them, is a famous example of a scale free network (Albert and Barabási, 2002, p. 52).

There also exists a vast amount of research on *Exponential random graph models* (Snijders et al., 2006; Goldenberg et al., 2010, p. 30) and *Gaussian random graph models*, where the absence of an edge in the graph implies uncorrelated

variables in the joint multivariate normal distribution of all data points or nodes (Lauritzen, 1996).

Perhaps two of the more significant papers on the topic of latent space graph models are Hoff et al. (2002) and Lyzinski et al. (2014). In Hoff et al. (2002) latent space positions of the nodes are inferred as parameters of a model; the distance of nodes in the latent space inversely affects the probability of an edge between these nodes in the model. Then, in Lyzinski et al. (2014) the class of *random dot product graph models* is defined; here the probability of an edge between nodes is a function of the dot product of their corresponding latent positions. Lyzinski et al. (2014) suggest inferring latent positions via the "adjacency spectral embedding," a mapping of the nodes to positions given by several reweighted adjacency eigenvectors, where the reweighting is equal to the square root of the corresponding eigenvalues. Both the stochastic blockmodel and the degree corrected stochastic blockmodel, which will be defined in the immediately following Section 1.2.4, can be shown to be special cases of random dot product graph models.

1.2.4 Stochastic blockmodels

In this section we begin by giving a brief history of the stochastic blockmodel (SBM) and then proceed to give a rigorous definition of the SBM parametrisation we will use in this thesis. Throughout Chapters 2 and 4 SBMs will be used to sample graphs from.

SBM literature review

The SBM, which is widely used in the networks literature, allows us to encode block structure in a random graph via different probabilities of edges within and between node-blocks. It is often accredited to Holland et al. (1983), who model edge data incorporating several different relation types, i.e., a model with several adjacency matrices on the same set of nodes. Interestingly they state that in the presence of multiple relation data adjacency matrices are often highly correlated and not much information is lost by aggregation. They define the original version of the SBM and then go on to define two variants of the

SBM called the pair-dependent SBM and the SBM with reciprocity. The aim of these variant models seems to be to accommodate the increased likelihood of undirected edges in social networks. The models are fit using maximum likelihood estimates obtained in closed form from a Bernoulli or multinomial distributional assumption on the edges. They propose a test on the standard stochastic blockmodel to test the hypothesis that the reciprocation or number of undirected edges between blocks is adequately explained by the model.

Another major step in SBM history was the introduction of the degree-corrected SBM in Karrer and Newman (2011). They noticed that the standard SBM does not allow for a broad enough degree variation and hence introduce an additional n parameters encoding the expected degree of the n vertices in the network. Karrer and Newman (2011) find the fit of the degree corrected SBM to outperform the standard SBM in several synthetic and real network data examples. In their work they allow for multiedges and self-loops and hence assume the number of edges between two vertices to be Poisson distributed. They argue that in the limit of large sparse graphs, "there is essentially no difference" between their Poisson version of the SBM and the standard SBM. The Poisson assumption significantly simplifies working out the distribution of network degrees in comparison to the standard SBMs, where the degrees follow a so called Poisson Binomial distribution. In Butler and Stephens (2017) two approximations to the distribution of the degrees in a standard SBM are discussed.

Sengupta and Chen (2018) further build up on the degree-corrected SBM by proposing a popularity-adjusted blockmodel, in which the expected degree parameter from Karrer and Newman (2011) is split into K parameters representing the popularity of a given node in the K different communities in the SBM.

A summary of the different SBM versions, recent advances and fundamental limits for community detection in SBMs can be found in Abbe (2017). Lei (2016) propose a goodness-of-fit test for the stochastic blockmodel.

Remark 1.13. A team at Microsoft Research worked on the limit of a convergent sequence of graphs and called this limit a *graphon*. The theory is outlined in the book by Lovász (2012). The SBM can also be described by a graphon in the limit. ◁

SBM parametrisation

The definition and parametrisation below is adapted from Lei and Rinaldo (2015).

Definition 1.14. Consider a graph with node set $\{v_1, \ldots, v_n\}$. Split this node set into K disjoint blocks denoted $\mathcal{B}_1, \ldots, \mathcal{B}_K$. We encode block membership of the nodes via a membership matrix $M \in \{0,1\}^{n \times K}$, where $M_{i,j} = 1$ if $v_i \in \mathcal{B}_j$ and $M_{ij} = 0$ otherwise. Finally, we fix the probability of edges between blocks to be constant and collect these probabilities in a probability matrix $P \in [0,1]^{K \times K}$, i.e., for nodes $v_i \in \mathcal{B}_l$ and $v_j \in \mathcal{B}_m$ the probability of an edge between v_i and v_j is equal to $P_{l,m}$. $\triangleleft$

Hence, the parameters of the SBM are $M \in \{0,1\}^{n \times K}$ and $P \in [0,1]^{K \times K}$, where the number of nodes $n \in \mathbb{N}$ and the number of clusters $K \in \mathbb{N}$ are implicitly defined via the dimensions of M. We simulate graphs from this model by fixing these parameters and then sampling edges from Bernoulli trials. The Bernoulli parameter of the trial corresponding to the edge connecting v_i to v_j is given by entry (i,j) of the matrix $B_A = MPM^T - \text{diag}(MPM^T)$, where the subtraction of the diagonal matrix ensures that the sampled network will not contain any self-loops.

Note that in van Mieghem (2011, p. 23) a membership matrix M, such as the one described here, encoding a given partition of the node set, is referred to as characteristic matrix and community matrix.

Since we work with undirected graphs, we sample symmetric adjacency matrices from this model by sampling the upper triangular half of the adjacency matrix from the SBM and then equating the lower triangular part of the adjacency matrix to the transpose of the upper triangle.

Example 1.15. To give the reader a better sense of what SBM samples look like, we display samples from two different SBMs in Figure 1.2.

For both SBMs used in Figure 1.2, $n = 30$, $K = 3$ and the number of nodes per block is equal to 10, i.e., $M_{1,1} = \ldots = M_{10,1} = M_{11,2} = \ldots = M_{20,2} = M_{21,3} = \ldots = M_{30,3} = 1$ and all other entries of M equal 0. The matrix of edge probabilities P is composed of only two values. On the diagonal we have p_w encoding the probability of edges within the different blocks to be the same for

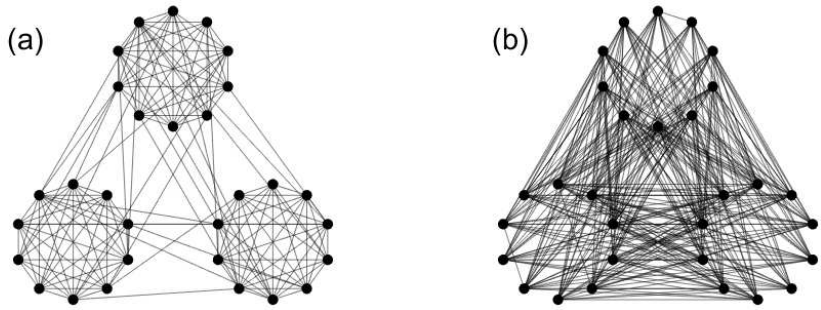

Figure 1.2: Networks sampled from two different SBMs.

all blocks. Off-diagonal we have p_b to encode the probability of edges between nodes in different blocks,

$$P = \begin{pmatrix} p_w & p_b & p_b \\ p_b & p_w & p_b \\ p_b & p_b & p_w \end{pmatrix}.$$

The network in Figures 1.2(a) is sampled from an SBM, where $(p_b, p_w) = (0.1, 0.9)$. While the network in Figures 1.2(b) is sampled from an SBM, where $(p_b, p_w) = (0.9, 0.1)$. The networks in Figures 1.2(a) and (b) display homophilic and heterophilic cluster structure, respectively. These concepts will be introduced in Section 1.3.2 and 1.3.3. ◁

1.2.5 Graph datasets

Real data examples of observed or inferred networks stem from many areas of application. A good overview of the many sciences in which network datasets and spectral network methods can be found, is given in Cvetkovic and Gutman (2011). In Appendix A.1 we display a review of literature standard graph datasets, which we conducted towards the beginning of this PhD project. Based on this review we established Zachary's karate dataset as our main, literature standard, real data example on which we illustrate results produced throughout the project. As established in the review in Appendix A.1, this dataset is a common benchmark in the literature, since it lends itself to clustering; see, e.g., Bickel and Sarkar (2016), Chen and Hero (2015), Fortunato (2010) and Karrer and Newman (2011). The remainder of this section will be a brief discussion of the history and properties of Zachary's karate dataset.

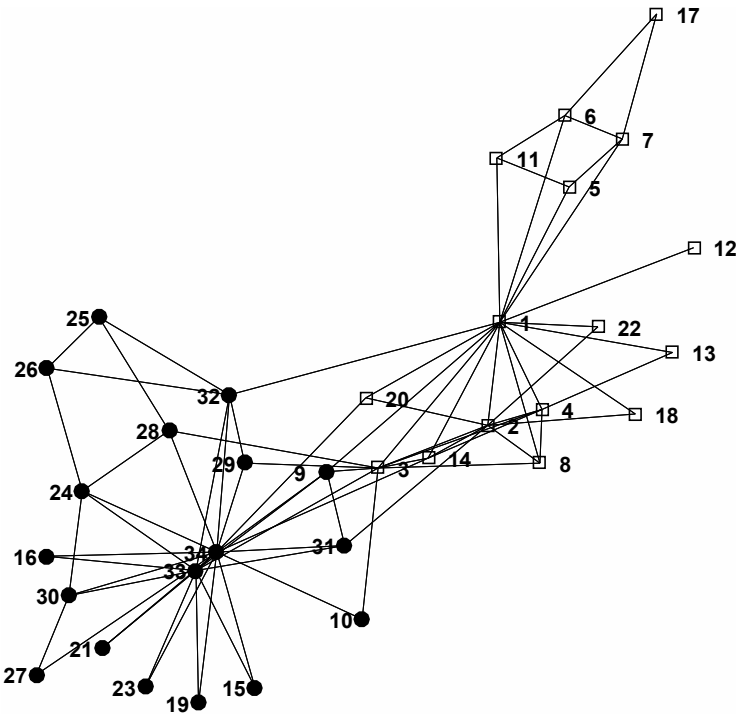

Figure 1.3: Zachary's karate social network. The filled circular nodes and the square unfilled nodes mark the two clusters present in this dataset. The numbering of nodes is as recorded by Zachary (1977).

The karate dataset was originally recorded and analysed by Zachary (1977). Our standard version of the karate dataset was obtained from Batagelj and Mrvar (2006). It contains two square matrices with 34 entries. We work with the first of the two matrices, a square matrix with 34 entries, '`ZACHE`,' which is the adjacency matrix of a social network describing presence of interaction in a university karate club. The second matrix, called '`ZACHC`,' containing "the relative strength of association" measured by the number of situations during which interactions occurred, will not be used by us.

In Figure 1.3 we plot the karate social network. Eventually, the karate club split up over an internal argument between the instructor (member number 1) and the club's president (member number 34) over the price of karate lessons and members in the dataset chose which of the two resulting groups they wanted to join. This eventual split is suggested as one of two ground truth clusterings of this network in Zachary (1977, p. 468), who calls it "CLUB AFTER SPLIT FROM DATA". The other proposed ground truth clustering

is according to social factions which were recorded by Zachary (1977, p. 468), who named this clustering "FACTION MEMBERSHIP FROM DATA". These two possible ground truths only differ in the classification of member number 9, who is said to be difficult to cluster in Zachary (1977, p. 464). Member number 9 was 3 weeks away from his black belt test at the time of the split and was only able to absolve this test in the social faction he did not support. Therefore, member number 9 ended up joining the social faction, he did not support, in the eventual split. We chose the social faction membership as ground truth clustering (displayed in Figure 1.3), since we thought that the social structure represented by the factions was the better motivated clustering of a network where edges describe social interaction.

Correctly identifying two clusters, with respect to either of the two ground truth clusterings recorded in Zachary (1977), comes with several challenges which are discussed in the clustering literature. Several more nodes, besides member 9, are identified as difficult to cluster correctly. Fortunato (2010) in his summary of community detection on graphs states that members 3, 9 and 10 are commonly misclustered. Karrer and Newman (2011) miscluster member 10 when fittting a degree corrected SBM and also state this member to be often misplaced by clustering algorithms. In Figure 1.3 the nodes corresponding to these club members can be seen to be connected to members of both clusters.

1.3 Spectrum of the adjacency matrix

In this section we list the known eigenproperties of the adjacency matrix (Section 1.3.1) and introduce homophilic cluster structure via a variant of Cheegers inequality (Section 1.3.2), as well as heterophilic cluster structure via an example (Section 1.3.3).

1.3.1 Eigenproperties of the adjacency matrix

We denote the eigenvalues of the adjacency matrix A by $\mu_1 \geq \mu_2 \geq ... \geq \mu_n$. Note that in this ordering we consider the largest eigenvalues to be the first. We begin by listing known eigenproperties of the adjacency matrix in List 1.16.

This list highlights the many connections the adjacency eigenproperties have to their underlying graph structure.

List 1.16.

(1) A graph is connected if and only if μ_1 occurs with multiplicity one in the adjacency spectrum with a corresponding eigenvector with non-zero entries, which are all of the same sign (Stanić, 2015, p. 11).

(2) For d-regular graphs, $\mu_1 = d$ with corresponding eigenvector $u_1 = \mathbf{1}_n/\sqrt{n}$, where $\mathbf{1}_n$ is a column vector with n entries all equal to 1 (Hoory et al., 2006, p. 453). The adjacency matrix of a d-regular graph with several connected components can be arranged to be block diagonal. The characteristic equation of a block diagonal matrix factorises into polynomials corresponding to the different blocks or connected components (Bernstein, 2009, p. 291). Therefore, for d-regular graph, the multiplicity r of the first eigenvalue $\mu_1 = d$ equals the number of the connected components $\mathcal{S}_k$ in the graph. The corresponding eigenvectors for the r-dimensional eigenspace can be chosen to be $\mathbf{1}_{\mathcal{S}_k}$, $k \in \{1, 2, ..., r\}$, where $\mathbf{1}_{\mathcal{S}_k}$ are indicator vectors with $(\mathbf{1}_{\mathcal{S}_k})_i = 1$ if $v_i \in \mathcal{S}_k$ and $(\mathbf{1}_{\mathcal{S}_k})_i = 0$ otherwise.

(3) As shown in Appendix A.2, from the Gershgorin theorem it immediately follows that $|\mu_i| \leq d_{\max}$ for $i \in \{1, 2, ..., n\}$ (van Mieghem, 2011, p. 29). In van Mieghem (2011, p. 29) the following additional bound on the adjacency eigenvalues is given, $\mu_i \in \left[-\sqrt{\frac{2(n-1)}{n}|V|}, \sqrt{\frac{2(n-1)}{n}|V|}\right]$ for $i \in \{1, 2, ..., n\}$.

(4) The only rational eigenvalues of the adjacency matrix are integers, e.g., "$\frac{3}{4}$ will never be an eigenvalue of A" (van Mieghem, 2011, p. 29).

(5) Let G be a connected graph with diameter Δ. Then, the adjacency matrix A of G has at least $\Delta + 1$ distinct eigenvalues (Brouwer and Haemers, 2011, p. 5).

(6) Since the adjacency matrix A for general graphs satisfies $\mathrm{tr}\,(A) = 0$, its eigenvalues sum to zero, i.e., $\sum_{i=1}^{n} \mu_i = 0$ (Brouwer and Haemers, 2011, p. 3; Butler, 2008, p. 4).

(7) The number of edges in a graph is equal to $1/2 \sum_{i=1}^{n} \mu_i^2$ (van Mieghem, 2011, p. 30) and the number of triangles is equal to $1/6 \sum_{i=1}^{n} \mu_i^3$ (van Mieghem, 2011, p. 31).

(8) The largest adjacency eigenvalue is bounded below by the average degree, i.e., $\mu_1 \geq \frac{2|E|}{n}$ (van Mieghem, 2011, p. 46; Zumstein, 2005, p. 14).

(9) There exists a constant c such that for every d-regular graph with n vertices of diameter Δ, $\mu_2 \geq 2\sqrt{d-1}(1 - c/\Delta^2)$ (Hoory et al., 2006, p. 484).

(10) A graph is bipartite if and only if its adjacency spectrum is symmetric around 0 (Stanić, 2015, p. 12). Furthermore, the adjacency matrix eigenvectors of bipartite graphs satisfy the following structure, if $\binom{x}{y}$ is an eigenvector with corresponding eigenvalue μ then $\binom{x}{-y}$ is also an eigenvector with corresponding eigenvalue $-\mu$, where the eigenvectors have been ordered according to the bipartite partition (Brouwer and Haemers, 2011, p. 6).

(11) A graph is regular and connected if and only if there exists a polynomial p such that $J = p(A)$, where J is the matrix of all ones (van Mieghem, 2011, p. 44).

(12) $(A^k)_{i,j}$ is the number of walks of length k starting at vertex i and ending at vertex j (Butler, 2008, p. 4). ◁

Another area of work of great importance stemming from the field of random matrix theory, mentioned in Hoory et al. (2006), Bordenave and Lelarge (2010) and many other papers, is the so called *Wigner's semicircle law*. Wigner's semicircle law provides the limit distribution of the eigenvalues of the adjacency matrix of any graph as the size of the matrix tends to infinity under the condition that elements of the adjacency matrix are sampled independently from distributions with finite moments. This theorem does not apply to d-regular graphs as independently sampling the elements of the adjacency matrix does not necessarily produce matrices, where all row sums are equal to d. For d-regular graphs the asymptotic distribution of the eigenvalues has also been derived under the condition, that the number of *cycles*, paths which start and

end in the same vertex, grow slower than linearly with n and can also be found in Hoory et al. (2006).

1.3.2 Homophilic cluster structure

In this section we discuss the relation of the expansion ratio with the adjacency spectrum and how it motivates the use of the largest adjacency eigenvalues to indicate the number of clusters present in a network. Then, we discuss the expander mixing lemma and how it reveals that cluster structure is encoded by eigenvalues at both ends of the adjacency spectrum.

Moving forward we define *eigengaps* of the adjacency matrix, $\mathcal{M}_i = \mu_i - \mu_{i+1}$, to be differences of successive eigenvalues. The first eigengap $\mathcal{M}_1 = \mu_1 - \mu_2$ is of particular significance and is sometimes referred to as the spectral gap of the adjacency spectrum (Hoory et al., 2006, p. 454). In Theorem 1.17 the expansion ration defined in Definition 1.12 is related to the adjacency spectrum via the first adjacency eigengap.

Theorem 1.17. (Hoory et al., 2006, p. 454; Stanić, 2015, p. 111) Let G be a d-regular graph with spectrum $d = \mu_1 \geq \mu_2 \geq ... \geq \mu_n$. Then,

$$\frac{d - \mu_2}{2} \leq h(G) \leq \sqrt{2d(d - \mu_2)}. \qquad \triangleleft$$

In the interpretation of this theorem it will be helpful to define *homophilic clusters* as clusters where nodes within a cluster have many connections to each other and have fewer connections to the rest of the graph.

> Homophilic clusters obtained from the end of the spectrum containing the larger eigenvalues, with many edges within the clusters and few between, correspond to parts of the graph similar in terms of edge density to cliques (see Definition 1.4). An example of a network displaying homophilic cluster structure can be found in Figure 1.2(a).

Remark 1.18. From Theorem 1.17 we see that a large spectral gap indicates a high expansion ratio. From the definition of the expansion ratio, Definition 1.12, we see that hence, a large spectral gap indicates that any small

vertex set has a large boundary set. Therefore for d-regular graphs, Theorem 1.17 motivates using small second adjacency eigengaps to indicate the presence of two large homophilic clusters in the graph. We will discuss equivalent and more general results relating the expansion ratio to the Laplacian spectra in Sections 1.4.2 and 1.5.3. ◁

Another interesting measure in Hoory et al. (2006) is $\max(|\mu_2|, |\mu_n|)$, i.e., the largest eigenvalue in absolute value of the spectrum other than μ_1. It is used in the now following Expander Mixing Lemma 1.20. In Hoory et al. (2006) this lemma is accredited to Alon and Chung (1988). Alon and Chung (1988) use it in their work in a constructive proof showing the existence of paths of a certain length in a general graph after removing a specified number of edges or vertices from it. The application named in their work is the study of fault tolerant arrays, where vertices in the graph represent processing elements and edges represent communication links. We have not used the version showed in Hoory et al. (2006) or Alon and Chung (1988). We reformulated the Expander Mixing Lemma in terms of the density of two disjoint vertex subsets.

Definition 1.19. Given a pair of vertex-disjoint vertex sets $\mathcal{S}_1$ and $\mathcal{S}_2$. We denote the set of edges incident to a vertex in $\mathcal{S}_1$ and $\mathcal{S}_2$ by $E(S_1, S_2)$. Then the *density* $\rho(\mathcal{S}_1, \mathcal{S}_2)$ of the pair $(\mathcal{S}_1, \mathcal{S}_2)$ is defined to be

$$\rho(\mathcal{S}_1, \mathcal{S}_2) = \frac{|E(\mathcal{S}_1, \mathcal{S}_2)|}{|\mathcal{S}_1||\mathcal{S}_2|}. \qquad \triangleleft$$

Now we are able to see that the Expander Mixing Lemma provides a bound for the difference between the density of any two vertex sets $\mathcal{S}_1$ and $\mathcal{S}_2$ in G and the expected number of edges between the two sets in a random d-regular graph with edge density d/n.

Lemma 1.20. *Expander Mixing Lemma* Let G be a d-regular graph with n vertices. Then, for all $\mathcal{S}_1, \mathcal{S}_2 \subseteq V$ we have

$$\left|\rho(\mathcal{S}_1, \mathcal{S}_2) - \frac{d}{n}\right| \leq \frac{\max(|\mu_2|, |\mu_n|)}{\sqrt{|\mathcal{S}_1||\mathcal{S}_2|}}. \qquad \triangleleft$$

This theorem further reinforces the earlier statement from Remark 1.18. A second eigenvalue close to $\mu_1 = d$, i.e., a small spectral gap, means that it

is possible to identify clear-cut clusters as the edge density of vertex subset pairs is allowed to have greater variation between all pairs of vertex sets. The fact that $|\mu_n|$ appears to also be informative to the regularity of the density of vertex sets in a graph is discussed in the next Section 1.3.3.

1.3.3 Heterophilic cluster structure

The value $\max(|\mu_2|, |\mu_n|)$ in Lemma 1.20 gives importance to both ends of the adjacency spectrum and hence validates the idea that the negative end of the spectrum also has significance in measuring the regularity of a graph and in identifying large clusters. This thought is substantiated in the block-model literature (Rohe et al., 2011), where so called '*heterophilic*' clusters are described and found by looking at the eigenvectors corresponding to the largest eigenvalues of the normalised Laplacian. These heterophilic clusters have fewer connections within than between clusters and are exemplified by romantic relationship graphs where the genders are the two biggest clusters identifiable.

Heterophilic clusters obtained from the end of the spectrum containing the smaller adjacency eigenvalues, with few edges within the clusters and many between, correspond to parts of the graph similar in terms of edge density to independent sets (see Definition 1.5). An example of a network displaying heterophilic cluster structure can be found in Figure 1.2(b).

In Example 1.21 the eigenproperties of adjacency matrices corresponding to bipartite graphs and the heterophilic cluster structure are visualised on a cycle on 8 nodes.

Example 1.21. In this example, we work with a d-regular bipartite graph, where the vertex sets forming the bipartite partition, i.e., the two heterophilic clusters with no edges within and all edges between sets, can be recovered from the eigenvector corresponding to the largest negative eigenvalue, which is equal to $-d$ as can be seen in Points (2) and (10) of List 1.16. In Figure 1.4 we plot the eigenvalues and eigenvectors of the adjacency matrix corresponding to the 2-regular bipartite graph C_8, which is plotted in Figure 1.1(a) in Section 1.2.

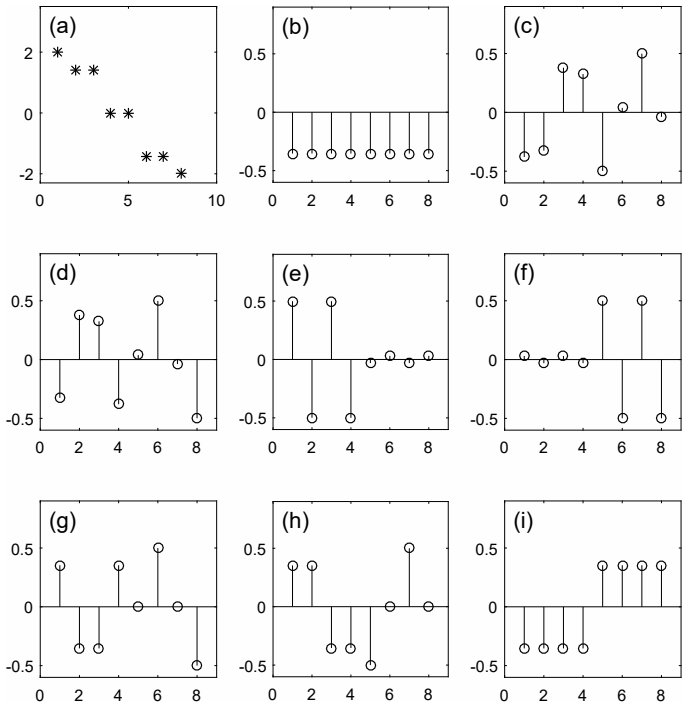

Figure 1.4: Adjacency eigenvalues and eigenvectors of C_8: (a) Eigenvalues of the adjacency matrix corresponding to C_8, where the x-axis displays eigenvalue indices and the y-axis shows the eigenvalue values. (b)–(i) adjacency eigenvectors corresponding to C_8, where (b) contains the eigenvector corresponding to the largest adjacency eigenvalue, (c) the adjacency eigenvector corresponding to the second largest eigenvalue and so on. The x-axis in plots (b)–(i) shows the eigenvector entry indices, while the y-axis shows the eigenvector entry values.

In Figure 1.4(a) we are nicely able to observe the symmetry of the adjacency spectrum around 0 (List 1.16(10)). The first eigenvalue having multiplicity one indicates that the graph is connected (List 1.16(2)). We furthermore see in Figure 1.4(i) that the bipartite partition sets of C_8 , i.e., vertices with indices $\{1, 2, 3, 4\}$ and $\{5, 6, 7, 8\}$, can be recovered perfectly from the 8^{th} eigenvector. Since there is no part of the graph, which has many connections within and less to the remainder of the graph, the second eigenvector in Figure 1.4(b) does not indicate a clear partition.

We can observe that eigenvectors 1 and 8 satisfy the structure described in List 1.16(10), while the other eigenvectors do not. However, it can be confirmed that the eigenvectors displayed in Figure 1.4 form a valid eigenbasis. This issue arises due to the three eigenvalues with multiplicity 2. Since the adjacency matrix is symmetric, the eigenspaces corresponding to the repeated eigenvalues

have dimension 2 and any basis of these spaces produces a valid eigenvector set. Hence, via the structure described in List 1.16(10) we would be able to generate two additional, equally valid, sets of eigenvectors for this adjacency matrix, which would be different to the one displayed in Figure 1.4. ◁

1.4 Spectrum of the unnormalised graph Laplacian

Moving away from the adjacency matrix A, we now present the unnormalised graph Laplacian, which is equal to the difference of the degree and the adjacency matrix, $L = D - A$. In Section 1.4.1, we start this section off by deriving the quadratic form of the unnormalised graph Laplacian in (1.4), which leads on to List 1.22 of the known eigenproperties of the unnormalised graph Laplacian. Then in Section 1.4.2 we review the Cheeger inequality and its recent extension for the unnormalised Laplacian. In Section 1.4.3 we discuss the mapping of the vertex set to locations determined by a subset of unnormalised Laplacian eigenvectors; the key point to understand in this section is displayed in Equation (1.5), where the Laplacian eigenvector entries are shown to be perturbed averages of the entries in their neighbourhood.

1.4.1 Eigenproperties of the unnormalised Laplacian

Recall from Section 1.1 that there is a rough correspondence between the larger end of the spectrum of the adjacency matrix A and the lower end of the spectrum of the unnormalised graph Laplacian L. Therefore, we order the eigenvalues of the unnormalised graph Laplacian in opposite order to the ones of the adjacency matrix and denote them by $\lambda_1 \leq \lambda_2 \leq ... \leq \lambda_n$.

Two early papers working with the unnormalised graph Laplacian not even referring to it as such are by Donath and Hoffman (1973) and Fiedler (1973). Donath and Hoffman (1973) prove bounds for the number of edges being cut in a k-partition of a graph in terms of eigenvalues of the unnormalised graph Laplacian. Fiedler (1973) emphasises the importance of the second eigenvalue of the unnormalised graph Laplacian, henceforth often referred to as the Fiedler

eigenvalue in several theoretical results. His results enforce the statements made about the spectral gap of the adjacency matrix in Remark 1.18. Note that measures like the spectral gap of the adjacency matrix and the Laplacian of a graph will be loosely related for general graphs in Section 2.10.1. Fiedler most notably relates the second eigenvalue to the minimal number of vertices or edges, which need to be removed in order for the graph to lose connectivity, and proves it to be non-decreasing as edges are added to a graph. In the context of clustering his work can hence be understood to illustrate the significance of the second eigenvector for partitioning a graph in two pieces or defining two clusters in the graph as mentioned in von Luxburg (2007) and in Remark 1.18 of Section 1.3.

In the following we derive the quadratic form of the unnormalised graph Laplacian. Especially notable when talking about quadratic forms and Laplacian eigenvectors is the book by Biyikoğu et al. (2007). The entire book is a theoretical approach to the eigenvectors of the graph Laplacian. Amongst other things they work on the connection between the eigenfunctions of the continuous Laplace-Beltrami operators on manifolds and the eigenvectors of the graph Laplacian on graphs. Viewing the Laplacian as a discretisation of the Laplace–Beltrami operator is done in several other publications as for example Belkin and Niyogi (2003).

$$
\begin{aligned}
L &= D - A, \\
\Rightarrow L\mathbf{f} &= D\mathbf{f} - A\mathbf{f}, \\
\Rightarrow (L\mathbf{f})_i &= d_i f_i - \sum_{j=1}^{n} a_{ij} f_j = \sum_{j=1}^{n} a_{ij} f_i - \sum_{j=1}^{n} a_{ij} f_j = \sum_{j=1}^{n} a_{ij}(f_i - f_j), \qquad (1.3) \\
\Rightarrow \mathbf{f}^T L\mathbf{f} &= \sum_{i=1}^{n} f_i \sum_{j=1}^{n} a_{ij}(f_i - f_j) = \sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}(f_i^2 - f_j f_i) \\
&= \sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}\frac{1}{2}(f_i^2 + f_j^2 - 2f_j f_i) = \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}(f_i - f_j)^2. \qquad (1.4)
\end{aligned}
$$

The sum in Equation (1.3) only includes terms in the neighbourhood of vertex v_i. Hence applying the graph Laplacian to a signal $\mathbf{f}$ produces a vector where each term corresponding to a certain vertex v_i is a weighted difference of the signal at vertex v_i and the signal in its neighbourhood.

Equation (1.4) is commonly found in the literature and is sometimes referred to as the total variation of a graph signal or Dirchlet energy (Mateos et al., 2019, p. 19). It is understood to be a smoothness measure of a signal over a given graph. In the context of the continuous operator analogue of the unnormalised Laplacian, Shuman et al. (2013, p. 87) derive the quadratic form in Equation (1.4) by taking the Euclidean norm of the gradient of a graph signal. Dong et al. (2016, p. 6164), Mateos et al. (2019, p. 26) and Dong et al. (2019, p. 51) show how Equation (1.4) can be incorporated into the cost function, when learning a graph, in order to impose signal smoothness over the learned graph.

The result in Equation (1.4) is also presented in von Luxburg (2007). From it follow two key properties of L mentioned in a multitude of papers such as von Luxburg (2007), Spielman (2007) and others.

List 1.22.

(1) L is positive semi-definite since all edge weights are non-negative.

(2) If the graph consists of a single connected component the smallest eigenvalue of L is 0 and the corresponding eigenvector is $\mathbf{1} = (1, ..., 1)^T$ otherwise the eigenvectors are as described in item (3) of this list.

We can furthermore add to this list that:

(3) The multiplicity of the 0 eigenvalue is equal to the number of connected components in G and the corresponding eigenvectors are indicator vectors establishing which vertex is element of which connected component (von Luxburg, 2007).

(4) For undirected graphs L is symmetric (it is the sum of two symmetric matrices).

(5) The spectrum of L is contained in $[0, 2d_{\max}]$ (van Mieghem, 2011, p. 68). As shown in Appendix A.2, this immediately follows from the Gershgorin theorem. van Mieghem (2011, p. 72) adds to this that the largest unnormalised Laplacian eigenvalue is bounded above by n, $\lambda_n \leq n$.

(6) The only rational eigenvalues of the adjacency matrix and unnormalised Laplacian are integers. "For example, $\frac{3}{4}$ will never be an eigenvalue of [...] L." (van Mieghem, 2011, p. 29)

(7) Let G be a connected graph with diameter Δ. Then, the unnormalised Laplacian matrix L of G has at least $\Delta+1$ distinct eigenvalues (Brouwer and Haemers, 2011, p. 5).

(8) The Laplacian spectrum of a graph and its complement are related as follows $\lambda_{n-j+1}(\overline{G}) = n - \lambda_j(G)$ and the corresponding eigenvectors are identical (van Mieghem, 2011, p. 72) .

(9) The sum of eigenvalues equals twice the number of edges, $\sum_{i=1}^{n} \lambda_i = 2|E|$ (Brouwer and Haemers, 2011, p. 4; van Mieghem, 2011, p. 68).

(10) The sum of the j largest unnormalised Laplacian eigenvalues is bounded below by the sum of the j largest degrees, $\sum_{i=1}^{j} d_{(i)} \leq \sum_{i=n-j+1}^{n} \lambda_i$ for $j \in \{1, \ldots, n\}$, where $d_{(i)}$ denotes the i^{th} largest degree. The proof follows from Schur's Majorisation Theorem (Zumstein, 2005, p. 44; van Mieghem, 2011, p. 68)

(11) For a connected graph the j^{th} largest unnormalised Laplacian eigenvalue can be lower-bounded by a function of the j^{th} largest degree, $\lambda_{n-j+1} \geq d_{(j)} - j + 2$ for $j \in \{1, \ldots, n-1\}$, where $d_{(j)}$ denotes the j^{th} largest degree (Brouwer and Haemers, 2008).

(12) The maximum eigenvalue of the unnormalised Laplacian matrix λ_n is bounded below by the maximum degree $d_{\max}$, i.e., $\lambda_n \geq d_{\max}$ (Zumstein, 2005, p. 18).

(13) There exist the following two upper bounds on the second smallest unnormalised Laplacian eigenvalue. Firstly, $\lambda_2 \leq \frac{d_u + d_v}{2}$ for any $u \neq v$ (van Mieghem, 2011, p. 81). And secondly, $\lambda_2 \leq \frac{n}{n-1} d_{\min}$. This bound is attained on the complete graph and the inequality can be sharpened even further to produce, $\lambda_2 \leq d_{\min}$ (van Mieghem, 2011, p. 82).

(14) The number of all possible spanning trees, denoted ξ, of a graph with n vertices is equal to the product of all Laplacian eigenvalues except for the first eigenvalue divided by n, i.e., $\xi = \frac{1}{n} \prod_{j=2}^{n} \lambda_j$. In van Mieghem (2011, p. 76) ξ is referred to as the complexity of a graph and in Butler (2014, p. 19) this result is referred to as Kirchoff's matrix tree theorem (Godsil and Royle, 2001, p. 284; van Mieghem, 2011, p. 77; Butler, 2014, p. 19). ◁

1.4.2 Cheeger inequality

In this section we discuss the relation of the Laplacian spectrum to the expansion ratio as was done for the adjacency spectrum in Section 1.3.2. Then we show a recent extension of this relationship and discuss its implication on using the Laplacian eigenvalues to indicate the number of homophilic clusters in a graph.

In Theorem 1.23 we relate the second smallest Laplacian eigenvalue to the expansion ratio of a graph which was defined in Definition 1.12. Note that unlike the equivalent result for the adjacency spectrum in Theorem 1.17, which only applied to d-regular graphs, the result in Theorem 1.23 applies to general graphs.

Theorem 1.23. *Cheeger inequality* (Stanić, 2015, p. 186; van Mieghem, 2011, p. 95) Let G be a connected graph with $n > 3$. Then,

$$\frac{h(G)^2}{2d_{\max}} \leq \lambda_2 \leq 2h(G). \qquad \triangleleft$$

The first half of the inequality in Theorem 1.23 follows from the provided reference by spotting that $h(G) \leq \sqrt{2d_{\max}\lambda_2 - \lambda_2^2} \leq \sqrt{2d_{\max}\lambda_2}$ and rearranging this inequality to give the one in Theorem 1.23.

From Theorem 1.23 it follows that the family of expander graphs can also be defined based on the second smallest eigenvalues of the Laplacian matrix instead of the expansion ratio by rescaling the constant lower-bounding the expansion ratio. Furthermore, following the argument in Remark 1.18, which was made in the context of the adjacency matrix, Theorem 1.23 motivates the use of the second smallest Laplacian eigenvalue to indicate the presence of 2 large homophilic clusters in the network.

Rather recently the relation of the expansion ratio to the Laplacian spectrum has been extended. In Theorem 1.24 we present this extension together with the definition of the generalised expansion ratio.

Theorem 1.24. (Lee et al., 2014) For every graph G and every $r \in \mathbb{N}$, let

$h_r(G)$ denote the r-way expansion constant defined to be,

$$h_r(G) = \min_{\mathcal{S}_1,\mathcal{S}_2,\ldots,\mathcal{S}_r} \max_{1\leq j\leq r} \frac{|\partial \mathcal{S}_j|}{|\mathcal{S}_j|},$$

where the subsets $\mathcal{S}_1, \mathcal{S}_2, \ldots \mathcal{S}_r \subseteq V$ are nonempty and disjoint; but, not necessarily exhaustive in V. Then, we have

$$\frac{h_r(G)^2}{O(r^4)} \leq \lambda_r \leq 2h_r(G). \qquad \triangleleft$$

Note that Theorem 1.23 is a special case of Theorem 1.24 with $r = 2$.

As stated in Lee et al. (2014), Theorem 1.24 acts as theoretical justification for clustering networks using the eigenvectors corresponding to the r smallest unnormalised Laplacian eigenvalues. It extends the insight from Theorem 1.23, which justified the use of the second eigenvalue to indicate the presence of 2 homophilic clusters in the network, to the r^{th} smallest Laplacian eigenvalue indicating the presence of r homophilic clusters in the network.

1.4.3 Using Laplacian eigenvectors for graph drawing

We begin this section by giving a graph theoretic definition of graph embeddings and then noting that, in line with much of the literature, we use a more relaxed version of the term graph embedding in this thesis. Then, we briefly review an argument from Spielman (2007), which motivates using the Laplacian eigenvectors for the purpose of graph embeddings.

Definition 1.25. (Bondy and Murty, 2008) An *embedding* of a graph $G = (V, E)$ on a surface T is a function taking each vertex v_i of G to a point $q(v_i)$ of T, and each edge v_iv_j of G, between any two vertices v_i and v_j, to an arc in T, with endpoints $q(v_i)$ and $q(v_j)$, in such a way that the only intersection between the points and arcs in the surface are those corresponding to incidences between edges and vertices of G. $\triangleleft$

Graphs which are embeddable in the plane are called *planar* graphs.

Example 1.26. To visualise the concept of planar graphs we go back to Figures 1.1(d) and 1.1(c) from Section 1.2. From Figure 1.1(d) we immediately

that the graph C_5 is planar as none of the edges intersect except at nodes. However, we see that in Figure 1.1(c) several edges intersect and it turns out that there exists no way to draw K_5 in the plane such that edges only intersect at nodes; hence K_5 is not a planar graph. Interestingly, it is has been shown that any graph can be embedded into $\mathbb{R}^3$ (Cohen et al., 1995). ◁

Remark 1.27. In the literature, e.g., von Luxburg (2007), Mateos et al. (2019, p. 25) and Spielman (2007), embeddings of graphs are often understood to refer to any map from a graph to some space, i.e., edges are allowed to intersect in the image of this map. In this thesis use this more relaxed definition of graph embeddings rather than the graph theoretic definition given in Definition 1.25. ◁

Spielman (2007) outlines how the eigenvectors corresponding to the smallest unnormalised Laplacian eigenvalues provide a good mapping of a graph into lower dimensional spaces. The reasoning follows from the eigenvalue eigenvector equations,

$$\begin{aligned}
(L\ \chi_2)_i &= \lambda_2\ \chi_{2i} \\
d_i\chi_{2i} - \sum_{j=1}^{n} a_{ij}\chi_{2j} &= \lambda_2\chi_{2i} \\
(d_i - \lambda_2)\chi_{2i} &= \sum_{j=1}^{n} a_{ij}\chi_{2j} \\
\Rightarrow \qquad \chi_{2i} &= \frac{1}{d_i - \lambda_2} \sum_{j:v_j\in\mathcal{N}(v_i)} a_{ij}\chi_{2j}, \qquad (1.5)
\end{aligned}$$

where $\mathcal{N}(v_i)$ denotes the neighbourhood of vertex v_i in the graph. When the mapping from the i^{th} vertex to (χ_{2i}, χ_{3i}) is considered, then, Equation (1.5) shows that each vertex is mapped to a location which is formed as an average, disregarding λ_2 for a second, of the vertex locations in its neighbourhood. λ_2 and λ_3 can be interpreted as perturbations of this mapping of vertices to the average location of the vertices in their graph neighbourhood. This mapping produces graphs which appear to be close to planar. The spectral clustering algorithm operates in the space spanned by the eigenvectors corresponding to the smallest unnormalised Laplacian eigenvalues. In fact it runs a k-means algorithm on the image of the mapping, which maps the i^{th} vertex to $(\chi_{1i}, ..., \chi_{ki})$.

Hence, it exploits exactly the close to planar view provided by this mapping as it is possible to spot small cuts in this space. We will use the spectral clustering algorithm in Section 2.12 to visualise the impact of the representation matrix choice on inference drawn from graphical analysis. This embedding of graphs utilised in the spectral clustering algorithm is commonly used in the latent space modelling literature introduced in Section 1.2.3 to estimate latent vertex positions.

A similar argument is made in Shuman et al. (2013); they interpret the Laplacian eigenvectors as signals on graphs, by taking each component of the eigenvector to be a function value at the corresponding vertex in the graph. They then define the number of zero crossings of an eigenvector on a graph to be the number of edges, which connect a vertex with a positive eigenvector entry to a vertex with a negative eigenvector entry. Then they count these zero crossings for each Laplacian eigenvector on a random signal graph and observe that larger eigenvalues correspond to eigenvectors with a larger number of zero crossings. From this they are able to derive the notion of frequency for the eigenvalues as the eigenvectors corresponding to small eigenvalues vary slowly across the graph and hence correspond to low frequency signals. This observation is explained by the eigenvalue eigenvector rational described by Spielman (2007), which was that the eigenvector values are averages of their neighbouring eigenvector values only perturbed by their corresponding eigenvalue. As this perturbation grows the smoothness of the eigenvector on the graph is perturbed by a larger amount and hence the eigenvector turns into a higher frequency signal. This connection of the Laplacian spectrum to the notion of frequency of a signal nicely justifies the use of the word spectrum for the set of eigenvalues. Viewing the Laplacian eigenvalues to correspond to signal frequencies and the corresponding eigenvectors as graph frequency components is a fundamental piece of theory in the graph signal processing literature and can also be found in Sandryhaila and Moura (2014, p. 84), Deri and Moura (2017, p. 785), Mateos et al. (2019, p. 19) amongst many others. In Ortega et al. (2018, pp. 814-5) the frequency interpretation of the eigenvalues and eigenvectors is applied to all three representation matrices.

Remark 1.28. This frequency interpretation of the Laplacian eigenvectors can be used as one explanation for us recovering homophilic clusters from eigenvectors corresponding to eigenvalues on one extreme end of the spectrum

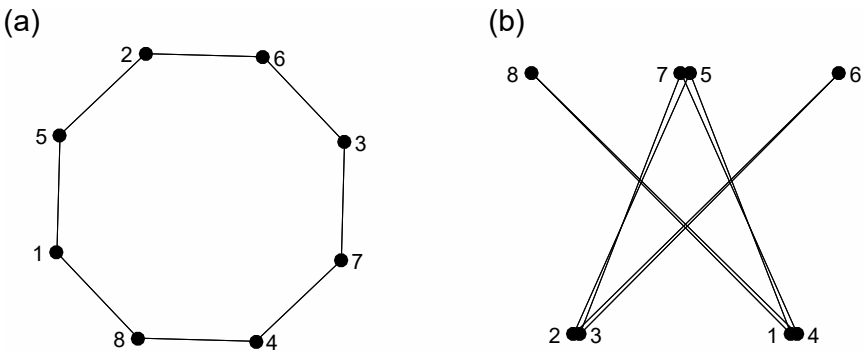

Figure 1.5: Plotting C_8 according to different eigenvectors of the unnormalised Laplacian: In (a) we display the mapping of the vertex set of C_8 to (χ_2, χ_3) and in we show the mapping of the vertex set of C_8 to (χ_7, χ_8).

and heterophilic clusters from eigenvectors corresponding to the other extreme end of the spectrum. Laplacian eigenvectors corresponding to small eigenvalues, i.e., low frequencies, vary slowly across the connection pattern of the graph and therefore well connected groups in the network have similar eigenvector values, which leads to us recovering homophilic cluster structures (discussed in Section 1.3.2). Contrariwise, Laplacian eigenvectors corresponding to large eigenvalues, i.e., high frequencies, vary rapidly across the connection pattern of the graph and therefore well connected groups in the network have dissimilar eigenvector values, which leads to us recovering the less well connected heterophilic cluster structures (discussed in Section 1.3.3). ◁

In Appendix A.3 we reproduce Figure 1 in von Luxburg (2007) and in doing so visualise the concept of graph embeddings and the frequency interpretation of the unnormalised Laplacian eigenvalues.

The interpretation of Equation (1.5) by Spielman (2007) leads on to the much wider topic of graph drawing. In Example 1.29 we explore the drawing of graphs according to their Laplacian eigenvectors.

Example 1.29. In Figure 1.5 we show two plots of the graph C_8. In Figure 1.5(a) we display the mapping of the vertex set to (χ_2, χ_3) and in Figure 1.5(b) the mapping of the vertex set to (χ_7, χ_8), which are the eigenvectors of L corresponding to its largest eigenvalues.

We see that Figure 1.5(a) provides a planar view of the graph and that from Figure 1.5(b) it is obvious that C_8 is a bipartite graph. These plots support

the clustering theory discussed Sections 1.3.2 and 1.3.3 and Remark 1.28, that clustering in the space of the eigenvectors corresponding to the smallest eigenvalues provides homophilic clusters with many edges within and few between and hence that we cut a minimal number of edges in order to disconnect the two clusters. While looking at the bipartite graph in Figure 1.5(b) enforces the argument, that clustering in the space of the eigenvectors corresponding to the largest eigenvalues provides heterophilic clusters with few edges within and many between and hence that we cut a large number of edges in order to disconnect the two clusters. Visually it also supports the fact stated in Section 1.2, that cycle graphs on an even number of vertices are bipartite graphs. ◁

The paper by Koren (2005) justifies this spectral approach to graph drawing by studying three different viewpoints of the problem of graph drawing. One of these viewpoints is equivalent to the point we take from Spielman (2007) in Equation (1.5). MATLAB has confirmed to us, in answer to a question we posted on `Mathworks.com`, that the theory laid out in Koren (2005) was used when their 'subspace' plotting routine was built.

We want to end on noting that there is a multitude of other approaches to graph drawing some of which are more heuristic and visual than this one. These usually place greater emphasis on the well separateness of nodes and edges. In terms of clustering this is not such an important issue. However, these drawing approaches help visual understanding of graphs, especially for graphs with a small number of nodes. For example, the graphs in Section 1.2 Figure 1.1(a) and Figure 1.5(b) both show the "bipartite view" of C_8. Figure 1.1(a) provides a clearer view of the graph.

1.5 Spectrum of the normalised graph Laplacian

We begin this section by defining the two normalised Laplacians and discussing their relationship in Section 1.5.1. Then we derive their quadratic forms as far as possible and give a list of their eigenproperties in Section 1.5.2. In Section 1.5.3 we review a version of Cheeger's inequality and the expander mixing

lemma which applies to the normalised Laplacians.

Commonly in the literature there are two normalised Laplacians, which are considered. They both follow from the unnormalised graph Laplacian by multiplying by the inverse of the degree matrix as follows,

$$L_{rw} = D^{-1}L; \qquad L_{sym} = D^{-1/2}LD^{-1/2}.$$

The book by Chung (1997) with the title "Spectral Graph Theory" relates the spectrum of L_{sym} to several graph invariants "in a way that other definitions (such as the eigenvalues of the adjacency matrices) often fail to do". In her proofs Chung makes heavy use of the *Rayleigh quotient* of L_{sym} defined as

$$\frac{\mathbf{f}^T L_{sym}\mathbf{f}}{\mathbf{f}^T\mathbf{f}},$$

where $\mathbf{f}$ is a signal on a graph. The minimum of this quotient (minimised over the space orthogonal to the one spanned by all eigenvectors corresponding to the $r-1$ smallest eigenvalues) is equal to the r^{th} smallest eigenvalue. Later on in this section we will present some of the results Chung (1997) derived.

1.5.1 The relationship of the two normalised Laplacians

We will begin our analysis by investigating how the spectra of the two normalised Laplacians are related to each other. First off, we define the concept of similarity as it can be found in Horn and Johnson (1985).

Definition 1.30. A matrix $B \in \mathbb{C}^{n\times n}$ is said to be *similar* to a matrix $A \in \mathbb{C}^{n\times n}$ if there exists a nonsingular matrix $S \in \mathbb{C}^{n\times n}$ such that

$$B = S^{-1}AS.$$

The transformation $A \to S^{-1}AS$ is called a *similarity transformation* by the *similarity matrix* S. ◁

From the functional definitions of the two normalised Laplacians it can be seen that L_{rw} is similar to L_{sym} via the similarity matrix $D^{1/2}$. Results for the

spectrum and the eigenvectors of similar matrices can be found in Horn and Johnson (1985, pp. 45, 60). We summarise them as follows:

Proposition 1.31. If $A, B \in \mathbb{C}^{n\times n}$ and if B is similar to A via S, then they have the same eigenvalues, counting multiplicity. Further, if $x \in \mathbb{C}^n$ is an eigenvector corresponding to an eigenvalue β of B, then Sx is an eigenvector of A corresponding to eigenvalue β. $\triangleleft$

Hence, L_{rw} and L_{sym} share eigenvalues and their eigenvectors are related via a transform by their similarity matrix $D^{1/2}$, i.e., if ζ is an eigenvector of L_{rw}, then, $D^{1/2}\zeta$ is an eigenvector of L_{sym}. We denote their shared spectrum by $\eta_1 \leq ... \leq \eta_n$. Note that the spectrum follows the eigenvalue indexing set out for the unnormalised Laplacian L in Section 1.4.

The existence of a real set of eigenvectors of L_{rw} and numerical issues in the calculation of its eigenvalues are discussed in Appendix A.4.

In Chapter 2 we work with the eigenvalues of the representation matrices. Since the two normalised Laplacians have a common eigenvalue set, the work in Chapter 2 applies to both normalised Laplacians. In Chapters 3 and 4 we work with the eigenvectors of the representation matrices. The results in Chapter 3 require symmetric matrices. Therefore, the work in Chapters 3 and 4 is only applicable to the eigenvectors of L_{sym}.

1.5.2 Eigenproperties of the normalised Laplacians

We begin by deriving the quadratic form of L_{sym} in order to get a better understanding of its interaction with a signal on a graph.

$$
\begin{aligned}
L_{sym} &= D^{-1/2} L D^{-1/2} = I - D^{-1/2} A D^{-1/2}, \\
\Rightarrow L_{sym}\mathbf{f} &= \mathbf{f} - D^{-1/2} A D^{-1/2}\mathbf{f}, \\
\Rightarrow (L_{sym}\mathbf{f})_i &= f_i - \sum_{j=1}^{n} \frac{1}{\sqrt{d_i d_j}} a_{ij} f_j = \frac{1}{\sqrt{d_i}} \left(\frac{d_i}{\sqrt{d_i}} f_i - \sum_{j=1}^{n} \frac{1}{\sqrt{d_j}} a_{ij} f_j \right) \\
&= \frac{1}{\sqrt{d_i}} \left(\sum_{j=1}^{n} \frac{1}{\sqrt{d_i}} a_{ij} f_i - \sum_{j=1}^{n} \frac{1}{\sqrt{d_j}} a_{ij} f_j \right) \\
&= \frac{1}{\sqrt{d_i}} \sum_{j=1}^{n} a_{ij} \left(\frac{f_i}{\sqrt{d_i}} - \frac{f_j}{\sqrt{d_j}} \right), \qquad (1.6) \\
\Rightarrow \mathbf{f}^T L_{sym}\mathbf{f} &= \sum_{i=1}^{n} f_i \frac{1}{\sqrt{d_i}} \sum_{j=1}^{n} a_{ij} \left(\frac{f_i}{\sqrt{d_i}} - \frac{f_j}{\sqrt{d_j}} \right) = \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \left(\frac{f_i^2}{d_i} - \frac{f_j f_i}{\sqrt{d_i d_j}} \right) \\
&= \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \left(\frac{1}{2} \frac{f_i^2}{d_i} + \frac{1}{2} \frac{f_j^2}{d_j} - \frac{f_j f_i}{\sqrt{d_i d_j}} \right) \\
&= \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} a_{ij} \left(\frac{f_i}{\sqrt{d_i}} - \frac{f_j}{\sqrt{d_j}} \right)^2. \qquad (1.7)
\end{aligned}
$$

From Equation (1.6) we deduce that applying L_{sym} to a signal at vertex v_i produces a weighted difference of the signal at v_i and the signal in its neighbourhood $\mathcal{N}(v_i)$. Here the weighting is by the edge weights and each signal is normalised by the square root of the corresponding vertex degrees and then separately by the square root of the vertex degree at v_i. The normalisation by vertex degrees is more straight forward in (1.8) (the equivalent form of (1.6) for L_{rw}), where each signal is normalised by the vertex degree at which it occurs.

$$
(L_{rw}\mathbf{f})_i = \sum_{j=1}^{n} a_{ij} \left(\frac{f_i}{d_i} - \frac{f_j}{d_j} \right). \qquad (1.8)
$$

We are not able to reach a quadratic sum form from $\mathbf{f}^T L_{rw}\mathbf{f}$; we cannot progress further than: $\mathbf{f}^T L_{rw}\mathbf{f} = \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}\left(\frac{f_i^2}{d_i} - \frac{f_i f_j}{d_j}\right)$. In contrast, for L_{sym} we reached a nice quadratic form in Equation (1.7), where now the normalisation by vertex degrees is less confused than it was in (1.6) as each signal is normalised by the square root of the vertex degree at which it occurs.

In order to get more familiar with these normalised Laplacians we have compiled a list of their properties below.

List 1.32.

(1) For undirected graphs L_{sym} is symmetric positive semi-definite as follows from Equation (1.7). That L_{rw} is also positive semi-definite follows from the fact that it shares its eigenvalues with L_{sym}.

(2) The smallest eigenvalue of L_{rw} and L_{sym} is equal to 0. Its multiplicity r is equal to the number of connected components in G and the corresponding eigenvectors are $\mathbf{1}_{\mathcal{S}_k}$ for L_{rw} and $D^{1/2}\mathbf{1}_{\mathcal{S}_k}$ for L_{sym}, where $\mathbf{1}_{\mathcal{S}_k}$ is the indicator vector of the connected component $\mathcal{S}_k$ as defined in Section 1.3 (von Luxburg, 2007).

(3) The spectrum of L_{sym} and L_{rw} is contained in $[0, 2]$ (van Mieghem, 2011, p. 64). As shown in Appendix A.2, this immediately follows from the Gershgorin theorem.

(4) The multiplicity of the eigenvalue 2 is equal to the number of bipartite connected components of G with at least two vertices (Butler and Chung, 2017, p. 47-3) .

(5) The unnormalised Laplacians have at least $\Delta + 1$ distinct eigenvalues (Butler and Chung, 2017, p. 47-3).

(6) A graph without vertices of degree 0 is bipartite if and only if the spectrum of the normalised Laplacians is symmetric around 1 (Butler and Chung, 2017, p. 47-3).

(7) $\sum_{i=1}^{n} \eta_i \leq n$ with equality holding if and only if G has no vertices of degree 0 (Chung, 1997, p. 6).

(8) For $n \geq 2$, $\eta_2 \leq \frac{n}{n-1}$ with equality holding if and only if G is the complete graph on n vertices. Also, for a graph without vertices of degree 0, we have $\eta_n \geq \frac{n}{n-1}$(Chung, 1997, p. 6).

(9) For a graph which is not a complete graph we have $\eta_2 \leq 1$ (Chung, 1997, p. 7).

(10) $\eta_n = 2$ if and only if a connected component of G is bipartite and non-trivial (Chung, 1997, p. 7). ◁

As we saw in Point (3) of List 1.32 the support of the eigenvalues of all normalised Laplacians is equal to $[0, 2]$ and does not depend on the graphs degree distribution. *Hence, the bounds on the eigenvalues in List 1.32 are only functions of the number of nodes n and are also independent of the degree distribution.* This independence of the degree distribution enables the comparison of spectra of general graphs on the same scale. Note that this was not the case for the adjacency matrix and the unnormalised graph Laplacian, where by the Gershgorin Theorem the support of the eigenvalues is equal to $[-d_{\max}, d_{\max}]$ and $[0, 2d_{\max}]$, respectively, and hence dependent on the degree distribution. That this dependence on the degree distribution also carries through to the bounds on the eigenvalues can be seen in Lists 1.16 and 1.22.

The normalised graph Laplacian eigenvalue's independence of the degree distribution carries over to the eigenvectors only for L_{rw}. As we saw in List 1.32 Point (2) the eigenvectors of L_{sym} do depend on the degree distribution. This dependence is the argument von Luxburg (2007) gives on why L_{rw} is to be preferred over L_{sym} when using the normalised Laplacians for clustering on graphs.

1.5.3 A variation of the Cheeger inequality and the expander mixing lemma

For the adjacency and unnormalised Laplacian spectra we presented their relations to the expansion ratio of a graph in Sections 1.3.2 and 1.4.2. In this section we present the relation of Butler's variant of the expansion ratio to the normalised Laplacian spectrum and also present a version of the expander mixing lemma for the normalised Laplacian spectrum.

In Theorem 1.33 we present the relation of Butler's expansion ratio defined in (1.2) to the normalised Laplacian spectrum.

Theorem 1.33. (Butler and Chung, 2017, p. 47-11) Let G be a weighted,

connected but not complete graph, then,

$$1 - \sqrt{1 - (h'(G)^2)} < \eta_2 \leq 2h'(G). \qquad \triangleleft$$

As discussed in Section 1.2.2, $h'(G)$ is not defined using the number of neighbours of a given vertex, but rather how well a given vertex is connected to these neighbours by taking the degree of the vertices into consideration. This variation does not change the interpretation that Theorem 1.33 can be used to motivate using the second smallest normalised Laplacian eigenvalue to indicate the presence of two large homophilic clusters in the graph.

Butler (2014) also produces a version of the expander mixing lemma which applies to the normalised Laplacian spectrum.

Lemma 1.34. (Butler, 2014, p. 24) For any graph G, for all $\mathcal{S}_1, \mathcal{S}_2 \subseteq V$ we have,

$$\left| E(\mathcal{S}_1, \mathcal{S}_2) - \frac{\text{vol}(\mathcal{S}_1)\ \text{vol}(\mathcal{S}_2)}{\text{vol}(V)} \right| \leq \max(|1 - \eta_2|, |1 - \eta_n|) \sqrt{\text{vol}(\mathcal{S}_1)\ \text{vol}(\mathcal{S}_2)}. \qquad \triangleleft$$

Note that unlike its adjacency matrix counterpart (Lemma 1.20), which only applied to d-regular graphs, Lemma 1.34 holds for general graphs. As in Section 1.3.2 eigenvalues on both ends of the spectrum can be seen to determine the overall regularity of the connectivity of the graph in Lemma 1.34.

1.6 The interaction of graph operations and the representation spectra

In this section we briefly review the interaction of several operations on graphs with the spectra of the graphs. The operations we review are the addition and removal of connected components in Section 1.6.1, the subgraph relation in Section 1.6.2 and the relation via edge contraction in Section 1.6.3. In Chapter 2 we make use of the relation of the representation spectra via the addition and removal of connected components.

We point readers interested in the effect of the operations of motif doubling,

graph splitting and joining and their effect on the representation spectra to Banerjee and Jost (2008), where their interactions with the normalised Laplacian spectrum are discussed. These operations seem to be particularly related to the eigenvalue 1 in the normalised Laplacian spectrum, which will play a prominent role in the upcoming Section 1.7.

1.6.1 Removal and addition of connected components

In this section we review the interaction of the representation spectra with the addition and removal of connected components to the graph. Connected components in a graph were defined in Definition 1.7. We begin by citing a lemma by Chung (1997), which applies to the normalised Laplacian. Then, we discuss its proof and observe that it extends to the other representation matrices treated in this thesis.

Lemma 1.35. (Chung, 1997, p. 7) The normalised Laplacian spectrum of a graph is the union of the normalised Laplacian spectra of its connected components, i.e., the union of the sets of normalised Laplacian eigenvalues corresponding to the connected components. ◁

Chung (1997) does not give a proof. However it trivially follows. Any of the three representation matrices of a graph containing several connected components can be arranged to be block diagonal by definition of connected components. Now the argument is completed by the fact that the characteristic equation of a block diagonal matrix factorises into polynomials corresponding to the different blocks or connected components (Bernstein, 2009, p. 291). This proof can be applied to all three representation spectra. Hence, we have the following result.

Lemma 1.36. The representation matrix spectrum of a graph is the union of the representation matrix spectra of its connected components, i.e., the union of the sets of eigenvalues of the representation matrices corresponding to the connected components. ◁

Lemma 1.36 will allow us to explain the monotonicity observed in our eigenvalue bounds in Section 2.7.

1.6.2 The interlacing theorems

The interlacing theorems can be used to relate the representation matrix eigenvalues of a graph and its subgraphs. One of the first occurrences of the interlacing theorems in the context of graphs together with a rich discussion of its many uses can be found in Haemers (1995). We begin the discussion by formally defining the subgraph relationship in the following.

Definition 1.37. (Bondy and Murty, 2008, p. 40) Let G be a graph with vertex set V and edge set E. Then a *subgraph* G' of G is formed either by removing a subset of vertices from V and removing all edges incident to the removed vertices from E or by removing a subset of edges from E. ◁

Now we are able to state one of the interlacing theorems.

Theorem 1.38. (Zumstein, 2005, p. 39) Let G' be a subgraph of a graph G, where one vertex has been removed. Then,

$$\mu_1(G) \leq \mu_1(G') \leq \mu_2(G) \leq \mu_2(G') \leq \ldots \leq \mu_{n-1}(G') \leq \mu_n(G). \qquad \triangleleft$$

Therefore, the adjacency spectra of a graph and its subgraphs are closely related. This inequality generalises to cases where several vertices or edges are deleted and to all three representation matrices. A good overview of the many interlacing theorems in the context of graphs is given in Porto and Allem (2017). Instances of this theorem in the context of graphs are also discussed in Stanić (2015, p. 11), van Mieghem (2011, p. 246), Godsil and Royle (2001, p. 193) and Johnson and Saiago (2018, p. 2).

Note that the interlacing theorem of the unnormalised Laplacian is a much stronger version of the result by Fiedler (1973), discussed in the introduction of Section 1.4.1, that the second eigenvalue of L is non-decreasing as edges are added.

1.6.3 Edge contractions

In this section we discuss how the contraction of edges interact with the normalised Laplacian spectrum. We begin by defining the contraction of an edge

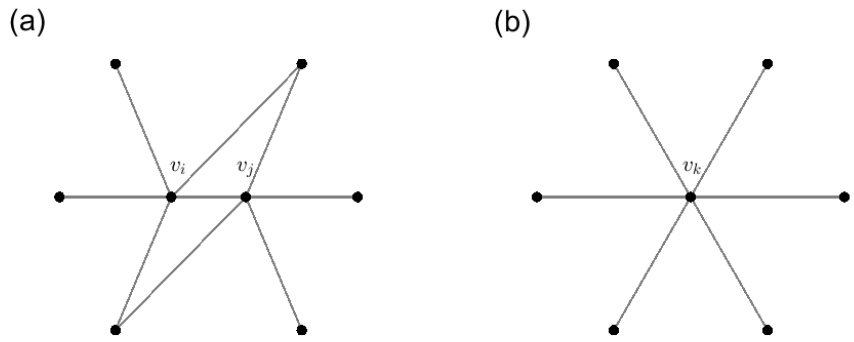

Figure 1.6: Example of an edge contraction: in (a) a graph is displayed. In (b) we display the graph arising from the graph in plot (a) when the edge between vertices v_i and v_j is contracted.

and giving an example of an edge contraction to visualise the concept. Then, we present a Lemma by Chung (1997), which describes the effect of edge contractions on the second eigenvalue of the normalised Laplacian.

From Bondy and Murty (2008) we know that *contracting an edge* between vertices v_i and v_j means removing vertices v_i and v_j and adding vertex v_k the neighbourhood of which is equal to $\mathcal{N}(v_i) \cup \mathcal{N}(v_j)$, where multiple edges are removed.

Example 1.39. An example of such a contraction is visualised in Figure 1.6. Here the edge between the vertices v_i and v_j in the left image is contracted. This leads to the graph in the right image, where vertices v_i and v_j have been removed, vertex v_k has been introduced and connected to all vertices in the neighbourhood of v_i and v_j and multiple edges have been removed. ◁

In Lemma 1.40 we present the main results of this subsection.

Lemma 1.40. (Chung, 1997, p. 13) Let $\eta_2(G)$ denote the second smallest normalised Laplacian eigenvalue of the graph G. If H is formed by edge con-

tractions from a graph G, then,

$$\eta_2(G) \leq \eta_2(H). \qquad \triangleleft$$

As was discussed in Section 1.5.3 a larger second eigenvalue corresponds to a more easily identified clustering into two components of the graph. It is surprising that contacting an edge seems to strengthen the evidence in the normalised Laplacian spectrum for the presence of a homophilic clustering into two clusters.

1.7 Twin vertices

In this section we formally define twin vertices in a graph and present Lemma 1.42, where the eigenstructure arising from the presence of twin vertices in a graph is shown. Finally, we present potential generalisations of Lemma 1.42.

We begin by formally defining twin vertices in a graph in Definition 1.41.

Definition 1.41. (Zumstein, 2005, p. 32) Two vertices v_i, v_j are called *twins* if for all other vertices v_k, v_k is connected to v_i if and only if v_k is connected to v_j. $\triangleleft$

In Bondy and Murty (2008, p. 15) twin vertices are referred to as similar vertices in the context of graph automorphisms. In Lemma 1.42 we discuss the eigenstructure of the representation matrices, which is implied by the presence of adjacent and nonadjacenct twin vertices in a graph.

Lemma 1.42. (Zumstein, 2005, p. 32) Let v_i, v_j be vertices in a graph G. We define the (Faria) vector $w(i,j)$ as

$$w(i,j)_k = \begin{cases} 1, & \text{if } k = i, \\ -1, & \text{if } k = j, \\ 0, & \text{otherwise.} \end{cases}$$

(1) Suppose v_i and v_j are nonadjacent twin vertices. Then, $\mu = 0$ is an eigenvalue of the adjacency matrix, $\lambda = d_i = d_j$ is an eigenvalue of the

unnormalised Laplacian and $\eta = 1$ is an eigenvalue of the normalised Laplacian L_{sym}, each with corresponding eigenvector $w(i, j)$.

(2) Suppose v_i and v_j are adjacent twin vertices. Then, $\mu = -1$ is an eigenvalue of the adjacency matrix, $\lambda = d_i + 1 = d_j + 1$ is an eigenvalue of the unnormalised Laplacian and $\eta = (d_i + 1)/d_i = (d_j + 1)/d_j$ is an eigenvalue of the normalised Laplacian L_{sym}, each with corresponding eigenvector $w(i, j)$. ◁

In Zumstein (2005) it is discussed how via Lemma 1.42 a complete set of eigenvalues and eigenvectors of complete graphs and stars can be obtained. Note that the adjacency and normalised Laplacian eigenvalue structure described in Lemma 1.42 can also be found in Butler and Chung (2017, p. 47-3). Teke and Vaidyanathan (2017) also work with the result in in Lemma 1.42 in their investigation of sparse eigenvector representation of graphs; they refer to the Faria eigenvectors as 2-sparse eigenvectors.

Remark 1.43. It is interesting to note that in the case of nonadjacent twin vertices in Lemma 1.42, the resulting eigenvalues of A and the normalised Laplacians are constant, while the resulting eigenvalue of L is dependent on the degree of the twin vertices. Therefore, the presence of twin vertices of very low or high degree might be manifested in eigenvalues on either extreme of the spectrum of the unnormalised Laplacian and therefore might misinform cluster structure inference, since a separation of single vertices into clusters might not be seen as an informative clustering. ◁

There exist several potential generalisations of Lemma 1.42 to larger sets of vertices sharing neighbours. First and foremost Zumstein (2005) goes on to prove the existence of an eigenvector in the spectrum of the symmetric normalised Laplacian, which is only nonzero on entries corresponding to vertex sets with an equal number of neighbourhoods in any given vertex subset in the graph. Another related idea can be found in van Mieghem (2011, p. 22), where he observes that an automorphism on a graph produces multiple eigenvectors belonging to the same eigenvalue. These eigenvectors are related via a permutation matrix, which acts as the permutation on the node set defined by the automorphism on the graph. Butler (2014, p. 24) and Godsil and Royle (2001, p. 196) discuss "equitable partitions" of the vertex set of a graph, which

is defined to be a partition $\mathcal{S}_1, \ldots, \mathcal{S}_r$ such that the number of neighbours in $\mathcal{S}_j$ of any vertex v in $\mathcal{S}_i$ is constant. Equitable partitions can be formed from orbits of graph automorphisms. Comparing these related, potential generalisations of Lemma 1.42 could lead to a result explaining how more global 'similar' structures, such as several homophilic or heterophilic clusters, are encoded in the eigenvectors of the representation matrices.

In Chapter 2 Remark 2.6 we observe twin vertices in the karate dataset, which was introduced in Section 1.2.5, and the resulting eigenvalue structure of the representation matrices. In Chapter 3 Example 3.3 we display several Faria eigenvectors of the adjacency matrix corresponding to a complete graph on 4 nodes.

1.8 Existing work on relating the representation matrix spectra

In this section we review the existing work on relating the representation matrix spectra. We begin with a curious result by Butler (2008) in Remark 1.44 and then move on to summarise results relating the representation matrix spectra via inequalities in Lemma 1.47.

In Remark 1.44 we quote a result from Butler (2008), where he relates the number of adjacency and normalised Laplacian eigenvalues within a certain portion of their spectral support.

Remark 1.44. Butler (2008) observes that it follows from Sylvester's Law of Inertia (not stated in this thesis) that the number of negative eigenvalues of the adjacency matrix is equal to the number of eigenvalues of the normalised Laplacians greater than 1, and the number of positive adjacency eigenvalues is equal to the number of eigenvalues of the normalised Laplacians smaller than 1. ◁

Now we turn to inequality relations of individual eigenvalues of the representation matrices. Little work has been done on relating the spectra of the

representation matrices of general graphs. All spectral relations of the representation matrices, which we found in the literature, can be proved using the linear algebra results stated in Theorems 1.45 and 1.46.

Theorem 1.45. *Weyl's inequality* (So, 1994) Let Φ and Ψ be $n \times n$ Hermitian matrices with eigenvalues $\phi_1 \geq \ldots \geq \phi_n$ and $\psi_1 \geq \ldots \geq \psi_n$, respectively. Denote the eigenvalues of their sum $\Phi + \Psi$ by $\tau_1 \geq \ldots \geq \tau_n$. Then,

$$\tau_{i+j-1} \leq \phi_i + \psi_j, \qquad \text{for } 1 \leq i, j;\ i + j - 1 \leq n \tag{1.9}$$

and

$$\phi_i + \psi_j \leq \tau_{i+j-n}, \qquad \text{for } 1 \leq i, j;\ i + j - n \leq n. \tag{1.10}$$

◁

Theorem 1.46. (So, 1994; Kolotilina, 2005) Let Φ and Ψ be positive semi-definite $n \times n$ Hermitian matrices with eigenvalues $\phi_1 \geq \ldots \geq \phi_n$ and $\psi_1 \geq \ldots \geq \psi_n$, respectively. Denote the eigenvalues of their product $\Phi\Psi$ by $\tau_1 \geq \ldots \geq \tau_n$. Then,

$$\tau_{i+j-1} \leq \phi_i \psi_j, \qquad \text{for } 1 \leq i, j;\ i + j - 1 \leq n \tag{1.11}$$

and

$$\phi_i \psi_j \leq \tau_{i+j-n}, \qquad \text{for } 1 \leq i, j;\ i + j - n \leq n. \tag{1.12}$$

◁

Lemma 1.47 summarises the relationships of the representation spectra, which we found in the literature. In its proof we show that all of these relationships can be proved using Theorems 1.45 and 1.46.

Lemma 1.47. Let G be a graph on n vertices. Then, for any $k \in \{1, \ldots, n\}$,

(1) $d_{\min} - \lambda_k \leq \mu_k \leq d_{\max} - \lambda_k$ (van Mieghem, 2011, p. 71). Note the eigenvalue ordering in this result follows our convention.

(2) $d_{(k)} - \mu_1 \leq \lambda_{n-k+1} \leq d_{(k)} - \mu_n$, where $d_{(k)}$ is the k^{th} largest degree in the graph (van Mieghem, 2011, p. 71).

(3) $d_{\max} - \mu_1 \leq \lambda_n \leq d_{\max} - \mu_n$ (Zumstein, 2005, p. 22).

(4) $d_{\min}\eta_k \leq \lambda_k \leq d_{\max}\eta_k$ (Zumstein, 2005, p. 22; Butler and Chung, 2017, p. 47-2).

Proof. Inequalities (1), (2) and (3) are proved using Theorem 1.45. We let $\Phi = D$, $\Psi = -A$ and hence, $\Phi + \Psi = L$. Therefore, $\phi_k = d_{(k)}$, $\psi_k = -\mu_{n-k+1}$, where the reordering of the indices is due to the negative sign in front of A reordering the magnitudes of the adjacency spectrum, and $\tau_k = \lambda_{n-k+1}$, where the reordering of the spectra is due to us ordering the unnormalised Laplacian spectrum in increasing order rather than in decreasing order as is the case for the eigenvalues in Theorem 1.45.

Now, (1) follows from choosing $i = 1$ and $k = n - j + 1$ in (1.9) and $i = n$ and $k = n - j + 1$ in (1.10). (2) follows from choosing $j = 1$ and $k = i$ in (1.9) and $j = n$ and $k = i$ in (1.10). (3) follows from (2) by choosing $k = 1$.

Inequality (4) is proved using Theorem 1.46. We let $\Phi = D^{-1}$, $\Psi = L$ and hence, $\Phi\Psi = L_{rw}$. Therefore, $\phi_k = 1/d_{(n-k+1)}$, where the reordering of the indices is due to the inverse of the diagonal matrix D reordering the magnitudes of the degrees, $\psi_k = -\lambda_{n-k+1}$ and $\tau_k = \eta_{n-k+1}$, where the reordering of the latter two spectra is due to us ordering the Laplacian spectra in increasing order rather than in decreasing order as is the case for the eigenvalues in Theorem 1.46.

Now, (4) follows from choosing $i = 1$ and $k = n - j + 1$ in (1.11) and $i = n$ and $k = n - j + 1$ in (1.12).

□

van Mieghem (2011, p. 71) states that Inequalities (1) and (2) in Lemma 1.47 can be proved via Weyl's inequality. While, Zumstein (2005) proves Inequalities (3) and (4) of Lemma 1.47 by making use of the Rayleigh quotient and the Courant-Fisher Theorem.

2

Comparing Spectra of Representation Matrices

An early version of the work in this chapter was published on arXiv in Lutzeyer and Walden (2017) and is cited in Bohannon et al. (2019), Dankulov et al. (2019) and Clark and Macdonald (2018). Parts of this chapter are a pre-copyedited version of a contribution published in Complex Networks and Their Applications VIII edited by Hocine Cherifi, Sabrina Gaito, José Fernendo Mendes, Esteban Moro and Luis Mateus Rocha published by Springer, Cham in Switzerland. The definitive authenticated version is available online via `https://doi.org/10.1007/978-3-030-36683-4_16`.

2.1 Introduction

In Chapter 1 we outlined the known spectral (eigenvalue) properties of the different representation matrices, i.e., the adjacency matrix A, the unnormalised Laplacian L and the normalised Laplacians L_{rw} and L_{sym}. In this chapter we compare the spectral properties of the graph representation matrices. As will be discussed in Section 2.1.1, a direct spectral comparison is not sensible. Therefore, the comparison is made pairwise by applying an affine transform-

ation to one of the representation matrices, which enables comparison whilst preserving certain key properties such as normalised eigengaps. Bounds are given on the eigenvalue differences thus found, which depend on the minimum and maximum degree of the graph. The monotonicity of the bounds and the structure of the graphs are related. The bounds are illustrated on a real social network graph, and on three model graphs, and good tightness properties are observed.

In the second part of this work, we will focus on an important function of the spectra of practical use, the eigengaps. The significance of the representation matrix eigengaps has already become apparent via the three variants of Cheeger's inequality discussed in Sections 1.3.2, 1.4.2 and 1.5.3. In the spectral clustering algorithm, eigengaps of the representation spectra are used to determine the number of clusters present in a network (von Luxburg, 2007). We extend the methodology used in the comparison of the representation matrix eigenvalues to study the normalised eigengap differences and derive bounds, which again turn out to be in terms of the graph's degree extremes.

It is found that if the degree extreme difference is large, different choices of representation matrix may give rise to disparate inference drawn from network analysis; smaller degree extreme differences result in consistent inference, whatever the choice of representation matrix.

We consider why we transform the representation spectra for comparison in Section 2.1.1, the nature of the transformation parameter choice in Section 2.1.2 and the uses of deriving a bound on the transformed eigenvalue differences in Section 2.1.3.

2.1.1 Motivation of the affine transformations

The aim of this work is to compare the spectra of the graph representation matrices A, L and L_{rw} or equivalently L_{sym}. A direct comparison of the observed spectra is difficult. Primarily this is because for two of three comparisons to be made, the ordering of the eigenvalues is reversed, i.e., as mentioned earlier in Section 1.1, there is a rough correspondence of the *larger* end of the adjacency matrix spectrum to the the *smaller* end of the Laplacian spectra and vice versa for the other ends. Furthermore, as discussed in Chapter 1, the sup-

ports of the three representation matrix spectra are different (van Mieghem, 2011, pp. 29, 64, 68), which makes it difficult to compare the spectra. An illustration of these complicating factors is given in plots (a), (b) and (c) of Figure 2.1, for the karate dataset graph introduced in Section 1.2.5.

For an insightful comparison it makes sense to relocate and scale the spectra via an affine transformation. Affine transformations preserve or reverse eigenvalue ordering and preserve eigengap sizes (differences between successive eigenvalues) relative to the spectral support of the representation matrices. Importantly, in the spectral clustering algorithms, eigengaps are used to determine the number of clusters present in the network, while the eigenvalue ordering, an even more fundamental property, is used to identify which eigenvectors to consider von Luxburg (2007). Therefore, when comparing the impact of the choice of the representation matrices on graphical analysis, we need to preserve or reverse eigenvalue ordering and preserve relative eigengap size. Further details on the properties of the affine transformations are given in Section 2.5.

In Appendix B.1 we show the exact transformation between representation matrix spectra, mapping one spectrum precisely onto another, to be either non-existent, if repeated eigenvalues in one spectrum correspond to unequal eigenvalues in another spectrum, or to be polynomial of degree smaller or equal than n. It is important to note that the exact mapping of one spectrum of the representation matrices onto the spectrum of another is not the aim of this work. In the mapping process the different information in the different representation matrix spectra, such as the eigenvalue ordering and the eigengaps, is lost. Therefore, higher order transformations are not suitable for a meaningful comparison of useful information, such as the eigengaps, carried in the representation spectra.

2.1.2 Motivation of the parameter choice

All three representation matrices are related through the degree matrix D. For d-regular graphs the representation matrices are already related by affine transformations and hence their eigenvalues are related by the same affine transformations. For general — non-d-regular graphs — we make parameter

choices for the affine transformations such that eigenvalue differences, after affine transformation, can be bounded above in terms of the elements of D. We find that these general parameter choices agree with the natural existing transformation parameters for d-regular graphs.

2.1.3 Motivation for calculating a bound on the eigenvalue differences

The motivation for bounding the eigenvalue difference is twofold. On the one hand, the structure of the bound and its dependency on graph parameters allow us to infer for which graphs we are able to see larger differences in the representation spectra. On the other hand, the bounds allow us to add a sensible scale to observed eigenvalues; the bound represents the theoretically maximal distance the eigenvalues and eigengaps can differ by and therefore we are able to compare observed differences to the maximal possible differences.

In Sections 2.2, 2.3 and 2.4, bounds are provided for the eigenvalue difference between any two of the three representation matrices, A, L and L_{rw}. The properties of the transformation used in each comparison are elaborated in Section 2.5. The analysis of the bound values over all graphs partitioned by their degree extremes is given in Sections 2.6 and 2.7. Section 2.8 displays a proof of concept by applying the bounds to graphs arising from a social network and three further model examples, and finds tightness of two of the bounds and the third bound is observed to be close to tight on our examples. In Section 2.9 we present an alternative approach to choosing the transformation parameters and highlight drawbacks of this alternative. In Section 2.10 the bound on normalised eigengap differences for each of the pairs of representation matrices is derived. These bounds are illustrated via examples with varying degree extreme difference, on a graph sampled from a stochastic blockmodel (SBM) and on a social network in Section 2.11. In Section 2.12 the potential impact of the representation matrix choice on inference drawn from graphical analysis is visualised by applying the spectral clustering algorithm to examples from Section 2.11. Finally, in Section 2.13 we show another use of our eigenvalue bounds by bounding the number of all possible spanning trees of a graph using the adjacency spectrum.

2.2 Bounding the individual eigenvalue difference for A and L

We begin by defining the spectral norm on matrices. Since A, L and L_{sym} are all real and symmetric, we pay particular attention to real, symmetric matrices, noting again that L_{sym} and L_{rw} have the same eigenvalues.

Definition 2.1. The spectral norm, or Euclidean 2-norm, of a matrix B, is defined to be the largest singular value of B, i.e., $\|B\|_2 \stackrel{\text{def}}{=}$[max eigenvalue of $\left(B^T B\right)]^{1/2}$. For any real, symmetric matrix, singular values and eigenvalues coincide in absolute value. Hence, we can express the spectral norm of a real, symmetric matrix Φ, with ordered eigenvalues $\phi_1, \ldots, \phi_n$, as $\|\Phi\|_2 = \max\left(|\phi_1|, |\phi_n|\right)$. ◁

Using the spectral norm, we will now compare the spectra of the adjacency matrix A and the unnormalised Laplacian L.

Theorem 2.2. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Let the eigenvalues of the corresponding adjacency matrix A and unnormalised Laplacian matrix L be denoted by $\mu_n \leq \cdots \leq \mu_1$ and $\lambda_1 \leq \cdots \leq \lambda_n$, respectively. Then, for all $i \in \{1, \cdots, n\}$,

$$|f_1(\mu_i) - \lambda_i| \leq \frac{d_{\max} - d_{\min}}{2} \stackrel{\text{def}}{=} e(A, L), \tag{2.1}$$

where $f_1(\mu_i) = c_1(A, L)\mu_i + c_0(A, L)$ with $c_1(A, L) = -1$ and $c_0(A, L) = (d_{\max} + d_{\min})/2$.

Proof. From Horn and Johnson (1991, p. 385) it follows that $f_1(A) = c_1(A, L)$ $A + c_0(A, L)I$ has eigenvalues of the form,

$$f_1(\mu_i) = c_1(A, L)\mu_i + c_0(A, L). \tag{2.2}$$

Now, we proceed to analyse the error made by approximating λ_i by $f_1(\mu_i)$. Here we view the affine transform $f_1(A)$ as a perturbation of L. Firstly, using inequality (4.3) in (Bai et al., 2000, p. 46)

$$\|f_1(A) - L\|_2 \geq |f_1(\mu_i) - \lambda_i| = |c_1(A, L)\mu_i + c_0(A, L) - \lambda_i|, \tag{2.3}$$

for $i \in \{1, 2, \ldots, n\}$ and $c_1(A, L) \leq 0$ (to ensure equal ordering in magnitude of the two compared spectra). In other words, the difference of corresponding eigenvalues of a symmetric matrix L and its 'perturbed' version $f_1(A)$ is bounded by the two norm of the perturbation $f_1(A) - L$. Now,

$$\|f_1(A) - L\|_2 \;=\; \|c_0(A, L)I - D + (1 + c_1(A, L))A\|_2 \,. \tag{2.4}$$

We are able to reexpress this norm in a simple form in terms of D by the choice $c_1(A, L) \stackrel{\text{def}}{=} -1$, i.e.,

$$\begin{aligned} \|f_1(A) - L\|_2 &= \|c_0(A, L)I - D\|_2 \\ &= \max\left(|c_0(A, L) - d_{\max}|\,, |c_0(A, L) - d_{\min}|\right). \end{aligned} \tag{2.5}$$

Equation (2.5) uses the fact that the eigenvalues of diagonal matrices are equal to their diagonal elements. As shown in Appendix B.2, (2.5) is analytically minimised by the choice $c_0(A, L) \stackrel{\text{def}}{=} (d_{\max} + d_{\min})/2$ to give,

$$\|f_1(A) - L\|_2 \;=\; \frac{d_{\max} - d_{\min}}{2}.$$

Putting (2.3) and (2.5) together, we obtain (2.1).

□

Remark 2.3. An alternative proof of Theorem 2.2 makes use of Weyl's inequality, as shown in Appendix B.3. ◁

For d-regular graphs, (2.1) gives $e(A, L) = 0$, so that $f_1(\mu_i) = \lambda_i$ and (2.2) then gives $\lambda_i = c_1(A, L)\mu_i + c_0(A, L) = d - \mu_i$, i.e., the eigenvalues are related by the required exact relation, as claimed in Section 2.1.2.

For general graphs, using (2.1), we can establish a rough correspondence, to within an affine transformation, between the eigenvalues of the adjacency matrix, A, and the unnormalised graph Laplacian, L, if the extremes of the degree sequence $d_{\max}$ and $d_{\min}$ are reasonably close.

2.3 Bounding the individual eigenvalue difference for L and L_{rw}

We now look at the spectral relationship between L and L_{rw}.

Theorem 2.4. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Let the eigenvalues of L and L_{rw} be denoted by $\lambda_1 \leq \cdots \leq \lambda_n$ and $\eta_1 \leq \cdots \leq \eta_n$, respectively. Then, for all $i \in \{1, \cdots, n\}$,

$$|f_2(\lambda_i) - \eta_i| \leq 2\,\frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}} \stackrel{\text{def}}{=} e\left(L, L_{rw}\right), \tag{2.6}$$

where $f_2(\lambda_i) = c_1(L, L_{rw})\lambda_i + c_0(A, L_{rw})$ with $c_1(L, L_{rw}) = 2/(d_{\max} + d_{\min})$ and $c_0(A, L_{rw}) = 0$.

Proof. As in the proof of Theorem 2.2, $f_2(L) = c_1(L, L_{rw})L + c_0(A, L_{rw})I$ has eigenvalues of the form,

$$f_2(\lambda_i) = c_1(L, L_{rw})\lambda_i + c_0(A, L_{rw}). \tag{2.7}$$

The following result is the equivalent of (2.3):

$$\|f_2(L) - L_{rw}\|_2 \geq |f_2(\lambda_i) - \eta_i\ | = |c_1(L, L_{rw})\lambda_i + c_0(A, L_{rw}) - \eta_i|\,, \tag{2.8}$$

for $i \in \{1, 2, \ldots, n\}$ and $c_1(L, L_{rw}) \geq 0$ (to ensure equal ordering in magnitude of the two compared spectra). In the steps that follow we use firstly that the eigenvalues of L_{rw} are all positive, and secondly that $\|L_{rw}\|_2 \leq 2$; as established in van Mieghem (2011, p. 64). Now,

$$\|f_2(L) - L_{rw}\|_2 \;\; = \;\; \|(c_1(L, L_{rw})D - I)\, L_{rw} + c_0(A, L_{rw})I\|_2\,.$$

The norm can be simplified by the choice $c_0(A, L_{rw}) = 0$, i.e.,

$$\begin{aligned} \|f_2(L) - L_{rw}\|_2 &= \|(c_1(L, L_{rw})D - I)\, L_{rw}\|_2 && (2.9)\\ &\leq \|L_{rw}\|_2\, \|c_1(L, L_{rw})D - I\|_2 \\ &\leq 2\max\left(|d_{\max}c_1(L, L_{rw}) - 1|\,, |d_{\min}c_1(L, L_{rw}) - 1|\right). && (2.10) \end{aligned}$$

As shown in Appendix B.4, this term is analytically minimised by the choice

$c_1(L, L_{rw}) \stackrel{\text{def}}{=} 2/(d_{\max} + d_{\min}) = 1/c_0(A, L)$ to give,

$$\|f_2(L) - L_{rw}\|_2 \;\leq\; 2\,\frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}}.$$

We thus obtain (2.6).

□

Note, since the multiplicative transformation parameter is nontrivial, $c_1(L, L_{rw}) = 2/(d_{\max} + d_{\min})$, we are unable to make use of Weyl's inequality for finding the bound (see Remark 2.3).

Asymptotically, as $d_{\max} \to \infty$ with fixed $d_{\min} > 0$, the bound $e\,(L, L_{rw})$ tends to 2. The restricted range of $d_{\min}$ is due to D^{-1}, the normalised Laplacian and consequently the bound $e(L, L_{rw})$ not being defined for $d_{\min} = 0$. As in Section 2.2, for d-regular graphs, where $d_{\max} = d_{\min} = d$, the bound equals zero. Hence, the spectra of $f_2(L)$ and L_{rw} are equal and from (2.7), $\eta_i = c_1(L, L_{rw})\lambda_i = \lambda_i/d$, i.e., the eigenvalues are related by the required exact relation, as claimed in Section 2.1.2.

Overall, the behaviour of this bound is similar to $e(A, L)$ obtained in Section 2.2. The smaller the difference of the degree sequence extremes $d_{\max}$ and $d_{\min}$, the closer the spectra of L and L_{rw} are related, as signified by a smaller bound on the eigenvalue differences.

2.4 Bounding the individual eigenvalue difference for A and L_{rw}

Now we move on to the final of our three possible relations: the spectra of A and L_{rw}.

Theorem 2.5. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Let the eigenvalues of A and L_{rw} be denoted by $\mu_n \leq \cdots \leq \mu_1$ and $\eta_1 \leq \cdots \leq \eta_n$, respectively. Then, for all $i \in \{1, \cdots, n\}$,

$$|f_3(\mu_i) - \eta_i| \leq \frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}} \stackrel{\text{def}}{=} e\,(A, L_{rw})\,, \tag{2.11}$$

where $f_3(\mu_i) = c_1(A, L_{rw})\mu_i + c_0(A, L_{rw})$ with $c_1(A, L_{rw}) = -2/(d_{\max} + d_{\min})$ and $c_0(A, L_{rw}) = 1$.

Proof. As in the proof of Theorems 2.2 and 2.4, $f_3(A) = c_1(A, L_{rw})A + c_0(A, L_{rw})I$ has eigenvalues of the form,

$$f_3(\mu_i) = c_1(A, L_{rw})\mu_i + c_0(A, L_{rw}). \tag{2.12}$$

We start with,

$$\|f_3(A) - L_{rw}\|_2 \geq |f_3(\mu_i) - \eta_i| = |c_1(A, L_{rw})\mu_i + c_0(A, L_{rw}) - \eta_i|, \tag{2.13}$$

for $i \in \{1, 2, \ldots, n\}$ and $c_1(A, L_{rw}) \leq 0$ (to ensure equal ordering in magnitude of the two compared spectra). Then,

$$\begin{aligned} \|f_3(A) - L_{rw}\|_2 &= \left\|c_1(A, L_{rw})A + c_0(A, L_{rw})I - I + D^{-1}A\right\|_2 \\ &= \left\|(c_1(A, L_{rw})D + I)D^{-1}A + c_0(A, L_{rw})I - I\right\|_2. \end{aligned} \tag{2.14}$$

At this stage we choose $c_0(A, L_{rw}) \stackrel{\text{def}}{=} 1$ as to simply the norm. Then,

$$\|f_3(A) - L_{rw}\|_2 \leq \|I + c_1(A, L_{rw})D\|_2 \left\|D^{-1}A\right\|_2 \tag{2.15}$$

$$= \max\left(\left|1 + d_{\max}c_1(A, L_{rw})\right|, \left|1 + d_{\min}c_1(A, L_{rw})\right|\right), \tag{2.16}$$

where in (2.15) we use the fact that $\|D^{-1}A\|_2 = 1$ (van Mieghem, 2011, p. 64). Finding $c_1(A, L_{rw})$ as the analytical minimum of this term is equivalent to finding $c_1(L, L_{rw})$ in (2.10) and therefore $c_1(A, L_{rw}) = -c_1(L, L_{rw}) \stackrel{\text{def}}{=} -2/(d_{\max} + d_{\min})$. Consequently,

$$\|f_3(A) - L_{rw}\|_2 \leq \frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}}.$$

Hence, (2.11) is obtained.

□

For d-regular graphs, $e(A, L_{rw})$ equals zero and hence the spectra of $f_3(A)$ and L_{rw} are equal and from (2.12), $\eta_i = 1 - c_1(L, L_{rw})\mu_i = 1 - (\mu_i/d)$, i.e., the spectra of L_{rw} and A are related by the required exact relation. For general

graphs, $e(A, L_{rw})$ is small for small degree extreme differences and hence the spectra of the two representation matrices exhibit smaller maximal differences, as was the case for $e(A, L)$ and $e(L, L_{rw})$, obtained in Sections 2.2 and 2.3, respectively.

2.5 Nature of transformations

Now we will visualise the derived transformations on the spectra from a real data set and analyse their properties. Through our analysis in Sections 2.2, 2.3 and 2.4 we have found the following transformation parameters for the affine transformations (2.2), (2.7) and (2.12),

$$f_1(\mu) = -\mu + \frac{d_{\max} + d_{\min}}{2}; \tag{2.17}$$

$$f_2(\lambda) = \frac{2}{d_{\max} + d_{\min}}\lambda; \tag{2.18}$$

$$f_3(\mu) = -\frac{2}{d_{\max} + d_{\min}}\mu + 1. \tag{2.19}$$

2.5.1 Karate data example

We first apply the transforms to the spectra of Zachary's karate dataset, which was introduced in Section 1.2.5 and will be further analysed in Sections 2.8 and 2.11.

The untransformed and transformed eigenvalues of the representation matrices A, L and L_{rw} of Zachary's karate dataset are shown in Figure 2.1. From Figures 2.1(a), (b) and (c) not much can be said about how the spectra compare. From Figures 2.1(d), (e) and (f) it can be observed that each pair of spectra of the karate dataset cover a similar range after transformation and that clearly the spectra of A and L_{rw} are the most similar of the three. It can also be seen that the largest eigengaps occur at opposing ends of the spectra when comparing A and L_{rw} with L. This observation is only possible from the plots including the transforms, i.e., Figures 2.1(d), (e) and (f). Clearly, the eigenvalue spectra become more comparable through utilising the proposed transformations.

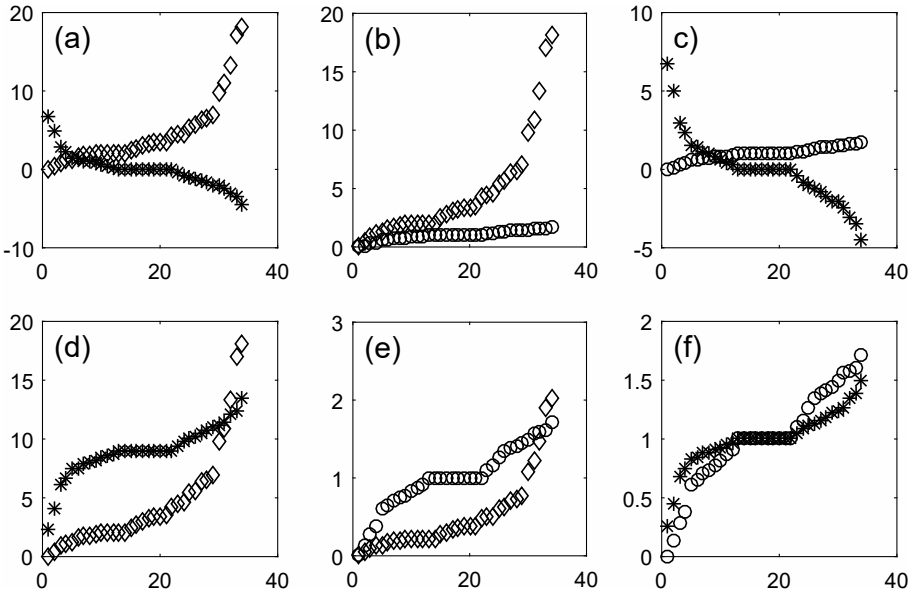

Figure 2.1: Transformation results for Zachary's karate dataset. The x-axes display eigenvalue indices, while the y-axes display eigenvalue values. First row: (a) eigenvalues μ of A (stars) and λ of L (diamonds); (b) eigenvalues λ of L and η of L_{rw} (circles); (c) eigenvalues μ of A and η of L_{rw}. Second row: the first of the two eigenvalue spectra are transformed by their respective transformations f_1, f_2 and f_3.

Remark 2.6. In Section 1.7 we described how vertices with identical neighbourhoods generate eigenvalues equal to 0 and 1 in the adjacency and normalised Laplacian spectrum, respectively. Interestingly, we spot several such vertices in the karate network in Figure 1.3, such as for example members 15, 16, 19 and 23, and several corresponding eigenvalues equal to 0 and 1 in Figure 2.1(c). It is very nice to see that the equivalence of these eigenvalues is highlighted by our transformation mapping them exactly onto each other in Figure 2.1(f). $\triangleleft$

In Chapter 1 we established that the location and size of eigengaps in all three representation spectra is linked to the evidence the different graph representation matrices encode of cluster structure in the underlying network. For practitioners it is of great value to consider the evidence of cluster structure encoded in *all three* representation spectra in an analysis in order to incorporate full information. Figure 2.1 serves as an example of how the comparison of the strength of evidence of the different cluster

structures is simplified by the use of our transformations, i.e., we recommend practitioners to compare the spectra of the representation matrices after having transformed one of the spectra by f_1, f_2 or f_3 instead of comparing untransformed spectra. The advantage of using our transformations specifically is that each transformation comes with a bound ($e(A, L)$, $e(L, L_{rw})$ and $e(A, L_{rw})$) on the maximal observable differences of the spectra. The bounds gives practitioners a sense for the scale of the differences they observe as we will see in the upcoming Section 2.8.

2.5.2 Transformation properties

All three transformations f_1, f_2 and f_3 are affine. The following simple observations are made for general affine transformations $f(x) = c_1 x + c_0$, where $c_1, c_0 \in \mathbb{R}, c_1 \neq 0$.

Ordering

Ordering is preserved for $c_1 > 0$, i.e., $f(\mu_1) \geq f(\mu_2)$ and reversed for $c_1 < 0$, i.e., $f(\mu_1) \leq f(\mu_2)$.

Eigengaps

Next we show that eigengaps, relative to the spectral support, are preserved. Assume that the domain of f is equal to the interval $[x_1, x_2]$. Hence f's image is equal to $[f(x_1), f(x_2)]$. Then normalised eigengaps are preserved:

$$\frac{\mu_1 - \mu_2}{x_2 - x_1} = \frac{c_1 (\mu_1 - \mu_2) + c_0 - c_0}{c_1 (x_2 - x_1) + c_0 - c_0} = \frac{f(\mu_1) - f(\mu_2)}{f(x_2) - f(x_1)}. \tag{2.20}$$

Mapping of spectral supports

As a final property we show the mappings for the spectral supports of the different representation matrices. The spectral supports of the three representation

matrices can be derived via Gershgorin's theorem as is done in Appendix A.2. The spectral supports are also given in van Mieghem (2011, pp. 29, 64, 68).

$$f_1 : [-d_{\max}, d_{\max}] \to \left[-\frac{d_{\max} - d_{\min}}{2}, \frac{3d_{\max} + d_{\min}}{2}\right] \tag{2.21}$$

$$f_2 : [0, 2d_{\max}] \to \left[0, \frac{4d_{\max}}{d_{\max} + d_{\min}}\right] \tag{2.22}$$

$$f_3 : [-d_{\max}, d_{\max}] \to \left[-\frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}}, \frac{3d_{\max} + d_{\min}}{d_{\max} + d_{\min}}\right]. \tag{2.23}$$

Remark 2.7. It is possible to choose the transformation parameters so that the transformation maps exactly to the spectral support of the target matrix. However, this results in a greater bound on the eigenvalue differences. Rather, since the mapped supports are of little consequence, we choose our transformation parameters according to the bound value, their treatment of eigenvalue ordering, and relative eigengap preservation. ◁

2.6 Relating the spectral bounds

We now compare the spectral bounds with each other and illustrate their relationship.

Firstly,

$$\begin{aligned} & e\left(A, L_{rw}\right) \leq e\left(L, L_{rw}\right) \\ \iff & \frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}} \leq 2\frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}}, \end{aligned} \tag{2.24}$$

which, obviously, always holds.

Next,

$$\begin{aligned} & e\left(A, L_{rw}\right) \leq e\left(A, L\right) \\ \iff & \frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}} \leq \frac{d_{\max} - d_{\min}}{2} \\ \iff & 2 \leq d_{\max} + d_{\min}. \end{aligned} \tag{2.25}$$

Table 2.1: Comparing bounds on eigenvalue differences of A, L and L_{rw}. The bound values are displayed as $(e(A,L), e(L,L_{rw}), e(A,L_{rw}))$.

		$d_{\min}$				
		0	1	2	3	4
$d_{\max}$	1	(0.5, ·, ·)	(0, 0, 0)	*	*	*
	2	(1, ·, ·)	(0.5, 0.67, 0.33)	(0, 0, 0)	*	*
	3	(1.5, ·, ·)	(1, 1, 0.5)	(0.5, 0.4, 0.2)	(0, 0, 0)	*
	4	(2, ·, ·)	(1.5, 1.2, 0.6)	(1, 0.67, 0.33)	(0.5, 0.29, 0.14)	(0, 0, 0)
	5	(2.5, ·, ·)	(2, 1.33, 0.67)	(1.5, 0.86, 0.43)	(1, 0.5, 0.25)	(0.5, 0.22,0.11)
	6	(3, ·, ·)	(2.5, 1.43, 0.71)	(2, 1, 0.5)	(1.5, 0.67, 0.33)	(1, 0.4, 0.2)

which applies to all graphs with edges.

Finally,

$$\begin{aligned} & e(L, L_{rw}) \leq e(A, L) \\ \iff & 2\frac{d_{\max} - d_{\min}}{d_{\max} + d_{\min}} \leq \frac{d_{\max} - d_{\min}}{2} \\ \iff & 4 \leq d_{\max} + d_{\min}. \end{aligned} \tag{2.26}$$

So for $d_{\max} + d_{\min} \geq 4$, $e(A,L) \geq e(L,L_{rw}) \geq e(A,L_{rw})$. The only other case of interest is attained by path graphs, i.e., for $d_{\min} = 1, d_{\max} = 2$, for which $e(L,L_{rw}) \geq e(A,L) \geq e(A,L_{rw})$. Interestingly, the maximal spectral differences between A and L_{rw} are always smaller than the two corresponding values for comparisons involving L.

For easier analysis of these inequalities, we display some sample values in Table 2.1. The first column of Table 2.1 only contains values of $e(A,L)$ as for $d_{\min} = 0$ the normalised Laplacian and hence $e(L,L_{rw})$ and $e(A,L_{rw})$ are not well-defined. In practice this is of little consequence since disconnected nodes are commonly removed from the dataset as a preprocessing step. On the diagonal of Table 2.1, where $d_{\max} = d_{\min}$, i.e., for the d-regular graphs, we find $e(A,L) = e(L,L_{rw}) = e(A,L_{rw}) = 0$ due to the direct spectral relation of the representation matrices.

Remark 2.8. It is interesting to note, that since the spectral support of neither A, nor L, is bounded, the bound on their eigenvalue difference is also not

bounded above. This does not apply to the other two bounds as we have $e(L, L_{rw}) \leq 2$ and $e(A, L_{rw}) \leq 1$. $\triangleleft$

2.7 Explaining the structure of Table 2.1

In the forthcoming analysis we will make use of the graph theoretic concepts of subgraphs and connected components in a graph, which were introduced in Definitions 1.37 and 1.7, respectively.

2.7.1 Partitioning

We will illustrate the class of unweighted graphs characterised by certain degree extremes $d_{\min}$ and $d_{\max}$ in order to visualise the class of graphs to which the bounds apply. This will allow us to explain the monotonicity of the three bounds in the sense of Table 2.1. We will denote the class of graphs with $d_{\min} = j$ and $d_{\max} = k$ by $\mathcal{C}_{j,k}$. For any graph, $d_{\min}$ and $d_{\max}$ are unique, and hence no graph is in two or more classes of $\{\mathcal{C}_{j,k}\}_{j=0}^{k}$. Therefore, the classes $\{\mathcal{C}_{j,k}\}_{j=0}^{k}$ are disjoint. Furthermore, since the degree extremes of a graph always exist, we have that $\{\mathcal{C}_{j,k}\}_{j=0}^{k}$ is a partition of the class of graphs with $d_{\max} = k$, (denoted by $\mathcal{C}_{\cdot,k}$ in the following).

For illustrative purposes, we consider the class of graphs $\mathcal{C}_{\cdot,2}$, i.e., the class of graphs to which the bounds in the second *row* of Table 2.1 apply. In Figure 2.2 we display the partition of $\mathcal{C}_{\cdot,2}$, with examples of elements in each class. All classes $\mathcal{C}_{j,k}$ are infinite in size, so only a few arbitrary sample elements are displayed here. The elements of each class are marked by dashed ellipses. Graph G_3 is an element of $\mathcal{C}_{0,2}$ and consists of five connected components. Graphs in $\mathcal{C}_{0,2}$ are denoted by G_a, graphs in $\mathcal{C}_{1,2}$ by H_a and graphs in $\mathcal{C}_{2,2}$ by I_a for $a \in \mathbb{Z}$.

Remark 2.9. The subgraph relation of two graphs does not impose a logical ordering of the classes $\{\mathcal{C}_{j,k}\}_{j=0}^{k}$. This can be seen in Figure 2.2 by observing that for example H_4 is a subgraph of G_3 and G_1 is a subgraph of H_3. Therefore, we make use of the addition and removal of connected components to relate graphs in the classes $\{\mathcal{C}_{j,k}\}_{j=0}^{k}$. $\triangleleft$

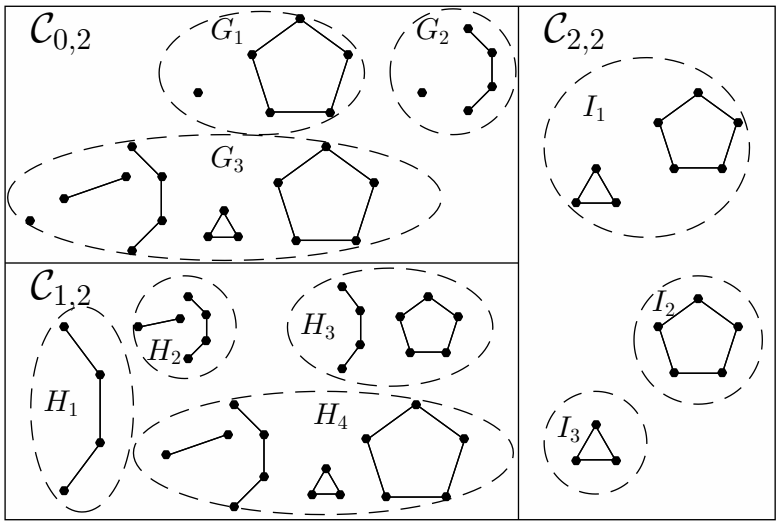

Figure 2.2: This figure illustrates the partition $\mathcal{C}_{.,2} = \mathcal{C}_{0,2} \cup \mathcal{C}_{1,2} \cup \mathcal{C}_{2,2}$. Examples of elements of the classes $\mathcal{C}_{0,2}$ containing all graphs with $d_{\min} = 0$ and $d_{\max} = 2$, $\mathcal{C}_{1,2}$, with element for which $d_{\min} = 1$ and $d_{\max} = 2$ and $\mathcal{C}_{2,2}$, the class of 2-regular graphs are shown.

2.7.2 Analysis of bound monotonicity in rows of Table 2.1

We now turn to a specific manipulation of graphs — addition and removal of connected components — which allows us to order the spectra of the graphs and hence understand the monotonicity observed in the bounds in rows of Table 2.1.

In Figure 2.2, we see that we can obtain G_3 from H_4 by adding a single disconnected vertex and that we can obtain H_4 from I_1 by adding the connected components, the line and the 2-complete component. In general, we can obtain a graph in $\mathcal{C}_{j,k}$ from a graph in $\mathcal{C}_{j+1,k}$, for all $j \leq k-1, k \in \mathbb{N}$, by adding to the graph one or more connected components in which all vertices are of degree greater or equal to j and smaller or equal to k, with at least one vertex attaining degree j.

Recall from Lemma 1.36 in Section 1.6.1, that the representation matrix spectrum of a graph is the union of the representation matrix spectra of its connected components.

Since all graphs in $\mathcal{C}_{j+1,k}$ can be extended to lie in $\mathcal{C}_{j,k}$ by adding one or

more connected components, it can be argued that all spectra of graphs in $\mathcal{C}_{j+1,k}$ are subsets of spectra of graphs in $\mathcal{C}_{j,k}$. Therefore, the support of the spectra of graphs in $\mathcal{C}_{j,k}$ must be larger or equal to the support of spectra of graphs in $\mathcal{C}_{j+1,k}$. Hence, we expect the spectral bounds, $e(\cdot,\cdot)$, we derived to be decreasing or constant with increasing $d_{\min}$ and constant $d_{\max}$. This phenomenon can be observed in Table 2.1 when traversing each *row*.

2.7.3 Analysis of bound monotonicity in columns of Table 2.1

In Section 2.7.2, when traversing the rows of Table 2.1, we were adding connected components to graphs in $\mathcal{C}_{j+1,k}$ to obtain graphs in $\mathcal{C}_{j,k}$. We observed decreasing bound size with increasing indices corresponding to $d_{\min}$. Here we find increasing bound size with increasing indices corresponding to $d_{\max}$: it seems sensible that increasing the support of the degree distribution should also increase the spectral support and hence the bounds. (The *degree distribution* is a probability distribution from which vertex degrees of a graph are sampled.)

We argue that any graph in $\mathcal{C}_{j,k}$, can be extended to be a graph in $\mathcal{C}_{j,k+1}$, by adding a connected component with all vertex degrees greater or equal to j and smaller or equal than $k+1$ with at least one node attaining degree $k+1$. Then the argument goes exactly as in Section 2.7.2, that spectra of graphs in $\mathcal{C}_{j,k}$ are subsets of spectra of graphs $\mathcal{C}_{j,k+1}$ and therefore the spectral support and hence the spectral bounds on $\mathcal{C}_{j,k+1}$ have to be greater or equal than the respective quantities for $\mathcal{C}_{j,k}$. So, any of the spectral bounds, $e(\cdot,\cdot)$, will be increasing or constant with increasing $d_{\max}$ and constant $d_{\min}$, as seen in the *columns* of Table 2.1.

2.8 Visualising the eigenvalue bounds on data

In this section we illustrate the bounds on the spectra of Zachary's karate dataset, which we already met in Section 2.5.1. We then proceed to explore the bounds on three different graphs with $d_{\min} = 1$ and $d_{\max} = 17$, choices

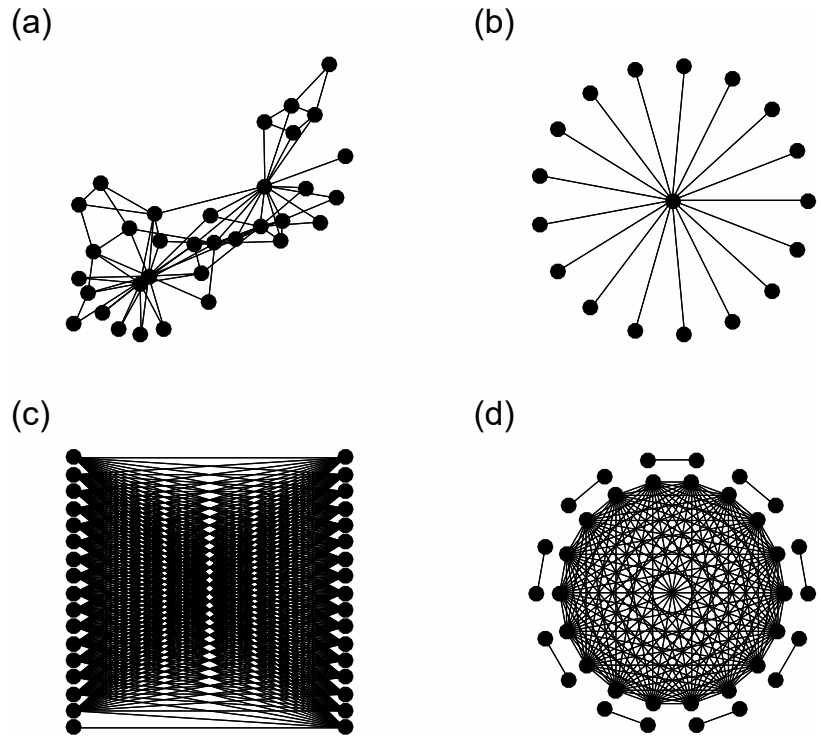

Figure 2.3: (a) graph given by the karate data set. (b) graph A, the star graph on 18 nodes. (c) graph B, a bipartite graph with degree sequence $\{1, \{16\}^{16}, \{17\}^{17}\}$ (see Section 2.8.2). (d) graph C, a graph consisting of 9 2-complete components and one 18-complete component.

appropriate for the karate dataset.

The graph defined in the karate data set is plotted in Figure 2.3(a) together with the three graph examples in Figures 2.3(b)-(d) analysed in Section 2.8.2.

2.8.1 Zachary's karate dataset

In Figure 2.4, we display a proof of concept of the bounds derived in Sections 2.2, 2.3 and 2.4. We display the transformed eigenvalues, together with the eigenvalues we compare to, and the derived bounds. The eigenvalue bounds are centred around the average value of each eigenvalue pair in order to highlight the maximal difference achievable by each individual eigenvalue pair under comparison. For the karate dataset the bound values $(e(A, L), e(L, L_{rw}), e(A, L_{rw}))$ are equal to $(8.00, 1.78, 0.89)$. The values in plot (a) are much larger than the ones in plots (b) and (c), as was to be expected

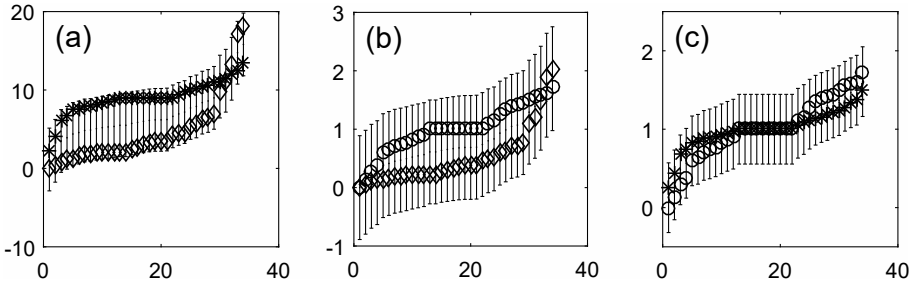

Figure 2.4: Eigenvalue bounds on the karate eigenvalues. The x-axes display eigenvalue indices, while the y-axes display eigenvalue values. In plot (a) we display the bound $e(A, L)$ via the intervals together with the transformed eigenvalues of the adjacency matrix $f_1(\mu)$ (stars) and the eigenvalues of the Laplacian λ (diamonds). (b) the bound $e(L, L_{rw})$ is displayed via the intervals, the diamonds correspond to the transformed Laplacian eigenvalues $f_2(\lambda)$ and the circles are the eigenvalues of the normalised graph Laplacian η. (c) the bound $e(A, L_{rw})$ is displayed via the intervals, the stars correspond to the transformed adjacency eigenvalues $f_3(\mu)$ and the circles are the normalised Laplacian eigenvalues η.

from Remark 2.8. The particular bounds displayed here are valid for all graphs with $d_{\min} = 1$ and $d_{\max} = 17$, i.e., all graphs in $\mathcal{C}_{1,17}$. The fact of the bounds being almost attained in plot (a), and not attained in plots (b) and (c), is more a consequence of the structure of the graph given by the karate data set than tightness and quality of the bounds. Since the three bounds $e(A, L)$, $e(L, L_{rw})$ and $e(A, L_{rw})$ apply to entire classes $\mathcal{C}_{j,k}$ at a time, we can only hope to achieve tightness on these classes — bounds being attained for some elements in $\mathcal{C}_{j,k}$ — and not on each individual element of them.

So for our particular social network, the karate dataset, firstly the bound being almost attained in plot (a) tells us that the spectra of A and L deviate almost as much as theoretically possible for a graph in $\mathcal{C}_{1,17}$. Secondly, from plot (c) we see that the spectra of A and L_{rw} are rather similar for the karate data set and could theoretically deviate significantly more for a different graph in $\mathcal{C}_{1,17}$ having degree extreme difference 16. The bounds give us insight into how particular representation matrix spectra behave for the karate dataset.

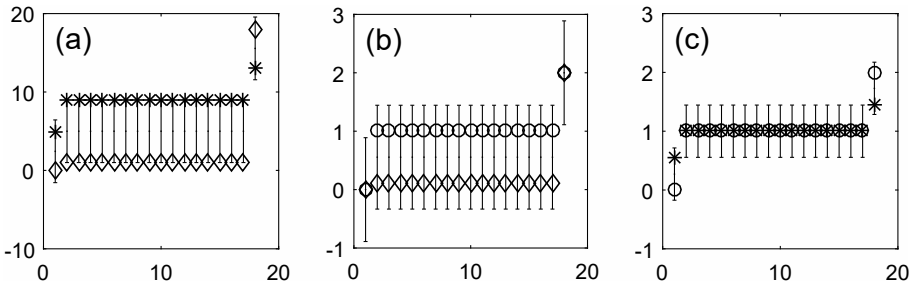

Figure 2.5: Spectra and transformed spectra with their bound for graph A, a star on 18 nodes, with degree sequence $\{\{1\}^{17}, 17\}$. The format follows Figure 2.4.

2.8.2 Three examples exploring the bounds for graphs in $\mathcal{C}_{1,17}$

We now proceed to explore the class of graphs with $d_{\min} = 1$ and $d_{\max} = 17$, $\mathcal{C}_{1,17}$, to which the bounds for the karate dataset displayed in Figure 2.4 apply. We will consider the star on 18 vertices, 'graph A,' a bipartite graph, 'graph B,' and a graph containing several 1- and 17-regular connected components, 'graph C.'

Graph A (Star graph)

Figure 2.5 displays the three bounds on the spectra of the star graph with degree sequence $\{\{1\}^{17}, 17\}$, where the notation $\{x\}^y$ is shorthand for the multiset consisting of y elements equal to x.

Especially notable, is that from Figure 2.5(a) the bound $e(A, L)$ can be seen to be tight: the spectra of A and L have maximal distance on individual eigenvalue level for 16 of the 18 eigenvalues. Further, for the 16 eigenvalues of maximal distance between spectra of A and L, the spectra of A and L_{rw} coincide, up to transformation.

Graph B (Bipartite graph)

We will now work with a bipartite graph on 34 nodes. The degree sequence of the bipartite graph is $\{1, \{16\}^{16}, \{17\}^{17}\}$. The graphs falling under the same spectral bound are only restricted in their degree sequence extremes, not by

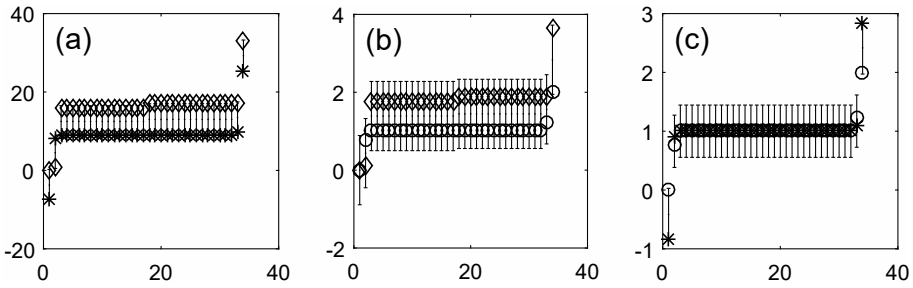

Figure 2.6: Spectra and transformed spectra with their bound for graph B, the bipartite graph on 34 nodes, with degree sequence $\{1, \{16\}^{16}, \{17\}^{17}\}$. The format follows Figure 2.4.

the number of nodes in the graphs.

In Figure 2.6(a), 15 of the eigenvalue pairs attain the maximal distance given by the bound and the remaining 19 eigenvalue pairs are close to attaining the maximal distance. Also, all but one of the eigenvalues of the graph Laplacian L are larger than the ones of the adjacency matrix A. Figure 2.5(a) displayed the opposite phenomenon, the eigenvalues of A being mostly larger than the eigenvalues of L. Another, very interesting, phenomenon can be spotted in the behaviour of the first two eigenvalues of A and L in Figure 2.6(a). To discuss this we need the following:

Definition 2.10. Let Φ and Ψ be real symmetric matrices with eigenvalues $\phi_1, \ldots, \phi_n$ and $\psi_1, \ldots, \psi_n$, respectively and let f be an affine transformation applied to a matrix spectrum. Then, in the context of a bound of the form $e(\Phi, \Psi) \geq |f(\phi_i) - \psi_i|$, we will use the term *maximal crossover* of eigenvalues to express that $f(\phi_i)$ attains one of the ends of the bound $e(\Phi, \Psi)$, $f(\phi_{i+1})$ attains the opposite end of the bound $e(\Phi, \Psi)$, and vice versa for ψ_i and ψ_{i+1}. $\triangleleft$

We see that the first two eigenvalues of A and L are very close to a maximal crossover, where $f = f_1$, with rounded distances $\lambda_1 - f_1(\mu_1) = 7.49$ and $\lambda_2 - f_1(\mu_2) = -7.06$ for a bound value $e(A, L) = 8$. The concept of a maximal crossover will be significant in Section 2.10 when bounding eigengaps.

Figure 2.6(b) also provides an interesting insight: the two last eigenvalues $f_2(\lambda_{34})$ and η_{34} attain the opposing ends of the bound $e(L, L_{rw})$, with a rounded distance between them equal to $f_2(\lambda_{34}) - \eta_{34} = 1.67$ on the bound value $e(L, L_{rw}) = 1.78$. Hence, $e(L, L_{rw})$ is close to tight on $\mathcal{C}_{1,17}$.

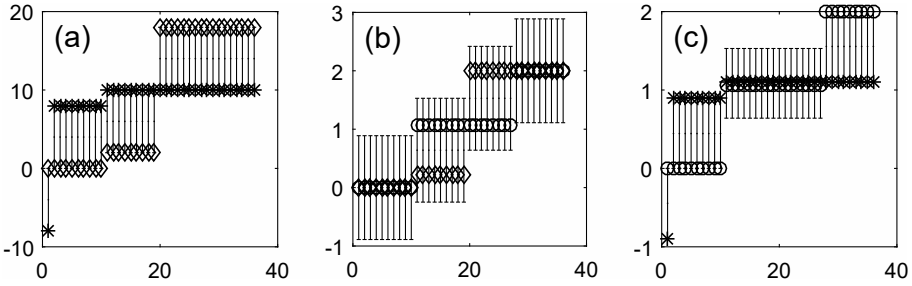

Figure 2.7: Spectra and transformed spectra together with their bound for graph C, consisting of 10 connected components, 9 of which are the complete graph on 2 nodes and one of which is the complete graph on 18 nodes. Therefore the degree sequence is equal to $\{\{1\}^{18}, \{17\}^{18}\}$. The format follows Figure 2.4.

Finally, we see in Figure 2.6(c) that the eigenvalues of A and L_{rw} mostly coincide up to transformation, as was the case for the eigenvalues in Figure 2.5(c). However, the first and last eigenvalue pairs can be seen to almost attain the edges of the bounds, where $|f_3(\mu_1) - \eta_1| = |f_3(\mu_{34}) - \eta_{34}| = 0.83$ on the bound value of $e(A, L_{rw}) = 0.89$. As with $e(L, L_{rw})$ in plot (b), this highlights that there is not much room to improve upon the bound $e(A, L_{rw})$ and indeed we will find it to be tight in the next example.

Graph C

As we saw in Section 2.6, $\mathcal{C}_{1,17}$ also contains graphs which consist of several connected components. In Figure 2.7 we hence consider the spectra and bounds corresponding to graph C, consisting of 9 pairs of vertices connected only to one another, i.e., 18 nodes of degree 1, and a complete connected component on 18 nodes.

In Figure 2.7(a), we spot a maximal crossover with $f = f_1$ between eigenvalues $f_1(\mu_1)$, λ_1 and $f_1(\mu_2)$, λ_2 and again for the pairs with indices 19 and 20.

From Figure 2.7(c) we find the bound $e(A, L_{rw})$ to be tight on 19 of the 36 eigenvalue pairs. We furthermore observe a maximal crossover with $f = f_3$ between eigenvalues $f_3(\mu_1)$, η_1 and $f_3(\mu_2)$, η_2. The eigenvalues in Figure 2.7(c) display greater differences than in Figures 2.5(c) and 2.6(c), where the majority of eigenvalues agreed up to transformation.

Remark 2.11. In the case of a graph consisting of more than one connected

component, as in Figure 2.7, it is beneficial to consider connected components separately. When this is done, the degree sequence extremes are decreased or stay equal and hence the transformation becomes more accurate or stays the same and the bounds decrease or stay equal (see Sections 2.7.2 and 2.7.3). For graph C underlying Figure 2.7, the spectra of the connected components can individually be exactly related as the connected components are 1- and 17-regular. ◁

2.9 An alternative approach via numerical optimisation

In this section we present an alternative approach to choosing the transformation parameters for the spectral comparison of the representation matrices. Instead of analytically choosing the transformation parameters as is done in Sections 2.2, 2.3 and 2.4, we can numerically determine the affine transformation parameters, which minimise the 2-norm of the matrix difference, e.g., $\|c_1(A, L)A + c_0(A, L)I - L\|_2$ in (2.3). This minimisation is a convex optimisation problem and can therefore, be readily implemented.

In Figure 2.8 we display the affine transformation parameters determined by the numerical optimisation of the matrix differences in the 2-norm on 100 graphs sampled from a SBM. Recall that SBMs were introduced in Section 1.2.4.

In the first row of plots in Figure 2.8 we see that our analytical transformation parameter choice for the comparison of A and L coincides with the numerically optimal parameters, which numerically optimise the 2-norm on the transformed matrix difference. Therefore, the two approaches are equivalent for the comparison of A and L for the sampled graphs. In the second and third row of Figure 2.8, we find the two approaches of choosing the transformation parameters produce unequal values. In the spectral comparison of the Laplacian matrices L and L_{rw} the numerically optimal multiplicative parameter is lower than the analytically chosen parameter, while the opposite trend can be observed for the additive parameters. This leads to an overall smaller bound on the spectral difference than $e(L, L_{rw})$ arising from our analytical parameter

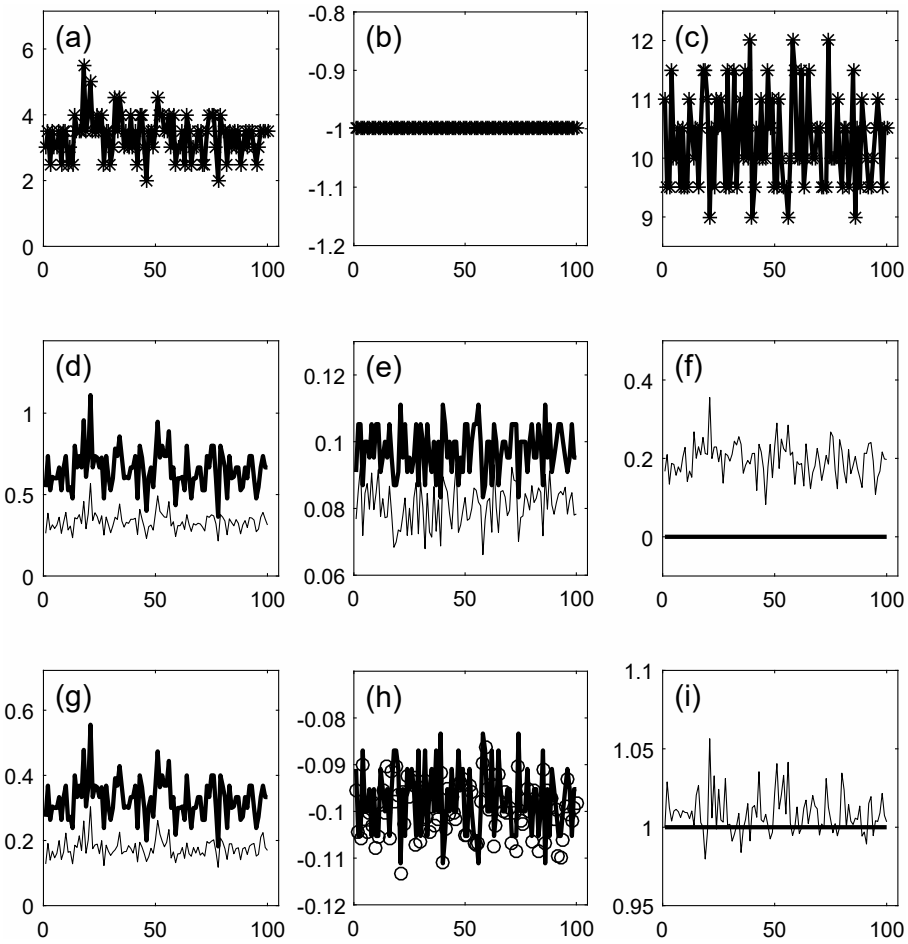

Figure 2.8: We compare the analytical and numerical parameter choice for the pairwise spectral comparison of the graph representation matrices of 100 draws from a SBM with parameters $(p_b, p_w) = (0.1, 0.9)$, $n = 30$ nodes, $K = 3$ and equally many nodes per block. In all plots the thin line, stars and circles correspond to values obtained via numerical optimisation and the thick black line corresponds to the values obtained via the analytical parameter choices made in Sections 2.2, 2.3 and 2.4. The x-axes display the indices of the SBM draws, while the y-axes display the different values obtained from the draws. First column: shows the bound values; second column: the multiplicative parameter of the affine transformation; third column: the additive parameter of the affine transformation. First row: comparison of A and L; second row: L and L_{rw}; third row: A and L_{rw}.

choice. Similarly, for the comparison of A and L_{rw} the numerical parameter choice varies slightly from the analytical parameter choice and produces lower bounds than $e(A, L_{rw})$.

Note these results do not speak to the tightness of our bounds since the matrix 2-norm, the object which is bounded, depends on the chosen transformation

parameters itself. The aim of this work is to choose transformation parameters which allow for a good comparison of the representation matrix spectra and *allow for analytical insight into their differences via D*; the aim is not primarily to minimise the spectral difference of the representation matrices in the 2-norm as is done by the numerical approach presented in this section. Crucially, analytical insight from the functional form of the parameters is not achieved by choosing the transformation parameters numerically, since we are unable to give a functional form for these parameter choices, which depends on graph constants such as the degree extremes of the underlying graph. Therefore, while smaller bounds could be obtained on the matrix differences from the numerical optimisation approach, the analytical bounds presented in this chapter have the additional advantage of being in terms of the degree extreme difference, which allows for significant theoretical insight into the spectral differences of the representation matrices. This insight would be lost if the numerical optimisation approach was taken.

2.10 Bounds on normalised eigengap differences

Following on from the work in Section 2.8, we examine the eigengaps corresponding to the three representation matrices in this section. We derive bounds on the normalised eigengap differences, where each eigengap is normalised by the spectral support of its corresponding representation matrix. Furthermore, we investigate under which conditions on the spectra the derived eigengap bounds are tight.

2.10.1 Normalised eigengap difference of A and L

Let $\mathcal{M}_i$ denote the i^{th} eigengap of A, $\mathcal{M}_i = \mu_i - \mu_{i+1}$ and $\mathcal{L}_i$ denote the i^{th} eigengap of L, $\mathcal{L}_i = \lambda_{i+1} - \lambda_i$ for $i \in \{1, 2, \ldots, n-1\}$. The spectral support of A and L is equal to $[-d_{\max}, d_{\max}]$ and $[0, 2d_{\max}]$, respectively (van Mieghem, 2011, pp. 29, 68), so the length of the support is $\ell(\mu) = \ell(\lambda) = 2d_{\max}$.

From (2.20) we see that the normalisation of transformed eigengaps by the transformed spectral support is equal to the normalisation of the untrans-

formed eigengaps by the untransformed spectral support, i.e.,

$$\left|\frac{f_1(\mu_{i+1}) - f_1(\mu_i)}{\ell(f_1(\mu))} - \frac{\mathcal{L}_i}{\ell(\lambda)}\right| = \left|\frac{\mathcal{M}_i}{\ell(\mu)} - \frac{\mathcal{L}_i}{\ell(\lambda)}\right|.$$

We will therefore start our analysis, and continue likewise in Sections 2.10.2 and 2.10.3, by considering the untransformed normalised eigengaps. Note that the eigengap normalisation by the spectral support of the corresponding representation matrix is crucial to be able to make a meaningful comparison of eigengap magnitudes. Furthermore, $\ell(\mu) = \ell(\lambda)$, so the normalisation of the two bounds separately is equivalent to normalising the entire difference by a single value. This will not be the case in Sections 2.10.2 and 2.10.3, where the lengths of the eigenvalue supports are different.

Theorem 2.12. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Then, for all $i \in \{1, \cdots, n\}$,

$$\left|\frac{\mathcal{M}_i}{2d_{\max}} - \frac{\mathcal{L}_i}{2d_{\max}}\right| \leq \frac{d_{\max} - d_{\min}}{2d_{\max}} \stackrel{\text{def}}{=} g(A, L).$$

Proof. The bound on the normalised eigengap difference of the adjacency matrix A and the unnormalised Laplacian L, can be derived as follows:

$$\begin{aligned}
\left|\frac{\mathcal{M}_i}{2d_{\max}} - \frac{\mathcal{L}_i}{2d_{\max}}\right| &= \frac{1}{2d_{\max}} \left|\mu_i - \mu_{i+1} - \lambda_{i+1} + \lambda_i\right| \\
&= \frac{1}{2d_{\max}} \left|\mu_i - \frac{d_{\max} + d_{\min}}{2} + \lambda_i - \mu_{i+1} + \frac{d_{\max} + d_{\min}}{2} - \lambda_{i+1}\right| \\
&= \frac{1}{2d_{\max}} \left|[\lambda_i - f_1(\mu_i)] + [f_1(\mu_{i+1}) - \lambda_{i+1}]\right| \qquad (2.27) \\
&\leq \frac{1}{2d_{\max}} 2e(A, L) \qquad (2.28) \\
&= \frac{d_{\max} - d_{\min}}{2d_{\max}},
\end{aligned}$$

where the triangle inequality is used to go from (2.27) to (2.28).

□

It follows from Equation (2.27), that the bound on the eigengaps is tight if and only if the individual mapped eigenvalue differences $\lambda_i - f_1(\mu_i)$ and

$\lambda_{i+1} - f_1(\mu_{i+1})$ have opposing signs and occupy the extremes of the bound $e(A, L)$. In Definition 2.10, such a situation is called a maximal crossover for general transformations f. We observed in Section 2.8, Figure 2.6(a), a situation very close to a maximal crossover with $f = f_1$. A maximal crossover is shown in Figure 2.7(a) for $f = f_1$; therefore, $g(A, L)$ is the optimal *constant* bound since it is tight for some cases in $\mathcal{C}_{1,17}$.

2.10.2 Normalised eigengap difference of L and L_{rw}

We proceed to derive a bound on the normalised eigengap difference of L and L_{rw}, denoted $g(L, L_{rw})$. Let $\mathcal{N}_i = \eta_{i+1} - \eta_i$ denote the i^{th} eigengap of L_{rw} for $i \in \{1, \ldots, n-1\}$. For L_{rw}, $\ell(\eta) = 2$ (van Mieghem, 2011, p. 64).

Theorem 2.13. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Then, for all $i \in \{1, \cdots, n\}$,

$$\left| \frac{\mathcal{L}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| \leq 2\frac{d_{\max} - d_{\min}}{d_{\max}} \stackrel{\text{def}}{=} g(L, L_{rw}).$$

Proof. The bound on the normalised eigengap difference of the unnormalised Laplacian L and the normalised Laplacian L_{rw}, is derived as follows:

$$\begin{aligned} \left| \frac{\mathcal{L}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| &= \frac{1}{2} \left| \frac{1}{d_{\max}} (\lambda_{i+1} - \lambda_i) + \eta_i - \eta_{i+1} \right| \\ &\leq \frac{1}{2} \left[\left| \frac{1}{d_{\max}} \lambda_{i+1} - \eta_{i+1} \right| + \left| \frac{1}{d_{\max}} \lambda_i - \eta_i \right| \right], \end{aligned} \tag{2.29}$$

from the triangle inequality. A comparison with (2.8), shows that these eigenvalue differences can be bounded by replacing $c_1(L, L_{rw})$ in (2.8) by $1/d_{\max}$ and then making this substitution into (2.10):

$$\begin{aligned} \left| \frac{\mathcal{L}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| &\leq \frac{1}{2} \left[4 \max \left(0, \left| \frac{d_{\min}}{d_{\max}} - 1 \right| \right) \right] \\ &= 2 \left(1 - \frac{d_{\min}}{d_{\max}} \right) \\ &= 2\frac{d_{\max} - d_{\min}}{d_{\max}}. \end{aligned}$$

□

We see from (2.29), that observing a maximal crossover with $f(\lambda) = \lambda/d_{\max}$ in the normalised spectra, would render $g(L, L_{rw})$ a tight bound.

2.10.3 Normalised eigengap difference of A and L_{rw}

Finally, we will derive the bound on the normalised eigengaps of A and L_{rw}, denoted by $g(A, L_{rw})$.

Theorem 2.14. Consider a graph G with degree extremes $d_{\min}$ and $d_{\max}$. Then, for all $i \in \{1, \cdots, n\}$,

$$\left| \frac{\mathcal{M}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| \leq \frac{d_{\max} - d_{\min}}{d_{\max}} \stackrel{\text{def}}{=} g(A, L_{rw}).$$

Proof. The bound on the normalised eigengap difference of the adjacency matrix A and the normalised Laplacian L_{rw}, is derived as follows:

$$\begin{aligned} \left| \frac{\mathcal{M}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| &= \frac{1}{2} \left| 1 - 1 + \frac{1}{d_{\max}} (\mu_i - \mu_{i+1}) + \eta_i - \eta_{i+1} \right| \\ &\leq \frac{1}{2} \left[\left| 1 - \frac{1}{d_{\max}} \mu_{i+1} - \eta_{i+1} \right| + \left| 1 - \frac{1}{d_{\max}} \mu_i - \eta_i \right| \right], \end{aligned}$$

via the triangle inequality. Again, as in the proof of Theorem 2.5, we chose the translation parameter equal to 1 to simplify the norm. Then, a comparison with (2.13), shows that these eigenvalue differences can be bounded by replacing $c_1(A, L_{rw})$ by $-1/d_{\max}$ and then substituting into (2.16) to give

$$\begin{aligned} \left| \frac{\mathcal{M}_i}{2d_{\max}} - \frac{\mathcal{N}_i}{2} \right| &\leq \max \left(0, \left| 1 - \frac{d_{\min}}{d_{\max}} \right| \right) \\ &= \frac{d_{\max} - d_{\min}}{d_{\max}}. \end{aligned}$$

□

Again, as in Sections 2.10.1 and 2.10.2, a maximal crossover with $f(\mu) = 1 - (\mu/d_{\max})$ in the normalised spectra corresponds to $g(A, L_{rw})$ being a tight bound.

2.11 Visualising the eigengap bounds on data

We observe our bounds on the eigengaps of graphs with varying degree extremes in Section 2.11.1, a graph sampled from a SBM in Section 2.11.2 and the karate data set in Section 2.11.3. We find that the eigengaps of the representation matrices are different, giving further evidence that the graphical structure captured in the different spectra differs significantly and in a structured manner. Furthermore, we make use of our bounds to put the magnitude of the observed eigengaps into the broader perspective of all graphs with corresponding degree extremes.

Throughout this section we interpret large eigengaps in the context of spectral clustering. Hence, large eigengaps are understood to indicate the number of clusters discovered by the different representation matrices, i.e., a large r^{th} eigengap is understood to indicate a clustering into r communities (von Luxburg, 2007, p. 410). Furthermore, the size of the eigengap is understood to represent relative strength of evidence for a certain clustering (von Luxburg, 2007, p. 407).

2.11.1 Eigengap bound examples for varying degree extremes

The functional form of our bounds implies that a growing degree extreme difference allows for a greater spectral difference between graph representation matrices. We visualise this effect by varying the degree maximum in graph C by replacing the 18-complete component by a k-complete component, the resulting graphs are denoted $C(k)$. The graph C we have considered so far will hence be referred to as graph $C(18)$ in the new notation.

The effects of changing $d_{\max}$ by varying the size of the k-complete component in graphs $C(k)$ with $k \in \{3, \ldots, 18\}$ are shown in Figure 2.9. Figures 2.9(a), (d) and (g) display plots of graphs $C(3)$, $C(10)$ and $C(18)$, respectively. Figures 2.9(b), (e) and (h) show the eigengap bounds on the adjacency and unnormalised Laplacian normalised eigengaps corresponding to graphs $C(3)$, $C(10)$ and $C(18)$, respectively. Figures 2.9(c), (f) and (i) show the eigengap bounds on the unnormalised and normalised Laplacian normalised eigengaps. Fig-

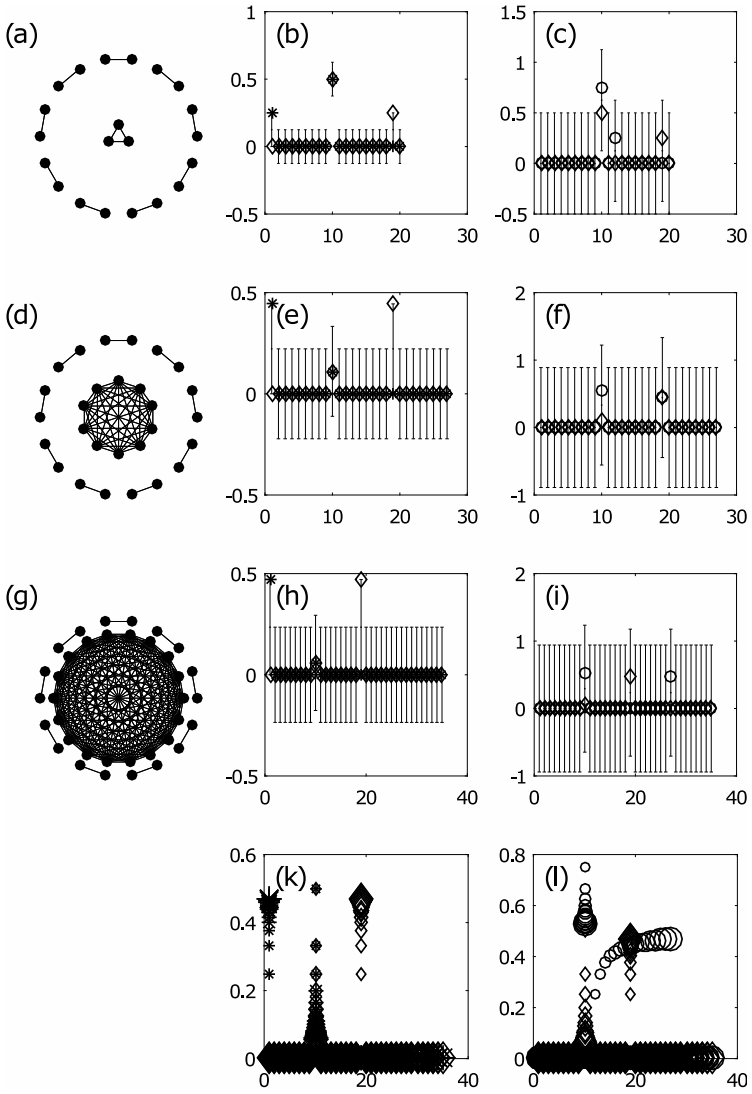

Figure 2.9: Effects of changing $d_{\max}$ by varying the size of the k-complete component in graphs $C(k)$ with $k \in \{3, \ldots, 18\}$. The x-axes display eigengap indices, while the y-axes display normalised eigengap values. (a), (d) and (g) contain plots of graphs $C(3)$, $C(10)$ and $C(18)$, respectively. (b), (e) and (h) bound $g(A, L)$ (shown via the intervals), together with the normalised eigengaps of the adjacency matrix $\mathcal{M}_i/\ell(\mu)$ (stars) and the normalised eigengaps of the Laplacian $\mathcal{L}_i/\ell(\lambda)$ (diamonds) corresponding to the graphs displayed at the start of their respective row of plots. (c), (f) and (i) bound $g(L, L_{rw})$ (shown via the intervals), normalised eigengaps of the Laplacian $\mathcal{L}_i/\ell(\lambda)$ (diamonds) and normalised eigengaps of the normalised Laplacian $\mathcal{N}_i/\ell(\eta)$ (circles). (k) normalised eigengaps of the adjacency matrices (stars) and the normalised eigengaps of the Laplacian (diamonds) corresponding to graphs $C(k)$ with $k \in \{3, \ldots, 18\}$. (l) normalised eigengaps of the Laplacian (diamonds) and the normalised eigengaps of the normalised Laplacian (circles) corresponding to graphs $C(k)$ with $k \in \{3, \ldots, 18\}$. In plots (k) and (l) the marker size of the normalised eigengaps increases with growing k.

ures 2.9(k) and (l) highlight the evolution of the eigengaps of the representation matrices as we vary k from 3 to 18; the eigengap marker size was chosen to increase with k to clearly show the evolution of the normalised eigengaps.

Note that the comparison of magnitudes of eigengaps corresponding to different representation matrices in Figure 2.9 is sensible since we normalised all eigengaps by their spectral support and hence have all eigengaps on the same scale.

While significant differences in the spectra of A and L were observed in Figure 2.7(a) of Section 2.8.2, here the eigengaps mostly agree (see Figure 2.9(h)). In Figure 2.9(k) the first eigengap of the adjacency matrix, the 10^{th} eigengap in both spectra and 19^{th} eigengap in the unnormalised Laplacian spectrum are noticeably larger than other eigengaps. The large first eigengap in the adjacency spectrum is understood as a measure of the connectedness of the graph (as discussed in Section 1.3.2) and therefore, is seen to measure the increasing connectedness in the graph as the size of the fully connected component grows; nothing about clustering is necessarily indicated by the presence of the large first eigengap. The large 10^{th} eigengap in both spectra and 19^{th} eigengap in the unnormalised Laplacian spectrum are understood to indicate clusterings into 10 and 19 clusters, respectively, where the 10 clusters are likely to be the 10 connected components present in the graphs. The spectra of A and L agree on the 10^{th} eigengap, which is found to decrease with increasing $d_{\max}$ in Figure 2.9(k). The 19^{th} eigengaps show clear disagreement between representation matrix spectra, and the bound $g(A, L)$ can be seen to be tight on all three displayed examples, graphs $C(3)$, $C(10)$ and $C(18)$. This is also true for the first eigengaps. The tightness of the bound here is due to the maximal crossover discussed in Section 2.10.1. Figure 2.9(k) shows the difference in the eigengaps of the two representation matrices growing with increasing degree extreme difference.

Turning now to Figures 2.9(c), (f) and (i) it is seen that the edges of the normalised eigengap bounds are not attained in Figure 2.9. However, as discussed in Section 2.8.1, this is not informative on the tightness of the bound, which is valid for *all elements* in the class of graphs $\mathcal{C}_{1,k}$; it is rather concerned with the particular structure of our examples. A feature of this comparison is the changing location of a notable eigengap, which is labelled as the $(k+9)^{\text{th}}$ norm-

alised eigengap present in the normalised Laplacian spectrum corresponding to graph $C(k)$, i.e., the number of clusters detected by the normalised Laplacian seems to depend on k. From Figure 2.9(l), as the degree extreme difference grows we find the difference of the $(k+9)^{\text{th}}$ eigengaps of the two Laplacians to grow. In all of our examples, the $(k+9)^{\text{th}}$ eigengap in the normalised Laplacian spectrum never overtakes the only other notable eigengap, the 10^{th}, in magnitude. So, in the normalised Laplacian spectrum, the evidence in favour of the clustering into the 10 connected components remains the strongest in all our examples. For L this is only the case for small degree extreme differences as the 10^{th} eigengap gets overtaken at $d_{\max} - d_{\min} = 3$ by the growing 19^{th} eigengap as seen in Figure 2.9(k).

It is very interesting to note that in this example a degree extreme difference larger than 3 is sufficient for the three representation matrix spectra to highlight different structures within the graph.

Remark 2.15. A further area of application for the bounds derived in this chapter arises in a big data context. On large datasets the calculation of all three representation matrix spectra will come at a significant cost. However, calculating the bounds on normalised eigengap difference is inexpensive as it only requires calculation of the degree extremes, which are likely to be required at several stages of the graph's analysis. Using our inexpensive bounds practitioners will be able to anticipate the maximally different results they could have obtained using a different representation matrix. In Figure 2.9(b) for example the 10^{th} eigengap of the matrices is clearly bounded away from 0 and knowledge of any one of the two spectra would suffice to say with certainty that the other spectrum will also contain a significantly larger 10^{th} eigengap than other eigengaps in the spectrum. $\triangleleft$

Remark 2.16. The examples of Figure 2.9 show that the eigengap spectra of the three representation matrices can deviate significantly for several eigengaps, while approximately agreeing for the majority. We find these differences to increase with growing degree extreme differences as suggested by the functional form of our bounds. Since the large eigengaps in particular are commonly used to inform graphical analysis, a few large deviations in eigengap spectra can lead to significantly different inference. $\triangleleft$

2.11.2 Eigengap bounds on a SBM sample

Recall that SBMs and their parametrisation were formally defined in Section 1.2.4.

Figure 2.10 contains eigengaps corresponding to the three representation matrices of a graph arising as a realisation of a SBM with parameters $n = 600$; $K = 6$; $M_{1,1} = \ldots = M_{100,1} = M_{101,2} = \ldots = M_{200,2} = M_{201,3} = \ldots = M_{300,3} = M_{301,4} = \ldots = M_{400,4} = M_{501,6} = \ldots = M_{600,6} = 1$ and all other entries of M equal 0;

$$P = \begin{pmatrix} 0.9 & 0.1 & 0.1 & 0 & 0 & 0 \\ 0.1 & 0.9 & 0.1 & 0 & 0 & 0 \\ 0.1 & 0.1 & 0.9 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.1 & 0.9 & 0.9 \\ 0 & 0 & 0 & 0.9 & 0.1 & 0.9 \\ 0 & 0 & 0 & 0.9 & 0.9 & 0.1 \end{pmatrix}.$$

Hence, we have a network of 6 blocks with 100 nodes per block. From observation of P, we recognise that all realisations from this model consist of at least *two connected components* since the probability of an edge between blocks $\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}$ and $\{\mathcal{B}_4, \mathcal{B}_5, \mathcal{B}_6\}$ is equal to 0.

In order to gain more insight into the block structure of this SBM we will consider different definitions of clusters or blocks. In Section 1.3.3 we described the so called '*heterophilic*' clusters, which have fewer connections within than between clusters. Our blockmodel includes such heterophilic blocks in $\{\mathcal{B}_4, \mathcal{B}_5, \mathcal{B}_6\}$. While blocks $\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}$ form clusters in the more traditional sense, first discussed in Section 1.3.2, with more edges within rather than between blocks, which we referred to as '*homophilic*' cluster structure.

In Figure 2.10 we find that all three spectra have non-zero normalised eigengaps with indices 2, 4 and 598. In addition, A has a large first normalised eigengap and L a large 301^{st} normalised eigengap. Hence, with exception of the large 301^{st} normalised eigengap in L's spectrum, in the context of the spectral clustering algorithms the eigengaps suggest searching for similar numbers of clusters.

From von Luxburg (2007, pp. 397–8) we know that the nonzero second eigengap following a zero first eigengap of the Laplacians in Figure 2.10 encodes the separation of the network into two connected components. It is interesting to see that all representation matrices put higher weight on the block structure encoded in the SBM which corresponds to the large 4^{th} and 598^{th} normalised

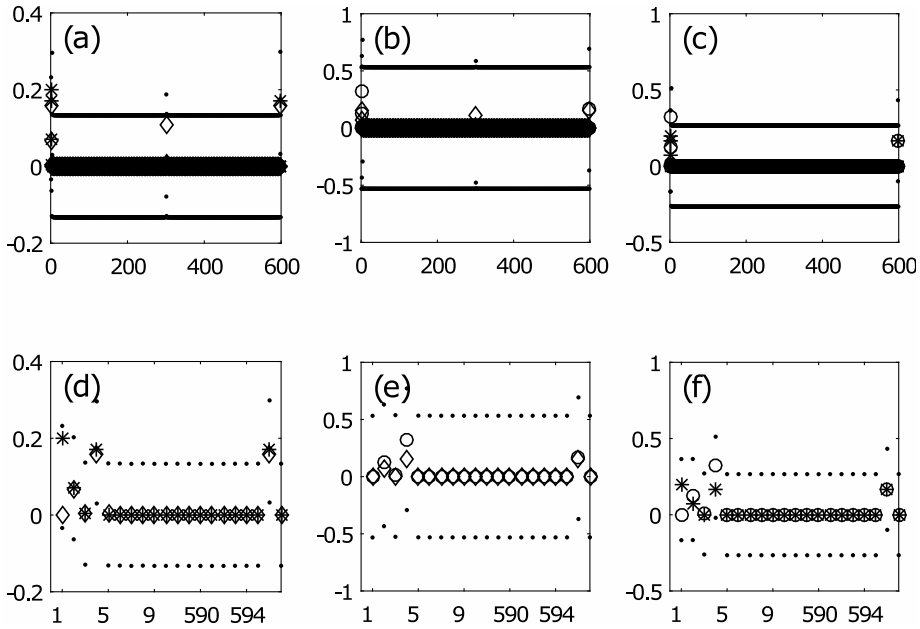

Figure 2.10: Plots (a), (b) and (c) display all normalised eigengaps and plots (d), (e) and (f) display the first and last 10 normalised eigengaps of a graph sampled from a SBM. The x-axes display eigengap indices, while the y-axes display normalised eigengap magnitudes. The stars correspond to eigengaps of A normalised by the spectral support of A, the diamonds display the corresponding quantity for L and the circles represent the corresponding quantity for L_{rw}. The solid black dots correspond to the three bounds, where $g(A, L)$ is displayed in plots (a) and (d), $g(L, L_{rw})$ is shown in plots (b) and (e) and $g(A, L_{rw})$ is displayed in plots (c) and (f).

eigengaps each suggesting the search for the block structure encoded mainly in one of the two connected components of our SBM. It is especially interesting that for A and L the 4^{th} and 598^{th} normalised eigengaps are approximately equal, while for L_{rw} the 4^{th} normalised eigengap is almost twice as large as the 598^{th} normalised eigengap. Both components have symmetric parameters and therefore it is highly interesting that the eigengaps in L_{rw} suggest stronger evidence for the homophilic cluster structure encoded by the 4^{th} eigengap rather than the heterophilic cluster structure encoded by the 598^{th} eigengap. A much larger simulation of SBMs with varying parameters, which is beyond the scope of this work, could establish whether L_{rw} does indeed have a structured preference for homophilic over heterophilic blocks.

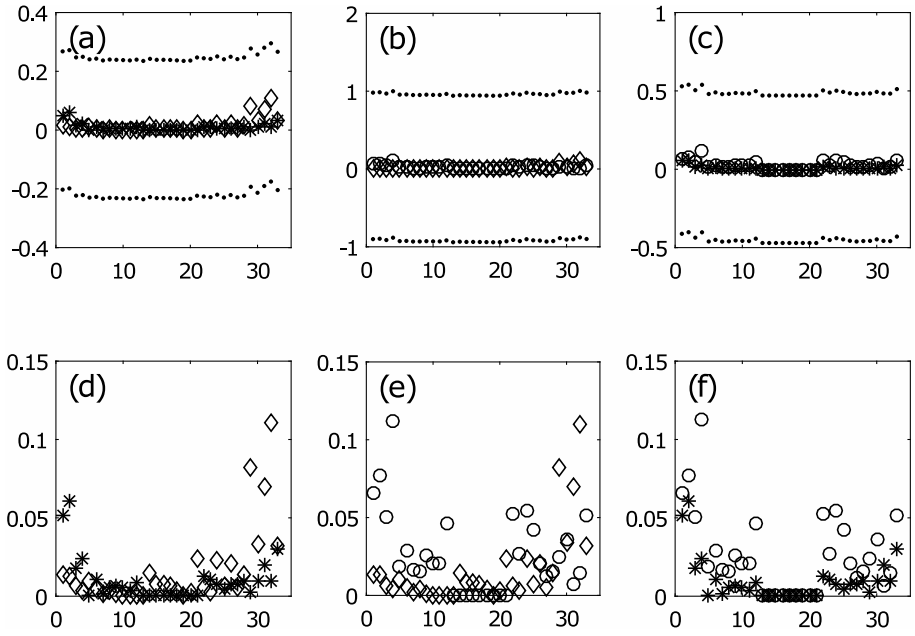

Figure 2.11: Plots (a), (b) and (c) display the karate dataset normalised eigengaps with eigengap bounds. Plots (d), (e) and (f) display the karate dataset normalised eigengaps without eigengap bounds. The x-axes display eigengap indices, while the y-axes display normalised eigengap magnitudes. The stars correspond to eigengaps of A normalised by the spectral support of A, the diamonds display the corresponding quantity for L and the circles represent the corresponding quantity for L_{rw}. The solid black dots correspond to the three bounds, where $g(A, L)$ is displayed in plot (a), $g(L, L_{rw})$ is shown in plot (b) and $g(A, L_{rw})$ is displayed in plot (c).

2.11.3 Eigengap bounds on the karate dataset

Figure 2.11 contains the eigengaps of the karate data set. From Figure 2.11 (a)-(c) we see that only a small proportion of our bounds is covered by the normalised eigengaps of the karate dataset indicating that there could be much larger differences in the eigengaps for a graph with equal degree extremes. In Figure 2.11 (d)-(f) we observe that the evidence found in the spectra is clearly in disagreement. The three largest normalised eigengaps of the three representation matrices in descending order are the $2^{\text{nd}}, 1^{\text{st}}$ and 33^{rd} for A, $32^{\text{nd}}, 29^{\text{th}}$ and 31^{st} for L and $4^{\text{th}}, 2^{\text{nd}}$ and 1^{st} for L_{rw}. Hence, in the context of spectral clustering, A and L_{rw} see strongest evidence for a homophilic clustering into 2 and 4 components, respectively, while L suggests clustering according to the eigengaps of the heterophilic end of the spectrum. Interestingly, we find many less 0-eigengaps for this real dataset when compared to the eigengaps

corresponding to the simulated graph in Figure 2.10. For the karate dataset the representation matrices suggest not only different numbers of clusters, but also different types (homophilic vs heterophilic).

Chen and Hero (2015) suggest that there is a deeper community structure in the karate dataset. They find the modularity of their clustering to increase as they moved from recovering 2 to 3 to 4 homophilic clusters. This insight from a modularity method is most closely represented by the eigengaps of the normalised Laplacian L_{rw}, which also suggests recovering $4, 2$ or 3 homophilic clusters.

2.12 Application to spectral clustering

In this section we illustrate the direct implication of the graph representation matrix choice on the outcome of spectral clustering in several examples, the eigengaps of which were studied in Section 2.11. We first review the spectral clustering algorithm itself. Then we cluster graph $C(18)$ according to the two large eigengaps which arose from the analysis in Section 2.11.1. Finally, we discuss a clustering of the karate data set according to all three representation matrices. We find the choice of representation matrix to have a significant impact on the clustering outcome.

2.12.1 Spectral clustering

The spectral clustering algorithm is comprehensively introduced and derived using several different approaches in von Luxburg (2007). In this work we use the simplest form of spectral clustering, which is to run the k-means algorithm in the space spanned by the first r eigenvectors of any of the three representation matrices. Large eigengaps are understood to indicate the number of clusters discovered by the different representation matrices, i.e., a large r^{th} eigengap is understood to indicate a clustering into r communities (von Luxburg, 2007, p. 410). This interpretation of the eigengaps can be motivated via the Cheeger inequalities, which we presented in Sections 1.3.2, 1.4.2 and 1.5.3.

Spectral clustering has had a big impact in the recent analysis of graphs, there currently is an active literature working on the theoretical guarantees which can be given for a clustering recovered in this way (Lei and Rinaldo, 2015; Rohe et al., 2011). The SBM has dominated the research. For their variant of the spectral clustering algorithm, Rohe et al. (2011) bound the number of misclustered nodes under the SBM. Lei and Rinaldo (2015) work in a sparser regime than Rohe et al. (2011), bounding the number of nodes for which clustering correctness cannot be guaranteed under the SBM, when using the spectral clustering algorithm. They also bound the number of misclustered nodes, under the degree corrected SBM, for the spherical k-median spectral clustering algorithm. Both results are derived for the adjacency matrix A. Lyzinski et al. (2014) manage to prove asymptotically almost sure error free clustering (under a few assumptions also on the minimal eigengap size) resulting from the spectral clustering on the SBM and the degree corrected SBM.

In practice, the k-means step in the spectral clustering algorithm produces different results depending on its randomly selected starting points according to the k-means++ algorithm in the MATLAB implementation. We have chosen to remedy this by varying the `Replicates` input of the k-means algorithm. In this section we set the number of replicates to 10 for all clusterings since we observed it to yield stable results for the graphs in this section. Hence, the clustering is repeated 10 times using new initial cluster centroid positions each time. The clustering with lowest within cluster sum to centroid distance is then returned as the final result.

2.12.2 Spectral clustering of graph $C(18)$

We found that all representation matrices recover the 10 connected components perfectly when applying the spectral clustering algorithm to their respective first 10 eigenvectors. Additionally to the large tenth eigengap, we observed a large 19^{th} eigengap in the *unnormalised* and a large 27^{th} eigengap in the *normalised* Laplacian spectrum for graph $C(18)$ in Section 2.11.1. At first sight it is unclear what structure is suggested by these eigengaps and hence we display the clustering according to the first 19 and 27 eigenvectors of the unnormalised and normalised Laplacians in Figures 2.12(a) and (b), respectively.

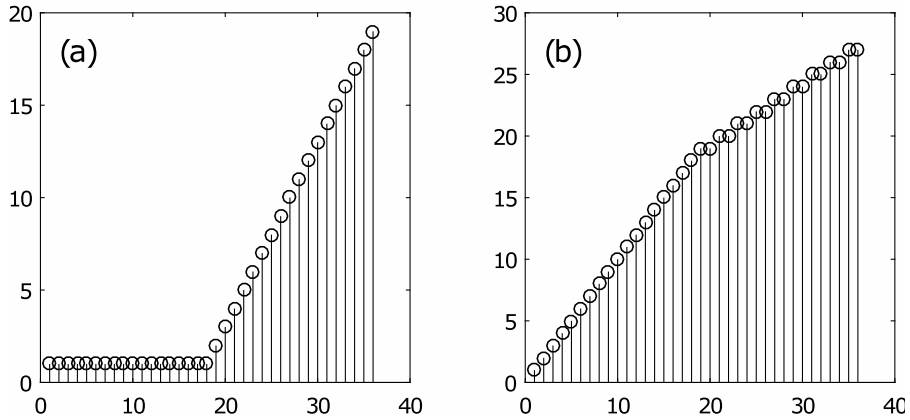

Figure 2.12: (a) and (b) display the spectral clustering of graph $C(18)$ according to the first 19 eigenvectors of the unnormalised Laplacian L and the first 27 eigenvectors of the normalised Laplacian, respectively. The x-axes display a numbering of the 36 nodes in the graph; numbers 1 to 18 correspond to nodes in the 18-complete component and nodes 19-36 make up the 9 2-complete components. The y-axes display cluster indices ranging from 1 to 19 in (a) and 1 to 27 in (b).

From Figure 2.12(a) we see that the large 19^{th} eigengap in the unnormalised Laplacian spectrum suggests a clustering where the 18-complete component is recovered as a single cluster and the 9 2-complete components are treated as 18 individual clusters. From Figure 2.12(b) we find that when clustering according to the first 27 eigenvectors of the normalised Laplacian we recover the 18-complete component as 18 clusters of one node each and each of the 2-complete components is exactly recovered. When varying $k \in \{1, \ldots, 18\}$, the number of clusters in this clustering depends on the size of the k-complete component since each node in it is taken to be a separate cluster. This clarifies why we observed the large $(k+9)^{\text{th}}$ eigengap in the normalised Laplacian spectrum as we varied k in Figure 2.9.

Essentially, the two clusterings presented in Figure 2.12 are complementary, in each case clustering all nodes of the same degree in one cluster per connected component and splitting the nodes of the remaining degree into separate clusters of single nodes.

The eigengaps, in conjunction with the spectral clustering algorithm, have revealed splits of the graphs, some of which might not have been anticipated.

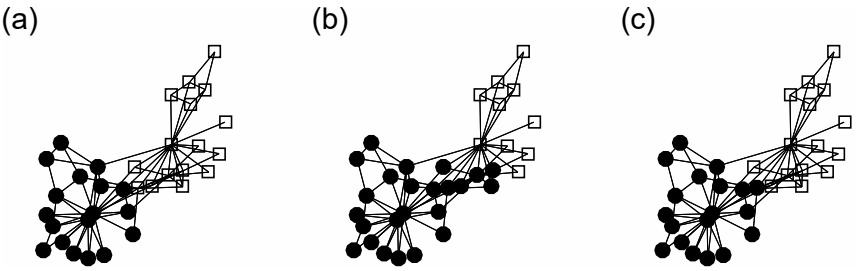

Figure 2.13: Spectral clustering of the karate graph according to the first two eigenvectors of the adjacency matrix A in (a), the unnormalised Laplacian L in (b) and the normalised Laplacian L_{rw} in (c).

2.12.3 Spectral clustering of the karate data set graph

To further illustrate the impact of the representation matrix choice, we display the spectral clustering according to the first two eigenvectors of each of the three representation matrices corresponding to the karate data set.

Since for the karate dataset, we have a reasonably large degree extreme difference, $d_{\max} - d_{\min} = 17$, we expect to see deviating results in the spectral clustering according to the different representation matrices.

At first sight, all clusterings displayed in Figure 2.13 seem sensible. The clustering according to the two normalised Laplacians agrees and hence only one of the two is displayed. The clustering according to the adjacency matrix extends the cluster marked with the unfilled, square nodes by one node, karate club member 3, in comparison to the clustering according to the normalised Laplacian. In contrast, the unnormalised Laplacian detects 5 nodes less (karate club members 2, 4, 8, 14 and 20) in the cluster marked by the unfilled square nodes, than does the normalised Laplacian.

We find the clustering according to the adjacency matrix A to agree with the ground truth clustering into social factions within the karate club as recorded by Zachary (1977). The normalised Laplacians misplace one out of 34 nodes, which is known to be difficult to cluster correctly in the literature (Fortunato, 2010). The unnormalised Laplacian however, misplaces 6 nodes, only one of which is known to us to be commonly misclustered. Hence, the unnormalised Laplacian clustering is clearly outperformed by the other two representation matrices when using the first two eigenvectors to find two communities in the

karate data set. In von Luxburg et al. (2008) the conditions under which spectral clustering using the unnormalised Laplacian converges are shown to be more restrictive than the conditions under which spectral clustering according to the normalised Laplacian converges. von Luxburg et al. (2008) hence advocate using the normalised Laplacian for spectral clustering over the unnormalised Laplacian. Our clustering results on graph $C(18)$ and the karate data set agree with this recommendation.

Remark 2.17. For our examples, the choice of representation matrix clearly impacts the results of cluster inference via spectral clustering. We suggest considering the degree extreme difference as a parameter in graphical analysis to infer the potential impact of the choice of representation matrix. ◁

2.13 Bounding the complexity of a graph in terms of the adjacency spectrum

Another potential use of our eigenvalue and eigengap bounds is the transfer of graph theoretic results from one representation spectrum to another. In this section we illustrate this use by studying the complexity of a graph.

Recall from Point (14) in List 1.22, that the complexity or number of all possible spanning trees of a graph G, denoted ξ, is related to the spectrum of the unnormalised Laplacian matrix as follows,

$$\xi = \frac{1}{n} \prod_{j=2}^{n} \lambda_j. \tag{2.30}$$

In his discussion of the complexity van Mieghem (2011, pp. 76-8) describes how for d-regular graphs the complexity can also be computed from the adjacency eigenvalues, due to the exact affine map relating eigenvalues of the two representation matrices.

$$\xi = \frac{1}{n} \prod_{j=2}^{n} (d - \mu_j).$$

For general graphs, our bound $e(A, L)$ can be used to provide an upper bound on the graph complexity using the adjacency matrix spectrum as is done in Theorem 2.18.

Theorem 2.18. Let G be a graph with n nodes and degree extremes $d_{\min}$ and $d_{\max}$. Then the G's complexity can be bounded as follows,

$$\xi \leq \frac{1}{n} \prod_{j=2}^{n} \left(\left| \frac{d_{\max} + d_{\min}}{2} - \mu_i \right| + \frac{d_{\max} - d_{\min}}{2} \right).$$

Proof. We begin by producing an upper bound on the eigenvalues of the unnormalised Laplcian using our bound $e(A, L)$ from Theorem 2.2.

$$\left| \frac{d_{\max} + d_{\min}}{2} - \mu_i - \lambda_i \right| \leq \frac{d_{\max} - d_{\min}}{2}$$

$$\Rightarrow \left| |\lambda_i| - \left| \frac{d_{\max} + d_{\min}}{2} - \mu_i \right| \right| \leq \frac{d_{\max} - d_{\min}}{2} \tag{2.31}$$

$$\Rightarrow \lambda_i - \left| \frac{d_{\max} + d_{\min}}{2} - \mu_i \right| \leq \frac{d_{\max} - d_{\min}}{2} \tag{2.32}$$

$$\Rightarrow \lambda_i \leq \left| \frac{d_{\max} + d_{\min}}{2} - \mu_i \right| + \frac{d_{\max} - d_{\min}}{2}, \tag{2.33}$$

where we use the reverse triangle inequality in (2.31) and use the fact that all Laplacian eigenvalues are positive, as shown in Appendix A.2, in (2.32).

Now by combining (2.30) with (2.33), we can bound the complexity of a general graph using the adjacency spectrum,

$$\xi \leq \frac{1}{n} \prod_{j=2}^{n} \left(\left| \frac{d_{\max} + d_{\min}}{2} - \mu_i \right| + \frac{d_{\max} - d_{\min}}{2} \right).$$

□

3

Extending the Davis–Kahan Theorem for Comparing Eigenvectors of Two Symmetric Matrices: Theory

An arXiv preprint of the work in this chapter can be accessed at Lutzeyer and Walden (2019a).

Having reviewed the eigenproperties of the representation matrices in Chapter 1 and compared their spectral properties in Chapter 2, we now move on to work on the eigenvectors of the representation matrices. In this chapter we extend the Davis–Kahan theorem, which, amongst other things, enables the comparison of the spaces spanned by the eigenvectors of the representation matrices.

More formally, the Davis–Kahan theorem can be used to bound the distance of the spaces spanned by the first r eigenvectors of any two symmetric matrices, given that the corresponding eigenvalues satisfy a certain structure. In this chapter, we observe that use can be made of matrix transformations to extend the applicability of the Davis–Kahan theorem by removing the straitjacket of such a highly constrained eigenvalue structure.

The value of the transformation step can be seen when thinking about the

many practical uses for such a comparison. Bounding the distance of the subspaces spanned by the eigenvectors of the graph representation matrices throws up some interesting issues, e.g., the largest eigenvalues of the adjacency matrix correspond to the smallest eigenvalues of the Laplacians, and an eigenvector comparison corresponds to comparing opposite ends of the eigenvalue spectrum; this renders the standard Davis–Kahan theorem inapplicable (due to the eigenvalue restrictions), while our new extended version incorporating a matrix transformation is highly apposite. A second related application is the comparison of spaces spanned by eigenvectors of the representation matrices to spaces spanned by the eigenvectors of their generating matrices in the case of the ubiquitous stochastic blockmodel (SBM); this can be used for consistency and rate of convergence studies for different methods based on the eigenvectors of the representation matrices in such a model (Cape et al., 2019; Eldridge et al., 2018; Rohe et al., 2011). A third example is the comparison of spaces spanned by the eigenvectors of the sample covariance matrix and its population covariance matrix in a spiked covariance model (Johnstone, 2001; Jung and Marron, 2009). In these second and third examples our new approach can deliver tighter bounds than the standard Davis–Kahan theorem, due to the matrix transform allowing a relaxation and utilisation of the eigenvalue structure imposed by the standard Davis–Kahan theorem.

To describe the novel contributions of this work, we need to introduce our standard notation. We denote the two symmetric matrices of interest by Φ and Ψ. The matrices formed from the eigenvectors of these matrices will be denoted by U and W, respectively. We work with two commonly considered unitarily invariant distance metrics to measure the distance of subspaces of $\mathbb{R}^n$. The considered metrics are $\|(I - UU^T)WW^T)\|_2$ and $\inf_{R\in O(r)} \|U - WR\|_F$, where $O(r)$ denotes the orthogonal group of $r \times r$ orthogonal matrices. Both of these metrics, which use the two norm or spectral norm and Frobenius norm respectively, are functionally related to the canonical angles between spaces and later we shall relate them via an inequality.

We summarise the content of this work in the following list.

1. For any two symmetric matrices Φ and Ψ, the standard Davis–Kahan theorem (Bhatia et al., 2013, pp. 211-2) enables bounding of the metric $\|(I-UU^T)WW^T)\|_2$, measuring the distance between spaces spanned by

eigenvectors of the matrices, given that their corresponding eigenvalues satisfy the standard structure specified in Requirements 3.14.

2. By applying a polynomial transformation to one of the matrices under comparison we leave $\|(I - UU^T)WW^T)\|_2$ unchanged, since the eigenvectors are unchanged, but can change the location of the corresponding eigenvalues (Theorem 3.20). This allows us to relax the necessary Davis–Kahan structure of the eigenvalues (Requirements 3.14) to a non-zero eigengap structure (Assumption 3.21). Furthermore, we find that since the Davis–Kahan bound (Theorem 3.13) depends on the matrices and eigenvalues under comparison, we can choose the polynomial transformation to not only relax the conditions on the eigenvalues, but also to reduce the Davis–Kahan bound value. Use of the metric $\inf_{R\in O(r)} \|U - WR\|_F$ is common (Lei and Rinaldo, 2015; Rohe et al., 2011; Tang and Priebe, 2018) and via Lemma 3.12 we are able to produce an upper bound for it using our sharpened version of the Davis–Kahan theorem as is done in Corollary 3.26.

In Section 3.1 we briefly discuss previous work on the Davis–Kahan theorem. In Section 3.2 we review two distance metrics on subspaces spanned by eigenvectors, and relate them via an inequality. In Section 3.3 we show that under a suitable assumption, we are always able to compare the spaces spanned by the first r eigenvectors of two symmetric matrices via the Davis–Kahan theorem; the bound involves a between-matrix eigengap. Section 3.4 presents an extension to the Davis–Kahan theorem which uses a polynomial matrix transformation of one of the symmetric matrices and also applies to the comparison of the spaces spanned by any two sets of consecutive and corresponding eigenvectors from two symmetric matrices. Assumptions on the spectra are developed. The special case of affine transforms is discussed in Section 3.5 and a proof of concept is given in Section 3.6.

3.1 Some background on the Davis–Kahan theorem

The Davis–Kahan theorem first appeared in Davis and Kahan (1970) and is a topic of current research, being of interest in, e.g., the analysis of spectral graph embedding methods (Tang and Priebe, 2018) and of principal component analysis of covariance matrices (Fan et al., 2013; Wang and Fan, 2017). von Luxburg (2007) states that the Davis–Kahan theorem forms the basis for the perturbation approach to spectral clustering of networks, where cluster structure is seen as a perturbed version of a structure with perfectly disjoint clusters. A new variant with a statistical flavour was recently published by Yu et al. (2015), where the denominator of the bound is reformulated to be purely in terms of one of the spectra under comparison.

Tightening the Davis–Kahan theorem using different norms, probabilistic methods and making mild assumptions on the structure of the perturbation or the underlying matrix has also been the subject of much recent research work. In their recent paper O'Rourke et al. (2018) use a probabilistic approach to sharpen the Davis–Kahan theorem using certain structural properties of the unperturbed matrix, such as having low rank. In Cape et al. (2019) the Davis–Kahan theorem is used as a coarse benchmark for their bound on the difference of spaces spanned by eigenvectors in the 2-to-infinity norm. Eldridge et al. (2018) also state that the tightness of the Davis–Kahan theorem is in many settings suboptimal and hence present their own tightened bound using the infinity norm. In contrast to these approaches, our generalised and tightened Davis–Kahan theorem applies to any pair of symmetric matrices.

As well as the tightness issues, a drawback of the Davis–Kahan theorem is that the choice of eigenvalue intervals might not be suitable for the two matrices we want to compare. For example in the particular context of one matrix being considered a perturbation of the other, von Luxburg (2007, p. 407) noted that "If the perturbation is too large or the eigengap is too small, we might not find a set S_1 such that both the first r eigenvalues of the perturbed and unperturbed versions of L are contained in S_1." Indeed, if we want to compare eigenvectors corresponding to eigenvalues lying on opposing ends of the spectra, as for example when comparing corresponding adjacency to graph

Laplacian eigenvectors, then, finding eigenvalue intervals such that the Davis–Kahan theorem can be applied is rarely possible. By applying a polynomial matrix transform to one of the matrices under comparison, we are able to show that for any two symmetric matrices, given nonzero eigengaps on either side of two contiguous sets of eigenvalues, we are able to find a valid set of eigenvalue intervals such that their corresponding eigenvectors can be bounded using a Davis–Kahan bound.

To summarise, the key innovation in this chapter is that the application of a matrix transformation allows us to broaden the class of cases in which the Davis–Kahan theorem applies by shifting the restrictive assumptions from the *eigenvalues* to the *transformation parameters*, with the latter being under our control. As a consequence the Davis–Kahan theorem can then be applied to the comparison of any two sets of consecutive and corresponding eigenvectors from two symmetric matrices (excluding degenerate examples).

3.2 Subspaces of $\mathbb{R}^n$ — spaces spanned by $r \leq n$ eigenvectors

We shall compare spaces spanned by eigenvectors of two symmetric matrices $\Phi, \Psi \in \mathbb{R}^{n \times n}$ with eigenvalues $\phi_1 \leq \phi_2 \leq \ldots \leq \phi_n$ and $\psi_1 \leq \psi_2 \leq \ldots \leq \psi_n$ and corresponding eigenvectors $\{u_1, u_2, \ldots, u_n\}$ and $\{w_1, w_2, \ldots, w_n\}$, respectively.

In fact we have reason to compare the subspaces of $\mathbb{R}^n$ spanned by $r \in \{1, \ldots, n\}$ consecutive eigenvectors of the two symmetric matrices. We define matrices $U_j = [u_{j+1}, \ldots, u_{j+r}]$ and $W_j = [w_{j+1}, \ldots, w_{j+r}]$ holding r consecutive eigenvectors of Φ and Ψ, respectively. Note that U_0 and W_0 correspond to the *first* r eigenvectors, i.e., the r eigenvectors corresponding to the r smallest eigenvalues. The value of $j \in \{0, 1, \ldots, n-r\}$ determines a particular contiguous block of r eigenvectors, i.e., j is an offset or shift parameter; if the value of j is irrelevant in a particular discussion then the parameter will be suppressed for brevity.

3.2.1 A comparison of eigenspaces and the space spanned by r eigenvectors

In this section we delineate the difference between eigenspaces and the space spanned by a set of eigenvectors. The understanding of these spaces and the difference between them is crucial to understand the object of inference of this work, where we are comparing the space spanned by a given set of eigenvectors as is standard in the literature.

Horn and Johnson (1985) define the eigenspace corresponding to an eigenvalue as follows.

Definition 3.1. Let $\Phi \in \mathbb{R}^{n\times n}$. For a given eigenvalue ϕ of Φ, the set of all vectors $u \in \mathbb{C}^n$ satisfying $\Phi u = \phi u$ is called the *eigenspace* of Φ corresponding to the eigenvalue ϕ. Note that every nonzero element of this eigenspace is an eigenvector of Φ corresponding to ϕ. ◁

Hence, linear combinations of eigenvectors corresponding to the same eigenvalue are also eigenvectors. But, linear combinations of eigenvectors corresponding to different eigenvalues are not as is shown in Remark 3.2.

Remark 3.2. Let u_i and u_j be two eigenvectors of a matrix, Φ, with corresponding eigenvalues ϕ_i and ϕ_j. Furthermore, let $x, y \in \mathbb{R}\backslash\{0\}$ such that $(xu_i + yu_j)$ is not equal to the vector of all zeros, then

$$\Phi(xu_i + yu_j) = \Phi xu_i + \Phi yu_j = \phi_i xu_i + \phi_j yu_j, \tag{3.1}$$

clearly, $(xu_i + yu_j)$ is an eigenvector of Φ if and only if $\phi_i = \phi_j$. Therefore, the nonzero linear combination of two eigenvectors, u_i and u_j, of a matrix, Φ, is itself an eigenvector of Φ if and only if u_i and u_j are eigenvectors corresponding to the same eigenvalue. ◁

In this work we will be comparing r consecutive eigenvectors of any two symmetric matrices, i.e., r eigenvectors corresponding to r consecutive eigenvalues ordered by magnitude, where $r \leq n$. This comparison will lead us to consider the subspace of $\mathbb{R}^n$ spanned by the r consecutive eigenvectors. Note, that the space spanned by the r consecutive eigenvectors of a matrix contains all linear combinations of the matrix's r consecutive eigenvectors whatever their

eigenvalue. The union of the eigenspaces corresponding to the r consecutive eigenvalues, counting multiplicities, contains only linear combinations of eigenvectors corresponding to the same eigenvalues as is discussed in Remark 3.2. Hence, in general, the space spanned by the r consecutive eigenvectors of a matrix is not equal to the union of the eigenspaces corresponding to the r consecutive eigenvalues, counting multiplicities. In fact, the union of eigenspaces corresponding to the r consecutive eigenvalues, counting multiplicities, is a subset of the space spanned by the r consecutive eigenvectors.

In Example 3.3 we show how several different eigenvectors sets can be chosen for a particular graph adjacency matrix and that the union of eigenspaces is a subset of the space spanned by the corresponding eigenvectors.

Example 3.3. Consider the complete graph on 4 nodes, denoted K_4, where every node is connected to all others in the graph. Note that K_4 is a 3-regular graph; its adjacency matrix takes the following form,

$$A_{K_4} = \begin{pmatrix} 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 1 \\ 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 \end{pmatrix}.$$

The eigenvalues of the adjacency matrix of a complete graph are known in general, in this case they are $\mu_1 = 3$ and $\mu_2 = \mu_3 = \mu_4 = -1$. The first eigenvector is a constant vector where all entries are equal, e.g., take $u_1 = \mathbf{1}$ to be the vector of all ones. The eigenvalue -1 has a corresponding three-dimensional eigenspace isomorphic to $\mathbb{R}^3$. The only restriction imposed by the eigenvalue eigenvector equations corresponding to the eigenvalue -1 is that the elements of each eigenvector should sum to 0. Hence, we can choose the following orthonormal set of eigenvectors of A_{K_4},

$$u_1 = \frac{1}{2}\,\mathbf{1}, \quad u_2 = \frac{1}{\sqrt{12}} \begin{pmatrix} 1 \\ -3 \\ 1 \\ 1 \end{pmatrix}, \quad u_3 = \frac{1}{\sqrt{6}} \begin{pmatrix} 1 \\ 0 \\ -2 \\ 1 \end{pmatrix}, \quad u_4 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}.$$

Alternatively, we can choose the eigenvectors u_2^*, u_3^* and u_4^* to be the non-orthogonal 2-sparse eigenvector set of K_4 mentioned in Teke and Vaidyanathan (2017, p. 5416). Recall that twin vertices and 2-sparse eigenvectors were introduced in Section 1.7.

$$u_1 = \frac{1}{2}\,\mathbf{1}, \quad u_2^* = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ 1 \\ -1 \\ 0 \end{pmatrix}, \quad u_3^* = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix}, \quad u_4^* = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ 0 \\ 0 \\ -1 \end{pmatrix}.$$

We denote the eigenspaces corresponding to the first and second eigenvalue of A_{K_4} as $\mathcal{E}_1$ and $\mathcal{E}_2$, respectively, i.e.,

$$\begin{aligned} \mathcal{E}_1 &= \operatorname{span}(u_1), \\ \mathcal{E}_2 &= \operatorname{span}(u_2, u_3, u_4) = \operatorname{span}(u_2^*, u_3^*, u_4^*). \end{aligned}$$

Note that both $\{u_1, u_2, u_3, u_4\}$ and $\{u_1, u_2^*, u_3^*, u_4^*\}$ span $\mathbb{R}^4$,

$$\mathbb{R}^4 = \operatorname{span}(u_1, u_2, u_3, u_4) = \operatorname{span}(u_1, u_2^*, u_3^*, u_4^*).$$

Hence, we are able to see that the union of the eigenspaces is unequal to the span of their corresponding eigenvector sets, in fact, the union of eigenspaces is a subset of the span of their corresponding eigenvectors, i.e.,

$$\mathcal{E}_1 \cup \mathcal{E}_2 \subset \mathbb{R}^4. \qquad \triangleleft$$

3.2.2 Groundwork

We now formally introduce the different spaces associated with matrices composed of orthogonal columns, which will be utilised when comparing the eigenvectors of symmetric matrices. Edelman et al. (1998) give a very nice introduction to these spaces. Below we summarise their definitions.

Definition 3.4. Let $1 \leq r \leq n$. We will be working with the following three matrix groups:

1. The *orthogonal group*, denoted $O(r)$, of $r \times r$ orthogonal matrices.
2. The *Stiefel manifold*, denoted $\mathbb{V}_{n,r}$, consisting of $n \times r$ matrices with orthonormal columns.
3. The *Grassmann manifold*, denoted $\mathbb{G}_{n,r}$, consisting of r-dimensional subspaces of $\mathbb{R}^n$. Elements of the Grassmann manifold are equivalence

classes of elements of the Stiefel manifold, where two elements of the Stiefel manifold are equivalent if their columns span the same subspace. ◁

An interesting observation in Edelman et al. (1998) is that both the Stiefel manifold and the Grassmann manifold can be expressed as quotient groups of the orthogonal group $O(n)$. The following Lemma helps us understand the equivalence classes formed on the Stiefel manifold to produce elements on the Grassmann manifold and is furthermore used to motivate our distance metric on eigenvector sets.

Lemma 3.5. If the columns of two elements $U, W \in \mathbb{V}_{n,r}$ span the same subspace of $\mathbb{R}^n$, then, there exists an orthogonal matrix $Q \in O(r)$ such that $U = WQ$. ◁

Our proof of Lemma 3.5 is shown in Appendix C.1. Note that two elements $U, W \in \mathbb{V}_{n,r}$, the columns of which span the same subspace of $\mathbb{R}^n$, correspond to a single element in the Grassmann manifold.

Definition 3.6. The orthogonal projector onto the subspace spanned by the eigenvectors $u_{j+1}, \ldots, u_{j+r}$ can be expressed in terms of $U_j = [u_{j+1}, \ldots, u_{j+r}]$ as $U_j U_j^T$ and the orthogonal projector onto the complementary subspace is equal to $(I - U_j U_j^T)$ (Meyer, 2000, p. 430). ◁

It is pointed out in Edelman et al. (1998, p. 319) that each element in the Grassmann manifold, $\mathbb{G}_{n,r}$ has a corresponding unique orthogonal projector onto its space of the form WW^T, where $W \in \mathbb{V}_{n,r}$.

3.2.3 Measuring distance between subspaces of $\mathbb{R}^n$

Now that we are familiar with the different spaces associated with the eigenvectors of symmetric matrices, we define two metrics on the spaces spanned by two eigenvector sets in Definitions 3.7 and 3.8. Then we introduce the notion of canonical angles between spaces in Definition 3.9 and relate both metrics to the canonical angles in Theorem 3.10 and Remark 3.11. We end by describing the relation of the two metrics in Lemma 3.12.

Definition 3.7. (Stewart and Sun, 1990, p. 95) Let $U, W \in \mathbb{V}_{n,r}$ be matrices with orthonormal columns which span r-dimensional subspaces of $\mathbb{R}^n$, $\mathcal{U}, \mathcal{W} \in \mathbb{G}_{n,r}$, respectively. We define the unitarily invariant metric on the distance of $\mathcal{U}$ and $\mathcal{W}$, denoted $\rho_1(\mathcal{U}, \mathcal{W})$, as,

$$\rho_1(\mathcal{U}, \mathcal{W}) = \inf_{R \in O(r)} \|U - WR\|_F. \tag{3.2}$$

$\triangleleft$

In Stewart and Sun (1990, p. 95) $\rho_1(\mathcal{U}, \mathcal{W})$ is motivated as a metric for the distance between subspaces of $\mathbb{R}^n$ by pointing out that $\rho_1(\mathcal{U}, \mathcal{W}) = 0$ when $\mathcal{U} = \mathcal{W}$. This follows directly from Lemma 3.5 by taking $R = Q$.

The second metric on subspaces is now defined.

Definition 3.8. (Stewart and Sun, 1990, p. 94) Let $U, W \in \mathbb{V}_{n,r}$ be matrices with orthonormal columns which span r-dimensional subspaces of $\mathbb{R}^n$, $\mathcal{U}, \mathcal{W} \in \mathbb{G}_{n,r}$, respectively. We define the unitarily invariant metric on the distance of $\mathcal{U}$ and $\mathcal{W}$, denoted $\rho_2(\mathcal{U}, \mathcal{W})$, as,

$$\rho_2(\mathcal{U}, \mathcal{W}) = \left\| UU^T \left(I - WW^T \right) \right\|_2. \qquad \triangleleft$$

Recall from Definition 3.4 that each different basis of a subspace corresponds to a Stiefel manifold element, while on the Grassmann manifold all bases of a subspace are represented by a single element. Note that the metric $\rho_1(\mathcal{U}, \mathcal{W})$ is directly dependent on Stiefel manifold elements U, W, while the metric $\rho_2(\mathcal{U}, \mathcal{W})$ is working with the projectors $UU^T, \left(I - WW^T \right)$. Since projectors are unique to their corresponding subspaces, they correspond to Grassmann manifold elements. In agreement with Deri and Moura (2017), we want to compare Grassmann manifold elements, i.e., entire subspaces rather than just individual bases of these spaces. The metric $\rho_1(\mathcal{U}, \mathcal{W})$ achieves this subspace comparison via the infimum over the orthogonal matrices ensuring that we are considering all bases of the spaces we want to compare. Therefore, both metrics are comparing Grassmann manifold elements. In Cape et al. (2019) the minimisation over all orthogonal matrices in (3.2) is described as enabling *basis alignment* of the two spaces under comparison.

Note that in an orthogonally invariant norm such as the Frobenius norm

the direct comparison of individual bases is not possible. The orthogonal invariance of the norm implies that the comparison of two elements on the Stiefel manifold corresponds to the comparison of a class of elements of the Stiefel manifold related via orthogonal transformations, i.e., $\|U - W\|_F = \|UQ' - WQ'\|_F$ for all $Q' \in O(r)$. Hence, without the minimisation over all orthogonal matrices in (3.2) we are simply considering transformations of both U and W by the same orthogonal transforms rather than by decoupled orthogonal transformations as achieved by the addition of the minimisation step.

When discussing the distance of subspaces of $\mathbb{R}^n$ the notion of canonical angles, as generalisation of angles between lines, is integral. We will therefore define canonical angles in Definition 3.9 and then discuss how both $\rho_1(\mathcal{U}, \mathcal{W})$ and $\rho_2(\mathcal{U}, \mathcal{W})$ are functionally related to the canonical angles. The following definition is adapted from Vu and Lei (2013).

Definition 3.9. Let $\mathcal{U}, \mathcal{W} \in \mathbb{G}_{n,r}$ be r-dimensional subspaces of $\mathbb{R}^n$ with orthogonal projectors UU^T and WW^T. Denote the singular values of $UU^T(I - WW^T)$ by $\gamma_1 \geq \ldots \geq \gamma_n$. The *canonical angles* between $\mathcal{U}$ and $\mathcal{W}$ are the numbers,

$$\theta_k(\mathcal{U}, \mathcal{W}) = \arcsin(\gamma_k).$$

In the literature the diagonal matrix $\Theta(\mathcal{U}, \mathcal{W}) = \operatorname{diag}(\theta_1(\mathcal{U}, \mathcal{W}), \ldots, \theta_n(\mathcal{U}, \mathcal{W}))$ is often considered. ◁

The following theorem states the functional relationship of $\rho_1(\mathcal{U}, \mathcal{W})$ to the canonical angles between $\mathcal{U}$ and $\mathcal{W}$.

Theorem 3.10. (Stewart and Sun, 1990, p. 95) If α_i is the cosine of the i^{th} canonical angle between $\mathcal{U}$ and $\mathcal{W}$, then,

$$\rho_1(\mathcal{U}, \mathcal{W}) = \sqrt{2 \sum_{i=1}^{n} (1 - \alpha_i)}. \tag{3.3}$$

◁

Theorem 3.10 will be of great help when calculating distances between eigenspaces as it provides an exact formula, which allows us to avoid the minimisation over all unitary matrices. The α_i's in Theorem 3.10 can be calculated

via Definition 3.9. Alternatively, Stewart and Sun (1990, p. 45) state that the cosines of the canonical angles, i.e., the α_i's, between $\mathcal{U}, \mathcal{W} \in \mathbb{G}_{n,r}$ are equal to the singular values of $W^T U$. Since $W^T U$ is only a $r \times r$ matrix, it is preferable to obtain the α_i's from $W^T U$, instead of the $n \times n$ projector $UU^T(I - WW^T)$ from Definition 3.9.

Remark 3.11. Canonical angles are defined via the singular values of the projector $UU^T\left(I - WW^T\right)$ in Definition 3.9. Hence, it trivially follows from the definition of the two norm that $\rho_2(\mathcal{U}, \mathcal{W})$ is equal to the sine of the largest canonical angle between $\mathcal{U}$ and $\mathcal{W}$. ◁

In Lemma 3.12 we relate the two distance metrics $\rho_1(\mathcal{U}, \mathcal{W})$ and $\rho_2(\mathcal{U}, \mathcal{W})$. Our proof of Lemma 3.12 below agrees with the first steps of the proof of Lemma 5.1 in Lei and Rinaldo (2015, p. 232) except for the addition of the $\min(r, n-r)$ term instead of r, since we do not want to exclude cases where $r > n - r$.

Lemma 3.12. Let $U, W \in \mathbb{V}_{n,r}$. Then, there exists a $Q \in O(r)$ such that,

$$\|U - WQ\|_F \leq \sqrt{2\min(r, n-r)}\left\|UU^T(I - WW^T)\right\|_2.$$

Proof. From (3.2) and Proposition 2.2 in Vu and Lei (2013, p. 2911) we have that,

$$\frac{1}{2}\left(\inf_{R \in O(r)} \|U - WR\|_F\right)^2 \leq \|\sin\Theta(\mathcal{U}, \mathcal{W})\|_F^2,$$

where the sine function is applied elementwise to the matrix $\Theta(\mathcal{U}, \mathcal{W})$.

Therefore, if we define Q to be the matrix for which the infimum over $O(r)$ is attained, then,

$$\begin{aligned}
\|U - WQ\|_F &\leq \sqrt{2}\,\|\sin\Theta(\mathcal{U}, \mathcal{W})\|_F \\
&= \sqrt{2}\left\|UU^T(I - WW^T)\right\|_F && (3.4) \\
&\leq \sqrt{2\min(r, n-r)}\left\|UU^T(I - WW^T)\right\|_2. && (3.5)
\end{aligned}$$

Here (3.4) follows from Equation (2.6) in Vu and Lei (2013, p. 2911). (3.5) makes use of the fact that for any matrix X the relation $\|X\|_F \leq \sqrt{\operatorname{rank}(X)}\,\|X\|_2$ holds (Bernstein, 2009, p. 628) and by applying the

following simplification,

$$\begin{aligned}\operatorname{rank}\left(UU^T(I-WW^T)\right) &= \min\left(\operatorname{rank}\left(UU^T\right), \operatorname{rank}\left(I-WW^T\right)\right)\\ &= \min\left(r, n-r\right).\end{aligned}$$

□

In Chapter 4, Figures 4.5 and 4.8, we will observe attained distances in both metrics $\rho_1(\mathcal{U},\mathcal{W})$ (calculated via Theorem 3.10) and $\rho_2(\mathcal{U},\mathcal{W})$ and the relationship of their magnitudes from Lemma 3.12.

3.3 Application of the Davis–Kahan theorem to the comparison of the spaces spanned by the first r eigenvectors

We firstly formally introduce the Davis–Kahan theorem in Section 3.3.1 and then in Section 3.3.2, we discuss its application to the comparison of the spaces spanned by the first r eigenvectors of two matrices and show that given non-zero r^{th} eigengaps in both spectra, this comparison can always be made.

3.3.1 The Davis–Kahan theorem

The Davis–Kahan theorem (Davis and Kahan, 1970) in the form given in Bhatia et al. (2013) is reexpressed in our Theorem 3.13, where we also give the orthogonal projections in terms of the matrix eigenvectors.

Theorem 3.13. *Davis–Kahan Theorem* (Bhatia et al., 2013, pp. 211–2) Let $\Phi, \Psi \in \mathbb{R}^{n\times n}$ be symmetric matrices. Let S_1 be an interval $[a,b]$ and S_2 be the complement in $\mathbb{R}$ of the interval $(a-\delta, b+\delta)$, i.e., the intervals S_1 and S_2 lie a distance $\delta > 0$ apart. Let the columns of matrix U be orthonormal eigenvectors corresponding to the eigenvalues of Φ contained in S_1 and W have its columns made up of orthonormal eigenvectors corresponding to the eigenvalues of Ψ

not contained in S_2. Then, for every unitarily invariant norm, denoted $\|\cdot\|$,

$$\left\|UU^T(I - WW^T)\right\| \leq \frac{1}{\delta}\|\Phi - \Psi\|. \tag{3.6}$$

◁

Note that the interval definitions are in terms of the parameter triplet (a, b, δ). From Theorem 3.13 we learn that for the comparison of eigenvectors $[u_{j+1}, \ldots, u_{j+r}]$ of Φ to $[w_{j+1}, \ldots, w_{j+r}]$ of Ψ we require the following conditions on the corresponding eigenvalues.

Requirement 3.14. For the Davis–Kahan theorem to apply to the comparison of the eigenvector matrices $U_j, W_j \in \mathbb{V}_{n,r}$, we need to choose Davis–Kahan intervals, $S_1 = [a, b], S_2 = \mathbb{R}\backslash(a - \delta, b + \delta)$ for some interval separation $\delta > 0$, such that, either

$$\begin{aligned} &\phi_{j+1}, \ldots, \phi_{j+r} \in S_1; \quad \phi_1, \ldots, \phi_j, \phi_{j+r+1}, \ldots, \phi_n \notin S_1; \\ &\psi_{j+1}, \ldots, \psi_{j+r} \notin S_2; \quad \psi_1, \ldots, \psi_j, \psi_{j+r+1}, \ldots, \psi_n \in S_2 \end{aligned}$$

or, by swapping S_1 and S_2,

$$\begin{aligned} &\psi_{j+1}, \ldots, \psi_{j+r} \in S_1; \quad \psi_1, \ldots, \psi_j, \psi_{j+r+1}, \ldots, \psi_n \notin S_1; \\ &\phi_{j+1}, \ldots, \phi_{j+r} \notin S_2; \quad \phi_1, \ldots, \phi_j, \phi_{j+r+1}, \ldots, \phi_n \in S_2. \end{aligned}$$ ◁

For example for several versions of the spectral clustering algorithm the second to r^{th} eigenvectors of the representation matrices are used to detect communities in networks (Riolo and Newman, 2014) and therefore, $j = 1$ is the appropriate offset in the eigenvector sets discussed in Requirements 3.14.

The Davis–Kahan theorem in Theorem 3.13 is slightly more powerful than the one stated in von Luxburg (2007). von Luxburg (2007) compares the eigenvectors corresponding to eigenvalues of the two matrices which fall within an interval S_1. The formulation in Theorem 3.13 allows eigenvalues corresponding to the eigenvectors under comparison to extend beyond S_1 as long as they do not enter S_2; the two intervals are separated by an interval of length δ.

3.3.2 Bounding the spaces spanned by the first r eigenvectors

Using the Davis–Kahan theorem stated in Theorem 3.13, we can obtain a bound on the distance of the spaces spanned by the first r eigenvectors of two symmetric matrices under the following very mild assumption.

Assumption 3.15. Assume matrices $\Phi, \Psi \in \mathbb{R}^{n\times n}$ to both have a non-zero r^{th} eigengap, i.e., $\phi_1 \leq \ldots \leq \phi_n$ and $\psi_1 \leq \ldots \leq \psi_n$ are such that $\phi_r \neq \phi_{r+1}$ and $\psi_r \neq \psi_{r+1}$. ◁

In the comparison of the spaces spanned by the first r eigenvectors it is natural to assume a non-zero r^{th} eigengap. If the r^{th} and $(r+1)^{\text{th}}$ eigenvalue are equal then their corresponding eigenvectors are shared. Therefore, an arbitrary choice would have to be made which of the basis elements of the (at least two-dimensional) eigenspace corresponding to the r^{th} eigenvalue should be considered in the eigenvector comparison and no meaningful comparison could be made.

Remark 3.16. Large eigengaps in representation spectra are commonly used to inform the number of eigenvectors which should be used in graphical analysis (von Luxburg, 2007). Therefore, it is usual to have a large r^{th} eigengap in the spectra under comparison. ◁

In Theorem 3.17 we demonstrate that, given Assumption 3.15, we are always able to compare the spaces spanned by the first r eigenvectors of two symmetric matrices using the Davis–Kahan theorem.

Theorem 3.17. Consider the matrices holding the eigenvectors corresponding to the r smallest eigenvalues of each matrix, namely $U_0 = [u_1, \ldots, u_r] \in \mathbb{V}_{n,r}$ and $W_0 = [w_1, \ldots, w_r] \in \mathbb{V}_{n,r}$. Suppose Assumption 3.15 holds for the matrix spectra under comparison.

Then, there exists a $Q \in O(r)$ such that,

$$\|U_0 - W_0 Q\|_F \leq \sqrt{2\min(r, n-r)}\frac{\|\Phi - \Psi\|_2}{\max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r)}. \tag{3.7}$$

Proof. By Lemma 3.12 there exists a $Q \in O(r)$ such that,

$$\|U_0 - W_0 Q\|_F \leq \sqrt{2\min(r, n-r)}\left\|U_0 U_0^T (I - W_0 W_0^T)\right\|_2 . \tag{3.8}$$

Now we bound the term $\left\|U_0 U_0^T (I - W_0 W_0^T)\right\|_2$ in (3.8) using the Davis–Kahan Theorem in the form of Theorem 3.13; in order for the Theorem to apply to our comparison of the first r eigenvectors of two matrices, i.e., for U to equal U_0 and for W to equal W_0, we need to find intervals such that the conditions identified in Requirements 3.14 are satisfied for $j = 0$. From the ordering of the eigenvalues it follows that the conditions in Requirements 3.14 simplify to finding intervals S_1 and S_2 such that the first r eigenvalues of one of the matrices are contained in S_1 – while the $(r+1)^{\text{th}}$ eigenvalue is not contained in S_1 – and the last $n - r$ eigenvalues of the other matrix are contained in S_2 – while the r^{th} eigenvalue is not contained in S_2.

Either of the two spectra can be used in the definition of S_1 and therefore, there exist two valid interval choices with parameters (a, b, δ) given by,

$$a_1 = \min(\phi_1, \psi_1), \quad b_1 = \phi_r, \quad \delta_1 = \psi_{r+1} - \phi_r; \tag{3.9}$$

$$a_2 = \min(\phi_1, \psi_1), \quad b_2 = \psi_r, \quad \delta_2 = \phi_{r+1} - \psi_r. \tag{3.10}$$

Parameter choice (3.9) leads to interval $S_1 = [\min(\phi_1, \psi_1), \phi_r]$ containing the first r eigenvalues of Φ and not containing ϕ_{r+1} by Assumption 3.15. Interval $S_2 = \mathbb{R}\backslash(\min(\phi_1, \psi_1) - \psi_{r+1} + \phi_r, \psi_{r+1})$ includes the last $n - r$ eigenvalues of Ψ and excludes ψ_r by Assumption 3.15.

Similarly, parameter choice (3.10) leads to interval $S_1 = [\min(\phi_1, \psi_1), \psi_r]$ containing the first r eigenvalues of Ψ and not containing ψ_{r+1} by Assumption 3.15. Interval $S_2 = \mathbb{R}\backslash(\min(\phi_1, \psi_1) - \phi_{r+1} + \psi_r, \phi_{r+1})$ includes the last $n - r$ eigenvalues of Φ and excludes ϕ_r by Assumption 3.15.

The parameter choice which results in the greater separation between the two chosen intervals leads to a smaller bound and is hence to be preferred. By using $\delta = \max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r)$ we ensure that we are working with the interval choice with greater separation.

Assumption 3.15 implies that $\delta = \max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r) > 0$, i.e., at least

one of the two parameter choices (3.9) and (3.10) yield valid Davis–Kahan intervals S_1 and S_2. To see this, note that Assumption 3.15 guarantees that the following inequalities are strict: $\phi_r < \phi_{r+1}$ and $\psi_r < \psi_{r+1}$. Now we either have $\phi_{r+1} > \psi_r$, which directly implies $\delta > 0$ or $\phi_{r+1} \leq \psi_r$, which together with the strict ordering inequalities, implies, $\phi_r < \phi_{r+1} \leq \psi_r < \psi_{r+1}$ and hence, $\phi_r < \psi_{r+1}$, which means $\delta > 0$.

We have demonstrated that we are always able to choose intervals, with separation $\max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r) > 0$, satisfying Requirements 3.14. Hence, continuing from (3.8), by the Davis–Kahan theorem,

$$\sqrt{2\min(r, n-r)}\left\|U_0 U_0^T \ (I - W_0 W_0^T)\right\|_2 \leq \sqrt{2\min(r, n-r)}\frac{\|\Phi - \Psi\|_2}{\max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r)},$$

so that (3.7) is obtained.

□

Remark 3.18. If $r = n$, it follows from Lemma 3.5, that $\|U_0 - W_0 Q\|_F = 0$. For this choice of r, the bound proven in Theorem 3.17 also equals zero since $\min(n, n-r) = 0$. When $r = n$, we have the issue of ϕ_{r+1} and ψ_{r+1} not being defined; in Yu et al. (2015, p. 317) this is addressed by defining $\phi_0 = \psi_0 = -\infty$ and $\phi_{n+1} = \psi_{n+1} = \infty$. In our case, any choice leading to a non-zero denominator in the bound will be sufficient to obtain the desired result. We hence find that for the degenerate parameter choice, $r = n$, the bound in Theorem 3.17 is tight. ◁

In Theorem 3.17 we bound the distance of the spaces spanned by the first r eigenvectors of any two symmetric matrices, (where the two matrices are required to have a nonzero r^{th} eigengap), *using only their spectra*. The denominator of the bound in (3.7) contains a very interesting quantity, namely $\max(\phi_{r+1} - \psi_r, \psi_{r+1} - \phi_r)$, which can be interpreted as a *between–matrix eigengap*.

3.4 Extension to polynomial mappings and non-zero offsets

We begin this section by defining polynomial matrix transformations and explore how they act on matrix spectra and eigenvectors in Section 3.4.1. Then we discuss the necessary assumptions on the matrix spectra (Section 3.4.2) and transformation parameters (Section 3.4.3) to guarantee the presence of valid Davis–Kahan intervals where one of the two matrices under comparison is transformed by a polynomial transformation. We are then able to generalise the Davis–Kahan Theorem 3.13 to the comparison of any r consecutive eigenvectors by considering the polynomial transformation of one of the matrices under comparison (Section 3.4.4).

3.4.1 Polynomial matrix transformations

Consider polynomial matrix transformations and their action on matrix spectra and eigenvectors.

Definition 3.19. (Horn and Johnson, 1985, p. 36) define the evaluation of a polynomial $p(x) = c_l x^l + c_{l-1} x^{l-1} + \ldots + c_1 x + c_0$ at a matrix Φ as

$$p(\Phi) = c_l \Phi^l + c_{l-1} \Phi^{l-1} + \ldots + c_1 \Phi + c_0 I. \qquad \triangleleft$$

Theorem 3.20. (Horn and Johnson, 1985) Let $p(\cdot)$ be a given polynomial. If ϕ is an eigenvalue of $\Phi \in \mathbb{R}^{n \times n}$, while u is an associated eigenvector, then $p(\phi)$ is an eigenvalue of the matrix $p(\Phi)$ and u is an eigenvector of $p(\Phi)$ associated with $p(\phi)$. $\triangleleft$

We therefore, are able to transform the *largest* adjacency eigenvalues to be comparable to the *smallest* Laplacian eigenvalues without altering their corresponding eigenvectors and hence keeping our object of inference, i.e., $\rho_1(\mathcal{U}, \mathcal{W})$, unchanged.

3.4.2 Assumptions on the spectra

In the proof of Theorem 3.17, we showed that Assumption 3.15 is sufficient to guarantee the presence of valid Davis–Kahan intervals in the case of the comparison of the spaces spanned by the first r eigenvectors of two spectra which are ordered and indexed in the same way; this corresponds to a zero offset ($j = 0$). We shall now consider the offset parameter j to be potentially greater than zero, and require non-zero j^{th} and $(j+r)^{\text{th}}$ eigengaps of the spectra under comparison.

Assumption 3.21. Assume, for given $1 \leq r \leq n,\ 0 \leq j \leq n-r$, the eigenvalues $\phi_1 \leq \ldots \leq \phi_n$ of $\Phi \in \mathbb{R}^{n\times n}$ and the eigenvalues $\psi_1 \leq \ldots \leq \psi_n$ of $\Psi \in \mathbb{R}^{n\times n}$ to have a nonzero j^{th} and $(j+r)^{\text{th}}$ eigengap, i.e., $\phi_{j+1} - \phi_j > 0$, $\phi_{j+r+1} - \phi_{j+r} > 0$, $\psi_{j+1} - \psi_j > 0$ and $\psi_{j+r+1} - \psi_{j+r} > 0$. ◁

3.4.3 Constraints on the polynomial transformation

Since the polynomial transformation can change the ordering of the eigenvalues, i.e., $\phi_i < \phi_j$ does not imply that $p(\phi_i) < p(\phi_j)$, we need to place further assumptions on the transformation parameters to ensure the presence of valid Davis–Kahan intervals. Consider a given choice of Davis–Kahan intervals, $S_1 = [a, b], S_2 = \mathbb{R}\backslash(a-\delta, b+\delta)$ for some interval separation δ. We will make use of the following two sets to refer to parts of the transformed spectrum, for given $1 \leq r \leq n,\ 0 \leq j \leq n-r$,

$$\mathcal{A}_1 = \Big\{ i \in \{1,\ldots,n\}\backslash\{j+1,\ldots,j+r\} : p(\phi_i) > b \Big\} \tag{3.11}$$

and

$$\mathcal{A}_2 = \Big\{ i \in \{1,\ldots,n\}\backslash\{j+1,\ldots,j+r\} : p(\phi_i) < a \Big\}. \tag{3.12}$$

Constraint 3.22. For given $1 \leq r \leq n,\ 0 \leq j \leq n-r$, let the transformation parameters of $p(\cdot)$ be chosen such that

$$\mathcal{A}_1 \cup \mathcal{A}_2 = \{1,\ldots,n\}\backslash\{j+1,\ldots,j+r\}. \tag{3.13}$$

◁

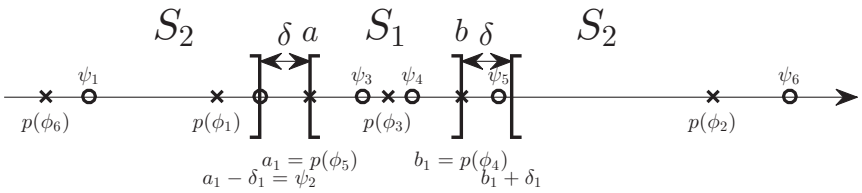

Figure 3.1: A possible eigenvalue configuration used to illustrate Davis–Kahan interval choice (3.14) and the necessary assumptions in order to guarantee its presence.

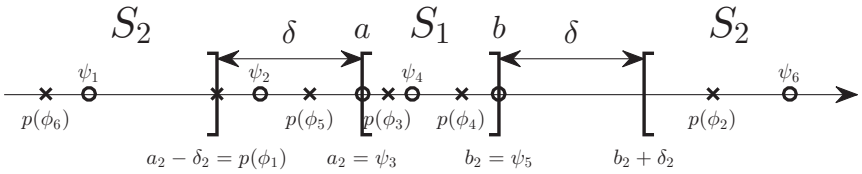

Figure 3.2: A possible eigenvalue configuration used to illustrate Davis–Kahan interval choice (3.15) and the necessary assumptions in order to guarantee its presence.

There are two possible Davis–Kahan interval choices. S_1 can either be defined based on eigenvalues in the spectrum of $p(\Phi)$ or Ψ. The two interval choices result in the following two interval parameter triplets (a, b, δ), for given $1 \leq r \leq n,\ 0 \leq j \leq n - r$,

$$
\begin{aligned}
a_1 &= \min_{i \in \{j+1,\ldots,j+r\}} p(\phi_i), \quad b_1 = \max_{i \in \{j+1,\ldots,j+r\}} p(\phi_i), \\
\delta_1 &= \min\left(\psi_{j+r+1} - b_1, a_1 - \psi_j\right); \qquad (3.14) \\
a_2 &= \psi_{j+1}, \quad b_2 = \psi_{j+r}, \\
\delta_2 &= \min\left[\min_{i \in \mathcal{A}_1} p(\phi_i) - b_2, a_2 - \max_{i \in \mathcal{A}_2} p(\phi_i)\right]. \qquad (3.15)
\end{aligned}
$$

Example 3.23. In Figure 3.1, we display an example of two spectra composed of 6 eigenvalues each, where the ψ_i's follow their natural ordering, but the transformed ϕ_i's have been given an arbitrary ordering and position for illustration purposes; in practice the choice of $p(\cdot)$ determines the ordering and position. We choose $r = 3, j = 2$, i.e., $U_2 = [u_3, u_4, u_5]$ and $W_2 = [w_3, w_4, w_5]$ are the eigenvector matrices to be compared. For this particular arbitrary set of eigenvalues, interval choice (3.14), is displayed in Figure 3.1, with $a_1 = p(\phi_5), b_1 = p(\phi_4)$ and $\delta_1 = p(\phi_5) - \psi_2$. In Figure 3.2 we illustrate interval choice (3.15), where $a_2 = \psi_3, b_2 = \psi_5$ and $\delta_2 = \psi_3 - p(\phi_1)$.

Interval choice (3.14) is less favourable than (3.15) since we can clearly observe that $\delta_2 > \delta_1$ by comparing Figures 3.2 and 3.1. The purpose of Figure 3.1 is mainly to demonstrate that both interval choices (3.14) and (3.15) are valid for the displayed set of eigenvalues. For simplicity we will proceed to discuss assumptions only in the context of Figure 3.2 and interval choice (3.15).

It can be deduced from Figure 3.2 that the eigenvalues in the illustration satisfy Assumption 3.21, i.e., $\phi_3 - \phi_2 > 0$, $\phi_6 - \phi_5 > 0$, $\psi_3 - \psi_2 > 0$ and $\psi_6 - \psi_5 > 0$. (For example, note that $p(\phi_3) \neq p(\phi_2) \implies \phi_3 > \phi_2$, because of the eigenvalue ordering, so $\phi_3 - \phi_2 > 0$.)

Without nonzero j^{th} and $(j+r)^{\text{th}}$ eigengaps for ϕ_i and ψ_i, no valid interval choice could be made, since any violation would immediately lead to a violation of the conditions (Requirements 3.14) for the presence of valid Davis–Kahan intervals.

For interval choice (3.15), Constraints 3.22 say that $\{p(\phi_1), p(\phi_2), p(\phi_6)\}$ must either be larger than $b_2 = \psi_5$ or smaller than $a_2 = \psi_3$. For the chosen comparison and transformation $p(\cdot)$, we find $\mathcal{A}_1 = \{2\}$ and $\mathcal{A}_2 = \{1, 6\}$ and hence, Constraints 3.22 is satisfied. ◁

As can be seen by observing ψ_5 in Figure 3.1 and $p(\phi_5)$ in Figure 3.2, for interval choices (3.14) and (3.15), the spectrum not explicitly used in the definition of S_1 can extend beyond the boundaries of S_1 on both sides. In Section 3.3 this issue could be avoided by defining the left boundary of S_1 to be $\min(\phi_1, \psi_1)$; here we introduce Constraints 3.24.

Constraint 3.24. A) For interval choice (3.14), let the transformation parameters of $p(\cdot)$ be chosen such that, for given $1 \leq r \leq n,\ 0 \leq j \leq n - r$,

$$\delta_1 > 0, \tag{3.16}$$

$$a_1 - \psi_{j+1} < \delta_1, \tag{3.17}$$

$$\psi_{j+r} - b_1 < \delta_1. \tag{3.18}$$

B) For interval choice (3.15), let the transformation parameters of $p(\cdot)$ be

chosen such that, for given $1 \leq r \leq n,\ 0 \leq j \leq n-r$,

$$\delta_2 > 0, \tag{3.19}$$

$$a_2 - \min_{i \in \{j+1,\ldots,j+r\}} p(\phi_i) < \delta_2, \tag{3.20}$$

$$\max_{i \in \{j+1,\ldots,j+r\}} p(\phi_i) - b_2 < \delta_2. \tag{3.21}$$

$\lhd$

In essence, in (3.17), (3.18), (3.20) and (3.21) we require that, if any of the eigenvalues with indices $(j+1), \ldots, (j+r)$ of the spectrum not explicitly used in the definition of S_1, fall outside of S_1, then they must not be further than δ_i away from the boundary of S_1, where $i \in \{1, 2\}$ depending on the interval choice we are considering. For interval choice (3.15) and the eigenvalues displayed in Figure 3.2, Constraints 3.24B) require $\delta_2 > 0$, $\psi_3 - p(\phi_5) < \psi_3 - p(\phi_1)$ and $p(\phi_4) - \psi_5 < \psi_3 - p(\phi_1)$.

3.4.4 The main theorem

Now we are able to extend the Davis–Kahan theorem to be applicable to any two symmetric matrices satisfying Assumption 3.21.

Theorem 3.25. Let $\Phi, \Psi \in \mathbb{R}^{n \times n}$ be symmetric matrices with eigenvalues $\phi_1 \leq \phi_2 \leq \ldots \leq \phi_n$ and $\psi_1 \leq \psi_2 \leq \ldots \leq \psi_n$ and corresponding eigenvectors $\{u_1, u_2, \ldots, u_n\}$ and $\{w_1, w_2, \ldots, w_n\}$, respectively. Let the matrices holding the eigenvectors corresponding to r consecutive eigenvalues of each matrix be denoted by $U_j = [u_{j+1}, \ldots, u_{j+r}] \in \mathbb{V}_{n,r}$ and $W_j = [w_{j+1}, \ldots, w_{j+r}] \in \mathbb{V}_{n,r}$, where $r, j \in \mathbb{N}$ with $1 \leq r \leq n,\ 0 \leq j \leq n-r$. Further, let Assumption 3.21 hold for the spectra of Φ and Ψ.

Let $p(\cdot)$ be a polynomial transformation satisfying Constraints 3.22 and 3.24A) ($\delta_i = \delta_1$, (3.14)) or Constraints 3.22 and 3.24B) ($\delta_i = \delta_2$, (3.15)) then, for every unitarily invariant norm, denoted $\|\cdot\|$,

$$\left\| U_j U_j^T (I - W_j W_j^T) \right\| \leq \frac{\|p(\Phi) - \Psi\|}{\delta_i}. \tag{3.22}$$

Proof. We bound the term $\left\|U_j U_j^T (I - W_j W_j^T)\right\|$ using the Davis–Kahan Theorem 3.13. From Theorem 3.20 we know that Φ and $p(\Phi)$ share eigenvectors and therefore, when bounding$\left\|U_j U_j^T (I - W_j W_j^T)\right\|$ we are able to apply the Davis–Kahan theorem to the spectrum of $p(\Phi)$ instead of the spectrum of Φ. For the Davis–Kahan theorem to apply we need to show that the interval requirements laid out in Requirements 3.14 are satisfied by interval choices (3.14) and (3.15). We begin by checking the four conditions laid out in Requirements 3.14 for interval choice (3.14):

Firstly, the condition $p(\phi_{j+1}), \ldots, p(\phi_{j+r}) \in S_1$ is always guaranteed by the definition of $S_1 = \left[\min_{i\in\{j+1,\ldots,j+r\}} p(\phi_i), \max_{i\in\{j+1,\ldots,j+r\}} p(\phi_i)\right]$ in interval choice (3.14).

Secondly, for interval choice (3.14), Constraints 3.22 imply that for all $l \in \{1, \ldots, n\}\backslash\{j+1, \ldots, j+r\}$ we have that either

$$p(\phi_l) < \min_{i\in\{j+1,\ldots,j+r\}} p(\phi_i) \qquad \text{or} \qquad p(\phi_l) > \max_{i\in\{j+1,\ldots,j+r\}} p(\phi_i). \tag{3.23}$$

This implies that

$$\max_{i\in\mathcal{A}_2} p(\phi_i) < \min_{i\in\{j+1,\ldots,j+r\}} p(\phi_i) = a_1$$

and

$$b_1 = \max_{i\in\{j+1,\ldots,j+r\}} p(\phi_i) < \min_{i\in\mathcal{A}_1} p(\phi_i).$$

Hence, Constraints 3.22 guarantee that $p(\phi_1), \ldots, p(\phi_j), p(\phi_{j+r+1}), \ldots,$ $p(\phi_n) \notin S_1$.

Next, the statements in Constraints 3.24A) can be rearranged to $\delta_1 > 0$, $\psi_{j+r} < \delta_1 + b_1$ and $a_1 - \delta_1 < \psi_{j+1}$. Since, $S_2 = \mathbb{R}\backslash(a_1 - \delta_1, b_1 + \delta_1)$, we find that $\psi_{j+1}, \ldots, \psi_{j+r} \notin S_2$ follows immediately from Constraints 3.24A).

Finally, from the definition of $\delta_1 = \min(\psi_{j+r+1} - b_1, a_1 - \psi_j)$ in interval choice (3.14) it follows that $\delta_1 \leq a_1 - \psi_j$, hence, $\psi_j \leq a_1 - \delta_1$ and similarly, $\delta_1 \leq \psi_{j+r+1} - b_1$, hence, $b_1 + \delta_1 \leq \psi_{j+r+1}$. Therefore, $\psi_1, \ldots, \psi_j, \psi_{j+r+1}, \ldots, \psi_n \in S_2$ follows immediately from the parameter choice (3.14). Therefore, under Constraints 3.22 and 3.24A) on the transformation parameters, interval choice (3.14) satisfies the conditions laid out in Requirements 3.14.

In the case of interval choice (3.15), Assumption 3.21, and Constraints 3.22 and 3.24B) suffice to guarantee Requirements 3.14:

Firstly, the condition $\psi_{j+1}, \ldots, \psi_{j+r} \in S_1$ is always guaranteed by the definition of $S_1 = [\psi_{j+1}, \psi_{j+r}]$ in interval choice (3.15).

Secondly, the eigengaps in the spectrum of Ψ implied by Assumption 3.21, specifically $\psi_{j+1} - \psi_j > 0$ and $\psi_{j+r+1} - \psi_{j+r} > 0$, guarantee that $\psi_1, \ldots, \psi_j$, $\psi_{j+r+1}, \ldots, \psi_n \notin S_1$.

Next, the statements in Constraints 3.24B) can be arranged to $\delta_2 > 0$, $\min_{i \in \{j+1,\ldots,j+r\}} p(\phi_i) > a_2 - \delta_2$ and $\max_{i \in \{j+1,\ldots,j+r\}} p(\phi_i) < b_2 + \delta_2$. Since, $S_2 = \mathbb{R} \backslash (a_2 - \delta_2, b_2 + \delta_2)$, we find that $p(\phi_{j+1}), \ldots, p(\phi_{j+r}) \notin S_2$ follows immediately from Constraints 3.24B).

Finally, from the definition of δ_2 in (3.15),

$$\delta_2 = \min\left(\min_{i \in \mathcal{A}_1} p(\phi_i) - b_2, a_2 - \max_{i \in \mathcal{A}_2} p(\phi_i)\right) \leq \min_{i \in \mathcal{A}_1} p(\phi_i) - b_2$$
$$\Rightarrow b_2 + \delta_2 \leq \min_{i \in \mathcal{A}_1} p(\phi_i), \tag{3.24}$$

and similarly,

$$\delta_2 = \min\left(\min_{i \in \mathcal{A}_1} p(\phi_i) - b_2, a_2 - \max_{i \in \mathcal{A}_2} p(\phi_i)\right) \leq a_2 - \max_{i \in \mathcal{A}_2} p(\phi_i)$$
$$\Rightarrow a_2 - \delta_2 \geq \max_{i \in \mathcal{A}_2} p(\phi_i). \tag{3.25}$$

Furthermore, Constraints 3.22 imply that for all $l \in \{1, \ldots, n\} \backslash \{j+1, \ldots, j+r\}$ either $p(\phi_l) \geq \min_{i \in \mathcal{A}_1} p(\phi_i)$ or $p(\phi_l) \leq \max_{i \in \mathcal{A}_2} p(\phi_i)$. This together with Equations (3.24) and (3.25) implies that for all $l \in \{1, \ldots, n\} \backslash \{j+1, \ldots, j+r\}$ either $p(\phi_l) \geq b_2 + \delta_2$ or $p(\phi_l) \leq a_2 - \delta_2$. Therefore, $p(\phi_1), \ldots, p(\phi_j)$, $p(\phi_{j+r+1}), \ldots, p(\phi_n) \in S_2$ follows from the definition of interval choice (3.15) together with Constraints 3.22.

By Constraints 3.24, $\delta_1 > 0$ and $\delta_2 > 0$. Therefore, all requirements of the Davis–Kahan theorem are satisfied. Hence, (3.22) holds, where, for (3.14), $\delta_i = \delta_1$, and for (3.15), $\delta_i = \delta_2$.

□

We extend our result in Theorem 3.25 to bound $\|U_j - W_j Q\|_F$, i.e., ρ_1, in Corollary 3.26. However, Theorem 3.25 is equally valid in its own right, if one prefers to work with a bound on the distance of the subspaces in the metric $\left\|U_j U_j^T (I - W_j W_j^T)\right\|_2$, i.e., ρ_2.

Corollary 3.26. Let $\Phi, \Psi, U_j, W_j, \delta_i$ and $p(\cdot)$ satisfy the conditions given in Theorem 3.25, then, there exists a $Q \in O(r)$ such that

$$\|U_j - W_j Q\|_F \leq \sqrt{2\min(r, n-r)} \frac{\|p(\Phi) - \Psi\|_2}{\delta_i}. \tag{3.26}$$

Proof. This follows by applying Lemma 3.12 to Theorem 3.25, noting that the matrix two-norm or spectral norm is unitarily invariant as is required in Theorem 3.25. □

Remark 3.27. For the parameter choice $j = 0$ and $p(x) = x$ for $x \in \mathbb{R}$, (3.22) yields the result in (3.7), where $\max_{i \in \mathcal{A}_1} p(\phi_i), \psi_0$ are defined to equal $-\infty$. Hence, Corollary 3.26 includes Theorem 3.17 as a special case, i.e., the standard Davis–Kahan theorem is a special case of our extended version, and consequently our bounds are guaranteed to be at least as tight as those given by the standard Davis–Kahan theorem. ◁

Analogously to Theorem 3.17, the bound in (3.26) *depends only on the spectra* of the two matrices.

Within the restrictions imposed by Constraints 3.22 and 3.24, we aim to choose the polynomial transformation such that it minimises the bound on the eigenvector difference. We therefore have to chose the polynomial transformation to minimise the numerator of our bound (3.22), the spectral norm of the matrix difference $\|p(\Phi) - \Psi\|_2$, and to maximise the bound's denominator, the maximal eigenvalue interval separation δ_1 or δ_2.

In practice, we recommend choosing transformations to separately minimise the two subproblems, posed by the two possible denominators, and then to work with the smaller of the two bounds in order to achieve an overall minimal bound.

The polynomial transformation of a matrix can also be interpreted as finding the ideal matrix in the sense of producing a minimal bound on the difference of the spaces spanned by the eigenvectors, while preserving the eigenvectors of the untransformed matrix.

3.5 Affine transforms

In this section we consider affine matrix transformations $f(\Phi) = c_1\Phi + c_0 I$, $(c_1, c_0 \in \mathbb{R})$, which are special cases of the polynomial matrix transformations $p(\Phi) = c_l\Phi^l + c_{l-1}\Phi^{l-1} + \ldots + c_1\Phi + c_0 I$ we worked with in Section 3.4.

We consider the cases $c_1 = 0, c_1 > 0$ and $c_1 < 0$ separately. For $c_1 = 0$ all information in the spectrum is lost, i.e., $f(\phi_i) = c_1\phi_i + c_0 = c_0$, and we cannot find intervals S_1 and S_2 such that the transformation parameters satisfy Constraints 3.22, so there exist no valid Davis–Kahan eigenvalue intervals. For $c_1 > 0$ the ordering of the eigenvalues is preserved in the transformed spectrum, while for $c_1 < 0$ the ordering of eigenvalues is reversed in the transformed spectrum. We treat these latter two cases separately throughout the rest of this section.

For affine transformations, the quantities for the $j^{\text{th}}, (j+1)^{\text{th}}, (j+r)^{\text{th}}$ and $(j+r+1)^{\text{th}}$ eigenvalues in the transformed spectrum used in the interval definitions (3.14) and (3.15) are displayed in Table 3.1. Note that all the quantities in Constraints 3.24 can be found simply by plugging in values from Table 3.1. We see for example, for an affine transformation satisfying Constraints 3.22 and 3.24 with $c_1 > 0$, interval choice (3.14) corresponds to the following Davis–Kahan intervals:

$$\begin{aligned} S_1 &= [f(\phi_{j+1}), f(\phi_{j+r})], \\ \delta_1 &= \min(\psi_{j+r+1} - f(\phi_{j+r}), f(\phi_{j+1}) - \psi_j), \\ S_2 &= \mathbb{R}\backslash(f(\phi_{j+1}) - \delta_1, f(\phi_{j+r}) + \delta_1). \end{aligned}$$

Note that the two interval choices (3.14) and (3.15) together with the distinction between affine transformations for $c_1 > 0$ and $c_1 < 0$ produce four different possible values for the Davis–Kahan interval separation δ, which are,

Table 3.1: Explicit form of the $j^{\text{th}}, (j+1)^{\text{th}}, (j+r)^{\text{th}}$ and $(j+r+1)^{\text{th}}$ eigenvalues in the spectrum of $f(\Phi)$, where $f(\phi_i) = c_1\phi_i + c_0$ is an affine transformation.

	$c > 0$	$c < 0$
$\max_{i \in \mathcal{A}_2} p(\phi_i)$	$f(\phi_j)$	$f(\phi_{j+r+1})$
$\min_{i \in \{j+1,\ldots,j+r\}} p(\phi_i)$	$f(\phi_{j+1})$	$f(\phi_{j+r})$
$\max_{i \in \{j+1,\ldots,j+r\}} p(\phi_i)$	$f(\phi_{j+r})$	$f(\phi_{j+1})$
$\min_{i \in \mathcal{A}_1} p(\phi_i)$	$f(\phi_{j+r+1})$	$f(\phi_j)$

from (3.14)

$$\delta_{1,+} = \min\left(\psi_{j+r+1} - f(\phi_{j+r}), f(\phi_{j+1}) - \psi_j\right), \tag{3.27}$$

$$\delta_{1,-} = \min\left(\psi_{j+r+1} - f(\phi_{j+1}), f(\phi_{j+r}) - \psi_j\right), \tag{3.28}$$

and, from (3.15),

$$\delta_{2,+} = \min\left(f(\phi_{j+r+1}) - \psi_{j+r}, \psi_{j+1} - f(\phi_j)\right), \tag{3.29}$$

$$\delta_{2,-} = \min\left(f(\phi_j) - \psi_{j+r}, \psi_{j+1} - f(\phi_{j+r+1})\right). \tag{3.30}$$

3.6 Proof of concept example

As a proof of concept we will apply Corollary 3.26 to the comparison of the spaces spanned by eigenvectors corresponding to the 3 smallest eigenvalues of the unnormalised and normalised Laplacian matrices of d-regular graphs.

In Figure 3.3 we visualise the attained distances, measured in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$ calculated using (3.3), of the spaces spanned by the three eigenvectors corresponding to the 3 smallest eigenvalues of L and L_{sym} for 25 randomly sampled 30-regular graphs with 300 nodes. In addition, we show bound values on $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$ produced via the standard Davis–Kahan theorem (right hand side of (3.7)) and our theorem with extended applicability (right hand side of (3.26)) using the simplest type of polynomial transformation, namely an affine transformation.

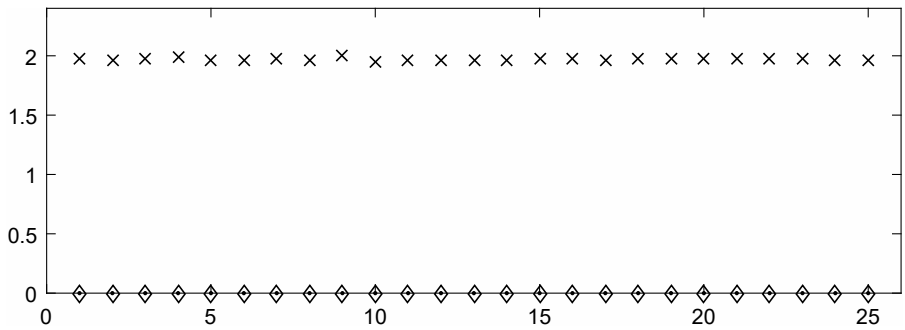

Figure 3.3: Dots represent the attained eigenvector distances, $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$, diamonds represent bound values calculated from Corollary 3.26, using an affine matrix transformation, and the 'x' symbols represent bound values calculated from Theorem 3.17. The x-axis displays the indices of the randomly drawn 30-regular graphs, while the y-axis shows magnitudes of the displayed values.

In the case of the graph Laplacian matrices we see that for d-regular graphs, the standard Davis–Kahan bound is non-zero in general. Contrariwise, for d-regular graphs, $L_{sym} = d^{-1}L$, i.e., there exists a choice of affine transformation parameters, namely $c_1 = d^{-1}, c_0 = 0$, which maps L exactly to L_{sym}, such that our extended Davis–Kahan bound is identically zero.

The algorithm, which will be discussed in Chapter 4, finds optimal transformation parameters in the case of affine transforms. The result shown in Figure 3.3 was obtained *automatically* using this algorithm, without any additional information. The equivalence of the spaces spanned by the first r eigenvectors was correctly identified by our extended Davis–Kahan bounds methodology.

In the above example both the standard Davis–Kahan theorem and our extended applicability version were able to be applied in comparing L and L_{sym}. In other cases, such as when comparing A and L, the standard Davis–Kahan theorem cannot even be applied since the required Davis–Kahan intervals of Requirements 3.14 do not exist; however, in this case our extended applicability version of the Davis–Kahan theorem can still be used to determine upper bounds. In fact, for d-regular graphs, (for which, e.g., $L_{sym} = I - d^{-1}A$), we find that our extended applicability bound identifies the equality of the subspaces spanned by the eigenvectors corresponding to the largest eigenvalues of A, and the spaces spanned by the eigenvectors corresponding to the smallest eigenvalues of the Laplacians L and L_{sym}.

4

Extending the Davis–Kahan Theorem for Comparing Eigenvectors of Two Symmetric Matrices: Computation and Applications

An arXiv preprint of the work in this chapter can be accessed at Lutzeyer and Walden (2019b).

4.1 Introduction

In Chapter 3 we introduced the extended Davis–Kahan theorem for comparing two sets of consecutive and corresponding eigenvectors from any two symmetric matrices.

The extension incorporated a polynomial transform of one of the matrices which allows a relaxation and utilisation of the eigenvalue structure imposed by the standard Davis–Kahan theorem. As a result the bounds determined by the extended theorem are always at least as tight as those from the standard

Davis–Kahan theorem.

Chapter 3 concentrated on the mathematics of the proposed approach. In this second part of the work, we turn our attention to computational issues, and also give some significant examples of applications of our extended Davis–Kahan theorem. The computational aspects are certainly challenging, and in this chapter we only give a full discussion for the case of affine transformations; however, as exemplified by the applications, the affine transformation can be very beneficial.

We classify the problem of finding optimal affine transformation parameters for the extended Davis–Kahan bound as a fractional program. Fractional programming problems can be transformed to be convex optimisation problems (see Sections 4.4.4 and 4.4.5) with a unique global solution as stated in Remark 4.6. Fractional programing seems to be a less-well known class of optimisation problems. For the history and recent advances in fractional programming see Frenk and Schaible (2008). An excellent overview of the solution and implementation of concave-convex fractional programs, the subclass of fractional programs which applies to our optimisation problem, is given in Isheden et al. (2012). In the wireless communication literature fractional programming approaches have recently found much use (Cheung and Hanzo, 2016; De Maio et al., 2011; Shen and Yu, 2018; Ye et al., 2018).

Once we have chosen a solution method for our optimisation problem we produce bounds in a range of different applications. In some of these applications the standard Davis–Kahan Theorem 3.13 does not apply since the eigenvalue structure is not satisfied, while in others bounds are available via Theorem 3.13 and our bound produced via Theorem 3.25 attains lower bound values. The matrices considered are graph representation matrices corresponding to graphs in a SBM and covariance matrices in a spiked covariance model, where the principal component analysis algorithm is well motivated.

The remainder of this chapter is structured as follows, in Section 4.2 we summarise our extended Davis–Kahan theorem proved in Chapter 3 using general polynomial transformations and then specialise to affine transformations. In Section 4.3 we present the problem of finding transformation parameters resulting in optimal bound values in our extended affine Davis–Kahan theorem. Furthermore, we prove that a trivial bound on the distance of the spaces spanned

by the compared eigenvectors is always outperformed for comparisons of eigenvectors corresponding to either the largest or smallest eigenvalues. In Section 4.4 we introduce concave-convex fractional programs, demonstrate that the problem from Section 4.3 can be brought into a fractional programming form and discuss two solution approaches, the Charnes–Cooper transformation and Dinkelbach's algorithm. In Section 4.5 we discuss the complications which arise when our solution approach is extended from affine to higher order polynomial transforms. Section 4.6 presents three significant examples in which our bound is computed. These examples are the comparison of the spaces spanned by the eigenvectors of the representation matrices in 4.6.4, of the representation matrices to their corresponding generating matrices in the SBM in Section 4.6.5 and of the sample and population covariance matrices in a spiked covariance model in Section 4.6.6.

4.2 Problem summary

4.2.1 Bounds

Let $\mathbb{V}_{n,r}$ denote the Stiefel manifold of $n \times r$ matrices with orthonormal columns. Let $\Phi, \Psi \in \mathbb{R}^{n \times n}$ be symmetric matrices with eigenvalues $\phi_1 \leq \phi_2 \leq \ldots \leq \phi_n$ and $\psi_1 \leq \psi_2 \leq \ldots \leq \psi_n$ and corresponding eigenvectors $\{u_1, u_2, \ldots, u_n\}$ and $\{w_1, w_2, \ldots, w_n\}$, respectively. For $r, j \in \mathbb{N}$ with $1 \leq r \leq n,\ 0 \leq j \leq n - r$, let the matrices holding the eigenvectors corresponding to r consecutive eigenvalues of each matrix be denoted by $U_j = [u_{j+1}, \ldots, u_{j+r}] \in \mathbb{V}_{n,r}$ and $W_j = [w_{j+1}, \ldots, w_{j+r}] \in \mathbb{V}_{n,r}$, the columns of which span the spaces $\mathcal{U}_j$ and $\mathcal{W}_j$, respectively. Under the stated conditions, it was shown in Chapter 3 that there exists a $Q \in O(r)$ such that

$$\begin{aligned} \|U_j - W_j Q\|_F &\leq \sqrt{2\min(r, n-r)}\, \left\|U_j U_j^T (I - W_j W_j^T)\right\|_2 \\ &\leq \sqrt{2\min(r, n-r)}\frac{\|p(\Phi) - \Psi\|_2}{\delta_i}, \end{aligned} \tag{4.1}$$

where $O(r)$ is the group of $r \times r$ orthogonal matrices, $p(\cdot)$ is a polynomial matrix transformation and $\delta_i,\ i \in \{1, 2\}$ are different values of the Davis–Kahan interval separation parameter corresponding to the two different Davis–Kahan

interval choices. The usual or standard Davis–Kahan bounds follow by taking $p(\Phi) = \Phi$ in the second inequality. In Chapter 3 we saw that the first norm is directly related to the metric $\rho_1(\mathcal{U}_j, \mathcal{W}_j) = \inf_{R\in O(r)} \|U_j - W_j R\|_F$, and the second relates to the metric $\rho_2(\mathcal{U}_j, \mathcal{W}_j) = \left\|U_j U_j^T \left(I - W_j W_j^T\right)\right\|_2$.

4.2.2 Affine transformations

In this chapter we focus on affine matrix transformations $f(\Phi) = c_1\Phi + c_0 I,\ (c_1, c_0 \in \mathbb{R})$ of one of the matrices under comparison (arbitrarily, Φ). So we set $p(\cdot) = f(\cdot)$ and for $1 \le r \le n,\ 0 \le j \le n - r$, from Chapter 3, the two Davis–Kahan interval choices take the form $S_1 = [a, b], S_2 = \mathbb{R}\backslash(a - \delta, b + \delta)$, with

$$
\begin{aligned}
a_1 &= \min_{i\in\{j+1,\ldots,j+r\}} f(\phi_i), \quad b_1 = \max_{i\in\{j+1,\ldots,j+r\}} f(\phi_i), \\
\delta_1 &= \min\left(\psi_{j+r+1} - b_1, a_1 - \psi_j\right); \qquad (4.2) \\
a_2 &= \psi_{j+1}, \quad b_2 = \psi_{j+r}, \\
\delta_2 &= \min\left[\min_{i\in\mathcal{A}_1} f(\phi_i) - b_2, a_2 - \max_{i\in\mathcal{A}_2} f(\phi_i)\right]. \qquad (4.3)
\end{aligned}
$$

The index sets $\mathcal{A}_1$ and $\mathcal{A}_2$ are given in Equations (3.11) and (3.12) in Chapter 3.

A main purpose of this chapter is to solve the problem of optimising the bound on the right-side of (4.1) when $p(\cdot) = f(\cdot)$, i.e., minimise

$$
\sqrt{2\min(r, n-r)}\frac{\|f(\Phi) - \Psi\|_2}{\delta_i}, \quad i \in \{1, 2\}, \qquad (4.4)
$$

subject to the associated constraints given in Chapter 3. For an affine transformation, Constraints 3.24 of Chapter 3 must be applied (Constraints 3.22 are subsumed). Constraints 3.24 take the form:

A) In the case of interval choice (4.2), let the transformation parameters of

$f(\cdot)$ be chosen such that, for given $1 \leq r \leq n,\ 0 \leq j \leq n-r$,

$$\begin{aligned} \delta_1 &> 0, \\ a_1 - \psi_{j+1} &< \delta_1, \\ \psi_{j+r} - b_1 &< \delta_1. \end{aligned}$$

B) For interval choice (4.3), let the transformation parameters of $f(\cdot)$ be chosen such that, for given $1 \leq r \leq n,\ 0 \leq j \leq n-r$,

$$\begin{aligned} \delta_2 &> 0, \\ a_2 - \min_{i \in \{j+1,\ldots,j+r\}} f(\phi_i) &< \delta_2, \\ \max_{i \in \{j+1,\ldots,j+r\}} f(\phi_i) - b_2 &< \delta_2. \end{aligned}$$

Also as pointed out in Section 3.5 in Chapter 3, in addition to the two Davis–Kahan interval choices which give δ_1, δ_2 in (4.4), there are also the possibilities $c_1 > 0$ and $c_1 < 0$ in the affine transform. So there are four different possible values for the Davis–Kahan interval separation δ, which are,

$$\delta_{1,+} = \min\left(\psi_{j+r+1} - c_1\phi_{j+r} - c_0, c_1\phi_{j+1} + c_0 - \psi_j\right), \tag{4.5}$$

$$\delta_{1,-} = \min\left(\psi_{j+r+1} - c_1\phi_{j+1} - c_0, c_1\phi_{j+r} + c_0 - \psi_j\right), \tag{4.6}$$

$$\delta_{2,+} = \min\left(c_1\phi_{j+r+1} + c_0 - \psi_{j+r}, \psi_{j+1} - c_1\phi_j - c_0\right), \tag{4.7}$$

$$\delta_{2,-} = \min\left(c_1\phi_j + c_0 - \psi_{j+r}, \psi_{j+1} - c_1\phi_{j+r+1} - c_0\right). \tag{4.8}$$

4.3 The bound as a numerical optimisation problem

In this section we frame the optimisation of the affine bounds (4.4) as constrained optimisation problems over the transformation parameters c_1, c_0. Solving the optimisation problems derived in this section results in the minimal bound under affine transformations on the distance of the spaces spanned by two sets of eigenvectors.

Davis–Kahan interval choices (4.2) and (4.3) impose a certain structure in-

volving different parameters. Based on the eigengaps of the two spectra one of these structures results in a more favourable interval configuration as was shown in Example 3.23. Therefore, interval choices (4.2) and (4.3) pose separate optimisation problems. Furthermore, the distinction between affine parametrisations with $c_1 > 0$ and $c_1 < 0$ results in different cost functions and constraints. These distinctions manifest as the four possible denominators in (4.4), namely (4.5)-(4.8). Hence, we have four optimisation subproblems, which need to be solved in order to obtain the *overall* optimal Davis–Kahan bound.

We study the solution of one in detail, and the rest follow analogously. In (4.9) we show the optimisation subproblem for interval choice (4.2) for affine transformations with $c_1 > 0$, i.e., where $\delta = \delta_{1,+}$ in (4.5). Then $a_1 = \min_{i\in\{j+1,\ldots,j+r\}} f(\phi_i) = f(\phi_{j+1}) = c_1\phi_{j+1} + c_0$ and $b_1 = \max_{i\in\{j+1,\ldots,j+r\}} f(\phi_i) = f(\phi_{j+r}) = c_1\phi_{j+r} + c_0$. (Here $a_1 = f(\phi_{j+1})$ and $b_1 = f(\phi_{j+r})$ are given in Table 3.1 of Chapter 3, but follow from the preservation of ordering of the eigenvalues for an affine transform with $c_1 > 0$). For interval choice (4.2) we need Constraints 3.24 A), which we apply in (4.9) via the last three rows of the "subject to" statement. From Corollary 3.26 of Chapter 3 the objective function to be minimised is $\|f(\Phi) - \Psi\|_2/\delta_{1,+}$. Hence, the first subproblem takes the form

$$\begin{aligned} \min_{c_1,c_0} \quad & \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta_{1,+}}, \\ \text{subject to} \quad & c_1 > 0, \\ & \delta_{1,+} > 0, \\ & \delta_{1,+} > \psi_{j+r} - c_1\phi_{j+r} - c_0, \\ & \delta_{1,+} > c_1\phi_{j+1} + c_0 - \psi_{j+1}. \end{aligned} \tag{4.9}$$

The remaining three subproblems follow a similar structure as (4.9), where δ equals (4.6), (4.7) or (4.8) and the values of the transformed spectrum are correspondingly taken from Table 3.1 of Chapter 3.

Remark 4.1. In the objective function in (4.9) we omit the constant $\sqrt{2\min(r, n-r)}$ present in (4.1), since it is inconsequential to the minimisation. However, solutions of (4.9) have to be multiplied by $\sqrt{2\min(r, n-r)}$ in

order to obtain valid bounds. ◁

In Proposition 4.2 we obtain a trivial upper bound on $\|U - WQ\|_F$ and demonstrate that, for comparisons of the first or the last r eigenvectors, solutions of (4.9) always approximate or improve upon this bound. Without considering the matrix transform, no such guarantees could be given.

Proposition 4.2.

1. Let $U_j, W_j \in \mathbb{V}_{n,r}$. Then, for any $j \geq 0$, there exists $Q \in O(r)$ such that,

$$\|U_j - W_j Q\|_F \leq \sqrt{2\min(r, n-r)}. \tag{4.10}$$

2. When comparing spaces spanned by the r eigenvectors corresponding to either the r largest or r smallest eigenvalues, the bound produced from (4.9) always approximates or improves upon (4.10).

Proof. We begin by proving part 1 of the proposition. From Lemma 3.12 of Chapter 3,

$$\|U_j - W_j Q\|_F \leq \sqrt{2\min(r, n-r)}\ \left\|U_j U_j^T (I - W_j W_j^T)\right\|_2.$$

Now, $U_j U_j^T$ and $(I - W_j W_j^T)$ are both projectors and hence all their eigenvalues are equal to either 0 or 1 (Bernstein, 2009, p. 358). Therefore,

$$\begin{aligned}\|U_j - W_j Q\|_F &\leq \sqrt{2\min(r, n-r)}\ \left\|U_j U_j^T\right\|_2 \left\|I - W_j W_j^T\right\|_2 \\ &\leq \sqrt{2\min(r, n-r)}.\end{aligned}$$

This concludes the proof of part 1.

In Appendix D.1 we show that, for comparisons of spaces spanned by the r eigenvectors corresponding to either the r largest or r smallest eigenvalues, the objective function in (4.9) approximates 1 as c_0 tends to either ∞ or $-\infty$. This result holds for all optimisation problems of the form (4.9), where $j = 0$ or $j = n - r$. As stated in Remark 4.1, multiplying the objective function by $\sqrt{2\min(r, n-r)}$ yields valid bound values. Hence, for $j \in \{0, n-r\}$ there exist large negative or positive values of c_0 such that the cost function (4.9) results in bounds approximating (4.10). For some problems of the form (4.9),

an appropriate choice of transformation parameters may result in a bound smaller than 1 as shown in Section 4.6. Therefore, for $j \in \{0, n-r\}$ the bound produced from a solution of (4.9) always approximates or improves upon (4.10).

□

Remark 4.3. When considering other than the first or last r eigenvectors, i.e., eigenvector comparisons with $1 \leq j \leq n-1-r$, we observe from Equations (4.5), (4.6), (4.7) and (4.8) that all four quantities $\delta_{1,+}, \delta_{1,-}, \delta_{2,+}$ and $\delta_{2,-}$ tend to $-\infty$ as $c_0 \to \pm\infty$. Hence, for eigenvector comparisons with $1 \leq j \leq n-1-r$, Constraints 3.24 are violated when $c_0 \to \pm\infty$, i.e., $c_0 \to \pm\infty$ lies outside of the feasible set of (4.9) and its related subproblems. Hence, for comparisons of spaces spanned by eigenvectors other than the first or last r, (ordered by corresponding eigenvalue magnitude), a similar statement to part 2 of Proposition 4.2 is not guaranteed. ◁

Remark 4.4. The trivial upper bound in (4.10) applies for all values of j. Suppose we are unable to find a set of affine transformation parameters such that $\|c_1\Phi + c_0 I - \Psi\|_2 < \delta$. This means there exists no affine transformation reducing the distance of the two matrices in the two norm to a value less than the distance of the relevant eigenvalues. In this case, the trivial bound in (4.10) should be used to bound the distance of the spaces spanned by the r consecutive eigenvectors of the two matrices instead of the bound resulting from the solutions of (4.9) and its related subproblems. ◁

4.4 Calculating the bound in practice using fractional programming

In this section we discuss how the optimisation subproblem (4.9) and its related subproblems can be efficiently solved using fractional programming theory.

First off, we briefly want to mention the dccp package (Shen et al., 2016) and how our subproblem does not fall within its class of problems. Optimisation problems with a cost function equal to the difference of convex functions, are solved in Shen et al. (2016). Shen et al. (2016) have published a python package with their paper draft which enables us to efficiently find the optimum of such

problems. However, our cost function in (4.9) cannot easily be brought into the form of the difference of two convex functions. On first sight, the $\log(\cdot)$-transformation of the cost function in (4.9) seems promising. However, we are unable to justify the convexity of an increasing, concave function, i.e., the $\log(\cdot)$ function, composed with a convex function, namely the $\|\cdot\|_2$-term in the numerator of the cost function in (4.9).

4.4.1 Fractional programming

Hence, we turn to fractional programming theory to solve our optimisation problems. Ratio optimisation problems are commonly called *fractional programs* (Frenk and Schaible, 2008). Hence, the optimisation problem (4.9) of choosing the affine transformation parameters $(c_1, c_0) \in \mathbb{R}^+ \times \mathbb{R}$ resulting in a minimal bound is a fractional program. We now describe the properties of fractional programs and their solutions. Then we transform (4.9) to fit the standard class of concave-convex fractional programs and discuss the implementation of its solution.

Firstly, we formally define fractional programs.

Definition 4.5. (Isheden et al., 2012) A general nonlinear *fractional program* has the form,

$$\begin{aligned} \max_{\mathbf{x}} \quad & \frac{g_1(\mathbf{x})}{g_2(\mathbf{x})}, \\ \text{subject to} \quad & \mathbf{x} \in \mathcal{F}, \end{aligned} \tag{4.11}$$

where $\mathcal{F} \subseteq \mathbb{R}^m$, $g_1, g_2 : \mathcal{F} \to \mathbb{R}$ and $g_2(\mathbf{x}) > 0$. Problem (4.11) is called a *concave-convex fractional program* if g_1 is concave, g_2 is convex, and $\mathcal{F}$ is a convex set; additionally $g_1(\mathbf{x}) \geq 0$ for $\mathbf{x} \in \mathcal{F}$ is required, unless g_2 is affine. $\lhd$

In Frenk and Schaible (2008), concave-convex fractional programs are referred to as concave fractional programs. For an excellent overview of the solution and implementation of concave-convex fractional programs we highly recommend Section III in Isheden et al. (2012). Isheden et al. (2012) focus on the application of optimising energy efficiency in wireless communication, which can, for example, be defined as the ratio of the amount of data transmitted and the energy consumed in the process.

Remark 4.6. For concave-convex fractional programs, a powerful and useful practical result is that any local maximum is a global maximum (Frenk and Schaible, 2008). ◁

When discussing the solution of concave-convex fractional programs the concept of equivalence of optimisation problems is essential.

Definition 4.7. Boyd and Vandenberghe (2004, p. 130) define two optimisation problems as *equivalent* if the solution of one problem can be readily obtained given the solution of the other problem and vice versa. ◁

Furthermore, we make use of the standard definition of the feasible set of an optimisation problem.

Definition 4.8. (Boyd and Vandenberghe, 2004, p. 127) The *feasible set* of an optimisation problem is equal to the set of points which satisfy all the constraints of the optimisation problem. ◁

4.4.2 Creating a concave-convex fractional program

Subproblem (4.9) can be transformed to fall into the class of concave-convex fractional programs. As pointed out in Schaible (1976),

$$\max_{\mathbf{x}\in\mathcal{F}} \left(\frac{g_1(\mathbf{x})}{g_2(\mathbf{x})}\right) = \frac{1}{\min_{\mathbf{x}\in\mathcal{F}} \left(\frac{g_2(\mathbf{x})}{g_1(\mathbf{x})}\right)}. \tag{4.12}$$

Using (4.12) we find that solving (4.9) is equivalent to solving,

$$\begin{aligned} \max_{c_1,c_0} \quad & \frac{\delta_{1,+}}{\|c_1\Phi + c_0 I - \Psi\|_2}, \\ \text{subject to} \quad & c_1 > 0, \\ & \delta_{1,+} > 0, \\ & \delta_{1,+} > \psi_{j+r} - c_1\phi_{j+r} - c_0, \\ & \delta_{1,+} > c_1\phi_{j+1} + c_0 - \psi_{j+1}. \end{aligned} \tag{4.13}$$

The bound value is found by transforming the solution of (4.13) according to (4.12).

We now demonstrate that the subproblems such as (4.13), are concave-convex fractional programs.

Lemma 4.9. The δ's in Equations (4.5), (4.6), (4.7) and (4.8) are all concave.

Proof. All the δ's in Equations (4.5)–(4.8) are equal to the minimum of two affine functions of the transformation parameters.

Let $f_1(x), f_2(x), x \in \mathbb{R}$ be affine functions. Then, $-f_1(x)$ and $-f_2(x)$ are still affine functions. Affine functions can be thought of as either convex or concave (Boyd and Vandenberghe, 2004, p. 67) and further the pointwise maximum of convex functions is convex (Boyd and Vandenberghe, 2004, p. 80). Hence, $\max(-f_1(x), -f_2(x))$ is a convex function.

If $g(x), x \in \mathbb{R}$ is a convex function, $-g(x)$ is concave (Boyd and Vandenberghe, 2004, p. 67). Therefore, $-\max(-f_1(x), -f_2(x))$ is concave. But $-\max(-f_1(x), -f_2(x)) = \min(f_1(x), f_2(x))$, is of the same form as the δ's, therefore δ is concave. □

All four subproblems share the denominator $\|c_1\Phi + c_0 I - \Psi\|_2$ which is easily shown to be convex via the triangle inequality. Furthermore, we require the denominator in (4.13) to be strictly positive. Since $\|c_1\Phi + c_0 I - \Psi\|_2$ is not affine we additionally require the numerator of (4.13) to be positive on its feasible set; this is ensured by the δ's being positive on this set. Hence, (4.13) and the remaining 3 subproblems are elements of the class of concave-convex fractional programs as in Definition 4.5.

4.4.3 Solving a concave-convex fractional program

Isheden et al. (2012) present several general approaches to solving concave-convex fractional programs and show that their optimality conditions are equivalent. We chose to discuss and implement the two most popular approaches.

In Section 4.4.4, we discuss the parameter-free approach where an equivalent convex problem is obtained through transformation of the optimisation parameters; the transformation used is commonly referred to as the Charnes-Cooper transformation. This transformed problem only needs to be solved once. In Cheung and Hanzo (2016) the Charnes–Cooper transformation is

used in the optimisation of the energy spectral efficiency of a communication network and in De Maio et al. (2011) the authors show that the maximum likelihood estimate of the steering direction of a signal for radar detection can be found by utilising the Charnes–Cooper transformation of a fractional programming problem.

In Section 4.4.5, we treat the parametric approach which introduces an additional parameter λ to obtain an equivalent problem, which is not jointly convex in (c_1, c_0) and λ. The equivalent problem, is however, convex in (c_1, c_0) and monotone in λ. Therefore, we iteratively solve the convex problem for (c_1, c_0) for a fixed λ and update λ using a Newton-Raphson step. This algorithm is accredited to Dinkelbach (1967). In Ye et al. (2018) minimisation of the system outage probability in a communication network using Dinkelbach's algorithm is discussed, using a closed form solution to the problem at each iteration.

For computational reasons we mainly utilise the parameter-free approach, i.e., the Charnes–Cooper transformation. This follows advice in Schaible (1981, 1983), who state that the iterative solution via Dinkelbach's algorithm is only to be preferred over the single Charnes–Cooper transformed problem, if the solution via Dinkelbach exploits the structure of the numerator and denominator of the fractional program which the Charnes–Cooper solution does not. For instance for quadratic fractional programs — fractional programs with a quadratic numerator and denominator and affine constraints — Dinkelbach's algorithm solves a quadratic program at every iteration, while the Charnes-Cooper transformation yields a concave problem. Therefore, if not many Dinkelbach iterations are necessary for convergence, then Dinkelbach's algorithm is to be preferred over the Charnes–Cooper approach for quadratic fractional programs. We find that for our problem (4.9) both the Charnes–Cooper transformation and Dinkelbach's algorithm solve convex or concave problems. Therefore, we prefer the Charnes–Cooper solution method. However, most importantly, the results of the two different approaches agree in our simulations, as would be anticipated from Isheden et al. (2012) who showed that the optimality conditions of the two approaches are equivalent, so in theory the results should indeed not vary.

4.4.4 The Charnes-Cooper transformation

In this section we discuss a variable transform originally proposed by Charnes and Cooper (1962) for linear fractional programs — fractional programs with an affine numerator and denominator and linear constraints. Schaible (1974) generalised the transformation to concave-convex fractional programs. Isheden et al. (2012) give a good recent summary of the transformation of concave-convex fractional problems. For the transformation of (4.13), appropriate transformation parameters are:

$$\begin{aligned} y_1 &= \frac{c_1}{\|c_1\Phi + c_0 I - \Psi\|_2}; \qquad y_2 = \frac{c_0}{\|c_1\Phi + c_0 I - \Psi\|_2}; \\ t &= \frac{1}{\|c_1\Phi + c_0 I - \Psi\|_2}. \end{aligned} \tag{4.14}$$

Let

$$\delta_{1,+}(t) = \min\left(t\psi_{j+r+1} - y_1\phi_{j+r} - y_2,\ y_1\phi_{j+1} + y_2 - t\psi_j\right).$$

Transforming (4.13) using the parameters in (4.14) we obtain the following convex optimisation problem:

$$\begin{aligned} \max_{y_1,y_2,t} \quad & \delta_{1,+}(t), \\ \text{subject to} \quad & t > 0, \\ & \|y_1\Phi + y_2 I - t\Psi\|_2 \leq 1, \\ & y_1 > 0, \\ & \delta_{1,+}(t) > 0, \\ & \delta_{1,+}(t) > t\psi_{j+r} - y_1\phi_{j+r} - y_2, \\ & \delta_{1,+}(t) > y_1\phi_{j+1} + y_2 - t\psi_{j+1}. \end{aligned} \tag{4.15}$$

In the original proposal of the transformation for linear fractional programs (Charnes and Cooper, 1962) the equality constraint $\|y_1\Phi + y_2 I - t\Psi\|_2 = 1$ was used. This constraint cannot be placed on concave-convex fractional problems, since convex optimisation problems can only have linear equality constraints (Boyd and Vandenberghe, 2004, p. 191). It is proved in Schaible (1974) that for concave-convex fractional programs the constraints $\|y_1\Phi + y_2 I - t\Psi\|_2 = 1$ and $\|y_1\Phi + y_2 I - t\Psi\|_2 \leq 1$ are equivalent. Therefore,

we work with the relaxed constraint $\|y_1\Phi + y_2 I - t\Psi\|_2 \leq 1$.

Note that (4.15) is not a linear program since the constraint $\|y_1\Phi + y_2 I - t\Psi\|_2 \leq 1$ contains a non-linear function of the parameters. The constraint is however convex; therefore, (4.15) is a convex optimisation problem.

We implement the 4 subproblems using the cvx package in MATLAB (Grant and Boyd, 2018, 2008). cvx does not accept strict inequalities. We therefore solve relaxed subproblems, where the strict inequalities are relaxed to include their boundaries, and then check whether the obtained solutions satisfy the strict inequalities. We have found this approach to work extremely well in practice with no convergence issues.

4.4.5 Dinkelbach's algorithm

Dinkelbach's algorithm was proposed in Dinkelbach (1967). Equivalent to (4.13) is the problem

$$\begin{aligned} \max_{c_1, c_0} \quad & \delta_{1,+} - \lambda \left\|c_1\Phi + c_0 I - \Psi\right\|_2, \\ \text{subject to} \quad & c_1 > 0, \\ & \delta_{1,+} > 0, \\ & \delta_{1,+} > \psi_{j+r} - c_1\phi_{j+r} - c_0, \\ & \delta_{1,+} > c_1\phi_{j+1} + c_0 - \psi_{j+1}. \end{aligned} \tag{4.16}$$

Dinkelbach (1967) state that the algorithm can be initialised at a feasible point (c_1, c_0), which is chosen such that the corresponding $\lambda = g_1(c_1, c_0)/g_2(c_1, c_0)$ is positive, or at $\lambda = 0$. When we initialise at $\lambda = 0$ then any feasible set of transformation parameters (c_1, c_0) can be chosen for the initialisation. Therefore, we choose to always initialise at $\lambda = 0$ and (c_1, c_0) to be equal to their respective optima from the Charnes–Cooper algorithm in the corresponding subproblem.

As with the Charnes–Cooper implementation we utilise relaxed subproblems, where the strict inequality constraints are relaxed to include their boundaries. Then we check and report if any of the strict inequality constraints in the

Dinkelbach implementation are violated. Our implementation of Dinkelbach's algorithm can be found in Algorithm 1. We have found the solution of Dinkelbach's algorithm to agree with the solution of the Charnes–Cooper algorithm in all cases we tested.

Dinkelbach's scheme was extremely useful for checking that our implementation of the Charnes–Cooper scheme was correct, but it offered no advantages over the latter, and was much slower.

Algorithm 1 Our implementation of Dinkelbach's algorithm

function DINKELBACH($\Phi, \Psi, \phi, \psi, j, r, \omega, \lambda_0, c_1^0, c_0^0$)
  $i = 1$
  $F_{\lambda_0} = \min(\psi_{j+r+1} - c_1^0\phi_{j+r} - c_0^0, c_1^0\phi_{j+1} + c_0^0 - \psi_j) - \lambda_0\|c_1^0\Phi + c_0^0 I - \Psi\|_2$
  while $|F_{\lambda_{i-1}}| > \omega$ **do**
    Solve (4.16) using λ_{i-1} as input and obtain (c_1^i, c_0^i)
    Check whether the inequality constrains are satisfied
    $F_{\lambda_{i-1}} = \min(\psi_{j+r+1} - c_1^i\phi_{j+r} - c_0^i, c_1^i\phi_{j+1} + c_0^i - \psi_j)$
      $- \lambda_{i-1}\left\|c_1^i\Phi + c_0^i I - \Psi\right\|_2$
    $\lambda_i = \frac{\min(\psi_{j+r+1} - c_1^i\phi_{j+r} - c_0^i, c_1^i\phi_{j+1} + c_0^i - \psi_j)}{\left\|c_1^i\Phi + c_0^i I - \Psi\right\|_2}$
    $i = i + 1$
  end while
  return $1/\lambda_i, c_1^i, c_0^i$
end function

4.5 Higher order polynomial transformations

In this section we briefly discuss the complications which arise when our solution approach is extended to higher order transformations of one of the matrices under comparison as was the case in Chapter 3, where we discussed the theory of polynomial transformations of Φ, $p(\Phi) = c_l\Phi^l + c_{l-1}\Phi^{l-1} + \ldots + c_1\Phi + c_0 I$. The affine transformations, which we concentrate on in this chapter, have the advantage of being monotone and hence the largest and smallest transformed eigenvalues in eigenvalue intervals are easily determined, e.g.,

$$\min_{i\in\{1,\ldots,r\}}(f(\phi_i)) = \begin{cases} c_1\phi_1 + c_0 & \text{for} \quad c_1 \geq 0, \\ c_1\phi_r + c_0 & \text{for} \quad c_1 < 0. \end{cases}$$

The overall minimal bound is thus found by considering 4 different optimisation problems, where affine transforms with $c_1 > 0$ and $c_1 < 0$ are treated separately.

By way of contrast, higher order transformations are not monotone in general and therefore the terms for the $j^{\text{th}}, (j+1)^{\text{th}}, (j+r)^{\text{th}}$ and $(j+r+1)^{\text{th}}$ eigenvalues in the transformed spectrum, namely $\max\limits_{i \in \mathcal{A}_2} p(\phi_i), \min_{i \in \{j+1,\ldots,j+r\}} p(\phi_i)$, $\max_{i \in \{j+1,\ldots,j+r\}} p(\phi_i)$ and $\min\limits_{i \in \mathcal{A}_1} p(\phi_i)$ cannot be simplified further; as we vary the transformation parameters $\{c_l, c_{l-1}, \ldots, c_1, c_0\}$ the eigenvalue at which the minima and maxima occur can change. As a result, it is *difficult to classify the resulting optimisation problem to fall within a certain class of solvable optimisation problems*. Hence, we restrict ourselves to the practical implementation of affine transformations in this work.

4.6 Visualising the bound values: three examples

Problem (4.9) and its solution via (4.13), (along with the three related optimisation subproblems), can be used to calculate bounds on the distance of the spaces spanned by eigenvectors of *any* two symmetric matrices satisfying Assumption 3.21 of Chapter 3. Therefore, we envisage that the affine transform could contribute tighter bounds in a range of fields where eigenvectors are used. We highlight three such applications. In Section 4.6.4 we study our bound on the distance of the spaces spanned by the eigenvectors of the three graph representation matrices. Then, in Section 4.6.5 we will apply the bound to the comparison of the representation matrices to their respective generating matrices in the SBM. Our final example application in Section 4.6.6 is in a principal component analysis setting, where we compare the space spanned by the eigenvectors of the sample covariance matrix and its corresponding population covariance matrix in a spiked covariance model.

4.6.1 The stochastic blockmodel

Throughout Sections 4.6.4 and 4.6.5 we generate networks from the SBM, which was introduced in Section 1.2.4. In this section we take the matrix of edge probabilities to be composed of only two values. On the diagonal we have p_w encoding the probability of edges within the different blocks to be the same for all blocks. Off-diagonal we have p_b to encode the probability of edges between nodes in different blocks. For example, for $K = 3$, P takes the form:

$$P = \begin{pmatrix} p_w & p_b & p_b \\ p_b & p_w & p_b \\ p_b & p_b & p_w \end{pmatrix}. \tag{4.17}$$

4.6.2 Scaling

For affine transforms, from (4.1) and (4.4) the bound of interest is the right-hand-side of

$$\begin{aligned} \|U_j - W_j Q\|_F &\leq \sqrt{2\min(r, n-r)}\, \left\|U_j U_j^T - (I - W_j W_j^T)\right\|_2 \\ &\leq \sqrt{2\min(r, n-r)} \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta_i}, \end{aligned}$$

where δ_i equals $\delta_{1,+}, \delta_{1,-}, \delta_{2,+}$ or $\delta_{2,-}$, from Equations (4.5)-(4.8), which generate the four different optimisation subproblems of the form (4.13).

In this section we divide all bound values and attained values by $\sqrt{2\min(r, n-r)}$, i.e., we consider instead,

$$\begin{aligned} \frac{\|U_j - W_j Q\|_F}{\sqrt{2\min(r, n-r)}} &\leq \left\|U_j U_j^T - (I - W_j W_j^T)\right\|_2 \\ &\leq \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta_i}, \end{aligned} \tag{4.18}$$

This rescaling has the advantage that, independent of n and r, our bound values are on the same relative scale. Furthermore, the trivial bound derived in Proposition 4.2 corresponds to the upper bound of 1 in all plots, rather than the value $\sqrt{2\min(r, n-r)}$, which would vary across the different simulations.

In what follows, (scaled) attained distance in the metric $\rho_1(\mathcal{U}_j, \mathcal{W}_j)$ refers to the quantity $1/\sqrt{2\min(r, n-r)}\inf_{R\in O(r)} \|U_j - W_j R\|_F$ and (scaled) attained

distance in the metric $\rho_2(\mathcal{U}_j, \mathcal{W}_j)$ refers to $\left\|U_j U_j^T - (I - W_j W_j^T)\right\|_2$. (For the first of these we recall from Chapter 3 that, when calculating distances, finding the matrix Q for which the infimum is attained can be avoided by the use of canonical angles.)

4.6.3 Eigenvectors of the representation matrices

In this section we give a brief review of the representation matrix eigenvectors.

Remark 4.10. For any symmetric matrix, we can find an orthogonal basis of $\mathbb{R}^n$ consisting of its eigenvectors (Abadir and Magnus, 2005, p. 179). Since A, L and L_{sym} are symmetric matrices, we can assume to always be working with their orthonormal sets of eigenvectors, which is the standard output of most eigenvector functions, as for example `eig()` in MATLAB. ◁

An orthogonal set of eigenvectors can easily be shown to exist if and *only if* the corresponding matrix is symmetric. Hence, since L_{rw} is not symmetric, L_{rw} does not have an orthogonal set of eigenvectors. Therefore, we will be using the symmetric normalised Laplacian L_{sym} instead of the random walk normalised Laplacian L_{rw} to visualise our results in Sections 4.6.4 and 4.6.5. Note that there exists results extending the Davis–Kahan theorem to singular value decompositions (Wedin, 1972). From initial inspection, it seems likely that our work here could be extended to also apply to the singular value decomposition of matrices. Therefore, it seems highly probable that similar bounds to the ones presented here could be derived for comparisons involving the singular vectors of L_{rw}.

Example 4.11 shows that the Laplacians and the adjacency matrix only share eigenvector sets for d-regular graphs and have unequal eigenvector sets for non-d-regular graphs.

Example 4.11. The constant vector $\mathbf{1} = (1, \ldots, 1)^T$ is an eigenvector of every graph Laplacian (van Mieghem, 2011, p. 64). Right multiplying the adjacency matrix A by $\mathbf{1}$ yields the vector of vertex degrees $(d_1, d_2, \ldots, d_n)^T$. It is easy to spot that this can satisfy the eigenvector equation

$$A\mathbf{1} = (d_1, d_2, \ldots, d_n)^T = \mu\mathbf{1}$$

if and only if the degrees of all vertices are equal. This example clearly shows that the set of eigenvectors of the Laplacian and the adjacency matrix can only be equal for d-regular graphs. ◁

4.6.4 Different pairs of representation matrices

In this section we calculate the bound on the distance of the spaces spanned by the eigenvectors of the representation matrices. Recall that the largest eigenvalues of the adjacency matrix correspond to the smallest eigenvalues of the Laplacians. Hence, in two of the three presented matrix comparisons we will compare spaces spanned by eigenvectors corresponding to eigenvalues on opposing ends of the eigenvalue spectrum. Therefore, in the majority of cases presented in this section the Davis–Kahan Theorem does not apply, while bound values can be obtained via our extended Davis–Kahan Theorem.

Throughout this section we consider SBMs with $K = 3$, where every block is composed of equally many nodes. The parameters identifying the compared eigenvectors are $j = 1$ and $r = 2$. This choice of j and r is motivated by the fact that the first eigenvector of the Laplacian matrices is a constant vector and is therefore not informative in the recovery of the blocks in the SBM. Hence, we are comparing spaces spanned by eigenvectors corresponding to the *second and third largest* eigenvalue of the adjacency matrix to the eigenvectors corresponding to the *second and third smallest* eigenvalues of the two graph Laplacians.

In Figures 4.1 (a) and (b) we plot the bound and attained values arising from the comparison of the spaces spanned by eigenvectors of the three representation matrices against the degree extremes of the corresponding graphs. We simulated 25 SBMs with equal parameters $n = 300, K = 3, (p_b, p_w) = (0.1, 0.6)$.

In Figure 4.1 (a) we observe that the bound grows with a growing degree extreme difference for the comparisons of L with A and L_{sym}, while the bound remains relatively constant across different degree extremes when A and L_{sym} are compared. The bound values for comparisons of L with A and L_{sym} are very close, differing only by very small amounts. The bound arising from the comparison of A and L_{sym} attains much lower values than the bound values arising from the other two comparisons. Hence, using an affine transformation

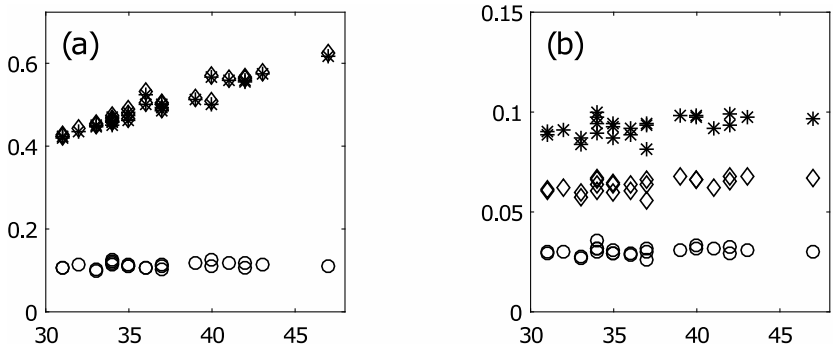

Figure 4.1: Different pairs of graph representation matrices. (a) Bound values plotted against the degree extreme differences (x-axis) of their corresponding network. The stars correspond to the comparison of spaces spanned by eigenvectors of A and L, diamonds correspond to the comparison of L and L_{sym} and circles correspond to the comparison of A and L_{sym}. (b) The attained distances in the metric $\rho_1(\mathcal{U}_1, \mathcal{W}_1)$ (see Chapter 3) are plotted against the degree extreme differences (x-axis) of their corresponding network. All values have been rescaled as discussed in the text.

we are able to obtain a very small bound on the difference of spaces spanned by the eigenvectors of A and L_{sym}, which suggests that they are very close. Not only are they close to each other, they also produce extremely similar bounds on the space spanned by the eigenvectors when individually compared to the eigenvectors of L.

In Figure 4.1 (b) we observe that the attained distances of the three comparisons remain rather constant across the different degree extreme differences.

The results in Figure 4.1 show that for the comparison of A and L_{sym} an affine transformation is sufficient to remove the dependence of the bound value on the degree extreme difference. In contrast the much higher bounds for the comparison of the eigenvectors of L with A and L_{sym} still depend on the degree extreme difference and the affine transformation was not sufficient to remove this dependence.

In Figure 4.2, we observe the effect of a growing number of network nodes n on our bound. For each value of $n \in \{30, 120, 210, 300\}$ we simulated 25 SBMs with equal parameters $K = 3, (p_b, p_w) = (0.1, 0.9)$. In all plots the four values of n are displayed on the x-axis. The first column of plots in Figure 4.2 displays boxplots of the bound values of the three possible pairwise comparisons, while the second and third columns display boxplots of the optimal affine

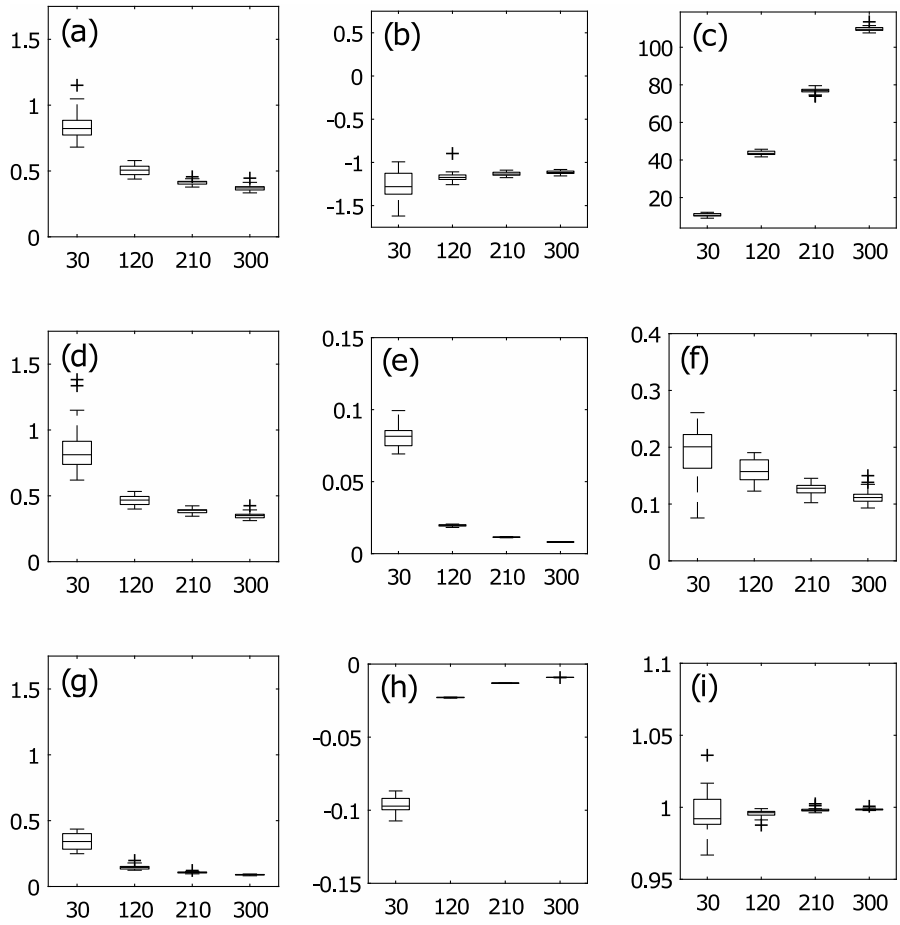

Figure 4.2: Different pairs of graph representation matrices and the effect of a growing number of nodes (x-axis). First column: boxplots of the bound values of the three possible pairwise comparisons; second column: boxplots of optimal transformation parameter c_1; third column: boxplots of optimal transformation parameter c_0. First row: comparison of A and L; second row: L and L_{sym}; third row: A and L_{sym}.

transformation parameters c_1 and c_0, respectively. The rows of plots show the comparison of spaces spanned by eigenvectors of A and L, (first), L and L_{sym}, (second) and A and L_{sym}, (third). The boxplots show the interquartile range (the distance of the 25% and 75% quartiles) as the boxes and highlight the median as lines in the middle of these boxes. The whiskers extend to the maximum observed value within a distance of 1.5 times the interquartile range from the box. All datapoints beyond the whiskers are marked using the symbol '+' in the plots. Just the 25 samples from each SBM parametrisation were sufficient to reveal general trends in the 9 plots.

From Figure 4.2 (a), (d) and (g) we immediately see that the bound values decrease as the number of nodes in the SBMs grows indicating that as n grows the differences between the spaces spanned by the representation matrix eigenvectors decrease. From Figure 4.2 (b), (c), (e), (f), (h) and (i) we observe the majority of optimal affine transformation parameters depend on n. For the comparison of A and L, in Figure 4.2(c), we find the optimal additive parameter c_0 to grow from roughly 10 for $n = 30$ to roughly 110 for $n = 300$, while the multiplicative parameter in Figure 4.2(b) remains almost constant for all values of n. Comparison of the Laplacian matrices L and L_{sym}, shows both transformation parameters to vary slightly with n, (Figure 4.2(e) and (f)). Finally, comparison of A and L_{sym}, shows the additive parameter c_0 to remain mostly constant with changing n (Figure 4.2(i)) while the multiplicative parameter c_1 grows with n (Figure 4.2(h)). The magnitude by which the transformation parameters change as n grows is clearly variable.

We were unable to run the simulation displayed in Figure 4.2 beyond $n = 300$ within reasonable computation time. The times for the simulations with the four different values of n were, $3, 21, 118$ and 472 minutes, respectively. For each value of n, 300 convex optimisation problems were solved since each bound on the three possible representation matrix comparisons was calculated for 25 different SBM samples per value of n and calculation of each bound involves the solution of 4 different convex optimisation problems. The main contribution to the computation time comes from setting up the optimisation problem in `cvx`. The issue being that the two norm constraint, $\|y_1\Phi + y_2 I - t\Psi\|_2 \leq 1$, in (4.15) is interpreted as having $2n^2 + n$ variables. Ideally, the convex optimisation implementation would be able to recognise that in fact only 3 variables are present in the constraint and would therefore be able to compute our bounds for much larger networks. However, for the purpose of this section it is satisfactory to consider networks with up to 300 nodes and therefore, we will not address this issue in the implementation of our optimisation procedure.

Remark 4.12. In all our simulations we have found the distance between the spaces spanned by the *eigenvectors* of A and L_{sym} to be significantly smaller than the distance arising from the remaining two pairwise comparisons of the representation matrices. In Section 2.6 the same behaviour was observed in our *spectral* bounds on the representation matrices. $\triangleleft$

4.6.5 Representation matrices and generating matrices

In this section we compare (i) spaces spanned by eigenvectors of the representation matrices to (ii) spaces spanned by eigenvectors of their corresponding generating matrices in the SBM $B_A = MPM^T - \text{diag}(MPM^T), B_L = \text{diag}(B_A\mathbf{1}_n) - B_A$ and $B_{L_{sym}} = \text{diag}(B_A\mathbf{1}_n)^{-1/2}B_L\text{diag}(B_A\mathbf{1}_n)^{-1/2}$. (Here $\mathbf{1}_n$ is a column vector of ones with n entries and the term $-\text{diag}(MPM^T)$ in the calculation of B_A ensures that our SBMs do not have self-loops.) It is natural to compare the spaces spanned by these eigenvectors, since consistency and rate of convergence of different methods, based on the eigenvectors of the representation matrices in a SBM setting, can be demonstrated (Cape et al., 2019; Eldridge et al., 2018; Rohe et al., 2011).

In Figures 4.3 (a) and (b) the bound and attained values from the comparison of the spaces spanned by eigenvectors of the three representation matrices to the eigenvectors of their respective generating matrices are plotted against the degree extremes of the corresponding graphs. In the adjacency matrix comparison, the eigenvectors corresponding to the three largest eigenvalues are compared, while for the Laplacians we concern ourselves with the eigenvectors corresponding to the three smallest eigenvalues. 25 realisations of a SBM with parameters $n = 210, K = 3, (p_b, p_w) = (0.1, 0.9)$ were simulated. In this comparison we included the first eigenvector of the matrices under comparison, i.e., we chose $j = 0$ and $r = 3$.

We see in Figure 4.3 that the ordering in magnitude of the attained values is reflected in the bound values, with the bounds on the unnormalised Laplacian comparison (L, B_L) taking the largest values. We observe that both the attained and the bound values in the comparison of the spaces spanned by eigenvectors of L and B_L seem to grow with growing degree extreme difference, which is not the case for the comparisons involving A and L_{sym}.

In Figure 4.4 we observe the effect of a growing number of network nodes n on our bound. For each value of $n \in \{30, 120, 210, 300\}$ we simulated 25 SBMs with parameters $K = 3, (p_b, p_w) = (0.1, 0.8)$. The transformation parameters are roughly centred around the parameters of the identity transformation, $f(x) = x$, i.e., $c_1 = 1$ and $c_0 = 0$. Interestingly, the variance of the additive parameter c_0 seems to be increasing with increasing n for the comparison

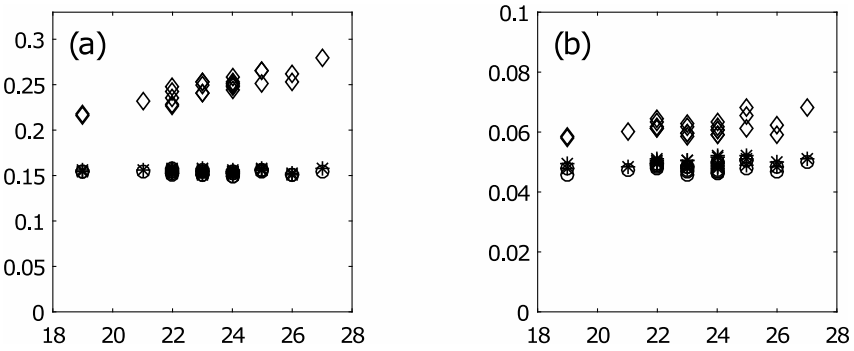

Figure 4.3: Comparison of the spaces spanned by eigenvectors of the representation matrices to those of their respective generating matrices, plotted against the degree extremes (x-axis) of their corresponding network. The stars correspond to the comparison of eigenvectors of A to B_A, diamonds correspond to the comparison of eigenvectors of L and B_L and circles correspond to the bound values arising from the comparison of L_{sym} and $B_{L_{sym}}$. (a) bound values, and (b) attained distances in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$, (see Chapter 3). The x-axes display the indices of the SBM samples. All displayed values are rescaled as discussed in the text.

(L, B_L). For all other displayed values in Figure 4.4 we find the variance of the observed values to decrease as the number of nodes in the network grows. As the number of nodes, n, in the SBM grows, the bound values of all three comparisons decrease.

For the spaces spanned by the leading eigenvectors of the representation matrices compared to their corresponding generating matrices, Figure 4.5 shows the usual Davis–Kahan bounds (Theorem 3.17 of Chapter 3), our sharpened bounds (Theorem 3.25 of Chapter 3 and (4.18)) and the attained values; all were standardised. Here 25 realisations of a SBM with parameters $n = 30, K = 3, (p_b, p_w) = (0.1, 0.6)$ were generated. In all three comparisons, our sharpened bound values improve on the usual Davis–Kahan bound values. In the case of the unnormalised Laplacian several of the usual Davis–Kahan bound values are greater than 1, therefore, even the trivial bound value of 1 (see Proposition 4.2 and Section 4.6.2) is tighter than the usual Davis-Kahan bound. In contrast, our bound produces values consistently lower than 1.

In addition to the attained distances of the spaces spanned by the eigenvectors in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$, we have shown the distance in the metric $\rho_2(\mathcal{U}_0, \mathcal{W}_0)$ in Figure 4.5, which is discussed in Chapter 3. As expected from Lemma 3.12 and Theorem 3.25 in Chapter 3, we find the attained values $\rho_2(\mathcal{U}_0, \mathcal{W}_0)$ to

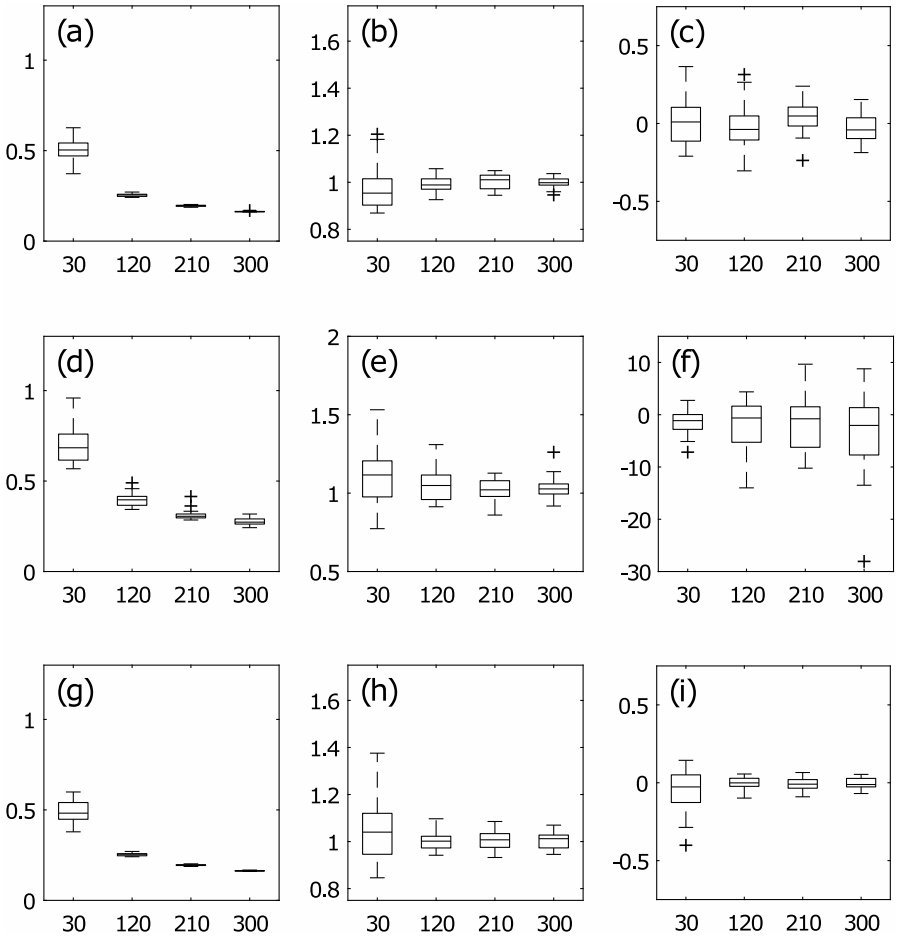

Figure 4.4: Representation matrices and generating matrices: the effect of a growing number of nodes (x-axis). First column: boxplots of the bound values; second column: boxplots of optimal transformation parameter c_1; third column: boxplots of optimal transformation parameter c_0. The first row is for the comparison of spaces spanned by eigenvectors of A and B_A, the second row, L and B_L, and the third row, L_{sym} and $B_{L_{sym}}$.

fall between the distances in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$ and our sharpened Davis–Kahan values. Of particular interest are the values attained in simulation number 10 in Figure 4.5(b), where we find the distance of the spaces spanned by the eigenvectors to come very close to 1 in the metric $\rho_2(\mathcal{U}_0, \mathcal{W}_0)$ and our sharpened bound to be very close to tight in this instance.

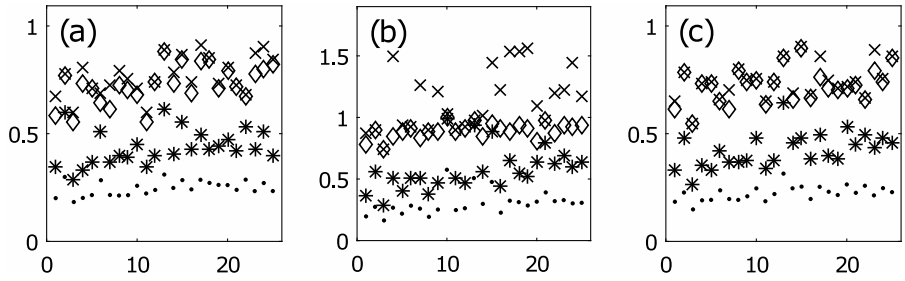

Figure 4.5: Representation matrices and generating matrices: comparison of spaces spanned by eigenvectors of (a) A and B_A, (b) L and B_L, (c) L_{sym} and $B_{L_{sym}}$. The x-axis shows SBM sample indices. The dots represent the attained distances in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$ and the stars represent attained distances in the metric $\rho_2(\mathcal{U}_0, \mathcal{W}_0)$ (see Chapter 3). The diamonds show our sharpened Davis–Kahan bound using an affine matrix transformation, and the x's represent usual Davis–Kahan bound values. All displayed values were rescaled as discussed in the text.

4.6.6 Sample covariance and population covariance matrices

Our final example of the application of our sharpened Davis–Kahan bound is in the setting of Principal Component Analysis (PCA). Consider $N \in \mathbb{N}$ independent, identically distributed samples $\{X_i\}$ from a multivariate normal distribution of dimension $p \in \mathbb{N}$, with mean $\mathbf{0}$ and covariance matrix Σ. Then let X be the $p \times N$ matrix with columns X_i with $i \in \{1, \ldots, N\}$. We denote the sample covariance matrix by $\hat{\Sigma} = XX^T/N$. When a low dimensional structure truly generates the covariance matrix, PCA is the correct tool to recover this low dimensional space. The standard PCA algorithm maps the data into the space spanned by the r eigenvectors corresponding to the largest eigenvalues of $\hat{\Sigma}$. In this setting it is of interest to study the convergence of the space spanned by these leading r eigenvectors of $\hat{\Sigma}$ to the leading eigenvectors spanning the true low dimensional covariance space of Σ. Using a so-called spiked covariance model, we will apply our bound to the spaces spanned by eigenvectors corresponding to the largest eigenvalues of Σ and $\hat{\Sigma}$.

Remark 4.13. The use of eigenvalues in the context of PCA of data and spectral clustering of graphs is different. In the context of PCA we consider the eigenvectors corresponding to the largest eigenvalues of the covariance matrix to maximise the variance of the data captured by our chosen eigenvectors. Contrariwise, when using one of the Laplacians in the spectral clustering al-

gorithm, we consider the eigenvectors corresponding to the smallest eigenvalues in order to capture low frequency, large neighbourhood effects on the graph as was discussed in Remark 1.28 in Section 1.4.3. Hence, in one application the variance is maximised and we are concerned with high frequency signals, while in the other application we are concerned with low frequency, low variance signals. ◁

The spiked covariance model was first introduced by Johnstone (2001), who described a phenomenon in real world data where the largest eigenvalues of the covariance matrix are separated by a large eigengap from the rest of the spectrum. Jung and Marron (2009) proved consistency of the first r eigenvectors of the estimated sample covariance towards the population covariance in the case of zero mean normally-distributed data and under certain conditions on the growth of the largest eigenvalues with growing dimensions p and N. The spiked covariance model has been found to be implied by the factor model (Fan et al., 2013; Shen et al., 2016; Wang and Fan, 2017), which models a multivariate time series as being driven by a few main factors. The factor model and consequently the spiked covariance model, find application to financial data.

In our parametrisation of the spiked covariance model $r \in \mathbb{N}$ determines the dimension of the low-dimensional latent space in which Σ is generated, $M \in \{0,1\}^{p\times r}$ encodes the latent dimension membership (similar to the SBM) and $P \in \mathbb{R}^{r\times r}$ encodes the correlations between latent dimensions. Σ has to be a valid covariance matrix, so must be symmetric positive semi-definite. As for the SBMs, we chose P to consist of only two values and take the form given in (4.17). We define the population covariance matrix as,

$$\Sigma = MPM^T + I.$$

By building our covariance matrix like this, we get a covariance matrix following the spiked covariance model, where r eigenvalues are significantly larger than the rest of the spectrum (Johnstone, 2001; Jung and Marron, 2009), the latter consisting of eigenvalues all equal to 1, as a result of adding I into the covariance structure.

In Figure 4.6(a) we plot our sharpened bound for the spaces spanned by the eigenvectors corresponding to the largest 3 eigenvalues of Σ and $\hat{\Sigma}$ and in Fig-

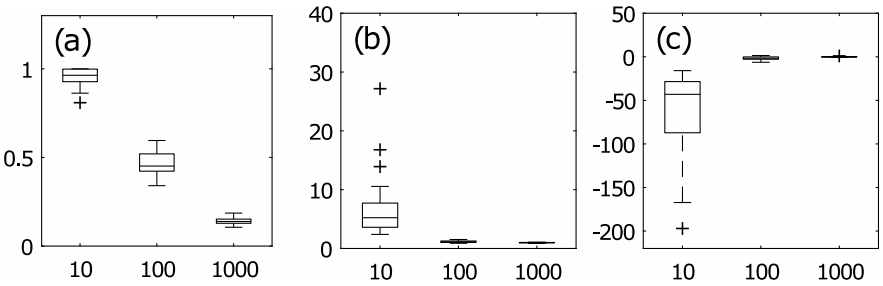

Figure 4.6: Effect of increasing sample size N (x-axis) for the space spanned by eigenvectors of the sample covariance matrix versus the population covariance matrix. (a) boxplots of bound values, (b) boxplots of optimal transformation parameter c_1, (c) boxplots of optimal transformation parameter c_0.

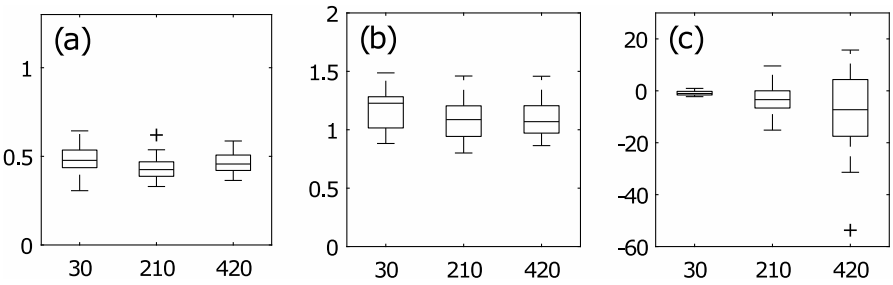

Figure 4.7: Effect of increasing dimension p (x-axis) for the space spanned by the eigenvectors of the sample covariance matrix versus population covariance matrix. (a) boxplots of bound values, (b) boxplots of optimal transformation parameter c_1, (c) boxplots of optimal transformation parameter c_0.

ure 4.6(b) and (c) the corresponding optimal transformation parameters c_1 and c_0, respectively. Here $N \in \{10, 100, 1000\}, p = 60, (p_b, p_w) = (0.2, 0.8), j = 0, r = 3$. In Figure 4.6(a) we see the bound decreases as the sample size N grows. This makes sense: more samples should improve the estimation performance of the sample covariance matrix and therefore the distance of the subspaces spanned by the first three eigenvectors of the sample and population covariance matrices should decrease. For $N = 10$ samples we find that in a few cases the bound value 1 is attained. This situation was theoretically discussed in Proposition 4.2. It is nice to see that we do indeed find the bound to converge to 1 with diverging transformation parameters in the worst case in practice. In plots (b) and (c) we find the transformation parameters to converge to the identity transformation as N grows.

The format of Figure 4.7 follows Figure 4.6 with the difference that we keep N

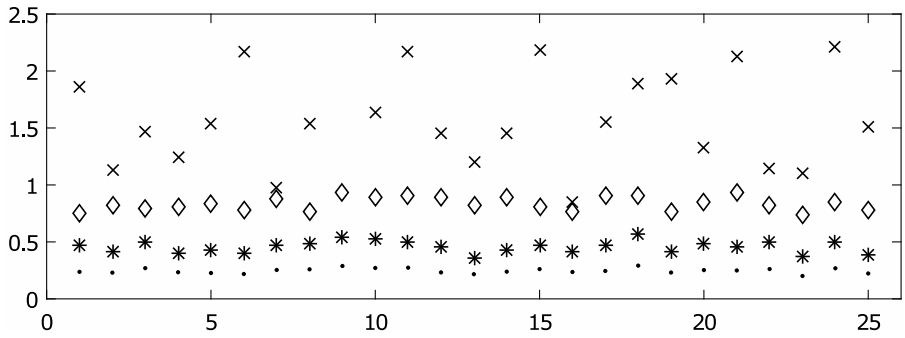

Figure 4.8: Spaces spanned by first three eigenvectors of the sample covariance matrix versus the population covariance matrix. The x-axis shows sample indices. The dots represent attained distances in the metric $\rho_1(\mathcal{U}_0, \mathcal{W}_0)$ and the stars represent attained distances in the metric $\rho_2(\mathcal{U}_0, \mathcal{W}_0)$ (see Chapter 3). The diamonds show our sharpened Davis–Kahan bound using an affine matrix transformation. The x's show the usual Davis–Kahan bounds. The x-axis displays the indices of the normal samples. All displayed values are rescaled as discussed in the text.

fixed at 100 and study the behaviour of our bound as p grows, $p \in \{30, 210, 420\}$. Interestingly, our bound remains fairly constant for the different values of p considered. The transformation parameters hover around the identity transformation $f(x) = x$ with the uncertainty in the additive parameter c_0 increasing as p increases.

Using $N = 100, p = 60, (p_b, p_w) = (0.4, 0.6), j = 0, r = 3$, and 25 simulations, Figure 4.8 shows that our extended and sharpened Davis–Kahan bounds improve upon the usual Davis–Kahan bounds by roughly a factor of 2. In Proposition 4.2 it was discussed that bound values above 1 are non-informative. In Figure 4.8 we observe our bound to consistently fall below 1, improving on the trivial bound; however, the usual Davis–Kahan bound attains values consistently above 1.

Conclusion

We have compared the eigenproperties of the graph representation matrices: the adjacency matrix A, the unnormalised graph Laplacian L and the normalised graph Laplacians L_{rw}, L_{sym} and found differences in their spectra and eigenvectors corresponding to general graphs.

Overall, we want to draw three conclusions from this thesis.

1) In Chapters 2 and 4 it became apparent that the degree extreme difference of the described network is clearly influential in the magnitude of the differences we observe in the representation matrix eigenproperties. *In examples with varying degree extremes, we saw that if the degree extreme difference is large, different choices of representation matrix may give rise to disparate inference drawn from network analysis; smaller degree extreme differences will result in consistent inference, whatever the choice of representation matrix.* Therefore, network analysts should be aware that the degree extreme difference is an indicator for potential differences in the analyses based on the different representation matrices.

2) Polynomial matrix transforms were established as a powerful tool in the comparison of the representation matrices, which enables significant insight into the differences of their eigenproperties. In Chapter 3 we also found them to play a crucial role in the tightening and extension of the Davis–Kahan theorem, which applies to any two symmetric matrices.

3) No representation matrix has turned out to strictly outperform another throughout the work presented here. Therefore, it is difficult to conclude this work by recommending a single representation matrix over the remaining representation matrices. However, we were able to observe that the eigenproperties of A and L_{sym} consistently displayed the smallest differences.

As a result of this work, we hope to have increased awareness about the importance of the choice of representation matrix for network analysis. As a concrete example of how our thesis work contributes to this issue, an alternative to considering all three representation matrices and their alternatives in each network analysis can be proposed, and consists of two steps: Firstly, pre-analysis one could make an informed representation matrix decision based on the eigenproperties we summarised in Chapter 1. As we saw the representation matrix eigenproperties carry different qualitative information about the underlying graph, such as for example, the number of spanning trees, the bipartiteness of the graph or the number of connected components. Depending on the aspects of a graph which are important in a particular analysis, one could choose the representation matrix where these aspects are most naturally encoded in the matrix eigenproperties. Secondly, post-analysis one should quantify maximal differences in the used eigenproperties via our bounds, which were derived in Chapters 2, 3 and 4. This second step can be seen as a robustness analysis of the obtained results and can help assess the impact of the initial representation matrix choice.

Specific, more detailed, conclusions from the work in the individual chapters follow now.

In *Chapter 1* clear differences in the strengths of the eigenproperties of the different representation matrices can be found in our lists of their eigenproperties. However, we also found similarities. A variant of Cheeger's inequality exists for all three spectra and the global clustering behaviour, where we obtain homophilic cluster structure and heterophilic cluster structure from opposite ends of the spectra, holds for all three representation matrices. Furthermore, we were able to show that all previously existing inequality spectral relations of the representation matrices can be proved via Weyl's inequality and a similar theorem which applies to the product of two matrices.

In *Chapter 2*, the degree extreme difference, $d_{\max} - d_{\min}$, was found to linearly upper bound the spectral differences of all three pairs of representation matrices when transforming one of the two spectra by an affine transformation. We explained the monotonicity found in the bounds by partitioning the class of graphs according to their degree extremes and considering the addition/deletion of connected components to/from the graph. Our bounds were

illustrated on several graphs, including Zachary's karate dataset; we found the bounds $e(A, L)$ and $e(A, L_{rw})$ to be tight and the bound $e(L, L_{rw})$ to be close to tight on our examples. The bounds were extended to bound normalised eigengaps of the representation matrices. In our examples in Section 2.11.1 we saw that a degree extreme difference larger than 3 was sufficient to change the inference drawn from the three representation matrix spectra. We further analysed the behaviour of the representation matrix eigengaps together with our bound on graphs with changing degree extremes, a graph sampled from a stochastic blockmodel with 600 nodes, where we observed homophilic and heterophilic cluster structure, and on the karate dataset. To visualise the impact of the representation matrix choice on the inference drawn from network analysis we ran the spectral clustering algorithm using all three representation matrices on a synthetic example graph and on the karate data set graph. Finally, we demonstrated how our eigenvalue bounds can be used to transfer graph theoretic results from one representation spectrum to another using the complexity of a graph as an example.

In *Chapter 3* Theorem 3.17 showed that Davis–Kahan bounds always exist for the comparison of spaces spanned by the first r eigenvectors of two symmetric matrices. Further, we proved an extended version of the Davis–Kahan theorem, which applies to the comparison of spaces spanned by any r consecutive corresponding eigenvectors of two symmetric matrices; Theorem 3.25 and Corollary 3.26 cover the metrics ρ_2 and ρ_1, respectively. The main tool in the extension of the theorem was the consideration of a polynomial transformation of one of the matrices under comparison. Our extended version *includes the original Davis–Kahan theorem as a special case* and consequently our bounds are guaranteed to be at least as tight as those given by the original Davis–Kahan theorem. The case of affine transformations as a class of polynomial transformations was discussed. A proof of concept example demonstrated that in the case of affine transformations our extended Davis–Kahan theorem, implemented as discussed in detail in Chapter 4, outperforms the conventional Davis–Kahan theorem by automatically recovering the exact relationship between the adjacency matrix and the unnormalised and normalised graph Laplacians in the case of d-regular graphs. One property of our extended Davis–Kahan theorem, shared with the original Davis–Kahan theorem, is that both matrices under comparison have to be known to calculate bounds if these are to be known

numerically rather than analytically. It is an open question how to optimally *estimate* transformation parameters for our extended Davis–Kahan theorem in situations where we do not have complete knowledge of the two matrices. The answer to this question is likely to be application specific since the amount of knowledge available about the two matrices and their difference varies in different applications.

In *Chapter 4* we demonstrated that via the Charnes–Cooper transformation the calculation of bounds from our extended Davis–Kahan theorem in the case of affine transformations is a convex optimisation problem. We furthermore, considered the extension of our implementation to higher order polynomial transformations and discussed its complications. We found that in comparisons amongst the representation matrices, and of representation matrices with their generating matrices (in a stochastic blockmodel setting), our fractional programming implementation of the affine Davis–Kahan bounds is superior to the standard Davis–Kahan bounds. The same was found when working with the eigenvectors of the sample and population covariance matrices in a spiked covariance model (for which the PCA algorithm is a well motivated analysis tool).

Bibliography

K. M. Abadir & J. R. Magnus, *Matrix algebra*, Cambridge, UK: Cambridge University Press, 2005.

E. Abbe, "Community detection and stochastic block models: Recent developments," *Journal of Machine Learning Research*, vol. 18, pp. 1–86, 2017.

R. Albert & A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, pp. 47–97, 2002.

N. Alon & F. R. K. Chung, "Explicit construction of linear sized tolerant networks," *Annals of Discrete Mathematics*, vol. 38, pp. 15–19, 1988.

M. Aouchiche & P. Hansen, "Two Laplacians for the distance matrix of a graph," *Linear Algebra and its Applications*, vol. 439, pp. 21–33, 2013.

Z. Bai, J. Demmel, J. Dongarra, A. Ruhe & H. van der Vorst, *Templates for the solution of algebraic eigenvalue problems: a practical guide*, Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2000.

F. Ball & P. Neal, "The asymptotic variance of the giant component of configuration model random graphs," *The Annals of Applied Probability*, vol. 27, pp. 1057–1092, 2017.

A. Banerjee & J. Jost, "On the spectrum of the normalized graph Laplacian," *Linear Algebra and its Applications*, vol. 428, pp. 3015–3022, 2008.

J. Banks, C. Moore, M. Newman & P. Zhang, "Community detection with the z-Laplacian," *Santa Fe Institute Working Paper,* 2014.

V. Batagelj & A. Mrvar, "Pajek datasets," http://vlado.fmf.uni-lj.si/pub/networks/data/, 2006 (Accessed: 2016-11-23).

M. Belkin & P. Niyogi, "Laplacian eigenmaps for dimensionality reduction and data representation," *Neural Computation*, vol. 15, pp. 1373–1396, 2003.

H. P. Benson, "Fractional programming with convex quadratic forms and functions," *European Journal of Operational Research*, vol. 173, pp. 351 – 369, 2006.

D. S. Bernstein, *Matrix mathematics: Theory, facts, and formulas (second edition)*, Princeton, NJ: Princeton University Press, 2009.

R. Bhatia, *Matrix analysis*, Berlin, Germany: Springer, 2013.

P. J. Bickel & P. Sarkar, "Hypothesis testing for automated community detection in networks," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 78, pp. 253–273, 2016.

T. Biyikoğu, J. Leydold & P. F. Stadler *Laplacian eigenvectors of graphs: Perron-Frobenius and Faber-Krahn type theorems*, Berlin, Germany: Springer, 2007.

A. W. Bohannon , B. M. Sadler & R. V. Balan, " A filtering framework for time-varying graph signals," In: L. Stankovic & E. Sejdic (eds.) *Vertex-frequency analysis of graph signals, signals and communication technology,* Cham, Switzerland: Springer, pp. 341-376, 2019.

J. Bondy & U. S. R. Murty, *Graph theory*, Berlin, Germany: Springer, 2008.

C. Bordenave & M. Lelarge, "Resolvent of large random graphs," *Random Structures & Algorithms*, vol. 37, pp. 332–352, 2010.

C. Bordenave, M. Lelarge & L. Massoulié, "Non-backtracking spectrum of random graphs: community detection and non-regular Ramanujan graphs," *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 1347–1357, 2015.

S. Boyd & L. Vandenberghe, *Convex optimization*, Cambridge, UK: Cambridge University Press, 2004.

A. E. Brouwer & W. H. Haemers, "A lower bound for the Laplacian eigenvalues of a graph — Proof of a conjecture by Guo," *Linear Algebra and its Applications* , vol. 429, pp. 2131–2135, 2008.

A. E. Brouwer & W. H. Haemers, *Spectra of graphs*, New York, NY: Springer, 2011.

T. Bühler & M. Hein, "Spectral clustering based on the graph p-Laplacian," In: *Proceedings of the 26th Annual International Conference on Machine Learning,* New York, NY: ACM, pp. 81–88, 2009.

S. Butler, *Eigenvalues and structures of graphs*, PhD thesis, University of California, San Diego, 2008.

S. Butler, "A gentle introduction to the normalized Laplacian," *The Bulletin of the International Linear Algebra Society IMAGE* , vol. 53, pp. 19 – 27, 2014.

S. Butler, " Algebraic aspects of the normalized Laplacian," In: A. Beveridge, J. Griggs, L. Hogben, G. Musiker & P. Tetali (eds.) *Recent Trends in Combinatorics*, Cham, Switzerland: Springer, pp. 295–315, 2016.

S. Butler & F. Chung, "Spectral graph theory," In: L. Hogben (ed) *Handbook of linear algebra (2nd edition)*, Boca Raton, FL: CRC Press, pp. 47/1—47/14, 2017.

K. Butler & M. A. Stephens, "The distribution of a sum of independent Binomial random variables," *Methodology and Computing in Applied Probability*, vol. 19, pp. 557 – 571, 2017.

J. Cape, M. Tang & C. E. Priebe, "Signal-plus-noise matrix models: eigenvector deviations and fluctuations," *Biometrika*, vol. 106, pp. 243 – 250, 2019.

A. Charnes & W. W. Cooper, "Programming with linear fractional functionals," *Naval Research Logistics Quarterly*, vol. 9, pp. 181 – 186, 1962.

P.-Y. Chen & A. O. Hero, "Deep community detection," *IEEE Transactions on Signal Processing*, vol. 63, pp. 5706–5719, 2015.

S. Chen, R. Varma, A. Sandryhaila & J. Kovacevic, "Discrete signal processing on graphs: Sampling theory," *IEEE Transactions on Signal Processing*, vol. 63, pp. 6510–6523, 2015.

S. Chen, A. Sandryhaila, J. M. F. Moura & J. Kovacevic, "Signal recovery on graphs: Variation minimization," *IEEE Transactions on Signal Processing*, vol. 63, pp. 4609–4624, 2015.

K. T. K. Cheung & L. Hanzo, "Distributed energy spectral efficiency optimization for partial/full interference alignment in multi-user multi-relay multi-cell MIMO systems," *IEEE Transactions on Signal Processing*, vol. 64, pp. 882–896, 2016.

F. R. K. Chung, *Spectral graph theory*, Providence, R.I.: American Mathematical Society, 1997.

R. A. Clark & M. Macdonald, "Eigenvector-based community detection for identifying information hubs in neuronal networks," *BioRxiv (doi: 10.1101/457143)*, 2018.

R. F. Cohen, P. Eades, T. Lin & F. Ruskey, "Three-dimensional graph drawing," In: R. Tamassia & I. G. Tollis (eds.) *Graph drawing,* Berlin, Germany: Springer, pp. 1–11, 1995.

D. Cvetkovic, R. Rowlinson & S. Simic, *Eigenspaces of graphs*, Cambridge, UK: Cambridge University Press, 1997.

D. Cvetkovic & I. Gutman, "Applications of graph spectra: An introduction to the literature," *Zbornik Radova*, vol. 14, pp. 9–34, 2011.

M. M. Dankulov, B. Tadić & R. Melnik, "Spectral properties of hyperbolic nanonetworks with tunable aggregation of simplexes," *Physical Review E*, vol. 100, pp. 012309, 2019.

C. Davis & W. M. Kahan, "The rotation of eigenvectors by a perturbation. III," *SIAM Journal on Numerical Analysis*, vol. 7, pp. 1–46, 1970.

A. Decelle, F. Krzakala, C. Moore & L. Zdeborova, "Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications," *Physical Review E*, vol. 84, pp. 066106, 2011.

J. A. Deri & J. M. F. Moura, "Spectral projector-based graph Fourier transforms," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, pp. 785–795, 2017.

W. Dinkelbach, "On nonlinear fractional programming," *Management Science*, vol. 13, pp. 492–498, 1967.

W. E. Donath & A. J. Hoffman, "Lower bounds for the partitioning of graphs," *IBM Journal of Research and Development*, vol. 17, pp. 420–425, 1973.

X. Dong, D. Thanou, P. Frossard & P. Vandergheynst, "Learning Laplacian matrix in smooth graph signal presentations," *IEEE Transactions on Signal Processing* , vol. 64, pp. 6160–6173, 2016.

X. Dong, D. Thanou, M. Rabbat & P. Frossard, "Learning graphs from data: A signal representation perspective," *IEEE Signal Processing Magazine*, vol. 36, pp. 44–63, 2019.

A. Edelman, T. A. Arias & S. T. Smith, "The geometry of algorithms with orthogonality constraints," *SIAM Journal on Matrix Analysis and Applications*, vol. 20, pp. 303–353, 1998.

M. Eichler, "Testing nonparametric and semiparametric hypotheses in vector stationary processes," *Journal of Multivariate Analysis*, vol. 99, pp. 968–1009, 2008.

J. Eldridge, M. Belkin & Y. Wang, "Unperturbed: spectral analysis beyond Davis–Kahan," *Proceedings of Algorithmic Learning Theory*, vol. 83, pp. 321–358, 2018.

J. P. Fairbanks, D. A. Bader & G. D. Sanders, "Spectral partitioning with blends of eigenvectors," *Journal of Complex Networks*, vol. 5, pp. 551–580, 2017.

J. Fan, Y. Liao & M. Mincheva, "Large covariance estimation by thresholding principal orthogonal complements," *Journal of the Royal Statistical Society B*, vol. 75, pp. 603–680, 2013.

M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal*, vol. 23, pp. 298–305, 1973.

S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, pp. 75–174, 2010.

H. Frenk & S. Schaible, "Fractional programming," In: C. Floudas & P. Pardalos (eds.) *Encyclopedia of optimization,* Boston, MA: Springer, pp. 1080–1091, 2008.

C. Godsil & G. Royle, *Algebraic graph theory*, Springer, 2001.

A. Goldenberg, A. X. Zheng, S. E. Fienberg & E. M. Airoldi, "A survey of statistical network models," *Foundations and Trends in Machine Learning*, vol. 2, pp. 129–233, 2010.

M. Grant & S. Boyd, "Graph implementations for nonsmooth convex programs," In: V. Blondel, S. Boyd & H. Kimura (eds.) *Recent advances in learning and control,* Boston, MA: Springer, pp. 95–110, 2008.

M. Grant & S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, 2018, (Accessed: 2018-09).

W. H. Haemers, "Interlacing eigenvalues and graphs," *Linear Algebra and its Applications*, vol. 226-228, pp. 593 - 616, 1995.

D. K. Hammond, P. Vandergheynst & R. Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, pp. 129–150, 2011.

P. D. Hoff, A. E. Raftery & M. S. Handcock, "Latent space approaches to social network analysis," *Journal of the American Statistical Association*, vol. 97, pp. 1090–1098, 2002.

P. W. Holland, K. B. Laskey & S. Leinhardt, "Stochastic blockmodels: First steps," *Social Networks*, vol. 5, pp. 109–137, 1983.

S. Hoory, N. Linial & A. Wigderson, "Expander graphs and their applications," *Bulletin of the American Mathematical Society*, vol. 43, pp. 439–561, 2006.

R. A. Horn & C. R. Johnson, *Matrix analysis*, Cambridge, UK: Cambridge University Press, 1985.

R. A. Horn & C. R. Johnson, *Topics in matrix analysis*, Cambridge UK: Cambridge University Press, 1991.

T. Ibaraki, "Parametric approaches to fractional programs," *Mathematical Programming*, vol. 26, pp. 345 – 362, 1983.

C. Isheden, Z. Chong, E. Jorswieck & G. Fettweis, "Framework for link-level energy efficiency optimization with informed transmitter," *IEEE Transactions on Wireless Communications*, vol. 11, pp. 2946–2957, 2012.

C. R. Johnson & C. M. Saiago, *Eigenvalues, multiplicities and graphs*, Cambridge, UK: Cambridge University Press, 2018.

I. M. Johnstone, "On the distribution of the largest eigenvalue in principal components analysis," *The Annals of Statistics*, vol. 29, pp. 295–327, 2001.

S. Jung & J. S. Marron, "PCA consistency in high dimension low sample size context," *The Annals of Statistics*, vol. 37, pp. 4104–4130, 2009.

B. Karrer & M. E. J. Newman, "Stochastic blockmodels and community structure in networks," *Physical Review E*, vol. 83, pp. 016107, 2011.

L. Y. Kolotilina, "The stengthened versions of the additive and multiplicative Weyl inequalitites," *Journal of Mathematical Sciences*, vol. 127, pp. 1976–1987, 2005.

Y. Koren, "Drawing graphs by eigenvectors: theory and practice," *Computers & Mathematics with Applications*, vol. 49, pp. 1867–1888, 2005.

F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová & P. Zhang, "Spectral redemption in clustering sparse networks," *Proceedings of the National Academy of Sciences*, vol. 110, pp. 20935–20940, 2013.

S. Kumar, J. Ying, J. V. de M. Cardoso & D. P. Palomar, "A unified framework for structured graph learning via spectral constraints," *arXiv:1904.09792 [stat.ML]*, 2019.

E. Lagunas, A. G. Marques, S. Chatzinotas & B. Ottersten, "Graph similarity based on graph Fourier distances," *IEEE 2018 26th European Signal Processing Conference (EUSIPCO)*, pp. 877–881, 2018.

S. L. Lauritzen, *Graphical models*, Oxford, UK: Oxford University Press, 1996.

J. R. Lee, S. O. Gharan & L. Trevisan, "Multiway Spectral partitioning and higher-order Cheeger inequalities," *Journal of the ACM*, vol. 61, pp. 37, 2014.

J. Lei & A. Rinaldo, "Consistency of spectral clustering in stochastic block models," *The Annals of Statistics*, vol. 43, pp. 215–237, 2015.

J. Lei, "A goodness-of-fit test for stochastic block models," *The Annals of Statistics*, vol. 44, pp. 401–424, 2016.

Z. Li, H. Chen & B. Yang, "Reparametrized stochastic block model adaptive and heterogeneous degree block distributions," *IEEE Access*, vol. 6, pp. 37615–37626, 2018.

L. Lovász, *Large networks and graph limits*, Providence, R.I.: American Mathematical Society, 2012.

A. Lubotzky, "Expander graphs in pure and applied mathematics," *Bulletin ofmerican Mathematical Society*, vol. 49, pp. 113–162, 2012.

J. F. Lutzeyer & A. T. Walden, "Comparing graph spectra of Adjacency and Laplacian matrices," *arXiv:1712.03769 [stat.ME]*, 2017.

J. F. Lutzeyer & A. T. Walden, "Extending the Davis-Kahan theorem for comparing eigenvectors of two symmetric matrices I: Theory," *arXiv:1908.03462 [math.ST]*, 2019.

J. F. Lutzeyer & A. T. Walden, "Extending the Davis-Kahan theorem for comparing eigenvectors of two symmetric matrices II: Computation and applications," *arXiv:1908.03465 [math.ST]*, 2019.

V. Lyzinski, D. L. Sussman, M. Tang, A. Athreya & C. E. Priebe, "Perfect clustering for stochastic blockmodel graphs via adjacency spectral embedding," *Electronic Journal of Statistics*, vol. 8, pp. 2905–2922, 2014.

A. De Maio, Y. Huang, D. P. Palomar, S. Zhang & A. Farina, "Fractional QCQP with applications in ML steering direction estimation for radar detection," *IEEE Transactions on Signal Processing*, vol. 59, pp. 172–185, 2011.

G. Mateos, S. Segarra, A. G. Marques & A. Ribeiro, "Connecting the dots," *IEEE Signal Processing Magazine*, vol. 36, pp. 16–43, 2019.

C. D. Meyer, *Matrix analysis and applied linear algebra*, Philadelphia, PA: SIAM, 2000.

M. Molloy & B. Reed, "A critical point for random graphs with a given degree sequence," *Random Structures & Algorithms*, vol. 6, pp. 161–180, 1995.

S. K. Narang& A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *IEEE Transactions on Signal Processing*, vol. 60, pp. 2786–2799, 2012.

M. E. J. Newman, S. H. Strogatz & D. J. Watts, "Random graphs with arbitrary degree distributions and their applications," *Physical Review E*, vol. 64, pp. 026118, 2001.

M. E. J. Newman & M. Girvan, "Finding and evaluating community structure in networks," *Physical Review E*, vol. 69, pp. 026113, 2004.

M. E. J. Newman, *Networks*, Oxford, UK: Oxford University Press, 2018.

M. A. Nunes, M. I. Knight & G. P. Nason, "Modelling and prediction of time series arising on a graph," In: A. Antoniadis, J.-M. Poggi & X. Brossat (eds.) *In Modelling and Stochastic Learning for Forecasting in High Dimensions*, Cham, Switzerland: Springer, pp. 183–192, 2015.

S. C. Olhede & P. J. Wolfe, "Order statistics of observed network degrees," *arXiv:1210.4377 [stat.ME]*, 2012.

S. C. Olhede & P. J. Wolfe, "Degree-based network models," *arXiv:1211.6537 [stat.ST]*, 2013.

S. O'Rourke, V. Vu & K. Wang, "Random perturbation of low rank matrices: Improving classical bounds," *Linear Algebra and its Applications*, vol. 540, pp. 26–59, 2018.

A. Ortega, P. Frossard, J. Kovacevic, J. M. F. Moura & P. Vandergheynst, "Graph signal processing: Overview, challenges, and applications," *Proceedings of the IEEE*, vol. 106, pp. 808–828, 2018.

R. Paffenroth, P. C. Du Troit L. L. Scharf, & A. P. Jayasumana, "Distributed pattern detection in cyber networks," *Cyber Sensing 2012*, vol. 8408, pp. 84080J, 2012.

Y. Park, C. E. Priebe & A. Youssef, "Anomaly detection in time series of graphs using fusion of graph invariants," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, pp. 67–75, 2013.

B. Pincombe, "Anomaly detection in time series of graphs using ARMA processes," *Australian Society for Operation Research (ASOR) Bulletin*, vol. 24, pp. 2, 2005.

I. Pitas, *Graph-based social media analysis*, Boca Raton, FL: CRC Press, 2016.

G. Porto & L. E. Allem, "Eigenvalue interlacing in graphs," *Proceeding Series of the Brazilian Society of Applied and Computational Mathematics*, vol. 5, pp. 010232, 2017.

H. Qiu, F. Han, H. Liu & B. Caffo, "Joint estimation of multiple graphical models from high dimensional time series," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 78, pp. 487–504, 2016.

M. A. Riolo & M. E. J. Newman, "First-principles multiway spectral partitioning of graphs," *Journal of Complex Networks*, vol. 2, pp. 121–140, 2014.

K. Rohe, S. Chatterjee & B. Yu, "Spectral clustering and the high-dimensional stochastic blockmodel," *The Annals of Statistics*, vol. 39, pp. 1878–1915, 2011.

A. Sandryhaila & J. M. F. Moura "Discrete signal processing on graphs," *IEEE Transactions on Signal Processing*, vol. 61, pp. 1644–1656, 2013.

A. Sandryhaila & J. M. F. Moura "Big data analysis with signal processing on graphs: Representation and processing of massive data sets with irregular structure," *IEEE Signal Processing Magazine*, vol. 31, pp. 80–90, 2014.

S. Schaible, "Parameter-free convex equivalent and dual programs of fractional programming problems," *Zeitschrift für Operations Research*, vol. 18, pp. 187–196, 1974.

S. Schaible, "Minimization of ratios," *Journal of Optimization Theory and Applications*, vol. 19, pp. 347–352, 1976.

S. Schaible, "Fractional programming: Applications and algorithms," *European Journal of Operations Research*, vol. 7, pp. 111–120, 1981.

S. Schaible, "Fractional programming," *Zeitschrift für Operations Research*, vol. 27, pp. 39–54, 1983.

S. Sengupta & Y. Chen, "A block model for node popularity in networks with community structure," *J. R. Statist. Soc. B*, vol. 80, pp. 365–386, 2018.

X. Shen, S. Diamond, Y. Gu & S. Boyd, "Disciplined convex-concave programming," *arXiv:1604.02639 [math.OC]*, 2016.

D. Shen, H. Shen & J. S. Marron, "A general framework for consistency of principal component analysis," *Journal of Machine Learning Research*, vol. 17, pp. 1–34, 2016.

K. Shen & W. Yu, "Fractional programming for communication systems—Part I: Power control and beamforming," *IEEE Transactions on Signal Processing*, vol. 66, pp. 2616–2630, 2018.

D. Shuman, S. K. Narang, P. Frossard, A. Ortega & P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Processing Magazine*, vol. 30, pp. 83–98, 2013.

R. Singh, A. Chakraborty & B. S. Manoj, "On Spectral Analysis of Node Centrality," *2016 IEEE International Conference on Advanced Networks and Telecommunications Systems (ANTS)*, pp. 1–5, 2016.

T. A. B. Snijders, P. E. Pattinson, G. L. Robins & M. S. Handcock, "New specifications for exponential random graph models," *Sociological Methodology*, vol. 36, pp. 99–153, 2006.

W. So, "Commutativity and spectra of Hermitian matrices," *Linear Algebra and its Applications*, vol. 212, pp. 121-129, 1994.

D. A. Spielman, "Spectral graph theory and its applications," *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pp. 29-38, 2007.

L. Spricer & T. Britton, "The configuration model for partially directed graphs," *Journal of Statistical Physics*, vol. 161, pp. 965-985, 2015.

G. W. Stewart & J. Sun, *Matrix perturbation theory*, San Diego, CA: Academic Press Inc., 1990.

Z. Stanić, *Inequalities for graph eigenvalues*, Cambridge, UK: Cambridge University Press, 2015.

M. Tang & C. E. Priebe, "Limit theorems for eigenvectors of the normalized Laplacian for random graphs," *The Annals of Statistics*, vol. 46, pp. 2360–2415, 2018.

O. Teke & P. P. Vaidyanathan, "Uncertainty principles and sparse eigenvectors of graphs," *IEEE Transactions on Signal Processing*, vol. 65, pp. 5406–5420, 2017.

N. Tremblay & P. Borgnat, "Graph wavelets for multiscale community mining," *IEEE Signal Processing Magazine*, vol. 62, pp. 5227–5239, 2014.

P. van Mieghem, *Graph spectra for complex networks*, Cambridge, UK: Cambridge University Press, 2011.

U. von Luxburg, "A tutorial on spectral clustering," *Statistics and Computing*, vol. 17, pp. 395–416, 2007.

U. von Luxburg, M. Belkin & O. Bousquet, "Consistency of spectral clustering," *The Annals of Statistics*, vol. 36, pp. 555–586, 2008.

V. Q. Vu & J. Lei, "Minimax sparse principal subspace estimation in high dimensions," *The Annals of Statistics*, vol. 41, pp. 2905–2947, 2013.

H. Wang, M. Tang, C. Priebe & Y. Park, "Inference in time series of graphs using locality statistics," *2013 IEEE Global Conference on Signal and Information Processing*, pp. 471–474, 2013.

W. Wang & J. Fan, "Asymptotics of empirical eigenstructure for high dimensional spiked covariance," *The Annals of Statistics*, vol. 45, pp. 1342–1374, 2017.

J. Wang, V. D. Calhoun, J. M. Stephen, T. W. Wilson & Y. Wang, "Integration of network topological features and graph Fourier transform for fMRI data analysis," *ISBI 2018*, pp. 92–96, 2018.

P.-A. Wedin, "Perturbation bounds in connection with singular value decomposition," *BIT Numerical Mathematics*, vol. 12, pp. 99–111, 1972.

Y. Ye, Y. Li, Z. Wang, X. Chu & H. Zhang, "Dynamic asymmetric power splitting scheme for SWIPT-based two-way multiplicative AF relaying," *IEEE Signal Processing Letters*, vol. 25, pp. 1014–1018, 2018.

Y. Yu, T. Wang & R. J. Samworth, "A useful variant of the Davis–Kahan theorem for statisticians," *Biometrika*, vol. 102, pp. 315–323, 2015.

W. W. Zachary, "An information flow model for conflict and fission in small groups," *Journal of Anthropological Research*, vol. 33, pp. 452–473, 1977.

H. Zenil, N. A. Kiani & J. Tegnér, "Quantifying loss of information in network-based dimensionality reduction techniques," *Journal of Complex Networks*, vol. 4, pp. 342–362, 2016.

P. Zumstein, *Comparison of spectral methods through the Adjacency matrix and the Laplacian of a graph*, PhD thesis, ETH Zürich, 2005.

A

Appendix for Chapter 1

A.1 A collection of real data sets

In Table A.1 we show several datasets discovered in our initial survey of the graphs and networks literature together with their properties. These properties are whether the graph is undirected, the dimensions of the dataset, whether nodes in the data set have geographic locations and a few references using the datasets. The survey shown in Table A.1 formed the basis of the decision to use the karate dataset as our literature standard, real dataset throughout the thesis: it clearly acts as a common example for clustering algorithms. In Table A.1 we left cells empty, where the answer is unknown to us.

In Table A.2 we show the datasets, which were used in more than one publication we surveyed. In Table A.2 it can be seen which of these more commonly used datasets often co-occur.

A.1. A collection of real data sets

Table A.1: Commonly used graph data sets and their properties.

Name	Undirected	Dimensions	Geo. loc.s	References and Comments
karate club (UCI)	Yes	34 nodes	No	Karrer and Newman (2011) Chen and Hero (2015) Bickel and Sarkar (2016) Fortunato (2010) clusters known
Sampson's monastery (UCI)	No	4 networks, 18 nodes	No	Hoff et al. (2002) Goldenberg et al. (2010) clusters known
Florentine families (UCI)	Yes	16 nodes	No	Hoff et al. (2002)
political blogs	No	1490 nodes	No	Chen et al. (2015) Karrer and Newman (2011) Chen et al. (2015) clusters known
Abilene Internet 2	Yes	25 networks, 28 nodes	Yes	Paffenroth et al. (2012)
Enron email	No	158 nodes	No	Goldenberg et al. (2010) Park et al. (2013) Wang et al. (2013)
acceleration signals	No	930 acceleration signals	No	Chen et al. (2015) 8-nearest neighbour graph
temperature data	Yes	150 nodes	Yes	Chen et al. (2015) Chen et al. (2015) Sandryhaila and Moura (2014) 8-nearest neighbour graph
Minnesota roads	Yes	2642 nodes	Yes	Shuman et al. (2013) Hammond et al. (2011) Narang and Ortega (2012)
Jester dataset 1	Yes	73421 nodes	No	Chen et al. (2015) 8-nearest neighbour graph
social interactions primary school	Yes	242 nodes	No	Tremblay and Borgnat (2014)
classroom data	Yes	27 nodes	No	Hoff et al. (2002)
dolphin social network	Yes	62 nodes	No	Chen and Hero (2015) Fortunato (2010) clusters known
facebook ego models 2012	Yes	9 networks		Bickel and Sarkar (2016) 1-hop friendship graph social cycles known
political book Newmann 2006	Yes	105 nodes	No	Bickel and Sarkar (2016)
mobile phone communications network	Yes	2600000 nodes		Fortunato (2010) many hierarchies in clusters
NIPS paper co-authorship	Yes	2037 nodes	No	Goldenberg et al. (2010)
1,2-hop authorship graphs	Yes	variable	No	Chen and Hero (2015)
Cerebral Cortex MRI	Yes	998 nodes	Yes	Shuman et al. (2013) Hammond et al. (2011)
ADHD-200 fMRI	Yes	973 networks, 264 nodes	Yes	Qiu et al. (2016)
Add Health HIV study	Yes	90118 initial participants	Yes	Goldenberg et al. (2010)
Framingham obesity study	Yes	5209 nodes		Goldenberg et al. (2010)
UK weather stations	Yes	102 nodes	Yes	Nunes et al. (2015) minimal spanning tree
protein interaction network in budding yeast	Yes	about 6000 nodes	No	Goldenberg et al. (2010)
protein-protein interaction network of rat proteome	Yes	313 nodes	No	Fortunato (2010)

Table A.2: Commonly used graph data sets and a subset of the publications using them as examples.

	karate club	online blogs	Enron email	temperature data	Minnesota roads	dolphin social network	Cerebral Cortex MRI	Sampson's monastery
Goldenberg et al. (2010)			✓					✓
Chen et al. (2015)		✓		✓				
Chen et al. (2015)		✓		✓				
Sandryhaila and Moura (2014)				✓				
Karrer and Newman (2011)	✓	✓						
Shuman et al. (2013)					✓		✓	
Hammond et al. (2011)					✓		✓	
Narang and Ortega (2012)					✓			
Chen and Hero (2015)	✓					✓		
Bickel and Sarkar (2016)	✓							
Fortunato (2010)	✓					✓		
Park et al. (2013)			✓					
Wang et al. (2013)			✓					
Hoff et al. (2002)								✓

A.2 Spectral support of the representation matrices

In this appendix we derive the spectral support of the representation matrices using the Gershgorin circle theorem.

Theorem A.1. *Gershgorin circle theorem* (Bernstein, 2009, p. 293) Let $B \in \mathbb{C}^{n\times n}$ with $(j,k)^{\text{th}}$-element b_{jk} and eigenvalue set $\{\beta_1, \ldots, \beta_n\}$. Furthermore, let $R_j^{(B)} = \sum_{\substack{k=1 \\ k\neq j}}^{n} |b_{jk}|$ be the sum of absolute values of the off-diagonal elements in row j. Then, for $i \in \{1, \ldots, n\}$

$$\beta_i \in \bigcup_{j=1}^{n} \left\{ s \in \mathbb{C} : |s - b_{jj}| \leq R_j^{(B)} \right\}. \qquad \triangleleft$$

Furthermore, we have that the eigenvalues of a real symmetric matrix are real. Therefore the union of discs in which the eigenvalues of the representation matrices are contained according to Theorem A.1 is restricted to intervals on the real line.

For the *adjacency matrix* A we have that $a_{jj} = 0$ for $j \in \{1, \ldots, n\}$. Furthermore $R_j^{(A)} = \sum_{\substack{k=1 \\ k\neq j}}^{n} |a_{jk}| \leq d_{\max}$ for $j \in \{1, \ldots, n\}$ following from the definition of $d_{\max}$. By applying the Gershgorin theorem we find that for $i \in \{1, \ldots, n\}$

$$\mu_i \in [-d_{\max}, d_{\max}].$$

The *unnormalised Laplacian* $L = D - A$ has diagonal entries $l_{jj} = \sum_{k=1}^{n} a_{jk} \leq d_{\max}$ for $j \in \{1, \ldots, n\}$. Furthermore $R_j^{(L)} = \sum_{\substack{k=1 \\ k \neq j}}^{n} |l_{jk}| = \sum_{k=1}^{n} a_{jk} \leq d_{\max}$ for $j \in \{1, \ldots, n\}$ by using the fact that the diagonal elements of A are equal to 0. We hence have that $l_{jj} = R_j^{(L)} \leq d_{\max}$. By applying the Gershgorin theorem we find that for $i \in \{1, \ldots, n\}$

$$\lambda_i \in [0, 2d_{\max}].$$

Finally the *normalised Laplacian* $L_{rw} = I - D^{-1}A$ has diagonal entries $(L_{rw})_{jj} = 1$ for $j \in \{1, \ldots, n\}$ using that fact that the adjacency matrix has a zero diagonal. Furthermore $R_j^{(L_{rw})} = \sum_{\substack{k=1 \\ k \neq j}}^{n} |(L_{rw})_{jk}| = \sum_{\substack{k=1 \\ k \neq j}}^{n} |\frac{a_{jk}}{d_j}| = \frac{1}{d_j} \sum_{\substack{k=1 \\ k \neq j}}^{n} |a_{jk}| = 1$ for $j \in \{1, \ldots, n\}$ by definition of the degree matrix. We hence have that $(L_{rw})_{jj} = R_j^{(L_{rw})} = 1$. By applying the Gershgorin theorem we find that for $i \in \{1, \ldots, n\}$

$$\eta_i \in [0, 2].$$

A.3 Smooth eigenvectors and their zero crossings

In this appendix we want to draw the readers attention to the issue of eigenvector smoothness, which we encountered when reproducing Fig. 1 in von Luxburg (2007). A smooth eigenvector plot is produced by choosing a specific ordering of eigenvector elements and manipulating the distance of the elements on the x-axis. We explain the link between these plots and the notion of frequencies corresponding to eigenvalues. Furthermore, we assess the impact of these findings on the inference drawn from the eigenvectors in our work.

von Luxburg (2007) simulates graphs by sampling from a mixture of four normal distributions with means 2, 4, 6 and 8. We choose the standard deviation of all four components to be equal to 0.2 by inspection of the sample histogram in Fig. 1 in von Luxburg (2007). We proceed to sample 200 points from this mixture. A histogram of our sample is shown in Figure A.1.

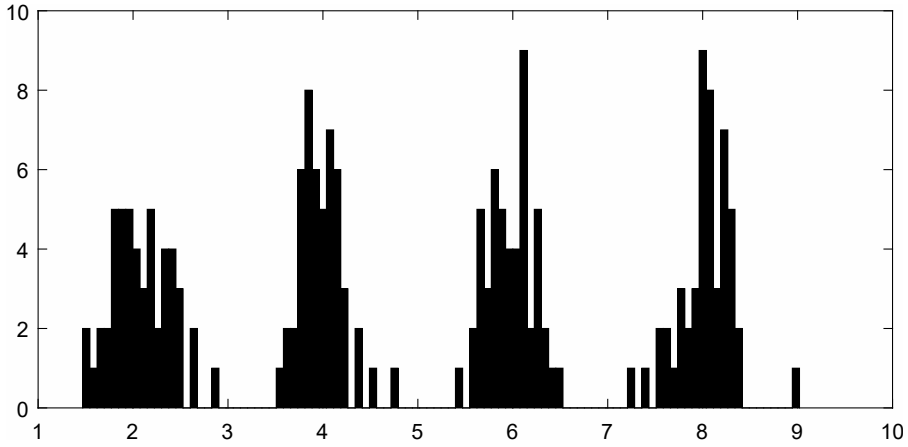

Figure A.1: Histogram of the data sample.

As in von Luxburg (2007) we produce a 10-nearest neighbours graph and a fully connected graph, which were both introduced in Section 1.2.3, from the sampled data. In both graphs we weight the edges using the Gaussian similarity function $s(x_i, x_j) = \exp(-(x_i - x_j)^2/2)$, where x_i and x_j are sampled data values from the Gaussian mixture distribution. We then calculate the unnormalised Laplacian L and the normalised Laplacian L_{rw} together with their spectra and eigenvectors. In Figure A.2 we plot the eigenvector values against their corresponding indices, as one usually would. The eigenvectors appear noisy and do not match the ones displayed in von Luxburg (2007). Based on these plots it is difficult to decide whether there is any signal in these eigenvectors. For the 10-nearest neighbour graph, displayed in the first two rows of Figure A.2, we find that the first four eigenvectors are all either strictly non-negative or non-positive.

von Luxburg (2007) states that she is plotting eigenvector elements against the sampled data values they correspond to. This represents a reordering of the eigenvector values and hence a relabelling of the graph vertices. In calculation it can quickly be shown that permuting any graph representation matrix, e.g. A, and its eigenvectors κ by the same permutation matrix P results in valid

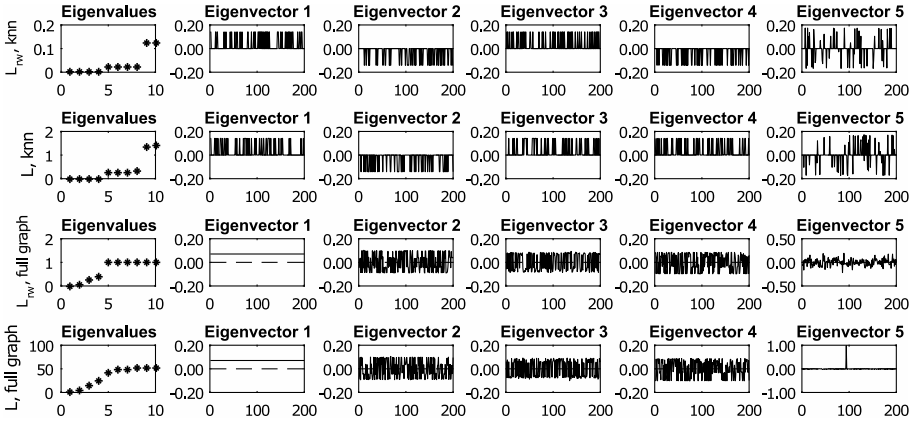

Figure A.2: The eigenvectors are plotted in the order in which they were sampled and are plotted against their corresponding indices $\{1, 2, \cdots, 200\}$. The format of the figure is chosen to match Fig. 1 in von Luxburg (2007).

permuted eigenvectors corresponding to the permuted representation matrix.

$$\begin{aligned} A\kappa &= \mu\kappa \\ PA\kappa &= \mu P\kappa \\ PAP^T P\kappa &= \mu P\kappa \end{aligned}$$

In Figure A.3 we plot the eigenvectors reordered according to the magnitude of their corresponding sample values, i.e., we plot $P\kappa$ instead of κ as was the case in Figure A.2.

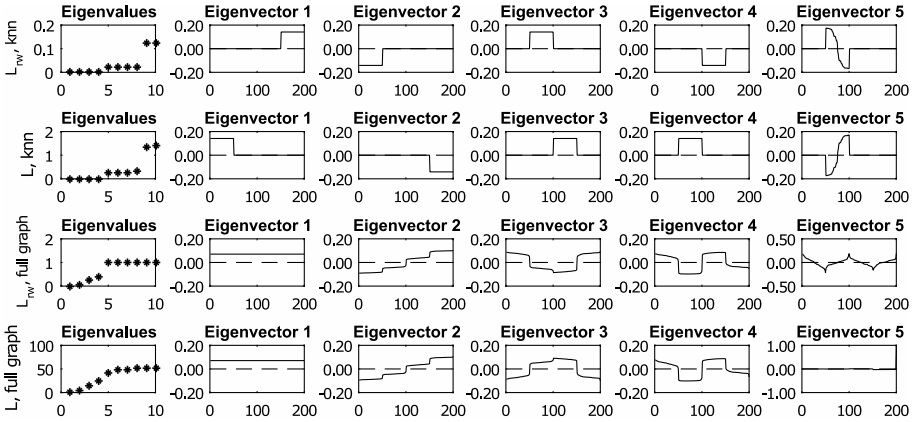

Figure A.3: The eigenvectors are reordered according to the magnitude of the sampled data value they correspond to. These reordered eigenvectors are plotted against the indices $\{1, 2, \cdots, 200\}$. The format of the figure is chosen to match Fig. 1 in von Luxburg (2007).

From Figure A.3 it is clearly visible that the eigenvectors have a structure corresponding to the magnitude of the sampled data. However, the eigenvectors still don't match the ones displayed in Figure 1 in von Luxburg (2007). The eigenvectors corresponding to the 10-nearest neighbour graph appear, as the theory suggests (List 1.22(3)), as indicator functions of the four connected components.

In addition to reordering the eigenvector elements the x-axis indices are also remapped in von Luxburg's plot. The map applied to the indices is $i \to x_i$, where x_i is the i^{th} sample point. The distance on the x-axis between any two eigenvector values will be equal to the distance of their corresponding sampled data points. In Figure A.4 we plot the eigenvector values against their corresponding data values, which results in eigenvector plots very similar to the ones in von Luxburg (2007).

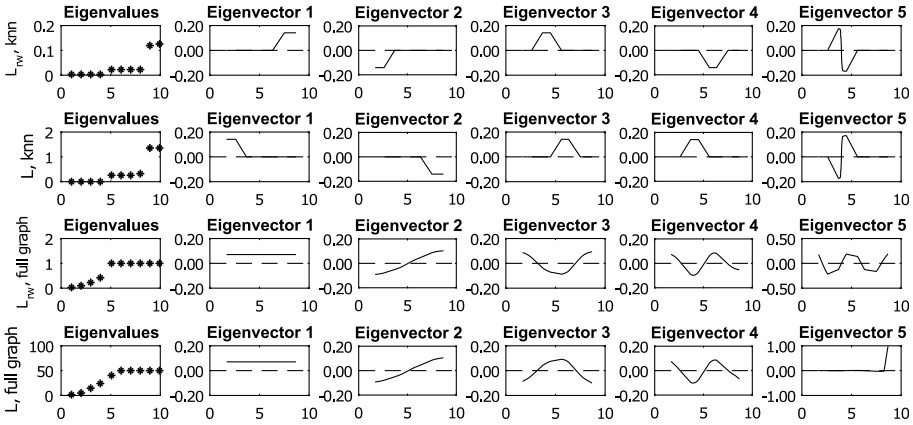

Figure A.4: The eigenvectors are reordered according to the magnitude of the sampled data value they correspond to. These reordered eigenvectors are plotted against their corresponding sample value. The format of the figure is chosen to match Fig. 1 in von Luxburg (2007).

In Figure A.4 the eigenvectors corresponding to the 10-nearest neighbour graph do not appear to be equal to indicator functions. This is a sampling issue. Since we have not sampled any data points in between the connected components, plotting the eigenvectors against the sampled data, results in a linear interpolation of the eigenvector values in unsampled regions. For the 10-nearest neighbour graph the plots in Figure A.3 gave better insight into the eigenvector structure. Contrariwise, the first four, smooth eigenvectors corresponding to the fully connected graph in Figure A.4 give better insight into the

graph structure. They resemble polynomials of increasing degree, where the number of zero crossings of the polynomial corresponding to eigenvector v_i is equal to $i - 1$. We have already met the concept of eigenvalues corresponding to certain signal frequencies in the eigenvectors across graphs in Section 1.4.3 and in the work of Shuman et al. (2013). This frequency argument explains why the ordering of the vertices according to their sampled data values produces such meaningful eigenvectors. In both the 10-nearest neighbour and the fully connected graph the strength, and in the nearest neighbour graph even the presence, of edges is dependent only on the difference in magnitude of the sampled values. *Hence, the ordering according to sample magnitudes represents a sensible embedding of the graph structure on a one dimensional axis.*

At first it seems surprising and worrying that the labelling of vertices has such an impact on the appearance of the eigenvectors. However, the labelling does not affect the outcome of spectral clustering and most other spectral algorithms. The eigenvector elements remain unchanged when their ordering is switched and hence the class assigned by a k-means algorithm run on the eigenvector elements remains unchanged across relabelling of the vertices.

We conclude that when generating graphs in this way there exists a sensible ordering of vertices, which gives the eigenvectors a smooth appearance. The result of inference on the eigenvectors remains unchanged by the relabelling.

A.4 The existence of a real eigenbasis of L_{rw}

When calculating the spectrum and eigenvectors of L_{rw} one can encounter complex values. This is surprising since L_{rw} is similar to the real and symmetric L_{sym}. To explain our surprise and investigate this issue further will be the topic of this appendix. We suggest several ways of obtaining the eigenvectors of L_{rw}, deduce a result on the existence of real eigenvectors for matrices in equivalence classes defined by similarity and discuss our findings in the context of the work by von Luxburg (2007).

Eigenvalues of Hermitian matrices are real Horn and Johnson (1985, p. 38). Since L_{sym} is Hermitian we know that its eigenvalues are real and so are the eigenvalues of L_{rw}, via their similarity relationship. Further, if $B \in \mathbb{C}^{n\times n}$ is real then there exists a real eigenvector associated to each real eigenvalue. This can be confirmed by noticing that real and imaginary parts of any complex eigenvector are themselves eigenvectors of B.

Therefore, there exists a set of real eigenvectors corresponding to L_{sym}. Since $D^{1/2}$ has real entries we find - via the similarity relationship - that L_{rw} also has a set of real eigenvectors corresponding to its eigenvalues.

In practice, if real eigenvectors of L_{rw} are obtained via the similarity relationship or by taking the real part of eigenvectors, obtained via the inbuilt MATLAB algorithm, then it is important to rescale the eigenvectors to have length 1. If the corresponding eigenvalue has geometric multiplicity 1, i.e., its corresponding eigenspace has dimension 1, then the rescaled eigenvectors of length 1 are all equal up to sign.

In our computations we found cases in which the eigenvalues of L_{rw} had very small imaginary parts of order 10^{-15}. We now know that the spectrum of L_{rw} must be real and hence we were dealing with an imaginary error in the real eigenvalues. We now draw attention to a fact highlighted on Horn and Johnson (1985, p. 37).

Remark A.2. A non-real complex eigenvalue of a matrix with real entries cannot have a real eigenvector associated to it. ◁

As is to be anticipated from Remark A.2, we got complex eigenvectors with much larger complex parts of order 10^{-2} associated with the complex eigenvalues. This initially confronted us with the problem of finding complex equivalents to the algorithms we were running on the eigenvectors. It is simpler to work with the real eigenvectors of L_{rw}, which we now know always exist due to the similarity relation between the two normalised Laplacians.

The argument made here can be extended to all matrices related via a similarity transform.

Proposition A.3. If a matrix B, which is not necessarily Hermitian, is similar to a Hermitian matrix Φ via a real similarity matrix, then there exists a set

of real eigenvectors for both matrices corresponding to their necessarily real eigenvalues. ◁

Horn and Johnson (1985) state that similarity is an equivalence relation on $\mathbb{C}^{n\times n}$, partitioning $\mathbb{C}^{n\times n}$ into disjoint equivalence classes. Hence, Proposition A.3 applies to all matrices in one of these equivalence classes containing one or more Hermitian matrix or matrices. This is a powerful result of practical relevance.

von Luxburg (2007) makes an argument for spectral clustering according to L_{rw} instead of L_{sym}. In the ideal clustering case of k connected components she finds the eigenvectors of L_{rw} to be indicator functions of the connected components. In this case the eigenvectors of L_{sym} are the indicator functions premultiplied by $D^{1/2}$, due to the similarity. If the graph contains vertices of very small degrees this might cause misclustering even in the ideal case of several connected components. Von Luxburg states that using L_{sym} does not have any numerical advantages over L_{rw}. Contrariwise, since L_{rw} is not Hermitian a complex error in the spectrum and the eigenvectors cannot be ruled out in computation. This does not make her argument invalid, but should be considered when choosing which of the two normalised Laplacians to work with.

B

Appendix for Chapter 2

B.1 Polynomial relation of the representation matrix spectra

In this appendix section we show how most matrix spectra can theoretically be exactly related via polynomial transforms.

The issue of mapping one set of representation matrix eigenvalues to another can be seen as a polynomial interpolation problem on the two sets of eigenvalues. Horn and Johnson (1991, p. 390) state that as long as there exists a function which maps one set of input points to the set of output points, the interpolation problem always has a unique solution and they also provide formulas to find this solution. Therefore, we find that as long as equal eigenvalues in the domain are not required to be mapped to unequal eigenvalues in the range, there exists a unique polynomial which maps one representation matrix spectrum to another.

In practice, the operation of finding the interpolating polynomial can be numerically unstable and therefore while the polynomial mapping from one spectrum to another theoretically exists, it cannot be readily obtained. For example, for the graph of the real karate data set the interpolating polynomial maps eigen-

values $\mu_{13}, \ldots, \mu_{22}$, which have a pairwise difference smaller than 10^{-14} onto eigenvalues of the unnormalised Laplacian which span an interval of width greater than 2. Obviously, finding the interpolating polynomial cannot be done without significant numerical error.

Furthermore, is it unclear how one would use the polynomial map between spectra for graphical inference. Relating the polynomial coefficients to vertex degrees could lead to great insight, but is hard to achieve. Finally, we are unaware of any results which would allow calculation of the transformation without calculating both representation spectra first, while our proposed affine transformation parameters only rely on the degree extremes which are readily available from the graph. Hence, while in theory, in some cases, there exists a more precise relation of the representation matrix spectra, our simple affine transformations have significant advantages.

B.2 Minimising the bound on the spectra of A and L

In Figure B.1 we illustrate the term $\max(|d_{\max} - c_0(A, L)|, |d_{\min} - c_0(A, L)|)$ as a function of $c_0(A, L)$.

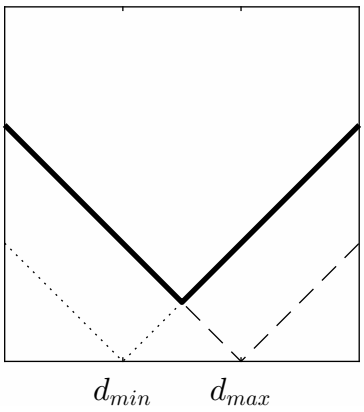

Figure B.1: This plot displays values of $c_0(A, L)$ on the x-axis. The dotted line corresponds to the function $|d_{\min} - c_0(A, L)|$, the dashed line plots the function $|d_{\max} - c_0(A, L)|$ and the solid line depicts the bound which is to be minimised $\max(|d_{\max} - c_0(A, L)|, |d_{\min} - c_0(A, L)|)$.

From Figure B.1 it is obvious that the bound is analytically minimised at the

intersection of the increasing part of the dotted line, $|d_{\min} - c_0(A, L)|$, and the decreasing part of the dashed line, $|d_{\max} - c_0(A, L)|$, which is calculated now.

$$\begin{aligned} c_0(A, L) - d_{\min} &= d_{\max} - c_0(A, L) \\ \Leftrightarrow c_0(A, L) &= \frac{d_{\max} + d_{\min}}{2}. \end{aligned}$$

B.3 Weyl's inequality and the bound of (2.1)

If we choose $c_1(A, L) = -1$ then $f_1(A)$ can be expressed in terms of L as follows,

$$f_1(A) = -A + c_0(A, L)I = L - D + c_0(A, L)I.$$

We denote the error between $f_1(A)$ and L by $\Omega = c_0(A, L)I - D$, and the error eigenvalues as follows: $\lambda_1(\Omega) \geq \lambda_2(\Omega) \geq \ldots \geq \lambda_n(\Omega)$. From Weyl's inequality in Theorem 1.45 with the substitutions, $\Phi = L$, $\Psi = \Omega$, $j = 1$ in (1.9) and $j = n$ in (1.10), it follows that for $i \in \{1, 2, \ldots, n\}$,

$$\lambda_i + \lambda_n(\Omega) \leq f_1(\mu_i) \leq \lambda_i + \lambda_1(\Omega).$$

So, in particular,

$$\begin{aligned} \lambda_n(\Omega) \leq f_1(\mu_i) - \lambda_i \leq \lambda_1(\Omega). \\ \Rightarrow c_0(A, L) - d_{\max} \leq f_1(\mu_i) - \lambda_i \leq c_0(A, L) - d_{\min}. \end{aligned}$$

Now, set $c_0(A, L) = (d_{\max} + d_{\min})/2$ to obtain the result of Theorem 2.2, which is a bound on the absolute difference.

This result extends the application of Weyl's inequality to the two representation matrix spectra, as is done in van Mieghem (2011, p. 71), by transforming one of the two spectra.

B.4 Minimising the bound on the spectra of L and L_{rw}

In Figure B.2 we illustrate the term $\max(|d_{\max}c_1(L, L_{rw}) - 1|, |d_{\min}c_1(L, L_{rw}) - 1|)$ as a function of $c_1(L, L_{rw})$.

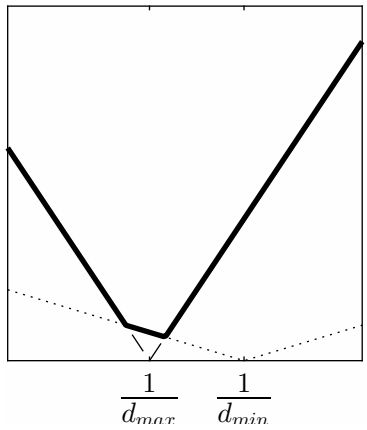

Figure B.2: This plot displays values of $c_1(L, L_{rw})$ on the x-axis. The dotted line corresponds to the function $|d_{\min}c_1(L, L_{rw}) - 1|$, the dashed line plots the function $|d_{\max}c_1(L, L_{rw}) - 1|$ and the solid line depicts the bound which is to be minimised $\max(|d_{\max}c_1(L, L_{rw}) - 1|, |d_{\min}c_1(L, L_{rw}) - 1|)$.

As is obvious from the illustration the bound is minimised at the intersection of the increasing part of the dashed line, $|d_{\max}c_1(L, L_{rw}) - 1|$, and decreasing part of the dotted line, $|d_{\min}c_1(L, L_{rw}) - 1|$. The location of this point will now be calculated.

$$\begin{aligned} d_{\max}c_1(L, L_{rw}) - 1 &= -\left(d_{\min}c_1(L, L_{rw}) - 1\right) \\ \left(d_{\max} + d_{\min}\right) c_1(L, L_{rw}) &= 2 \\ c_1(L, L_{rw}) &= \frac{2}{d_{\max} + d_{\min}}. \end{aligned}$$

C

Appendix for Chapter 3

C.1 Existence of Q — Proof of Lemma 3.5

Proof. Suppose $U, W \in \mathbb{V}_{n,r}$ such that the spaces spanned by their columns, $\mathcal{U}, \mathcal{W} \in \mathbb{G}_{n,r}$, coincide, i.e., $\mathcal{U} = \mathcal{W}$. We want to show that there exists an orthogonal matrix $Q \in O(r)$ such that $U = WQ$.

We begin by extending the bases of $\mathcal{U}$ formed by the columns of U and W to be bases of $\mathbb{R}^n$ by adding orthogonal columns $U_\perp, W_\perp \in \mathbb{V}_{n,n-r}$ using the Gram-Schmidt algorithm. By joining the columns we obtain the orthogonal matrices $U_n = [U, U_\perp]$ and $W_n = [W, W_\perp]$.

Consider $Q_n = W_n^T U_n$. Note, $W_n Q_n = W_n W_n^T U_n = U_n$. From $\mathcal{U} = \mathcal{W}$ it follows that their orthogonal complement spaces coincide, i.e., $\mathcal{U}^\perp = \mathcal{W}^\perp$. Since U and W span $\mathcal{U}$ and $U_\perp$ and $W_\perp$ span $\mathcal{U}^\perp$, we have by the definition of orthogonal complement spaces that $U^T W_\perp = W^T U_\perp = \mathbf{0}_{r,n-r}$, the matrix of all zeros. Therefore, Q_n is block diagonal,

$$\begin{aligned} Q_n = W_n^T U_n &= \begin{bmatrix} W^T \\ W_\perp^T \end{bmatrix} \begin{bmatrix} U & U_\perp \end{bmatrix} = \begin{bmatrix} W^T U & W^T U_\perp \\ W_\perp^T U & W_\perp^T U_\perp \end{bmatrix} \\ &= \begin{bmatrix} W^T U & \mathbf{0}_{r,n-r} \\ \mathbf{0}_{n-r,r} & W_\perp^T U_\perp \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} Q & \mathbf{0}_{r,n-r} \\ \mathbf{0}_{n-r,r} & Q_\perp \end{bmatrix}. \end{aligned}$$

Since U_n and W_n are orthogonal, Q_n is also orthogonal. Together the block diagonality and the orthogonality of Q_n imply orthogonality of Q. Finally, when only observing the top left $r \times r$ block of the equation $U_n = W_n Q_n$, we obtain $WW^T U = U$. Therefore, we have constructively shown the existence of an orthogonal matrix Q which satisfies $U = WQ$.

□

D

Appendix for Chapter 4

D.1 Proof of part 2 of Proposition 4.2

In this appendix we first work out the limit of bound (4.9) as either $c_0 \to \infty$ or $c_0 \to -\infty$ for comparisons with $j = 0$. Then we draw parallels to the comparisons involving the last r eigenvectors, i.e., the case where $j = n - r$.

We begin by producing a lower and an upper bound for the cost function in (4.9). Using the matrix triangle and reverse-triangle inequalities $\big|||\Phi|| - ||\Psi|||\big| \leq ||\Phi - \Psi||$ and $||\Phi + \Psi|| \leq ||\Phi|| + ||\Psi||$ on the numerator of the cost function, gives

$$\frac{||c_0| \, \|I\|_2 - \|-c_1\Phi + \Psi\|_2|}{\delta} \leq \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta} \leq \frac{|c_0| \, \|I\|_2 + \|c_1\Phi - \Psi\|_2}{\delta}. \tag{D.1}$$

But $\|I\|_2 = 1$ so henceforth this term will be omitted. Here δ is any of $\delta_{1,+}, \delta_{1,-}, \delta_{2,+}$ and $\delta_{2,-}$ corresponding to the four different subproblems which need to be solved in the affine case. For $j = 0$, we encounter the issue that ϕ_0 and ψ_0 are undefined and therefore, as is done in Yu et al. (2015, p. 317), we set them equal to $-\infty$. Then, (3.27), (3.28), (3.29) and (3.30) take the

following form,

$$\delta_{1,+} = \psi_{r+1} - c_1\phi_r - c_0, \tag{D.2}$$
$$\delta_{1,-} = \psi_{r+1} - c_1\phi_1 - c_0, \tag{D.3}$$
$$\delta_{2,+} = c_1\phi_{r+1} + c_0 - \psi_r, \tag{D.4}$$
$$\delta_{2,-} = \psi_1 - c_1\phi_{r+1} - c_0. \tag{D.5}$$

In order to work out the limit of the lower and upper bound in (D.1) as either $c_0 \to \infty$ or $c_0 \to -\infty$ we use l'Hopital's rule. For l'Hopital's rule to apply we require both the numerator and the denominator of the lower and upper bound in (D.1) to tend to 0 or ∞ as either $c_0 \to \infty$ or $c_0 \to -\infty$. We find that both numerators and their common denominators do indeed tend to ∞ as either $c_0 \to \infty$ or $c_0 \to -\infty$, i.e.,

$$\lim_{c_0 \to \pm\infty} ||c_0| - \|-c_1\Phi + \Psi\|_2| = \infty,$$
$$\lim_{c_0 \to \pm\infty} |c_0| + \|c_1\Phi - \Psi\|_2 = \infty,$$
$$\lim_{c_0 \to -\infty} \delta_{1,+} = \infty, \quad \lim_{c_0 \to -\infty} \delta_{1,-} = \infty, \quad \lim_{c_0 \to \infty} \delta_{2,+} = \infty, \quad \lim_{c_0 \to -\infty} \delta_{2,-} = \infty.$$

We start by evaluating the derivative with respect to c_0 of the numerator of both the lower and upper bound in (D.1) and of their denominators.

Consider firstly,

$$\frac{\partial}{\partial c_0} ||c_0| - \|c_1\Phi - \Psi\|_2| = \operatorname{sign}\left(|c_0| - \|c_1\Phi - \Psi\|_2\right) \operatorname{sign}(c_0) =$$
$$= \begin{cases} 1 & \text{for} \quad \|c_1\Phi - \Psi\|_2 < c_0, \\ -1 & \text{for} \quad 0 < c_0 < \|c_1\Phi - \Psi\|_2, \\ 1 & \text{for} \quad -\|c_1\Phi - \Psi\|_2 < c_0 < 0, \\ -1 & \text{for} \quad c_0 < -\|c_1\Phi - \Psi\|_2. \end{cases}$$

Next we note,

$$\frac{\partial}{\partial c_0}\left(|c_0| + \|c_1\Phi - \Psi\|_2\right) = \text{sign}(c_0) = \begin{cases} 1 & \text{for} \quad c_0 > 0, \\ -1 & \text{for} \quad c_0 < 0. \end{cases}$$

The derivative of the terms (D.2), (D.3), (D.4) and (D.5) follows trivially since they are linear functions of c_0:

$$\frac{\partial}{\partial c_0}\delta_{1,+} = -1; \quad \frac{\partial}{\partial c_0}\delta_{1,-} = -1; \quad \frac{\partial}{\partial c_0}\delta_{2,+} = +1; \quad \frac{\partial}{\partial c_0}\delta_{2,-} = -1.$$

For $\delta_{2,+}$, the limits of both the lower and the upper bound in (D.1) follow from l'Hopital's rule:

$$\lim_{c_0\to\infty} \frac{\frac{\partial}{\partial c_0}\left||c_0| - \|-c_1\Phi + \Psi\|_2\right|}{\frac{\partial}{\partial c_0}\delta_{2,+}} = 1 \quad \overset{\text{l'Hopital}}{\Longrightarrow} \lim_{c_0\to\infty} \frac{\left||c_0| - \|-c_1\Phi + \Psi\|_2\right|}{\delta_{2,+}} = 1,$$

$$\lim_{c_0\to\infty} \frac{\frac{\partial}{\partial c_0}\left(|c_0| + \|c_1\Phi - \Psi\|_2\right)}{\frac{\partial}{\partial c_0}\delta_{2,+}} = 1 \quad \overset{\text{l'Hopital}}{\Longrightarrow} \lim_{c_0\to\infty} \frac{\left(|c_0| + \|c_1\Phi - \Psi\|_2\right)}{\delta_{2,+}} = 1.$$

Since both the lower and the upper bound on (4.9) tend to 1 as $c_0 \to \infty$, the sandwich theorem says that

$$\lim_{c_0\to\infty} \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta_{2,+}} = 1.$$

We also use l'Hopital's rule and then the sandwich theorem to establish that for $\delta_{1,+}$,

$$\lim_{c_0\to-\infty} \frac{\|c_1\Phi + c_0 I - \Psi\|_2}{\delta_{1,+}} = 1,$$

and the same holds for $\delta_{1,-}$ and $\delta_{2,-}$.

For $j = n - r$, we set ϕ_{n+1} and ψ_{n+1} to equal ∞, as done in Yu et al. (2015, p. 317), and then quantities (3.27), (3.28), (3.29) and (3.30) corresponding to

the four different subproblems in the affine case are,

$$\begin{aligned}
\delta_{1,+} &= c_1\phi_{n-r+1} + c_0 - \psi_{n-r},\\
\delta_{1,-} &= c_1\phi_n + c_0 - \psi_{n-r},\\
\delta_{2,+} &= \psi_{n-r+1} - c_1\phi_{n-r} - c_0,\\
\delta_{2,-} &= c_1\phi_{n-r} + c_0 - \psi_n.
\end{aligned}$$

The above arguments can be applied to these values of δ without complications. Consequently, for $j = n - r$, the cost function in (4.9) tends to 1 as c_0 tends to either ∞ or $-\infty$.