CLUSTERING GRAPHS AND CONTINGENCY TABLES WITH SPECTRAL METHODS

Academic Doctoral Dissertation

Marianna Bolla
Institute of Mathematics
Budapest University of Technology and Economics

Budapest, 2016.

Contents

Introduction

Spectral clustering is a relatively new notion of the 1990s for methods that aim at finding clusters of data points or vertices of a graph by means of the eigenvalues and eigenvectors of a conveniently constructed matrix based on the data or graph. If one reads the tutorial [Lux] or looks for spectral clustering in Wikipedia, finds that the eigenvalues of a similarity matrix assigned to the data are used to perform dimension reduction, and also finds a collection of algorithms, which roughly give the following recipe: if you have data points, build a similarity graph on them; if you have the graph, take the adjacency or Laplacian matrix, occasionally some normalized versions of them, and for a given integer k, find the top or bottom k eigenvalues together with eigenvectors; then apply the k-means algorithm to the eigenvectors. They do not tell much about the choice of the best suitable matrix and the integer k; further about the relation between the so obtained clustering of the vertices and the measures characterizing a good clustering from the point of view of some reasonable requirements of those practitioners who really want to find groups in biological or social networks. Here we try to give answers to these questions in the form of precisely formulated theorems and practical considerations, while we use the tools of advanced linear algebra, multivariate statistics, and probability.

In his videotalk, Ravi Kannan (Microsoft Research, India) also pointed out the close relation between spectral clustering and statistics when making low-dimensional embedding of high-dimensional data. The abstract of his talk Clustering – Does Theory Help? (Simons Institute, Berkeley, December 9, 2013) says the following. "*Theoretical computer science has brought to bear powerful ideas to find nearly optimal clusterings, while statistics mixture models of data have been useful in understanding the structure of data and in developing clustering algorithms. However, in practice many heuristics (e.g., dimension reduction and the k-means algorithm) are widely used. The talk will describe some aspects of the theoretical computer science and statistics approaches, and attempt to answer the question: is there a happy marriage of these approaches with practice?*" Indeed, graph theoretical optimization problems are partly considered by theoretical computer scientists (from the point of view of algorithms and their computational complexity) and partly by statisticians (from the point of view of mixture models and parameter estimation). In this dissertation, I manage to use both approaches and reconcile them with the needs of the practitioners.

Spectral clustering is originated in *spectral graph theory*, which connects linear algebra and combinatorial graph theory. It was developed by [Biggs] and [Cv], later summarized in [Chu]. These books mainly consider relations between spectral and structural properties of graphs, and describe spectra of many well-known graphs. Since non-isomorphic graphs can have the same spectra, the eigenvectors are also needed to characterize their properties. Fiedler [Fid73] and Hoffman [Hof69, Hof70, Hof72] use the eigenvector, corresponding to the smallest positive Laplacian eigenvalue of a connected graph (the famous Fiedler vector), to find a bipartition of the vertices which approximates the *minimum cut problem*, see also Juhász and Mályusz [Juh-Mály]. From the two-clustering point of view, this eigenvector

becomes important when the corresponding eigenvalue is not separated from the trivial zero eigenvalue, but it is separated from the second smallest positive Laplacian eigenvalue. On the contrary, when there is a large spectral gap between the trivial zero and the smallest positive Laplacian eigenvalue (or equivalently, between the trivial 1 and the second largest positive eigenvalue of the transition probability matrix in the random walk view), there is no use of partitioning the vertices, the whole graph forms a highly connected cluster. This case has frequently been studied since Cheeger [Che], establishing a lot of equivalent or near equivalent advisable features of these graphs. There are many results about the relation between this gap and different kinds of expansion constants of the graph (see e.g., [Ho-Lin-Wid]), including random walk view of [Az-Gh, Dia-Str, Mei-Shi]. Roughly speaking, graphs with a large spectral gap are good expanders, the random walk goes through them very quickly with a high mixing rate and short commute time, see Lovász [Lov93] for an overview; they are also good magnifiers as their vertices have many neighbors; in other words, their vertex subsets have a large boundary compared to their volumes characterized by the isoperimetric number of Mohár [Moh88]; equivalently, they have high conductance (see Nash-Williams [Nash]) and show quasirandom properties discussed in Thomason [Thom87, Thom89], Bollobás [Bo], and Chung, Graham, Wilson [Chu-G-W, Chu-G]. For these favorable characteristics, they are indispensable in communication networks.

However, less attention has been paid to graphs with a small spectral gap, when several cases can occur: among others, the graph can be a bipartite expander of Alon [Alon] or its vertices can be divided into two sparsely connected clusters, but the clusters themselves can be good expanders (see [Le-Gh-Trev] and [Ng-Jo-We]). In case of several clusters of vertices the situation is even more complicated. The pairwise relations between the clusters and the within-cluster relations of the vertices of the same cluster show a great variety. Depending on the number and sign of the so-called *structural eigenvalues* of the *normalized modularity matrix*, defined in [Bol11c], we make inferences on the number of the underlying clusters and the type of connection between them. Furthermore, based on *spectral and singular value decompositions* (in the sequel, SD and SVD), we attempt to explore the structure of the actual data set, by treating graphs and contingency tables as statistical data, which methods are reminiscent of the classical and modern techniques of *multivariate statistical analysis*; see [Bol81, Bol-Tus85, Boletal98] for new algorithms and applications of SVD. In the case of large data sets, we also investigate random effects, and explore tendencies in the spectra and spectral subspaces when the sizes tend to infinity.

I started dealing with this topic at the end of the 1980s, when together with my PhD advisor, Gábor Tusnády, we used spectral methods for a binary clustering problem, where the underlying matrix turned out to be the generalization of the graph Laplacian to hypergraphs (this framework is throughly discussed in [Bol91]; however, hypergraphs will not be considered in the present dissertation). Then we defined Laplacian for multigraphs and edge-weighted graphs; further, we went beyond the expanders by investigating gaps within the spectrum and used eigenvectors corresponding to some structural eigenvalues to find clusters of vertices. We also considered minimum multiway cuts with different normalizations that were later called *ratio- and normalized cuts*. These results (preliminary version published in [Bol91]) appeared in [Bol93, Bol-Tus94]. Meanwhile, in the 1990s, spectral clustering became a fashionable area, a lot of papers in this topic appeared, sometimes redefining or modifying the above notions, using different notation, sometimes having a numerical flavor and suggesting algorithms without rigorous mathematical explanation. At the turn of the millennium, thanks to the spreading of the World Wide Web and the human genome project, a rush started to investigate evolving graphs, microarrays, and random situations different of the classical Erdős–Rényi one. László Lovász and coauthors considered convergent graph sequences and testable graph parameters, bringing statistical concepts into this

discrete area. Inspired by this, we investigated noisy graph and contingency table sequences, and testability of some balanced versions of the already defined minimum multiway cut densities in [Bol-Ko-Kr12]. In parallel, physicists introduced other measures characterizing community structure of networks, see Newman [New]. In [Bol11c] we also defined some penalized versions of the Newman–Girvan modularity that are related to regular cuts. These are present in situations when the network has clusters of vertices with a homogeneous information flow within or between them, i.e., clusters with small within- and between-cluster discrepancies.

The dissertation consists of three strongly intertwining chapters. In the first one, we introduce the optimization problems, in due course of which the graph based matrices emerge, and the *representation theorems* (see Theorems 1, 3, 9) follow with simple liner algebra. Based on these representation theorems, more complicated statements about the relation between multiway cuts and spectra can easily be proved in one direction, e.g., Theorem 2 and the first part of Theorem 4. In the more complicated other direction, in the second part of Theorem 4, we also state the converse statement that can be proved by means of *analysis of variance* considerations and using the so-called k-variance of the k-dimensional vertex representatives. In fact, this k-variance is the squared distance between the spectral subspace corresponding to the k bottom Laplacian eigenvalues and the subspace of the so-called *partition vectors* (they are stepwise constant with respect to the underlying k-partition of the vertices). In the $k = 2$ case we are able to directly estimate this 2-variance with the ratio of the two smallest positive eigenvalues, see Theorem 5. In Theorem 6, we state a finer version of the backward isoperimetric inequality in the general case of an edge-weighted graph. The representation theorems are generalized to rectangular arrays and joint distributions, see Theorems 10, 11, 12 (we follow Alfréd Rényi's setup for introducing the notion of maximal correlation), and discussed together with the sequential correlation maximization task of *correspondence analysis*. In this way, the biclustering and bifactorization problems are unified, and in this framework, the application of reproducing kernel Hilbert spaces becomes natural and well justified when we want to find non-linear separation between our data points (we will show how the so-called kernel trick forms the base of the fashionable image segmentation). Most notions and statements about multiway cuts are to be found in Chapter 1, together with introducing a unified and very general treatment for the factorization and classification of edge-weighted and directed graphs, or contingency tables. This setup will also provide the framework to prove testability of certain spectral subspaces in Chapter 3. Summarizing, in Chapter 1, we want to establish a common outline structure for the contents of each optimization problem and algorithm, with unified notation and principles. We will use this notation in the statements and proofs of the subsequent chapters. Therefore, occasionally, the notation here differs from that used in the paper where the original version was published. Since spectral clustering is a relatively new area, in the related bibliography, there is no unified wording and notation for most of these concepts, but we will compare our notation with others' used in the literature.

Most of the new results are proved in Chapter 2. Here we investigate noisy random graphs and contingency tables, of which the *generalized random graphs* of [McSh] are special cases. In Theorems 13 and 15, we give the spectral characterization of the generalized random graphs having a k-cluster structure: the noisy matrix has k structural eigenvalues with eigenvectors based on which the k-variance of the vertex representatives tends to 0 as the cluster sizes tend to infinity, roughly speaking, at the same rate, see Theorems 14 and 16. The results extend to the spectra and spectral subspaces of noisy rectangular arrays, see Theorems 19, 20, 21, and 22. Some results of Chapter 1 become relevant for large graphs and contingency tables only. Networks are modeled either by edge-weighted graphs or contingency tables, and usually subject to random errors due to their evolving

and flexible nature. Asymptotic properties of SD and SVD of the involved matrices are discussed when not only the number of the vertices or that of the rows and columns of the contingency table tend to infinity, but the cluster sizes also grow proportionally with them. Mostly, perturbation results for the SD and SVD of blown-up matrices burdened with a Wigner-type error matrix are investigated. If the increasing graph obeys a block model, then it will have as many structural (large absolute value) eigenvalues as the rank (k) of the underlying pattern matrix, and the representatives of the vertices form k well separated clusters in the representation, based on the corresponding eigenvectors. Special structures are collected in Table 2.1, which makes sense for growing graph sequences. Conversely, given a weight-matrix or rectangular array of nonnegative entries, we are looking for the underlying block-structure. In Theorems 18 and 23 we will show that under very general circumstances, the clusters of a large graph's vertices and those of the rows and columns of a large contingency table can be identified with high probability.

In this framework, so-called volume-regular cluster pairs are also considered with 'small' *within- and between-cluster discrepancies*. Here we prove the multiclass generalization of the expander mixing lemma and its converse. Roughly speaking, Theorem 25, proved in [Bol14b] estimates the k-way discrepancy of a contingency table in the k-clustering obtained by spectral clustering tools as $O(\sqrt{2k}\tilde{S}_k + s_k)$, where s_k is the k-th largest singular value of the normalized table, and $\tilde{S}_k$ is the squareroot of the sum of the weighted k-variances of the optimal $(k-1)$-dimensional row- and column-representatives. Note that, with subspace perturbation theory, this is small if there is a large gap between s_{k-1} and s_k. The analogue of this theorem for weighted graphs is Theorem 27, proved earlier in [Bol14a], that estimates the k-way discrepancy of an edge-weighted graph in the k-clustering obtained by spectral clustering tools as $O(\sqrt{2k}\tilde{S}_k + |\mu_k|)$, where μ_k is the k-th largest absolute value eigenvalues of the normalized modularity matrix, and $\tilde{S}_k$ is the squareroot of the weighted k-variance of the optimal $(k-1)$-dimensional vertex-representatives. Therefore, when we have a bipartite, biregular graph, up to a constant factor, Theorem 27 gives the same estimate for the discrepancy between the two independent vertex-sets as Lemma 3.2 of [Ev-Go-Lu]. In the converse direction, in Theorem 24, we estimate s_k by a (near zero) strictly increasing logarithmic function of the k-way discrepancy, see [Bol16]. We state this theorem for rectangular arrays of nonnegative entries, but similar results follow for edge-weighted and directed graphs too, see Theorems 26 and 28.

The message of Theorems 24 and 25 is that the k-way discrepancy, when it is 'small' enough, suppresses s_k. Conversely, s_k together with a 'small' enough $\tilde{S}_k$ also suppresses the k-way discrepancy. Using perturbation theory of spectral subspaces, in [Bol14a] (in the framework of edge-weighted graphs), we also discuss that a 'large' gap between s_{k-1} and s_k suppresses $\tilde{S}_k$. Therefore, if we want to find row–column cluster pairs of 'small' discrepancy, we must select a k such that there is a remarkable gap between s_{k-1} and s_k; further s_k is 'small' enough. Moreover, by using this k and the construction in the proof of the forward statement of Theorem 25, we are able to find these clusters with spectral clustering tools. It makes sense, for example, when we want to find clusters of genes and conditions simultaneously in microarrays so that genes of the same row-cluster would 'equally' influence conditions of the same column-cluster.

In Chapter 3, some theoretical applications of the results of Chapter 1 and 2, further their relation to testable graph parameters are discussed, together with some open questions. We will prove that the increasing noisy graph sequences of Chapter 2 converge in the sense of Borgs and coauthors [Borgsetal1] too. In Theorems 29 and 30 we prove that for any k, the leading k singular values and the corresponding eigen-subspaces of the normalized modularity matrix are testable, hence, the weighted k-variance is testable too, see [Bol14a]. Here we also present some parametric and nonparametric statistical methods to find the underlying

clusters of a given network. In fact, minimum multiway cuts or bicuts and maximum modularities of Chapter 1 are nonparametric statistics that are estimated by means of spectral methods. Algorithms for the representation based spectral clustering are consequences of Theorems 25 and 27. Recently, we have considered the spectral and discrepancy properties as so-called *generalized quasirandom properties* and have realized that the theorems discussed in Chapter 2 are able to prove implications between them. The relations between spectra, spectral subspaces, multiway discrepancies, and degree distribution of generalized random and quasirandom graphs (introduced in [Lov-Sos]) can be regarded as generalized quasirandom properties, since equivalences between them can also be proved for deterministic graph sequences, irrespective of stochastic models. However, the precise formulation and proof of Conjecture 1 is not ready yet.

Finally, homogeneous and inhomogeneous *probabilistic mixture models* are considered, and it is pointed out how the *EM algorithm* can be applied to them for the purpose of simultaneous clustering and parameter estimation when our data form a graph or contingency table. We chose the number of clusters based on the gaps in the normalized modularity spectrum. Starting with the spectral clusters, the EM iteration provided a fine-tuning of the clusters, together with multiscale evaluation (via parameters) of the vertices. At this point, statistical mixture models are combined with graph theoretical optimization.

Summarizing, I think that my main contributions to spectral clustering are the following:

1. I have extended the notion of the Laplacian and modularity matrix to hypergraphs and edge-weighted graphs. I have discussed the graph and contingency table based optimization problems in a unified way, in which framework estimates for different kinds of multiway cuts are obtained by the SD or SVD of the appropriately selected matrix. I have also extended the task to joint distributions; hence, generalized the problem of factor analysis or correspondence analysis, and justified the usage of reproducing kernel Hilbert spaces in image segmentation problems.

2. I have characterized spectra and spectral subspaces of generalized random graphs; extended the expander mixing lemma to the k-cluster case and proved its converse too. These theorems can give a hint for practitioners about the choice of the number of clusters. The original expander mixing lemma and its converse (for simple, regular graphs) treat the $k = 1$ case only, whereas the Szemerédi Regularity Lemma applies to the worst case scenario: even if there is not an underlying cluster structure, we can find cluster pairs with small discrepancy with an enormously large k (which does not depend on the number of vertices, it only depends on the discrepancy to be attained). I rather treat the intermediate case, and show that a moderate k suffices if our graph has a hidden k-cluster structure that can be revealed by spectral clustering tools. Under good clustering I generally understand clusters with small within- and between-cluster discrepancies.

3. I have proved that the normalized modularity spectra and spectral subspaces of noisy random graph sequences are testable in the sense of [Borgsetal1]. Based on this, I state so-called generalized quasirandom properties and can prove some implications between them, including spectra and the multiway discrepancy introduced for this purpose.

4. I have shown how to apply the EM algorithm to estimate the parameters of the homogeneous and inhomogeneous stochastic block models, together with finding the underlying clusters. Here the data form a graph or contingency table, which is not complete as the cluster memberships of the graph vertices or rows/columns of the table are missing. The initial clustering is obtained by spectral tools.

Note that all the numbered theorems are mine (occasionally with coauthors), whereas the theorems of other authors are not numbered.

Acknowledgements: I am indebted to my former PhD advisor, Gábor Tusnády with whom I have worked together in several projects and applied statistical methods in real-life problems. I am grateful to Katalin Friedl, András Krámli, László Lovász, Miklós Simonovits, and Vera T. Sós for valuable discussions on the graph spectra and convergence topics. I wish to express gratitude to persons whose books, papers, and lectures on spectral graph theory and multivariate statistics turned my interest to this field: F. Chung, D. M. Cvetkovic, M. Fiedler, A. J. Hoffman, B. Mohar, and C. R. Rao. I thank many colleagues for their useful comments on the related research and applications: László Babai, András Benczur, Endre Boros, Imre Csiszár, Villő Csiszár, Miklós Ferenczi, Zoltán Füredi, László Gerencsér, László Győrfi, Ferenc Juhász, Gergely Kiss, János Komlós, Vilmos Komornik, Tamás Lengyel, András Lukács, Péter Major, Katalin Marton, György Michaletzky, Dezső Miklós, Tamás F. Móri, Marcello Pelillo, Gábor Pete, Dénes Petz, Tomaz Pisanski, András Prékopa, András Recski, Lídia Rejtő, Tamás Rudas, András Simonovits, Tamás Szabados, Tamás Szántai, Domokos Szász, Gábor J. Székely, András Telcs, László Telegdi, Miklós Telek, Bálint Tóth, János Tóth, György Turán, Zsuzsa Vágó, and Katalin Vesztergombi.

Most of these results are contained in the book [Bol13], except Theorems 24, 26, 28 that I have managed to prove since then. Here I omit my earlier results and detailed description of former results of other authors. My publications related to spectral clustering are in the block [Bol93]–[Bol16] of the Bibliography. I am the only author of most of the papers in the last fifteen years, and as for the joint papers, I have the approval of my coauthoring colleagues. Occasionally, papers were coauthored with my PhD and BSM students, however, they basically helped in making simulations, figures, and smaller calculations, and not in the formation of the main theorems. Hereby, I enlist the names of all Hungarian and foreign students who have participated in this research (while writing their MS or PhD theses or taking student research courses): Andrea Bán, Erik Bodzsár, Brian Bullins, Sorathan Chaturapruek, Shiwen Chen, Calvin Cheng, Ahmed Elbanna, Dóra Farkas, Max Del Giudice, Edward Kim, Tamás Kói, Gábor Molnár–Sáska, László Nagy, Ildikó Priksz, Viktor Szabó, Zsolt Szabó, Cheng Wai Koo, and Joan Wang.

I am grateful to Károly Simon and Miklós Horváth for encouraging me to write this dissertation and letting me to go on sabbatical leaves when I managed to prove many of these results. The related research was partly supported by the Hungarian National Research Grants OTKA 76481 and OTKA-KTIA 77778; further, by the TÁMOP-4.2.2.B-10/1-2010-0009 and TÁMOP-4.2.2.C-11/1/KONV-2012-0001 projects that have successfully been closed.

Chapter 1

Multiway cuts and representations related to spectra

In this chapter, the optimization problems are introduced, in due course of which the graph based matrices emerge, and the representation theorems follow with some liner algebra. Based on these theorems, statements about the lower estimates of multiway cuts by means of spectra can easily be proved; the more complicated upper estimates include the eigensubspaces as well.

Most notions and statements about multiway cuts are to be found in this chapter, together with introducing a unified and very general treatment for the factorization and classification of edge-weighted and directed graphs, or contingency tables. This setup will also provide the framework to prove more sophisticated theorems in Chapters 2 and 3.

Graph spectra are used for almost 50 years to recover the structure of graphs. Different kinds of spectra are capable of finding multiway cuts corresponding to different optimization criteria. While eigenvalues give estimates for the objective functions of the discrete optimization problems, eigenvectors are used to find clusters of vertices via the k-means algorithm, which approximately solves the multiway cut problems with some normalization. The technique is often called *spectral relaxation*, but these methods are also reminiscent of some classical methods of *multivariate statistical analysis*, namely, principal component analysis (Pearson 1901, Hotelling 1933), factor analysis (Thrustone 1931, Thompson 1939), canonical correlation analysis (Hotelling 1936), and correspondence analysis (Benzécri et al. 1980, Greenacre 1984, and [Bol87b]). We also generalize the notion of representation for joint distributions, which technique justifies the way how non-linearities are treated by mapping the data into a feature space (reproducing kernel Hilbert space).

The technical proofs of most of the theorems in this chapter are omitted (they are to be found in the paper referred to in parentheses after the numbered theorems, the theorems of other authors are not numbered). However, some short proofs, demonstrating the representation or kernel techniques are presented.

1.1 Quadratic placement and multiway cut problems for graphs

Now, our data matrix corresponds to a graph. First, let $G = (V, E)$ be a simple graph on the vertex-set V and edge-set E with $|V| = n$ and $|E| \leq \binom{n}{2}$. Thus, the $|E| \times n$ data matrix $\boldsymbol{B}$

has 0-1 entries, the rows correspond to the edges, the columns to the vertices, and b_{ij} is 1 or 0 depending on whether the edge i contains the vertex j as an endpoint or not. The Gram-matrix $\boldsymbol{C} = \boldsymbol{B}^T\boldsymbol{B}$ is the non-centralized covariance matrix based on the data matrix $\boldsymbol{B}$, and is both positive definite and a Frobenius type matrix with nonnegative entries. Sometimes the matrix $\boldsymbol{C}$ is called *signless Laplacian* (see [Cv-Ro-Si]) and its eigenspaces are used to compare cospectral graphs. It is easy to see that $\boldsymbol{C} = \boldsymbol{D} + \boldsymbol{A}$, where $\boldsymbol{A} = (a_{ij})$ is the usual *adjacency matrix* of G, while $\boldsymbol{D}$ is the so-called degree-matrix, i.e., diagonal matrix, containing the vertex-degrees in its main diagonal. $\boldsymbol{A}$ being a Frobenius-type matrix, its maximum absolute value eigenvalue is positive, it is at most the maximum vertex-degree, and apart from the trivial case – when there are no edges at all – it is indefinite, as the sum of its eigenvalues, i.e., the trace of $\boldsymbol{A}$, is zero.

Instead of the positive definite matrix $\boldsymbol{C}$, for optimization purposes, as will be derived below, the *Laplacian matrix* $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A}$ is more suitable, which is positive semidefinite. This $\boldsymbol{L}$ is sometimes called combinatorial or difference Laplacian, whereas we will introduce the so-called *normalized Laplacian*, $\boldsymbol{L}_D = \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2}$ too. If our graph is regular, then $\boldsymbol{D} = d\boldsymbol{I}$ (where d is the common degree of the vertices and $\boldsymbol{I}$ is the identity matrix) and the eigenvalues of $\boldsymbol{C}$ and $\boldsymbol{L}$ are obtained from those of $\boldsymbol{A}$ by adding d to them or subtracting them from d, respectively.

There are other frequently used matrices, for example, the modularity and normalized modularity matrices preferred by physicists, latter one closely related to the normalized Laplacian, akin to the so-called *transition probability matrix* $\boldsymbol{D}^{-1}\boldsymbol{A}$ or the random walk Laplacian $\boldsymbol{I} - \boldsymbol{D}^{-1}\boldsymbol{A}$ (these matrices are not symmetric, still they have real eigenvalues). We will clarify in which situation which of these matrices is the best applicable. The whole story simplifies if we use edge-weighted graphs, and all these matrices come into existence naturally, while solving some minimum placement problems. Multiway cut problems also fit into this framework, since the optima of their objective functions are obtained by taking the above optima over so-called partition vectors (stepwise constant with respect to the hidden partition of the vertices), so they can easily be related to the Laplacian or normalized Laplacian spectra, whereas the precision of the estimates depends on the distance between the subspaces spanned by the corresponding eigenvectors and partition vectors. By an *analysis of variance* argument, this distance is the sum of the inner variances of the underlying clusters, the objective function of the *k-means algorithm*.

1.1.1 Representation of edge-weighted graphs

Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph, where $V = \{1, \ldots, n\}$ is the vertex-set and the $n \times n$ symmetric *edge-weight matrix* $\boldsymbol{W}$ has nonnegative real entries and zero diagonal. Here w_{ij} can be thought of as the similarity between vertices i and j, where 0 similarity means no connection (edge) at all. If G is a simple graph, then $\boldsymbol{W}$ is its adjacency matrix. Since $\boldsymbol{W}$ is symmetric, the weight of the edge between two vertices does not depend on its direction, i.e., our graph is *undirected*. We will mostly treat undirected graphs, except in Section 2.3.4, where a non-symmetric $\boldsymbol{W}$ corresponds to a directed edge-weighted graph.

The row-sums of $\boldsymbol{W}$ are

$$d_i = \sum_{j=1}^{n} w_{ij}, \quad i = 1, \ldots, n$$

which are called *generalized vertex-degrees* and collected in the main diagonal of the diagonal *degree-matrix* $\boldsymbol{D} = \mathrm{diag}(\mathbf{d})$, where $\mathbf{d} = (d_1, \ldots, d_n)^T$ is the so-called *degree-vector*. (Vectors are always columns, in this chapter have real coordinates, and ^{T} stands for the transposition.)

For a given integer $1 \le k \le n$, we are looking for k-dimensional representatives $\mathbf{r}_1, \ldots, \mathbf{r}_n \in \mathbb{R}^k$ of the vertices such that they minimize the objective function

$$Q_k = \sum_{i<j} w_{ij} \|\mathbf{r}_i - \mathbf{r}_j\|^2 \ge 0 \tag{1.1}$$

subject to

$$\sum_{i=1}^{n} \mathbf{r}_i \mathbf{r}_i^T = \boldsymbol{I}_k \tag{1.2}$$

where $\boldsymbol{I}_k$ is the $k \times k$ identity matrix. When minimized, the objective function Q_k favors k-dimensional placement of the vertices such that vertices connected with large-weight edges are close to each other. This is the base of many graph-drawing algorithms.

Let us put both the objective function and the constraint in a more favorable form. Denote by $\boldsymbol{X}$ the $n \times k$ matrix of rows $\mathbf{r}_1^T, \ldots, \mathbf{r}_n^T$. Let $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^n$ be the columns of $\boldsymbol{X}$, for which fact we use the notation $\boldsymbol{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$. Because of the constraint (1.2), the columns of $\boldsymbol{X}$ form an orthonormal system, hence, $\boldsymbol{X}$ is a *suborthogonal* matrix, and the constraint (1.2) can be formulated as $\boldsymbol{X}^T\boldsymbol{X} = \boldsymbol{I}_k$. With this notation, the objective function (1.1) is rewritten in the symmetrized form

$$\begin{aligned} Q_k &= \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij} \|\mathbf{r}_i - \mathbf{r}_j\|^2 = \sum_{i=1}^{n} d_i \|\mathbf{r}_i\|^2 - \sum_{i=1}^{n} \sum_{j=1}^{n} w_{ij} \mathbf{r}_i^T \mathbf{r}_j \\ &= \sum_{\ell=1}^{k} \mathbf{x}_\ell^T (\boldsymbol{D} - \boldsymbol{W}) \mathbf{x}_\ell = \mathrm{tr}[\boldsymbol{X}^T (\boldsymbol{D} - \boldsymbol{W}) \boldsymbol{X}]. \end{aligned} \tag{1.3}$$

Definition 1 *The matrix $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{W}$ is called the Laplacian corresponding to the edge-weighted graph $G = (V, \boldsymbol{W})$.*

For simple graphs, we get back the usual definition of the Laplacian, e.g., [Chu, Moh88].

The Laplacian is always positive semidefinite, since by (1.1), the quadratic form Q_1 is nonnegative; and it always has a zero eigenvalue, since its rows sum to zero. It can be shown, that the multiplicity of 0 as an eigenvalue of $\boldsymbol{L}$ is equal to the number of the connected components of $G = (V, \boldsymbol{W})$, i.e., the maximum number of disjoint subsets of V such that there are no edges connecting vertices of distinct subsets (no edge means an edge with zero weight). In terms of $\boldsymbol{W}$, the number of connected components of G is the maximum number of the diagonal blocks which can be achieved by the same permutation of the rows and columns of $\boldsymbol{W}$. Consequently, if G is connected, then 0 is a single eigenvalue with corresponding unit-norm eigenvector $\mathbf{u}_0 = \frac{1}{\sqrt{n}}\mathbf{1}$, where $\mathbf{1}$ denotes the all 1's vector. In the sequel, we will assume that G is *connected*, or equivalently, $\boldsymbol{W}$ is *irreducible*.

Since our objective function $Q_k = \mathrm{tr}[\boldsymbol{X}^T\boldsymbol{L}\boldsymbol{X}]$ is minimized under $\boldsymbol{X}^T\boldsymbol{X} = \boldsymbol{I}_k$, a simple linear algebra provides the following theorem.

Theorem 1 ([Bol-Tus94]) ***Representation theorem for edge-weighted graphs.*** *Let $G = (V, \boldsymbol{W})$ be a connected edge-weighted graph with Laplacian matrix $\boldsymbol{L}$. Let $0 = \lambda_0 < \lambda_1 \le \cdots \le \lambda_{n-1}$ be the eigenvalues of $\boldsymbol{L}$ with corresponding unit-norm eigenvectors $\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{n-1}$. Let $k < n$ be a positive integer such that $\lambda_{k-1} < \lambda_k$. Then the minimum of (1.1) subject to (1.2) is*

$$\sum_{i=0}^{k-1} \lambda_i = \sum_{i=1}^{k-1} \lambda_i$$

and it is attained with the optimal representatives $\mathbf{r}_1^*, \ldots, \mathbf{r}_n^*$, the transposes of which are row vectors of $\boldsymbol{X}^* = (\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{k-1})$.

The vectors $\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{k-1}$ are called *vector components* of the optimal representation. We remark the following.

- The dimension k does not play an important role here, the vector components can be included one after the other up to a k such that $\lambda_{k-1} < \lambda_k$.
- The eigenvectors can be arbitrarily chosen in the eigenspaces corresponding to possible multiple eigenvalues, under the orthogonality conditions. Further, the representatives can as well be rotated in $\mathbb{R}^k$. Indeed, nor the objective function, neither the constraint is changed if we use $\boldsymbol{R}\mathbf{r}_i$'s instead of $\mathbf{r}_i$'s, or equivalently, $\boldsymbol{X}\boldsymbol{R}$ instead of $\boldsymbol{X}$, where $\boldsymbol{R}$ is an arbitrary $k \times k$ orthogonal matrix.
- Since the eigenvector $\mathbf{u}_0$ has equal coordinates, the same first coordinates of the vertex representatives do not play an important role in the representation, especially when the representatives are used for clustering purposes. Therefore, $\mathbf{u}_0$ can be disregarded, and an optimal $(k-1)$-dimensional representation is performed based on the eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_{k-1}$.
- So far, we assumed that $\boldsymbol{W}$ has zero diagonal. We can as well see that in the presence of possible loops (some or all diagonal entries of $\boldsymbol{W}$ are positive) the objective function and the Laplacian remains the same, hence, Theorem 1 is applicable to this situation too.
- Representation of hypergraphs is discussed in [Bol92, Bol93]. In fact, the Laplacian of the hypergraph $H = (V, E)$ defined there is the same as the Laplacian of the edge-weighted graph $G = (V, \boldsymbol{W})$, with edge-weights

$$w_{ij} = \begin{cases} \sum_{e \in E} \sum_{i,j \in e} \frac{1}{|e|} & \text{if} \quad i \neq j \\ 0 & \text{if} \quad i = j, \end{cases}$$

  where $|e|$ is the number of vertices contained in the hyper-edge e. Therefore, hypergraphs will not be discussed here. Examples for spectra and representation of some well-known simple graphs are to be found in Section 1.1.3 of [Bol13].

1.1.2 Estimating minimum multiway cuts via spectral relaxation

Clusters (in other words, *modules* or *communities*) of graphs are typical (usually, loosely connected) subsets of vertices that can be identified, for example, with social groups or interacting enzymes in social or metabolic networks, respectively; they form special partition classes of the vertices. To measure the performance of a clustering, different kinds of multiway cuts are introduced and estimated by means of Laplacian spectra. The key motif of these estimations is that minima and maxima of the quadratic placement problems of Section 1.1 are attained on some appropriate eigenspaces of the Laplacian, while optimal multiway cuts are special values of the same quadratic objective function realized by step-vectors. Hence, the optimization problem, formulated in terms of the Laplacian eigenvectors, is the *continuous relaxation* of the underlying maximum or minimum multiway cut problem.

For a fixed integer $1 \leq k \leq n$, let $P_k = (V_1, \ldots, V_k)$ be a *proper k-partition* of the vertices, where the disjoint, non-empty vertex subsets $V_1, \ldots, V_k$ will be referred to as clusters or modules. Let $\mathcal{P}_k$ denote the set of all k-partitions. Optimization over $\mathcal{P}_k$ is usually NP-complete, except some special classes of graphs.

Definition 2 *The weighted cut between the non-empty vertex-subsets $U, T \subset V$ of the edge-weighted graph $G = (V, \boldsymbol{W})$ is*

$$w(U,T) = \sum_{i \in U} \sum_{j \in T} w_{ij}.$$

The minimum k-way cut of G is

$$\mathrm{mincut}_k(G) = \min_{P_k \in \mathcal{P}_k} \sum_{a=1}^{k-1} \sum_{b=a+1}^{k} w(V_a, V_b). \tag{1.4}$$

For a simple graph G, Fiedler [Fid73] called the quantity $\mathrm{mincut}_2(G)$ the edge-connectivity of G, because it is equal to the minimum number of edges that should be removed to make G disconnected. He used the notation $e(G)$ for the edge-connectivity of the simple graph G, and $v(G)$ for its vertex-connectivity (minimum how many vertices should be removed to make G disconnected). In his breakthrough papers [Fid72, Fid73], Fiedler proved that for any graph G on n vertices, that differs from the complete graph K_n, the relation

$$\lambda_1 \leq v(G) \leq e(G) \tag{1.5}$$

holds. In [Fid73], he also provided two lower estimates for λ_1 by $e(G)$:

$$\lambda_1 \geq 2e(G)(1 - \cos\frac{\pi}{n}) \tag{1.6}$$

and

$$\lambda_1 \geq C_1 e(G) - C_2 d_{\max}, \tag{1.7}$$

where $C_1 = 2(\cos\frac{\pi}{n} - \cos\frac{2\pi}{n})$, $C_2 = 2\cos\frac{\pi}{n}(1 - \cos\frac{\pi}{n})$, and $d_{\max} = \max_i d_i$ is the maximum vertex-degree. Compared to (1.5), this estimation makes sense in the $n \geq 3$ case. The bound of (1.7) is tighter than that of (1.6) if and only if $e(G) \geq \frac{1}{2} d_{\max}$. The two estimates are equal and sharp for the path graph P_n with $e(G) = 1$ and $\lambda_1 = 2(1 - \cos\frac{\pi}{n})$. The path graph can be split into two clusters by removing any of its edges, however, we would not state that it has two underlying clusters. The forthcoming ratio cut of P_n is minimized by removing the middle edge (for even n) or one of the middle edges (for odd n), thus, it provides balanced clusters.

Because of this two-sided relation between λ_1 and $e(G)$, the smallest positive Laplacian eigenvalue of a connected graph is able to detect the strength of its connectivity; therefore, Fiedler called λ_1 the *algebraic connectivity* of G. This relation between $\lambda_1(G)$ and $e(G)$ was also discovered by A. J. Hoffman [Hof70, Hof69], at the same time.

The proof of Fiedler gives us the following hint how to find the optimal 2-partition: the eigenvector $\mathbf{u}_1$ should be close to a step-vector over an appropriate 2-partition of the vertices. Note that because of its orthogonality to the vector $\mathbf{1}$, the vector $\mathbf{u}_1$ contains both positive and negative coordinates, and Juhász and Mályusz [Juh-Mály] separated the two clusters according to the signs. In the sequel, we will use the k-means algorithm for this purpose, in a more general setup. The vector $\mathbf{u}_1$ is frequently called *Fiedler-vector*.

Even in the simplest $k = 2$ case, the solution of the minimum cut problem is frequently attained by an uneven 2-partition, for example, if there is an almost isolated vertex (connected to few other vertices), it may form a cluster itself. To prevent this situation and rather find real-life loosely connected clusters, we require some balancing for the cluster sizes. For this purpose, in [Bol91, Bol-Tus94] (even in the preprint version) we defined a type of a weighted cut that, in addition, penalizes partitions with very unequal cluster sizes. This cut was later called *ratio cut*, see, e.g., [Hag-Kah].

Definition 3 *Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph and $P_k = (V_1, \ldots, V_k)$ a k-partition of its vertices. The k-way ratio cut of G corresponding to the k-partition P_k is*

$$g(P_k, G) = \sum_{a=1}^{k-1} \sum_{b=a+1}^{k} \left(\frac{1}{|V_a|} + \frac{1}{|V_b|} \right) w(V_a, V_b) = \sum_{a=1}^{k} \frac{w(V_a, \overline{V}_a)}{|V_a|}$$

and the minimum k-way ratio cut of G is

$$g_k(G) = \min_{P_k \in \mathcal{P}_k} g(P_k, G).$$

Assume that G is connected. Let $0 = \lambda_0 < \lambda_1 \leq \cdots \leq \lambda_{n-1}$ denote the eigenvalues of its Laplacian matrix $\boldsymbol{L}$ with corresponding unit-norm, pairwise orthogonal eigenvectors $\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{n-1}$. Namely, $\mathbf{u}_0 = \frac{1}{\sqrt{n}}\mathbf{1}$.

Theorem 2 ([Bol-Tus94]) *For the minimum k-way ratio cut of the connected edge-wighted graph $G = (V, \boldsymbol{W})$ the lower estimate*

$$g_k(G) \geq \sum_{i=1}^{k-1} \lambda_i$$

holds.

To illustrate the spectral relaxation technique, we describe the short proof here.

Proof. The k-partition P_k is uniquely determined by the $n \times k$ balanced partition matrix $\boldsymbol{Z}_k = (\mathbf{z}_1, \ldots, \mathbf{z}_k)$, where the a-th *balanced k-partition vector* $\mathbf{z}_a = (z_{1a}, \ldots, z_{na})^T$ is the following:

$$z_{ia} = \begin{cases} \frac{1}{\sqrt{|V_a|}} & \text{if} \quad i \in V_a \\ 0 & \text{otherwise.} \end{cases} \tag{1.8}$$

The matrix $\boldsymbol{Z}_k$ is trivially suborthogonal, and the set of balanced k-partition matrices is denoted by $\mathcal{Z}_k^B$. With the special representation in which the representatives $\tilde{\mathbf{r}}_1, \ldots, \tilde{\mathbf{r}}_n \in \mathbb{R}^k$ are row vectors of $\boldsymbol{Z}_k$, the ratio cut of $G = (V, \boldsymbol{W})$ corresponding to the k-partition P_k can be rewritten as

$$g(P_k, G) = \sum_{i=1}^{n-1} \sum_{j=i+1}^{n} w_{ij} \|\tilde{\mathbf{r}}_i - \tilde{\mathbf{r}}_j\|^2 = \sum_{a=1}^{k} \mathbf{z}_a^T \boldsymbol{L} \mathbf{z}_a = \mathrm{tr}(\boldsymbol{Z}_k^T \boldsymbol{L} \boldsymbol{Z}_k). \tag{1.9}$$

If we minimize it over balanced k-partition matrices $\boldsymbol{Z}_k \in \mathcal{Z}_k^B$, the so obtained minimum cannot go below the overall minimum $\sum_{i=0}^{k-1} \lambda_i$. This finishes the proof. □

Note that equality can be attained only in the $k = 1$ trivial case, otherwise the eigenvectors $\mathbf{u}_i$ $(i = 1, \ldots, k-1)$ cannot be partition vectors, since their coordinates sum to 0 because of the orthogonality to the $\mathbf{u}_0$ vector.

In the case of $k = 2$, in view of Theorem 2, $g_2(G)$ is bounded from below by λ_1, akin to the edge-connectivity of [Fid73]. The proof also suggests that the quality of the above estimation depends on, how close the k bottom eigenvectors of $\boldsymbol{L}$ are to partition vectors. The measure of the closeness of the involved subspaces is the *k-variance* of the k-dimensional vertex representatives $\mathbf{r}_1, \ldots, \mathbf{r}_n$ defined as

$$S_k^2(\mathbf{r}_1, \ldots, \mathbf{r}_n) = \min_{P_k \in \mathcal{P}_k} S_k^2(P_k; \mathbf{r}_1, \ldots, \mathbf{r}_n) = \min_{P_k = (V_1, \ldots, V_k)} \sum_{a=1}^{k} \sum_{j \in V_a} \|\mathbf{r}_j - \mathbf{c}_a\|^2 \tag{1.10}$$

where $\mathbf{c}_a = \frac{1}{|V_a|}\sum_{j\in V_a}\mathbf{r}_j$ is the center of cluster V_a $(a = 1,\dots,k)$. The minimum is obtained by the k-means algorithm. More precisely, we will apply the k-means algorithm for the optimal representatives, and if there is a large gap between λ_{k-1} and λ_k, we may expect that the optimum, given by the k-means algorithm is not far from that of the minimum k-way ratio cut. These issues are investigated in [Bol13] thoroughly, together with hypergraph cuts, here we do not discuss the details. We will rather give similar estimates for the normalized cut with the normalized Laplacian eigenvalues in the next section.

1.1.3 Normalized Laplacian and normalized cuts

Let $G = (V, \boldsymbol{W}, \boldsymbol{S})$ be a weighted graph on the vertex-set V $(|V| = n)$, where both the edges and vertices have nonnegative weights. The edge-weights are entries of $\boldsymbol{W}$ as in Section 1.1, whereas the diagonal matrix $\boldsymbol{S} = \mathrm{diag}(s_1,\dots,s_n)$ contains the positive vertex-weights in its main diagonal. Without loss of generality, we can assume that the entries in $\boldsymbol{W}$ and $\boldsymbol{S}$ both sum to 1. For the time being, the vertex-weights have nothing to do with the edge-weights. These individual weights are assigned to the vertices subjectively. For example, in a social network, the edge-weights are similarities between the vertices based on the strengths of their pairwise connections (like frequency of co-starring of actors), while vertex-weights embody the individual strengths of the vertices in the network (like the actors' individual abilities). This idea also appears in Borgs and coauthors [Borgsetal1], where they consider edge- and vertex-weighted graphs.

Now, we look for k-dimensional representatives $\mathbf{r}_1,\dots,\mathbf{r}_n$ of the vertices so that they minimize the objective function $Q_k = \sum_{i<j} w_{ij}\|\mathbf{r}_i - \mathbf{r}_j\|^2$ subject to $\sum_{i=1}^n s_i\mathbf{r}_i\mathbf{r}_i^T = \boldsymbol{I}_k$. With the notation and considerations of Section 1.1,

$$\begin{aligned}\min_{\sum_{i=1}^n s_i\mathbf{r}_i\mathbf{r}_i^T=\boldsymbol{I}_k} Q_k &= \min_{\boldsymbol{X}^T\boldsymbol{S}\boldsymbol{X}=\boldsymbol{I}_k}\mathrm{tr}(\boldsymbol{X}^T\boldsymbol{L}\boldsymbol{X})\\ &= \min_{\boldsymbol{X}^T\boldsymbol{S}\boldsymbol{X}=\boldsymbol{I}_k}\mathrm{tr}[(\boldsymbol{S}^{1/2}\boldsymbol{X})^T(\boldsymbol{S}^{-1/2}\boldsymbol{L}\boldsymbol{S}^{-1/2})(\boldsymbol{S}^{1/2}\boldsymbol{X})]\\ &= \sum_{i=0}^{k-1}\lambda_i(\boldsymbol{L}_S) = \sum_{i=1}^{k-1}\lambda_i(\boldsymbol{L}_S)\end{aligned}$$

where $\boldsymbol{L}_S = \boldsymbol{S}^{-1/2}\boldsymbol{L}\boldsymbol{S}^{-1/2}$ is the Laplacian normalized by $\boldsymbol{S}$, and because of the constraints, $\boldsymbol{S}^{1/2}\boldsymbol{X}$ is a suborthogonal matrix. Obviously, $\boldsymbol{L}_S$ is also positive semidefinite with eigenvalues $0 = \lambda_0(\boldsymbol{L}_S) \le \lambda_1(\boldsymbol{L}_S) \le \cdots \le \lambda_{n-1}(\boldsymbol{L}_S)$ and corresponding orthonormal eigenvectors $\mathbf{u}_0,\mathbf{u}_1,\dots,\mathbf{u}_{n-1}$. Furthermore, 0 is a single eigenvalue if and only if G is connected. The optimal k-dimensional representation is obtained by the row vectors of the matrix $\boldsymbol{S}^{-1/2}(\mathbf{u}_0,\mathbf{u}_1,\dots,\mathbf{u}_{k-1})$.

The special case, when the vertex-weights are the generalized degrees, that is $\boldsymbol{S} = \boldsymbol{D}$, has a distinguished importance.

Definition 4 *The matrix*

$$\boldsymbol{L}_D = \boldsymbol{D}^{-1/2}\boldsymbol{L}\boldsymbol{D}^{-1/2} = \boldsymbol{I}_n - \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2} = \boldsymbol{I}_n - \boldsymbol{W}_D$$

is called the normalized Laplacian of the edge-weighted graph $G = (V, \boldsymbol{W})$.

In the sequel, assume that $\sum_{i=1}^n\sum_{j=1}^n w_{ij} = 1$, which will be in accord with the joint distribution setup of Section 1.3. This is not a serious restriction, since neither the normalized Laplacian, nor the normalized cut to be introduced are affected by the scaling of the edge-weights. Some simple statements concerning the normalized Laplacian spectra:

- In Section 1.3 we will see that the eigenvalues of the *normalized edge-weight matrix* $\boldsymbol{W}_D = \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}$ are in the [-1,1] interval, since they are special correlations, the largest one being 1. Consequently, the eigenvalues of $\boldsymbol{L}_D$ are in the $[0, 2]$ interval. Let

$$0 = \lambda_0 \leq \lambda_1 \leq \cdots \leq \lambda_{n-1} \leq 2$$

  denote the spectrum of the normalized Laplacian $\boldsymbol{L}_D$.

- Trivially, 0 is a single eigenvalue of $\boldsymbol{L}_D$ if and only if G is connected (i.e., $\boldsymbol{W}$ is irreducible), and in this case, the corresponding unit-norm eigenvector is the $\sqrt{\mathbf{d}} := (\sqrt{d_1}, \ldots, \sqrt{d_n})^T$ vector. Furthermore, the normalized Laplacian spectrum of a disconnected graph is the union of those of its connected components.

- Since $\sum_{i=0}^{n-1} \lambda_i = \mathrm{tr}(\boldsymbol{L}_D) = n$, the following estimates hold:

$$\lambda_1 = \min_{i\in\{1,\ldots,n-1\}} \lambda_i \leq \frac{1}{n-1}\sum_{i=1}^{n-1} \lambda_i = \frac{n}{n-1} \leq \max_{i\in\{1,\ldots,n-1\}} \lambda_i = \lambda_{n-1}.$$

  Note that both of the above inequalities hold with equality at the same time, if and only if G is the complete graph K_n.

- For a simple graph G, which is not the complete graph K_n, $\lambda_1 \leq 1$ holds. Furthermore, $\lambda_1 = 1$ if and only if our graph is the complete k-partite graph $K_{n_1,\ldots,n_k}$ with some $1 < k < n$. This issue is discussed in [Boletal15] and in Section 1.1.5.

- Provided G is connected, 2 is an eigenvalue of it if and only if G is a bipartite graph.

- For the normalized Laplacian spectra of some well-known simple graphs, see Section 1.3 [Bol13].

Normalized Laplacian was used for spectral clustering in several papers (see, e.g., [Az-Gh, Mei-Shi]). These results are based on the observation that the SD of $\boldsymbol{L}_D$ solves the following quadratic placement problem.

Theorem 3 ([Bol-Tus94]) ***Representation theorem for edge- and vertex-weighted graphs*** *(when the vertex-weights are the generalized degrees). Let $G = (V, \boldsymbol{W})$ be a connected edge-weighted graph with normalized Laplacian $\boldsymbol{L}_D$. Let $0 = \lambda_0 < \lambda_1 \leq \cdots \leq \lambda_{n-1}$ be the eigenvalues of $\boldsymbol{L}_D$ with corresponding unit-norm eigenvectors $\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{n-1}$. Let $k < n$ be a positive integer such that $\lambda_{k-1} < \lambda_k$. Then the minimum of Q_{k-1} of (1.1) subject to*

$$\sum_{i=1}^{n} d_i \mathbf{r}_i \mathbf{r}_i^T = \boldsymbol{I}_{k-1} \quad \textit{and} \quad \sum_{i=1}^{n} d_i \mathbf{r}_i = \mathbf{0}$$

is $\sum_{i=1}^{k-1} \lambda_i$ and it is attained with the optimal $(k-1)$-dimensional representatives $\mathbf{r}_1^*, \ldots, \mathbf{r}_n^*$ the transposes of which are row vectors of $\boldsymbol{X}^* = \boldsymbol{D}^{-1/2}(\mathbf{u}_1, \ldots, \mathbf{u}_{k-1})$.

The vectors $\boldsymbol{D}^{-1/2}\mathbf{u}_1, \ldots, \boldsymbol{D}^{-1/2}\mathbf{u}_{k-1}$ are called *vector components* of the optimal representation. Here the second condition excludes the trivial $\boldsymbol{D}^{-1/2}\mathbf{u}_0 = \mathbf{1}$ vector.

Now we will use the normalized Laplacian matrix to find so-called minimum normalized cuts of edge-weighted graphs. Normalized cuts also favor balanced partitions, but the balancing is in terms of the cluster-volumes defined by the generalized degrees.

Definition 5 *Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph with generalized degrees $d_1, \ldots, d_n$ and assume that $\sum_{i=1}^n d_i = 1$. For the vertex-subset $U \subset V$ let $\mathrm{Vol}(U) = \sum_{i \in U} d_i$ denote the volume of U. The k-way normalized cut of G corresponding to the k-partition $P_k = (V_1, \ldots, V_k)$ of V is defined by*

$$\begin{aligned} f(P_k, G) &= \sum_{a=1}^{k-1} \sum_{b=a+1}^{k} \left(\frac{1}{\mathrm{Vol}(V_a)} + \frac{1}{\mathrm{Vol}(V_b)} \right) w(V_a, V_b) \\ &= \sum_{a=1}^{k} \frac{w(V_a, \overline{V}_a)}{\mathrm{Vol}(V_a)} = k - \sum_{a=1}^{k} \frac{w(V_a, V_a)}{\mathrm{Vol}(V_a)}. \end{aligned} \tag{1.11}$$

The minimum k-way normalized cut of G is

$$f_k(G) = \min_{P_k \in \mathcal{P}_k} f(P_k, G). \tag{1.12}$$

Apparently, $f_k(G)$ punishes k-partitions with 'many' inter-cluster edges of 'large' weights and with 'strongly' differing cluster volumes. The quantity $f_2(G)$ was introduced in [Moh89] for simple graphs and in [Mei-Shi] for edge-weighted graphs; further, for a general k, $f_k(G)$ is discussed in [Az-Gh] and [Bol-Mol02], where we called it k-density of G. Now, $f_k(G)$ will be related to the k smallest normalized Laplacian eigenvalues.

Theorem 4 ([Bol-Mol02]) *Assume that $G = (V, \boldsymbol{W})$ is connected and let $0 = \lambda_0 < \lambda_1 \leq \cdots \leq \lambda_{n-1} \leq 2$ denote the eigenvalues of its normalized Laplacian matrix. Then*

$$f_k(G) \geq \sum_{i=1}^{k-1} \lambda_i \tag{1.13}$$

and in the case when the optimal k-dimensional representatives of the vertices (see Theorem 3) can be classified into k well-separated clusters $V_1, \ldots, V_k$ in such a way that the maximum cluster diameter ε satisfies the relation $\varepsilon \leq \min\{1/\sqrt{2k}, \sqrt{2}\min_i \sqrt{p_i}\}$, where $p_i = \mathrm{Vol}(V_i)$, $i = 1, \ldots, k$, then

$$f_k(G) \leq c^2 \sum_{i=1}^{k-1} \lambda_i,$$

where $c = 1 + \varepsilon c'/(\sqrt{2} - \varepsilon c')$ and $c' = 1/\min_i \sqrt{p_i}$.

Note that the constant c of the upper estimate is greater than 1, and it is the closer to 1, the smaller ε is. The latter requirement is satisfied if there exists a 'very' well-separated k-partition of the k-dimensional Euclidean representatives. From Theorem 4 we can also conclude that the gap in the spectrum is a necessary but not a sufficient condition of a good classification. In addition, the Euclidean representatives should be well classified in the appropriate dimension. When the representatives $\mathbf{r}_1, \ldots, \mathbf{r}_n$ are endowed with the positive weights $d_1, \ldots, d_n$ ($\sum_{i=1}^n d_i = 1$ is assumed), then their *weighted k-variance* is defined as

$$\tilde{S}_k^2(\mathbf{r}_1, \ldots, \mathbf{r}_n) = \min_{P_k \in \mathcal{P}_k} \tilde{S}_k^2(P_k; \mathbf{r}_1, \ldots, \mathbf{r}_n) = \min_{P_k = (V_1, \ldots, V_k)} \sum_{a=1}^{k} \sum_{j \in V_a} d_j \|\mathbf{r}_j - \mathbf{c}_a\|^2 \tag{1.14}$$

where $\mathbf{c}_a = \frac{1}{\mathrm{Vol}(V_a)} \sum_{j \in V_a} d_j \mathbf{r}_j$ is the weighted center of cluster V_a ($a = 1, \ldots, k$). In the possession of k-dimensional representatives, we look for k clusters. Since the first coordinates of the representatives are coordinates of the vector $\boldsymbol{D}^{-1/2}\mathbf{u}_0 = \mathbf{1}$, they can as well be omitted.

In [Bol-Tus94], we directly estimated the weighted 2-variance of the optimal vertex representatives by the ratio of the two smallest positive normalized Laplacian eigenvalues (the proof is to be found in the cited papers).

Theorem 5 ([Bol-Tus94, Bol-Tus00]) *Let $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{n-1}$ be the eigenvalues of $\boldsymbol{L}_D$ with unit-norm eigenvectors $\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{n-1}$ and $\boldsymbol{D}$ be the diagonal degree-matrix. Then the 2-variance of the optimal 1-dimensional representatives $r_1^*, \ldots r_n^*$ (coordinates of the vector $\boldsymbol{D}^{-1/2}\mathbf{u}_1$) can be estimated from above as $\tilde{S}_2^2(r_1^*, \ldots, r_n^*) \leq \lambda_1/\lambda_2$.*

Theorem 5 indicates the following two-clustering property of the two smallest positive normalized Laplacian eigenvalues: the greater the gap between them, the better the optimal representatives of the vertices can be classified into two clusters. This fact, via Theorem 4, implies that the gap between the eigenvalues λ_1 and λ_2 of $\boldsymbol{L}_D$ is sufficient for the graph to have a small 2-way normalized cut. Note that this is not the usual spectral gap, which is between $\lambda_0 = 0$ and λ_1. For $k > 2$, the situation is more complicated and will be discussed in Chapter 2.

1.1.4 The isoperimetric number

For $k = 2$, the normalized cut is the symmetric version of the isoperimetric number (sometimes called Cheeger constant) introduced in the context of Riemannian manifolds and mathematical physics, see, e.g., [Che]. There is a wide literature of this topic together with expander graphs, see , e.g., [Alon, Ho-Lin-Wid, Moh88, Moh89] and [Chu], for a summary. We just discuss the most important relations of this topic to the normalized cut and clustering.

Definition 6 *Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph with generalized degrees $d_1, \ldots, d_n$ and assume that $\sum_{i=1}^n d_i = 1$. The isoperimetric number (Cheeger constant) of G is*

$$h(G) = \min_{\substack{\emptyset \neq U \subset V \\ \mathrm{Vol}(U) \leq \frac{1}{2}}} \frac{w(U, \overline{U})}{\mathrm{Vol}(U)}. \tag{1.15}$$

Since $\mathrm{Vol}(U)$ is the sum of the weights of edges emanating from U, while $w(U, \overline{U})$ is sum of the weights of those connecting U and $\overline{U}$, the relation $0 \leq h(G) \leq 1$ is trivial. Further, $h(G) = 0$ if and only if G is disconnected; therefore, only isoperimetric number of a connected graph is of interest. The isoperimetric number will later be considered as conditional probability, but first we investigate its relation to the smallest positive normalized Laplacian eigenvalue. Note that for simple graphs, $h(G)$ is not identical to the combinatorial isoperimetric number $i(G)$ which uses the cardinality of the subsets instead of their volumes in the denominator of (1.15), and hence, can exceed 1, see [Moh89] for details.

Intuitively, $h(G)$ is 'small' if 'few low-weight' edges connect together two disjoint vertex-subsets (forming a partition of the vertices) with 'not significantly' differing volumes; therefore, a 'small' $h(G)$ is an indication for a sparse cut of G. On the contrary, a 'large' $h(G)$ means that any vertex-subset of G has a large boundary compared to its volume, where the boundary of $U \subset V$ is the weighted cut between U and its complement in V. This is called good edge-expanding property of G, but here we do not want to give the exact definition of an expander graph which depends on many parameters and discussed in details (distinguishing between edge- and vertex-expansion) in many other places, see e.g. [Alon, Ho-Lin-Wid].

Now, a two-sided relation between $h(G)$ and the normalized Laplacian eigenvalue λ_1 is stated for edge-weighted graphs in the following theorem. Similar statements are proved in [Chu, Moh89] for simple graphs and in [Sin-Jer] for edge-weighted graphs, but without the upcoming improved upper bound.

Theorem 6 ([Bol-Mol04]) ***Improved Cheeger inequality.*** *Let $G = (V, \boldsymbol{W})$ be a connected edge-weighted graph with isoperimetric number $h(G)$, and let λ_1 denote the smallest positive eigenvalue of its normalized Laplacian $\boldsymbol{L}_D$. Then*

$$\frac{\lambda_1}{2} \leq h(G) \leq \min\{1, \sqrt{2\lambda_1}\}$$

always holds true. Furthermore, provided $\lambda_1 \leq 1$, the upper estimate improves to

$$h(G) \leq \sqrt{\lambda_1(2-\lambda_1)}.$$

Note that $\lambda_1 \leq 1$ is not a peculiar requirement. Indeed, Proposition 1 of the forthcoming Section 1.1.5 implies that $\lambda_1 \leq 1$ whenever G is not a soft-core weighted graph, i.e., it has at least one 0 weight.

An important application of the isoperimetric inequality is related in many aspects to random walks, see, e.g., [Chu, Lov93, Lux]. In fact, time-reversible Markov chains can be viewed as random walks on undirected, possibly edge-weighted graphs ($\boldsymbol{W}$ is symmetric). The walk can be described by a discrete time stochastic process $\xi_0, \xi_1, \ldots, \xi_t, \ldots$ with finite state space $\{1, \ldots, n\}$. The transition probabilities

$$\mathbb{P}(\xi_{t+1} = j \,|\, \xi_t = i) = \frac{w_{ij}}{d_i}$$

do not depend on t and are entries of the *transition probability matrix* $\boldsymbol{D}^{-1}\boldsymbol{W}$. The transition probability matrix is not symmetric, but its spectrum is the same as that of the symmetric matrix $\boldsymbol{W}_D$. Therefore, the transition probability matrix has real eigenvalues in the $[-1, 1]$ interval, they are the numbers $1 - \lambda_i$, where λ_i is the i-th largest eigenvalue of $\boldsymbol{L}_D$, and corresponding eigenvectors which are the vector components of the optimal representation of Theorem 3. Further, its largest eigenvalue is always 1 with corresponding eigendirection $\mathbf{1}$, and the multiplicity of 1 as an eigenvalue is equal to the number of the connected components of G. The random walk is *ergodic* if it has a unique stationary distribution. The necessary and sufficient condition of ergodicity is the *irreducibility* ($\lambda_1 > 0$) and *aperiodicity* ($\lambda_{n-1} < 2$). Therefore, the random walk on a connected and non-bipartite graph exhibits a unique stationary distribution which is just $\{d_1, \ldots, d_n\}$. The so-called *mixing rate* shows how rapidly the random walk converges to this stationary distribution. By the Cheeger inequality (Theorem 6) it follows that a relatively large λ_1 induces rapid mixing and short *cover time* which is the expected number of time to reach every vertex, see [Lov93].

We are rather interested in the case when λ_1 is near zero and the random walk cannot go through quickly the graph because of *bottlenecks* in it. Such a bottleneck can be the weighted cut between two disjoint vertex-subsets which gives the minimum in the definition of $h(G)$. More generally, if there are $k-1$ near zero eigenvalues of $\boldsymbol{L}_D$, then we may expect k clusters such that the random walk stays with high probability within the clusters and goes through between them with smaller probability.

In this direction, recent progress has been made by Luca Trevisan and coauthors [Gh-Trev, Le-Gh-Trev] who proved so-called *higher order Cheeger inequalities*. An important consequence of this type of results is summarized in the abstract of Luca Trevisan's talk (University of Illinois at Chicago, November 3, 2014): "*If λ_{k-1} is 'small' and λ_k is 'large', then the vertices can be partitioned into k subsets such that each subset defines a sparse cut and each subset induces an expanding subgraph. This points to a rigorous justification for the good performance of spectral clustering algorithms.*" Trevisan [Trev] also considered the upper end of the normalized Laplacian spectrum and its relation to the so-called bipartiteness ratio of the graph (called dual Cheeger inequality). However, their partitions do not exhaust necessarily the whole vertex-set.

In Chapter 2, we will use both ends of the normalized Laplacian spectrum and find cluster pairs of which 'sparse' cuts or 'dense' cuts (which do not necessarily exhaust V) of the aforementioned papers are special cases. Since it is a bit uncomfortable to treat eigenvalues near 0 or 2, we rather transform the normalized Laplacian matrix into a one with eigenvalues between -1 and 1, and consider its largest absolute value eigenvalues. For this purpose, we will introduce the normalized version of the forthcoming modularity matrix.

1.1.5 Modularity matrices and the Newman–Girvan modularity

The *modularity matrix* $\boldsymbol{M}$ was defined by Newman and Girvan [New-Gir, New] for simple graphs and naturally extends to edge-weighted graphs (see [Bol11c]) as

$$\boldsymbol{M} = \boldsymbol{W} - \mathbf{d}\mathbf{d}^T. \tag{1.16}$$

It is important that the edge-weight matrix $\boldsymbol{W}$ is normalized so that the sum of its entries is 1. The (i, j) entry of $\boldsymbol{M}$ just measures the deviation of w_{ij} (actual connection of vertices i and j) from $d_i d_j$ (their connection under independent attachment with the vertex-degrees as probabilities). It is easy to see that 0 is always an eigenvalue of $\boldsymbol{M}$ with corresponding eigendirection $\mathbf{1}$. However, it is not true that the modularity spectrum of a disconnected graph is the union of modularity spectra of its components, and sparse cuts are not related immediately to the eigenvalues of this modularity matrix. In case of simple graphs, $\boldsymbol{M}$ is usually indefinite, and it is negative semidefinite only for complete or complete multipartite graphs, see the forthcoming Theorem 8.

In [Bol11c], we also introduced the following normalized version of the modularity matrix, that will be intensively used in the discrepancy estimations of Chapter 2.

Definition 7 *Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph with the entries of $\boldsymbol{W}$ summing up to 1. The matrix*

$$\boldsymbol{M}_D = \boldsymbol{D}^{-1/2}\boldsymbol{M}\boldsymbol{D}^{-1/2} = \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2} - \sqrt{\mathbf{d}}\sqrt{\mathbf{d}}^T = \boldsymbol{W}_D - \sqrt{\mathbf{d}}\sqrt{\mathbf{d}}^T \tag{1.17}$$

is called normalized modularity matrix of G, where $\sqrt{\mathbf{d}} = (d_1, \dots, d_n)^T$.

It is easy to see that the eigenvalues of the normalized edge-weight matrix $\boldsymbol{W}_D = \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}$ are in the $[-1, 1]$ interval and the largest eigenvalue is always 1 with corresponding unit-norm eigenvector $\sqrt{\mathbf{d}}$. The only non-zero eigenvalue of the rank 1 term $\sqrt{\mathbf{d}}\sqrt{\mathbf{d}}^T$ is also 1 with the same eigenvector. Therefore, the spectrum of the matrix $\boldsymbol{M}_D$ is the same as the spectrum of $\boldsymbol{W}_D$, with the only exception that – due to the subtraction of the term $\sqrt{\mathbf{d}}\sqrt{\mathbf{d}}^T$ – the eigenvalue 1 of $\boldsymbol{W}_D$ becomes an eigenvalue 0 of $\boldsymbol{M}_D$ with eigenvector $\sqrt{\mathbf{d}}$. Hence, the spectrum of $\boldsymbol{M}_D$ is in $[-1, 1]$ and includes the 0.

These considerations give an exact relation between the normalized Laplacian and the normalized modularity spectrum. If the eigenvalues of $\boldsymbol{L}_D$ are $0 = \lambda_0 \le \lambda_1 \cdots \le \lambda_{n-1} \le 2$, then the spectrum of $\boldsymbol{M}_D$ consists of the numbers $\mu_i = 1 - \lambda_i$ $(i = 1, \dots, n-1)$ and $\mu_n = 0$ with corresponding eigenvector $\sqrt{\mathbf{d}}$. Further, the multiplicity of 0 is one more than the multiplicity of the eigenvalue 1 of $\boldsymbol{L}_D$. The multiplicity of 1 is one less than multiplicity of the eigenvalue 0 of $\boldsymbol{L}_D$; hence, 1 cannot be an eigenvalue of $\boldsymbol{M}_D$ if G is connected ($\boldsymbol{W}$ is irreducible).

In terms of the normalized modularity matrix, the minimization problem of Section 1.1.3

can be formulated as a maximization task in the following way.

$$\begin{aligned}
&\max_{\boldsymbol{X}^T\boldsymbol{D}\boldsymbol{X}=\boldsymbol{I}_{k-1}} \mathrm{tr}[(\boldsymbol{D}^{1/2}\boldsymbol{X})^T\boldsymbol{M}_D(\boldsymbol{D}^{1/2}\boldsymbol{X})]\\
=&\max_{\substack{\boldsymbol{X}^T\boldsymbol{D}\boldsymbol{X}=\boldsymbol{I}_{k-1}\\ \boldsymbol{X}^T\boldsymbol{D}\mathbf{1}=\mathbf{0}}} \mathrm{tr}[(\boldsymbol{D}^{1/2}\boldsymbol{X})^T(\boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2})(\boldsymbol{D}^{1/2}\boldsymbol{X})]\\
=&k-1-\min_{\substack{\boldsymbol{X}^T\boldsymbol{D}\boldsymbol{X}=\boldsymbol{I}_{k-1}\\ \boldsymbol{X}^T\boldsymbol{D}\mathbf{1}=\mathbf{0}}} \mathrm{tr}[(\boldsymbol{D}^{1/2}\boldsymbol{X})^T(\boldsymbol{I}_n-\boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2})(\boldsymbol{D}^{1/2}\boldsymbol{X})]\\
=&k-1-\min_{\substack{\sum_{i=1}^n d_i\mathbf{r}_i\mathbf{r}_i^T=\boldsymbol{I}_{k-1}\\ \sum_{i=1}^n d_i\mathbf{r}_i=\mathbf{0}}} Q_{k-1}.
\end{aligned}$$

The maximum is $k-1-\sum_{i=1}^{k-1}\lambda_i=\sum_{i=1}^{k-1}(1-\lambda_i)=\sum_{i=1}^{k-1}\mu_i$, that is the sum of the $k-1$ largest eigenvalues of $\boldsymbol{M}_D$.

The Newman–Girvan modularity directly focuses on modules of higher intra-community connections than expected based on the model of independent attachment of the vertices with probabilities proportional to their degrees. To maximize this modularity, many heuristic algorithms were recommended in the social network literature.

In [Bol11c], we extended the linear algebraic machinery developed for the Laplacian based spectral clustering to the modularity based community detection; further, introduced the notions of the balanced and normalized Newman–Girvan modularities. These considerations gave useful information on the choice of k and on the nature of the community structure in social networks.

Definition 8 *The Newman-Girvan modularity corresponding to the k-partition $P_k=(V_1,\ldots,V_k)$ of the vertex-set of the edge-weighted graph $G=(V,\boldsymbol{W})$, where the entries of $\boldsymbol{W}$ sum to 1, is*

$$M(P_k,G)=\sum_{a=1}^{k}\sum_{i,j\in V_a}(w_{ij}-d_id_j)=\sum_{a=1}^{k}[w(V_a,V_a)-\mathrm{Vol}^2(V_a)].$$

For given integer $1\le k\le n$, the k-module Newman-Girvan modularity of the edge-weighted graph G is

$$M_k(G)=\max_{P_k\in\mathcal{P}_k} M(P_k,G).$$

The entries d_id_j of the null-model matrix $\mathbf{d}\mathbf{d}^T$ correspond to the hypothesis of independence. In other words, under the null-hypothesis, vertices i and j are connected to each other independently, with probability d_id_j proportional (actually, because the sum of the weights is 1, equal) to their generalized degrees ($i,j=1,\ldots,n$). Hence, for given k, maximizing $M(P_k,G)$ is equivalent to looking for k modules of the vertices with intra-community connections higher than expected under the null-hypothesis.

We want to penalize partitions with clusters of extremely different sizes. To measure the size of cluster V_a, either the number of its vertices $|V_a|$ or its volume $\mathrm{Vol}(V_a)$ is used.

Definition 9 *The balanced Newman-Girvan modularity corresponding to the k-partition $P_k=(V_1,\ldots,V_k)$ of the vertex-set of $G=(V,\boldsymbol{W})$ ($\mathrm{Vol}(V)=1$) is*

$$BM(P_k,G)=\sum_{a=1}^{k}\frac{1}{|V_a|}\sum_{i,j\in V_a}(w_{ij}-d_id_j)=\sum_{a=1}^{k}\left[\frac{w(V_a,V_a)}{|V_a|}-\frac{\mathrm{Vol}^2(V_a)}{|V_a|}\right]$$

and the balanced k-module Newman-Girvan modularity of G is

$$BM_k(G)=\max_{P_k\in\mathcal{P}_k} BM(P_k,G).$$

Definition 10 *The normalized Newman-Girvan modularity corresponding to the k-partition $P_k = (V_1, \dots, V_k)$ of the vertex-set of $G = (V, \boldsymbol{W})$ ($\mathrm{Vol}(V) = 1$) is*

$$NM(P_k, G) = \sum_{a=1}^{k} \frac{1}{\mathrm{Vol}(V_a)} \sum_{i,j \in V_a} (w_{ij} - d_i d_j) = \sum_{a=1}^{k} \frac{w(V_a, V_a)}{\mathrm{Vol}(V_a)} - 1$$

and he normalized k-module Newman-Girvan modularity of G is

$$NM_k(G) = \max_{P_k \in \mathcal{P}_k} NM(P_k, G).$$

Here we used the fact that $\sum_{a=1}^{k} \mathrm{Vol}(V_a) = 1$. In view of (1.11), minimizing the normalized cut of G over k-partitions of its vertices is equivalent to maximizing the normalized Newman–Girvan modularity.

With similar techniques that we used in the previous sections, in [Bol11c], we maximized the balanced and the normalized Newman–Girvan modularities via minimizing the k-variance of the vertex representatives by choosing an appropriate representation, for which we used the unweighted and weighted k-means algorithm, respectively. Note that for the vertex representation, the eigenvectors were also multiplied with the squareroot of the absolute value of the corresponding eigenvalue.

We showed that both $BM_k(G)$ and $NM_k(G)$ is maximal for $k = pos$, where pos is the number of the positive eigenvalues of $\boldsymbol{M}$ and $\boldsymbol{M}_D$ (the two matrices have the same inertia). However, when n is 'large', it suffices to select a $k \leq pos$ such that there is a sudden gap between the $(k-1)$-th and k-th eigenvalues of $\boldsymbol{M}$ or $\boldsymbol{M}_D$, in decreasing order. For this k, we may say that there is a k-module *community structure* in the network: the clusters obtained have high intra- and low inter-cluster relations (higher and lower than expected in a random graph); for example, groups of strongly linked users or synopses of the brain, and groups of agents in strategic interaction networks following similar strategies when there are complementarities between them.

Likewise, both $BM_k(G)$ and $NM_k(G)$ is minimal for $k = neg$, where neg is the number of the negative eigenvalues of $\boldsymbol{M}$ and $\boldsymbol{M}_D$. However, when n is 'large', it suffices to select a $k \leq neg$ such that there is a sudden gap between the $k-1$-th and k-th eigenvalues of $\boldsymbol{M}$ or $\boldsymbol{M}_D$, in increasing order. For this k, we may say that there is a k-module *anticommunity structure* in the network: the clusters obtained have low intra- and high inter-cluster relations (lower and higher than expected in a random graph); for example, hub authorities or groups of agents in strategic interaction networks following similar strategies when there are substitute strategies between them (there are free-riders who do not buy the same goods as the neighbors, rather borrow those from them).

In [Bol11c], we showed some real-life examples and calculated the eigenvalues of $\boldsymbol{M}$ and $\boldsymbol{M}_D$ of some special structures that will be further investigated in Chapter 2 (see also Table 2.1). We experienced that the normalized modularity is best applicable for graphs which are far not regular (their generalized degrees differ significantly). In summary, the advantage of the modularity matrix versus the Laplacian is that here 0 is a natural divide, and the sign and the magnitude of the so-called structural eigenvalues (see Chapter 2) decide the type of the network modules: large positive eigenvalues of the modularity matrix are indications of a community, while large absolute value negative ones, of an anticommunity structure.

Together with Katalin Friedl and BSM students, in [Boletal15], we proved the following spectral properties of the modularity matrices. The statements are about the modularity spectra of complete and complete multipartite graphs, and those of their edge-weighted

analogues. A weighted graph is called *soft-core* if all its edge-weights are strictly positive (see, e.g., [Borgsetal1]). Likewise, we call a weighted graph *soft-core k-partite* with $2 \leq k \leq n$ clusters $V_1, \ldots, V_k$ (they form a proper partition of the vertices) if its edge-weights are

$$w_{ij} = \begin{cases} \text{positive} & \text{if} \quad c_i \neq c_j \\ 0 & \text{if} \quad c_i = c_j, \end{cases}$$

where c_i is the cluster membership of vertex i. Here the non-empty, disjoint vertex-subsets also form maximal independent sets of the vertices with zero-weighted edges within, and positively weighted edges between them. Note that from some intrinsic point of view (for example, from the point of view of the rank), only the position of the zeros are important, and not the exact values of the non-zero entries; see, e.g., [Shie-Pea] about the generic properties of matrices, which hold for every typical matrix (Lebesgue almost everywhere).

Proposition 1 ([Boletal15]) *If the connected weighted graph $G = (V, \boldsymbol{W})$ has an independent vertex-set of size $1 < k < n$, then its $\mu_{k-1} \geq 0$, where μ_i's are the eigenvalues of $\boldsymbol{M}_D$ in non-increasing order.*

We enclose the short proof so that to illustrate the application of the minimax principle.
Proof. Without loss of generality, assume that $w_{ij} = 0$ when $1 \leq i, j \leq k$. Since μ_{k-1} is the k-th largest eigenvalue (including the trivial $\mu_0 = 1$) of $\boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}$, the Courant–Fischer–Weyl minimax principle yields that

$$\mu_{k-1} = \max_{\substack{F \subset \mathbb{R}^n \\ \dim(F)=k}} \min_{\substack{\mathbf{x} \in F \\ \|\mathbf{x}\|=1}} \mathbf{x}^T \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}\mathbf{x}.$$

Therefore, to prove that $\mu_{k-1} \geq 0$, it suffices to find a k-dimensional subspace $F \subset \mathbb{R}^n$ such that $\min_{\substack{\mathbf{x} \in F \\ \|\mathbf{x}\|=1}} \mathbf{x}^T \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}\mathbf{x} = 0$. Set $F := \{\mathbf{x} : \mathbf{x} = (x_1, \ldots, x_k, 0, \ldots, 0) \in \mathbb{R}^n\}$. Clearly, for every $\mathbf{x} \in F$: $\mathbf{x}^T \boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}\mathbf{x} = 0$, and this also holds true for unit-norm $\mathbf{x}$'s. Therefore, the above minimum is also 0. This finishes the proof. □

By Proposition 1, the case $k = 2$ implies that $\mu_1 \geq 0$, or equivalently, $\lambda_1 \leq 1$ whenever G is not a soft-core weighted graph, i.e., it has at least one 0 weight. If this is the case, the improved Cheeger inequality of Theorem 6 is applicable.

Proposition 2 ([Boletal15]) *The modularity spectrum of the complete multipartite graph $K_{n_1,\ldots,n_k}$ consists of $k-1$ strictly negative eigenvalues and zero with multiplicity $n-k+1$.*

To prove the further statements, we will extensively use the following well-known characterization of the complete multipartite graphs (including the complete graphs): an unweighted connected graph is complete multipartite if and only if it has no three-vertex induced subgraph with exactly one edge. More generally, we are able to give a similar characterization of weighted soft-core multipartite graphs.

Lemma 1 ([Boletal15]) *A weighted graph is soft-core multipartite if and only if it has no triangle with exactly one positively weighted edge.*

Such a triangle is called *forbidden pattern.*

Theorem 7 ([Boletal15]) *If the connected weighted graph $G = (V, \boldsymbol{W})$ is not soft-core multipartite, then the largest eigenvalue of its normalized modularity matrix is strictly positive.*

We also enclose the short proof so that to illustrate the application of the representation technique and the minimax principle.

Proof. The largest eigenvalue μ_1 of $\boldsymbol{M}_D$ is the second largest eigenvalue of $\boldsymbol{W}_D$, whose largest eigenvalue is 1 with corresponding eigenvector $\sqrt{\mathbf{d}}$ (this is single if our graph is connected). Therefore, we think in terms of the two largest eigenvalues of $\boldsymbol{W}_D$. We can again assume that the first three vertices form the forbidden pattern and so, the upper left corner of this matrix looks like

$$\begin{pmatrix} 0 & \frac{w_{12}}{\sqrt{d_1 d_2}} & 0 \\ \frac{w_{21}}{\sqrt{d_1 d_2}} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

with $w_{12} = w_{21} > 0$.

Then the Courant–Fischer–Weyl minimax principle yields

$$\mu_1 = \max_{\substack{\|\mathbf{x}\|=1 \\ \mathbf{x}^T\sqrt{\mathbf{d}}=0}} \mathbf{x}^T \boldsymbol{D}^{-1/2} \boldsymbol{W} \boldsymbol{D}^{-1/2} \mathbf{x}.$$

Therefore, to prove that $\mu_1 > 0$, it suffices to find an $\mathbf{x} \in \mathbb{R}^n$ that satisfies conditions $\|\mathbf{x}\| = 1$, $\mathbf{x}^T\sqrt{\mathbf{d}} = 0$, and for which, $\mathbf{x}^T \boldsymbol{D}^{-1/2} \boldsymbol{W} \boldsymbol{D}^{-1/2} \mathbf{x} > 0$. (The unit norm condition can be relaxed here, because $\mathbf{x}$ can later be normalized, without changing the sign of the above quadratic form.)

Indeed, let us look for $\mathbf{x}$ of the form $\mathbf{x} = (x_1, x_2, x_3, 0, \ldots, 0)^T$ such that

$$\sqrt{d_1} x_1 + \sqrt{d_2} x_2 + \sqrt{d_3} x_3 = 0. \tag{1.18}$$

Then the inequality

$$\mathbf{x}^T \boldsymbol{D}^{-1/2} \boldsymbol{W} \boldsymbol{D}^{-1/2} \mathbf{x} = \frac{2 x_1 x_2 w_{12}}{\sqrt{d_1 d_2}} > 0$$

can be satisfied with any $\mathbf{x} = (x_1, x_2, x_3, 0, \ldots, 0)^T$ such that x_1 and x_2 are both positive or both negative, in which case, due to (1.18),

$$x_3 = -\frac{\sqrt{d_1} x_1 + \sqrt{d_2} x_2}{\sqrt{d_3}}$$

is a good choice, and will have the opposite sign. (Note that all the d_i's are positive, since we deal with connected weighted graphs.) This finishes the proof. □

Since, $\boldsymbol{M}$ and $\boldsymbol{M}_D$ have the same inertia, Theorem 7 together with Proposition 2 gives the following statement of equivalence.

Theorem 8 ([Boletal15]) *The modularity and normalized modularity matrix of a simple connected graph is negative semidefinite if and only if it is complete multipartite.*

Note that complete graphs are also understood, since they are complete multipartite with singleton clusters.

1.2 Biclustering of contingency tables

1.2.1 SVD of normalized contingency tables

Now, more generally, our underlying objects will be rectangular arrays of nonnegative entries. They may contain frequency counts for the joint distribution of two discrete random variables

taking on finitely many values (the values can as well be textual, in which case the variables are called categorical). For example, keyword–document matrices or microarrays are such. In microarrays, rows correspond to the genes and columns to different conditions, while the corresponding entries are expression levels of genes under specific conditions (a 0-1 matrix is a special case of it). Let $\boldsymbol{C}$ be a contingency table on row set $Row = \{1, \ldots, m\}$, column set $Col = \{1, \ldots, n\}$, where $\boldsymbol{C}$ is $m \times n$ rectangular matrix of nonnegative real entries c_{ij}'s. Without loss of generality, we can assume that there are no identically zero rows or columns (otherwise they are omitted). Here c_{ij} is some kind of association between the objects or categories corresponding to row i and column j, where 0 means no interaction at all. Usually, the entries of $\boldsymbol{C}$ are normalized, either with a uniform bound, say 1 (like probabilities), or the sum of the entries is 1 (reminiscent of a joint distribution). This normalization will have importance in Section 1.3. Here it has no relevance, since the normalized table to be introduced is invariant under scaling the entries of $\boldsymbol{C}$. Let the row-sums of $\boldsymbol{C}$ be

$$d_{row,i} = \sum_{j=1}^{n} c_{ij}, \quad i = 1, \ldots, m \tag{1.19}$$

and the column-sums

$$d_{col,j} = \sum_{i=1}^{m} c_{ij}, \quad j = 1, \ldots, n \tag{1.20}$$

which are collected in the main diagonal of the $m \times m$ diagonal matrix $\boldsymbol{D}_{row}$ and that of the $n \times n$ diagonal matrix $\boldsymbol{D}_{col}$, respectively.

For a given integer $1 \leq k \leq \min\{m, n\}$, we are looking for k-dimensional representatives $\mathbf{r}_1, \ldots, \mathbf{r}_m \in \mathbb{R}^k$ of the rows and $\mathbf{q}_1, \ldots, \mathbf{q}_n \in \mathbb{R}^k$ of the columns such that they minimize the objective function

$$Q_k = \sum_{i=1}^{m} \sum_{j=1}^{n} c_{ij} \|\mathbf{r}_i - \mathbf{q}_j\|^2 \tag{1.21}$$

subject to

$$\sum_{i=1}^{m} d_{row,i} \mathbf{r}_i \mathbf{r}_i^T = \boldsymbol{I}_k \quad \text{and} \quad \sum_{j=1}^{n} d_{col,j} \mathbf{q}_j \mathbf{q}_j^T = \boldsymbol{I}_k. \tag{1.22}$$

When minimized, the objective function Q_k favors k-dimensional placement of the rows and columns such that representatives of highly associated rows and columns are close to each other. This is equivalent to the problem of *correspondence analysis*.

Let us put both the objective function and the constraints in a more favorable form. Let $\boldsymbol{X}$ be the $m \times k$ matrix of rows $\mathbf{r}_1^T, \ldots, \mathbf{r}_m^T$, and $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^n$ denote the columns of $\boldsymbol{X}$, for which fact we use the notation $\boldsymbol{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$. Because of the constraint (1.22), the vectors $\boldsymbol{D}_{row}^{-1/2} \mathbf{x}_i$ $(i = 1, \ldots, k)$ form an orthonormal system, hence, $\boldsymbol{D}_{row}^{-1/2} \boldsymbol{X}$ is a suborthogonal matrix. Therefore, the first part of the constraint can be formulated as $\boldsymbol{X}^T \boldsymbol{D}_{row} \boldsymbol{X} = \boldsymbol{I}_k$. Likewise, let $\boldsymbol{Y}$ be the $n \times k$ matrix of rows $\mathbf{q}_1^T, \ldots, \mathbf{q}_n^T$, and $\boldsymbol{Y} := (\mathbf{y}_1, \ldots, \mathbf{y}_k)$. Hence, the second part of the constraint (1.22) can be formulated as $\boldsymbol{Y}^T \boldsymbol{D}_{col} \boldsymbol{Y} = \boldsymbol{I}_k$ and the matrix $\boldsymbol{D}_{col}^{-1/2} \boldsymbol{Y}$ is also suborthogonal.

With this notation, the objective function (1.21) is rewritten as

$$\begin{aligned} Q_k &= \sum_{i=1}^{m} d_{row,i}\|\mathbf{r}_i\|^2 + \sum_{j=1}^{n} d_{col,j}\|\mathbf{q}_j\|^2 - \sum_{i=1}^{m}\sum_{j=1}^{n} c_{ij}\mathbf{r}_i^T\mathbf{q}_j \\ &= \sum_{\ell=1}^{k} \mathbf{x}_\ell^T \boldsymbol{D}_{row}\mathbf{x}_\ell + \sum_{\ell=1}^{k} \mathbf{y}_\ell^T \boldsymbol{D}_{col}\mathbf{y}_\ell - \sum_{\ell=1}^{k} \mathbf{x}_\ell^T \boldsymbol{C}\mathbf{y}_\ell \\ &= \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{D}_{row}\boldsymbol{X}) + \mathrm{tr}(\boldsymbol{Y}^T\boldsymbol{D}_{col}\boldsymbol{Y}) - \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{C}\boldsymbol{Y}) \\ &= 2k - \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{C}\boldsymbol{Y}) = 2k - \mathrm{tr}[(\boldsymbol{D}_{row}^{1/2}\boldsymbol{X})^T\boldsymbol{C}_D(\boldsymbol{D}_{col}^{1/2}\boldsymbol{Y})], \end{aligned}$$

where the matrix $\boldsymbol{C}_D = \boldsymbol{D}_{row}^{-1/2}\boldsymbol{C}\boldsymbol{D}_{col}^{-1/2}$ is called *normalized contingency table*. Let

$$\boldsymbol{C}_D = \sum_{k=0}^{r-1} s_k \mathbf{v}_k \mathbf{u}_k^T \tag{1.23}$$

be the SVD of $\boldsymbol{C}_D$, where $r \leq \min\{m,n\}$ is the rank of $\boldsymbol{C}_D$, or equivalently (as there are no identically zero rows or columns), the rank of $\boldsymbol{C}$. Here $1 = s_0 \geq s_1 \geq \cdots \geq s_{r-1} > 0$ are the non-zero singular values of $\boldsymbol{C}_D$. They cannot exceed 1, since they are correlations (see Section 1.3). Furthermore, 1 is a single singular value if $\boldsymbol{C}_D$ (or equivalently, $\boldsymbol{C}$) is non-degenerate (or non-decomposable, with the wording of [Bol13]), i.e., when $\boldsymbol{C}\boldsymbol{C}^T$ (if $m \leq n$) or $\boldsymbol{C}^T\boldsymbol{C}$ (if $m > n$) is irreducible. In this case, $\mathbf{v}_0 = (\sqrt{d_{row,1}}, \ldots, \sqrt{d_{row,m}})^T$ and $\mathbf{u}_0 = (\sqrt{d_{col,1}}, \ldots, \sqrt{d_{col,n}})^T$ is the singular vector pair corresponding to $s_0 = 1$.

Note that the singular spectrum of a degenerate contingency table can be composed from the singular spectra of its non-degenerate parts, as well as their singular vector pairs. Therefore, in the future, the non-degenerate nature of the underlying contingency table will be assumed. With some simple linear algebra, the following can be proved.

Theorem 9 ([Bol14b]) ***Representation theorem for contingency tables.*** *Let $\boldsymbol{C}$ be a non-degenerate contingency table. Let $1 = s_0 > s_1 \geq \cdots \geq s_{r-1}$ denote the positive singular values of the normalized table $\boldsymbol{C}_D$ with unit-norm singular vector pairs $\mathbf{v}_i, \mathbf{u}_i$ $(i = 0, \ldots, r-1)$, and $k \leq r$ be a positive integer such that $s_{k-1} > s_k$. Then the minimum of (1.21) subject to (1.22) is $2k - \sum_{i=0}^{k-1} s_i$ and it is attained with the optimal k-dimensional row-representatives $\mathbf{r}_1^*, \ldots, \mathbf{r}_m^*$ and column-representatives $\mathbf{q}_1^*, \ldots, \mathbf{q}_n^*$ the transposes of which are row vectors of the matrices $\boldsymbol{X}^* = \boldsymbol{D}_{row}^{-1/2}(\mathbf{v}_0, \mathbf{v}_1, \ldots, \mathbf{v}_{k-1})$ and $\boldsymbol{Y}^* = \boldsymbol{D}_{col}^{-1/2}(\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{k-1})$, respectively.*

We remark the following.

- Provided 1 is a single singular value (when $\boldsymbol{C}$ is non-degenerate), the first columns of the matrices $\boldsymbol{X}^*$ and $\boldsymbol{Y}^*$ are $\boldsymbol{D}_{row}^{-1/2}\mathbf{v}_0$ and $\boldsymbol{D}_{col}^{-1/2}\mathbf{u}_0$, i.e., the constantly $\mathbf{1}$ vectors of $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively. Therefore they do not contribute to the separation of the representatives, and the k-dimensional representatives are in a $(k-1)$-dimensional hyperplane of $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively.

- Note that the dimension k does not play an important role here, the vector components can be included successively up to a k such that $s_{k-1} > s_k$. We remark that the singular vectors can arbitrarily be chosen in the isotropic subspaces corresponding to possible multiple singular values, under the orthogonality conditions.

- As for the joint distribution view (when the rows and columns belong to the categories of two categorical variables, see [Bol87b]), correspondence analysis uses the above $(k-1)$-dimensional row- and column-representatives for simultaneously plotting the row- and column-categories in $\mathbb{R}^{k-1}$ (with $k = 2, 3$ or 4 in most applications), and hence, the practitioner can draw conclusions from their mutual positions. Indeed, this representation has the following optimum properties: the closeness of categories of the same variable reflects the similarity between them, while the closeness of categories of the two different variables reflects their frequent simultaneous occurrence. For example, $\boldsymbol{C}$ being a microarray, the representatives of similar function genes as well as representatives of similar conditions are close to each other; likewise, representatives of genes that are responsible for a given condition are close to the representative of that condition.

- One frequently studied example of a rectangular array is the keyword–document matrix. Here the entries are associations between documents and words. Based on network data, the entry in the i-th row and j-th column is the relative frequency of word j in document i. Latent semantic indexing looks for real scores of the documents and keywords such that the score of a any document be proportional to the total scores of the keywords occurring in it, and vice versa, the score of any keyword be proportional to the total scores of the documents containing it. Not surprisingly, the solution is given by the SVD of the contingency table, where the document- and keyword-scores are the coordinates of the left and right singular vectors corresponding to its largest non-trivial singular value which gives the constant of proportionality. This idea can be generalized in the following way. We can think of the above relation between keywords and documents as the relation with respect to the most important topic (or context, or factor). After this, we are looking for another scoring with respect to the second topic, which is independent of the first one; and so on, up to k (where k is a positive integer not exceeding the rank of the table). This method is reminiscent of the principal component analysis, and in [Bol87b] we proved that the correspondence analysis indeed solves this factorization problem, together with spacial representations. The solution is given by the singular vector pairs corresponding to the k largest singular values of the table. The problem is also related to the Page-rank.

- In another view, a 0-1 contingency table can be considered as part of the adjacency matrix of a bipartite graph on vertex set $Row \cup Col$. However, it would be uncomfortable to always distinguish between these two types of vertices, I will rather use the framework of correspondence analysis, and formulate the statements in terms of the rows and columns.

1.2.2 Normalized bicuts of contingency tables

We are given an $m \times n$ contingency table $\boldsymbol{C}$ on row set Row and column set Col as introduced in the previous section. For a fixed integer k, $0 < k \leq r = \mathrm{rank}(\boldsymbol{C})$, we want to simultaneously partition the rows and columns of $\boldsymbol{C}$ into disjoint, nonempty subsets

$$Row = R_1 \cup \cdots \cup R_k, \quad Col = C_1 \cup \cdots \cup C_k$$

so that the cuts $c(R_a, C_b) = \sum_{i \in R_a} \sum_{j \in C_b} c_{ij}$, $a, b = 1, \ldots, k$ between the row-column cluster pairs be as homogeneous as possible.

Definition 11 *The normalized bicut of the contingency table $\boldsymbol{C}$ with respect to the k-partitions $P_{row} = (R_1, \ldots, R_k)$ and $P_{col} = (C_1, \ldots, C_k)$ of its rows and columns and the collection of signs σ is defined as follows:*

$$\nu_k(P_{row}, P_{col}, \sigma) = \sum_{a=1}^{k}\sum_{b=1}^{k}\left(\frac{1}{\mathrm{Vol}(R_a)} + \frac{1}{\mathrm{Vol}(C_b)} + \frac{2\sigma_{ab}\delta_{ab}}{\sqrt{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}}\right) c(R_a, C_b), \quad (1.24)$$

where

$$\mathrm{Vol}(R_a) = \sum_{i\in R_a} d_{row,i} = \sum_{i\in R_a}\sum_{j=1}^{n} c_{ij}, \quad \mathrm{Vol}(C_b) = \sum_{j\in C_b} d_{col,j} = \sum_{j\in C_b}\sum_{i=1}^{m} c_{ij}$$

are volumes of the clusters (also see formulas (1.19) and (1.20)), δ_{ab} is the Kronecker delta, and the sign σ_{ab} is equal to 1 or -1 (it only has relevance in the $a = b$ case), and $\sigma = (\sigma_{11}, \ldots, \sigma_{kk})$ is the collection of the relevant signs.

The normalized k-way bicut of the contingency table $\boldsymbol{C}$ is the minimum of (1.24) over all possible k-partitions $\mathcal{P}_{row,k}$ and $\mathcal{P}_{col,k}$ of its rows and columns, and over all possible collections of signs σ:

$$\nu_k(\boldsymbol{C}) = \min_{P_{row}, P_{col}, \sigma} \nu_k(P_{row}, P_{col}, \sigma).$$

Note that $\nu_k(\boldsymbol{C})$ penalizes row- and column clusters of extremely different volumes in the $a \neq b$ case, whereas in the $a = b$ case σ_{aa} moderates the balance between $\mathrm{Vol}(R_a)$ and $\mathrm{Vol}(C_a)$.

Theorem 10 ([Bol14b]) *Let $1 = s_0 \geq s_1 \geq \cdots \geq s_{r-1} > 0$ be the positive singular values of the normalized contingency table $\boldsymbol{C}_D = \boldsymbol{D}_{row}^{-1/2}\boldsymbol{C}\boldsymbol{D}_{col}^{-1/2}$ belonging to $\boldsymbol{C}$. Then for any positive integer $k \leq r$, such that $s_{k-1} > s_k$,*

$$\nu_k(\boldsymbol{C}) \geq 2k - \sum_{i=0}^{k-1} s_i.$$

Observe, that in the case of a symmetric table, we get the same result with the representation, based on the eigenvectors corresponding to the largest absolute value eigenvalues of the normalized modularity matrix. However, $\nu_k(P_{row}, P_{col}, \sigma)$ cannot always be directly related to the normalized cut, except in the following two special cases.

- When the $k-1$ largest absolute value eigenvalues of the normalized modularity matrix $\boldsymbol{M}_D$ are all positive, or equivalently, if the k smallest eigenvalues (including the zero) of the normalized Laplacian matrix are farther from 1 than any other eigenvalue which is greater than 1. In this case, the $k-1$ largest singular values (apart from the 1) of $\boldsymbol{C}_D$ are identical to the $k-1$ largest eigenvalues of $\boldsymbol{M}_D$, and the left and right singular vectors are identical to the corresponding eigenvector with the same orientation. Consequently, for the k-dimensional (in fact, $(k-1)$-dimensional) row- and column-representatives $\mathbf{r}_i = \mathbf{q}_i$ $(i = 1, \ldots, n = m)$ holds. With the choice $\sigma_{bb} = 1$ $(b = 1, \ldots, k)$, the corresponding $\nu_k(\boldsymbol{C})$ is twice the normalized cut of our weighted graph, where the weights of edges within the clusters do not count. In this special situation, the normalized bicut also favors k-partitions with low inter-cluster edge-densities (therefore, intra-cluster densities tend to be large, as they do not count in the objective function).

- When the $k-1$ largest absolute value eigenvalues of $\boldsymbol{M}_D$ are all negative, then $\mathbf{r}_i = -\mathbf{q}_i$ for all $(k-1)$-dimensional row and column representatives, and any (but only one) of them can be the corresponding vertex representative. Now $\nu_k(\boldsymbol{C})$, which is attained with the choice $\sigma_{bb} = -1$ $(b = 1, \ldots, k)$, differs from the normalized cut in that it also counts the edge-weights within the clusters. Indeed, in the $a = b$, $R_a = C_a = V_a$ case

$$\|\mathbf{r}_i - \mathbf{q}_j\|^2 = \frac{1}{\mathrm{Vol}(V_a)} + \frac{1}{\mathrm{Vol}(V_b)} + \frac{2}{\sqrt{\mathrm{Vol}(V_a)\mathrm{Vol}(V_b)}} = \frac{4}{\mathrm{Vol}(V_a)}$$

  if $i, j \in V_a$. Here, by minimizing the normalized k-way cut, rather a so-called anti-community structure (see Section 1.1.5) is detected in that $c(R_a, C_a) = c(V_a, V_a)$ is suppressed to compensate for the term $\frac{4}{\mathrm{Vol}(V_a)}$. This fact favors k-partitions of the vertices with low intra-cluster edge-densities.

In some real-life problems, e.g., clustering genes and conditions of microarrays, we rather want to find clusters of similarly functioning genes that equally (not especially weakly or strongly) influence conditions of the same cluster; this issue discussed in details in Chapter 2. Note that Dhillon [Dhil] also suggests a multipartition algorithm that runs the k-means algorithm simultaneously for the row- and column representatives, but not with our objective function behind it.

1.3 Representation of joint distributions

Now we will give an abstract description of the issues discussed in the previous sections in terms of two-variate distributions. With the help of joint distributions, representation can be discussed in a more general framework, of which graphs and contingency tables are special finite cases. This representation is reminiscent of the techniques of multivariate statistical analysis (principal component, canonical correlation, correspondence analysis). We will use some parts of this section in Chapter 3 when consider the continuous limit objects as kernels of integral operators taking conditional expectation with respect to the joint distribution.

1.3.1 General setup

Let (ξ, η) be a pair of real-valued random variables – neither of them being constant with probability 1 – defined over the product space $\mathcal{X} \times \mathcal{Y}$ having joint distribution $\mathbb{W}$ with marginal distributions $\mathbb{P}$ and $\mathbb{Q}$, respectively. Assume that the dependence between ξ and η is regular, i.e., their joint distribution $\mathbb{W}$ is absolutely continuous with respect to the product measure $\mathbb{P} \times \mathbb{Q}$, and let w denote the Radon–Nikodym derivative of $\mathbb{W}$ with respect to $\mathbb{P} \times \mathbb{Q}$, see Rényi [Reny59b] for details.

In the spirit of Breiman and Friedman [Bre-Fri], let $H = L^2(\xi)$ and $H' = L^2(\eta)$ be the set of random variables which are functions of ξ and η and have zero expectation and finite variance with respect to $\mathbb{P}$ and $\mathbb{Q}$, respectively. Both H and H' are Hilbert spaces with the covariance as inner product; further, they are embedded as subspaces into the L^2-space defined likewise by the (ξ, η) pair over the product space.

1.3.2 Integral operators

The following linear operators taking conditional expectation between the two margins (with respect to the joint distribution) will play a crucial role in the future, see also [Bre-Fri]. They

are integral operators defined as follows:

$$P_{\mathcal{X}} : H' \to H, \quad \psi = P_{\mathcal{X}}\phi = \mathbb{E}(\phi \,|\, \xi), \quad \psi(x) = \int_{\mathcal{Y}} w(x,y)\phi(y)\,\mathbb{Q}(dy)$$

and

$$P_{\mathcal{Y}} : H \to H', \quad \phi = P_{\mathcal{Y}}\psi = \mathbb{E}(\psi \,|\, \eta), \quad \phi(y) = \int_{\mathcal{X}} w(x,y)\psi(x)\,\mathbb{P}(dx).$$

It is easy to see that $P_{\mathcal{X}}^* = P_{\mathcal{Y}}$ and vice versa, because of the relation

$$\langle P_{\mathcal{X}}\phi, \psi\rangle_H = \langle P_{\mathcal{Y}}\psi, \phi\rangle_{H'} = \mathrm{Cov}_{\mathbb{W}}(\psi, \phi), \tag{1.25}$$

where $\langle ., . \rangle$ denotes the (real) inner product, while $\mathrm{Cov}_{\mathbb{W}}$ is the so-called *covariance function* with respect to the joint distribution $\mathbb{W}$, defined as

$$\mathrm{Cov}_{\mathbb{W}}(\psi, \phi) = \int_{\mathcal{X}\times\mathcal{Y}} \psi(x)\phi(y)\mathbb{W}(dx, dy) = \int_{\mathcal{X}}\int_{\mathcal{Y}} \psi(x)\phi(y)w(x,y)\mathbb{Q}(dy)\mathbb{P}(dx).$$

Assume that

$$\int_{\mathcal{X}}\int_{\mathcal{Y}} w^2(x,y)\mathbb{Q}(dy)\mathbb{P}(dx) < \infty. \tag{1.26}$$

Under this condition $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ are Hilbert–Schmidt operators, and therefore compact, with SVD-like decomposition

$$P_{\mathcal{X}} = \sum_{i=1}^{\infty} s_i \langle ., \phi_i\rangle_{H'}\psi_i, \quad P_{\mathcal{Y}} = \sum_{i=1}^{\infty} s_i \langle ., \psi_i\rangle_{H}\phi_i \tag{1.27}$$

where for the 'singular values' $1 > s_1 \geq s_2 \geq \cdots \geq 0$ holds (with the only point of accumulation 0 if they are countably infinitely many), since the operators $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ are in fact orthogonal projections from one margin onto the other, but the projections are restricted to the subspaces H and H', respectively (see [Bol87a] for details). The function pairs ψ_i, ϕ_i can be chosen (even in case of multiple singular values) so that $\{\psi_i\}_{i=1}^{\infty} \subset H$ and $\{\phi_i\}_{i=1}^{\infty} \subset H'$ form complete orthonormal systems. Since $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ are Hilbert–Schmidt operators, $\sum_{i=1}^{\infty} s_i^2 < \infty$. The decompositions of (1.27) are essentially unique (apart from function pairs corresponding to multiple singular values). We remark that denoting by ψ_0 and ϕ_0 the constantly 1 random variables, $\mathbb{E}(\phi_0|\xi) = \psi_0$ and $\mathbb{E}(\psi_0|\eta) = \phi_0$, however, this pair is not considered as a function pair with singular value $s_0 = 1$, since they have no zero expectation. Consequently, we will subtract 1 from the kernel, but with this new kernel, $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ will define the same integral operators. (This resembles the role of the normalized modularity matrix.)

Especially, if $\mathbb{W}$ is symmetric (H and H' are isomorphic in terms of the distributions too), then in view of (1.25), $P_{\mathcal{X}} = P_{\mathcal{Y}}$ is a selfadjoint linear operator, since

$$\langle P_{\mathcal{X}}\phi, \psi\rangle_H = \mathrm{Cov}_{\mathbb{W}}(\phi, \psi) = \mathrm{Cov}_{\mathbb{W}}(\psi, \phi) = \langle P_{\mathcal{Y}}\psi, \phi\rangle_{H'}.$$

The SD of $P_{\mathcal{X}} : H' \to H$ is

$$P_{\mathcal{X}} = \sum_{i=1}^{\infty} \lambda_i \langle ., \psi_i'\rangle_{H'}\psi_i.$$

Here for the eigenvalues, $|\lambda_i| \leq 1$ holds, and the eigenvalue–eigenfunction equation looks like

$$P_{\mathcal{X}}\psi_i' = \lambda_i\psi_i$$

where ψ_i and ψ_i' are identically distributed, whereas their joint distribution is $\mathbb{W}$ ($i = 1, 2, \ldots$).

1.3.3 Maximal correlation and optimal representations

From now on, we will intensively use analogues of separation theorems for the singular values and eigenvalues for matrices, see [Rao]. In view of these, the SVD (1.27) gives the solution of the following task of *maximal correlation*, posed by Gebelein [Geb] and Rényi [Reny59a]. We are looking for $\psi \in H$, $\phi \in H'$ such that their correlation is maximum with respect to their joint distribution $\mathbb{W}$. Using (1.25) and separation theorems,

$$\max_{\psi\in H\ \phi\in H'} \mathrm{Corr}_{\mathbb{W}}(\psi,\phi) = \max_{\|\psi\|=\|\phi\|=1} \mathrm{Cov}_{\mathbb{W}}(\psi,\phi) = s_1$$

and it is attained on the non-trivial ψ_1, ϕ_1 pair. In the finite, symmetric case, maximal correlation is related to some conditional probabilities in [Bol-Mol02].

The maximal correlation task is equivalent to the following:

$$\min_{\|\psi\|=\|\phi\|=1} \|\psi-\phi\|^2 = \min_{\|\psi\|=\|\phi\|=1} (\|\psi\|^2 + \|\phi\|^2 - 2\mathrm{Cov}_{\mathbb{W}}(\psi,\phi)) = 2(1-s_1). \tag{1.28}$$

Correspondence analysis ([Benz, Green], and [Bol87b]) is on the one hand, a special case of the problem of maximal correlation being $\mathcal{X}$ and $\mathcal{Y}$ finite sets, but on the other hand, it is a generalization in that we are successively finding maximal correlations under some orthogonality constraints.

The product space is now an $m \times n$ contingency table with row set $\mathcal{X} = \{1,\dots,m\}$ and column set $\mathcal{Y} = \{1,\dots,n\}$, whereas the entries $w_{ij} \geq 0$ ($\sum_{i=1}^m \sum_{j=1}^n w_{ij} = 1$) embody the joint distribution over the product space, with row-sums $p_1,\dots,p_m$ and column-sums $q_1,\dots,q_n$ as marginal distributions.

Hence, the effect of $P_{\mathcal{X}} : H' \to H$, $P_{\mathcal{X}}\phi = \psi$ is the following:

$$\psi(i) = \frac{1}{p_i}\sum_{j=1}^n w_{ij}\phi(j) = \sum_{j=1}^n \frac{w_{ij}}{p_i q_j}\phi(j)q_j, \quad i = 1,\dots,m. \tag{1.29}$$

Therefore, $P_{\mathcal{X}}$ is an integral operator with kernel $K_{ij} = \frac{w_{ij}}{p_i q_j}$ (instead of integration, we have summation with respect to the marginal measure $\mathbb{Q}$).

Consider the SVD

$$P_{\mathcal{X}} = \sum_{k=1}^{r-1} s_k \langle .,\phi_k\rangle_{H'}\psi_k,$$

where r is the now finite rank of the contingency table ($r \leq \min\{n,m\}$). The singular value $s_0 = 1$ with the trivial $\psi_0 = 1$, $\phi_0 = 1$ factor pair is disregarded as their expectation is 1 with respect to the $\mathbb{P}$- and $\mathbb{Q}$-measures, respectively; therefore, the summation starts from 1. If we used the kernel $K_{ij} - 1$, we could eliminate the trivial factors. Assume that there is no other singular value 1, i.e., our contingency table is non-degenerate. Then, by the orthogonality, the subsequent left- and right-hand side singular functions have zero expectation with respect to the $\mathbb{P}$- and $\mathbb{Q}$-measures, and they solve the following successive maximal correlation problem. For $k = 1,\dots,r-1$, in step k we want to find $\max \mathrm{Corr}_{\mathbb{W}}(\psi,\phi)$ subject to

$$\mathrm{Var}_{\mathbb{P}}(\psi) = \mathrm{Var}_{\mathbb{Q}}(\phi) = 1, \quad \mathrm{Cov}_{\mathbb{P}}(\psi,\psi_i) = \mathrm{Cov}_{\mathbb{Q}}(\phi,\phi_i) = 0, \quad i = 0,1,\dots,k-1.$$

(Note that the condition for $i = 0$ is equivalent to $\mathbb{E}_{\mathbb{P}}(\psi) = \mathbb{E}_{\mathbb{Q}}(\phi) = 0$.) By [Bol87b], the maximum is s_k and it is attained on the ψ_k, ϕ_k pair.

Now, we are able to define the joint representation of the general Hilbert-spaces H, H' – introduced in Subsection 1.3.1 – with respect to the joint measure $\mathbb{W}$ in the following way.

Definition 12 *We say that the pair* $(\mathbf{X}, \mathbf{Y})$ *of* k*-dimensional random vectors with components in* H *and* H'*, respectively, form a* k*-dimensional representation of the product space endowed with the measure* $\mathbb{W}$ *if* $\mathbb{E}_{\mathbb{P}}\mathbf{X}\mathbf{X}^T = \boldsymbol{I}_k$ *and* $\mathbb{E}_{\mathbb{Q}}\mathbf{Y}\mathbf{Y}^T = \boldsymbol{I}_k$ *(i.e., the components of* $\mathbf{X}$ *and* $\mathbf{Y}$ *are uncorrelated with zero expectation and unit variance, respectively); and the joint distribution of* X_i *and* Y_i *is* $\mathbb{W}$ $(i = 1, \dots, k)$*. Further, the cost of this representation is defined as*

$$Q_k(\mathbf{X}, \mathbf{Y}) = \mathbb{E}_{\mathbb{W}}\|\mathbf{X} - \mathbf{Y}\|^2.$$

The pair $(\mathbf{X}^*, \mathbf{Y}^*)$ *is an optimal representation if it minimizes the above cost.*

Analogously to the finite case, the following representation theorem was proved in [Bol13].

Theorem 11 ([Bol13]) ***Representation theorem for joint distributions.*** *Let* $\mathbb{W}$ *be a joint distribution with marginal distributions* $\mathbb{P}$ *and* $\mathbb{Q}$*. Assume that among the singular values of the conditional expectation operator* $P_{\mathcal{X}} : H' \to H$ *(see (1.27)) there are at least* k *positive ones and denote by* $1 > s_1 \geq s_2 \geq \cdots \geq s_k > 0$ *the largest ones. The minimum cost of a* k*-dimensional representation is* $2\sum_{i=1}^{k}(1 - s_i)$ *and it is attained with* $\mathbf{X}^* = (\psi_1, \dots, \psi_k)$ *and* $\mathbf{Y}^* = (\phi_1, \dots, \phi_k)$*, where* ψ_i, ϕ_i *is the function pair corresponding to the singular value* s_i $(i = 1, \dots, k)$*.*

We remark that when $\mathcal{X}$ and $\mathcal{Y}$ are finite sets, the solution corresponds to the SVD of the normalized contingency table. Though, this matrix seemingly does not have the same normalization as the kernel, our numerical algorithm for the SVD of a rectangular matrix is capable to find orthogonal singular vectors in the usual Euclidean norm, which corresponds to the Lebesgue measure and not to the $\mathbb{P}$- or $\mathbb{Q}$-measures. This is why, in correspondence analysis, we use the SVD of the matrix $\boldsymbol{C}_D$, and back-transform the singular vectors so that to get the representatives. Observe that if we have a non-degenerate contingency table, then $s_i < 1$ $(i = 1, \dots, k)$, therefore the minimum cost is strictly positive.

In the symmetric case, we can also define a representation. Now the $\mathbf{X}, \mathbf{X}'$ pair is identically distributed, but usually not independent; they are connected with the symmetric joint measure $\mathbb{W}$.

Definition 13 *We say that the* k*-dimensional random vector* $\mathbf{X}$ *with components in* H *forms a* k*-dimensional representation of the product space* $H \times H'$ *(*H *and* H' *are isomorphic) endowed with the symmetric measure* $\mathbb{W}$ *(and marginal measure* $\mathbb{P}$*) if* $\mathbb{E}_{\mathbb{P}}\mathbf{X}\mathbf{X}^T = \boldsymbol{I}_k$ *(i.e., the components of* $\mathbf{X}$ *are uncorrelated with zero expectation and unit variance). Further, the cost of this representation is defined as*

$$Q_k(\mathbf{X}) = \mathbb{E}_{\mathbb{W}}\|\mathbf{X} - \mathbf{X}'\|^2,$$

where $\mathbf{X}$ *and* $\mathbf{X}'$ *are identically distributed and the joint distribution of* X_i *and* X_i' *is* $\mathbb{W}$ $(i = 1, \dots, k)$*. The random vector* $\mathbf{X}^*$ *is an optimal representation if it minimizes the above cost.*

Theorem 12 ([Bol13]) ***Representation theorem for symmetric joint distributions****. Let* $\mathbb{W}$ *be a symmetric joint distribution with marginal* $\mathbb{P}$*. Assume that among the eigenvalues of the conditional expectation operator* $P_{\mathcal{X}} : H' \to H$ *(*H *and* H' *are isomorphic) there are at least* k *positive ones and denote by* $1 > \lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k > 0$ *the largest ones. Then the minimum cost of a* k*-dimensional representation is* $2\sum_{i=1}^{k}(1 - \lambda_i)$ *and it is attained by* $\mathbf{X}^* = (\psi_1, \dots, \psi_k)$ *where* ψ_i *is the eigenfunction corresponding to the eigenvalue* λ_i $(i = 1, \dots k)$*.*

In the case of a finite $\mathcal{X}$ (vertex set of an edge-weighted graph), we have a weighted graph with edge-weights w_{ij} $(\sum_{i=1}^{n}\sum_{j=1}^{n} w_{ij} = 1)$. The operator $P_{\mathcal{X}}$ deprived of the trivial

factor corresponds to its normalized modularity matrix with eigenvalues in the [-1,1] interval (1 cannot be an eigenvalue if the underlying graph is connected), and eigenfunctions which are the transformed eigenvectors. As the numerical algorithm gives an orthonormal set of the eigenvectors in Euclidean norm, some back-transformation is needed to get uncorrelated components with unit variance, therefore we use the normalized modularity matrix instead of the kernel $K_{ij} = \frac{w_{ij}}{d_i d_j}$ expected from (1.29), where $d_i = \sum_{j \in \mathcal{X}} w_{ij}$ is the generalized degree of vertex i ($i \in \mathcal{X}$). We remark that the above formula for the kernel corresponds to the so-called copula transformation of the joint distribution $\mathbb{W}$ into the unit square. This idea appears when vertex- and edge-weighted graphs are transformed into step-functions over $[0,1] \times [0,1]$, see the definition of graphons (e.g., [Borgsetal1]). This transformation can be performed in the non-symmetric and non-finite cases too. Also observe that neither the kernel nor the contingency table or graph is changed under measure preserving transformations of $\mathcal{X}$ or $\mathcal{Y}$, by the theory of exchangeable sequences and arrays. In particular, the labeling of the vertices or rows/columns is immaterial here.

In the framework of joint distributions, the *Cheeger constant* $h(G)$ can be viewed as a conditional probability and related to the symmetric maximal correlation in the following way. The weight matrix $\boldsymbol{W}$ (with sum of its entries 1) defines a discrete symmetric joint distribution $\mathbb{W}$ with the same marginals $\mathbb{D} = \{d_1, \ldots, d_n\}$. Let H denote the Hilbert space of $V \to \mathbb{R}$ random variables taking on at most n different values with probabilities $d_1, \ldots, d_n$, and having zero expectation and finite variance. Let us take two identically distributed (i.d.) copies $\psi, \psi' \in H$ with joint distribution $\mathbb{W}$. Then, obviously,

$$h(G) = \min_{\substack{B \subset \mathbb{R} \text{ Borel-set} \\ \psi, \psi' \in H \text{ i.d.} \\ \mathbb{P}_{\mathbb{D}}(\psi \in B) \le 1/2}} \mathbb{P}_{\mathbb{W}}(\psi' \in \overline{B} | \psi \in B).$$

The *symmetric maximal correlation*, defined by the symmetric joint distribution $\mathbb{W}$, is the following (it was introduced in [Bol-Mol02]):

$$r_1 = \max_{\psi, \psi' \in H \text{ i.d.}} \mathrm{Corr}_{\mathbb{W}}(\psi, \psi') = \max_{\substack{\psi, \psi' \in H \text{ i.d.} \\ \mathrm{Var}_{\mathbb{D}} \psi = 1}} \mathrm{Cov}_{\mathbb{W}}(\psi, \psi').$$

In view of (1.28) and Theorem 12, $r_1 = 1 - \lambda_1$, provided $\lambda_1 \le 1$.

With this notation, the result of Theorem 6 can be written in the equivalent form as follows.

Proposition 3 ([Bol-Mol02]) *Let $\mathbb{W}$ be the symmetric joint distribution of two discrete random variables taking on at most n different values, where the joint probabilities of $\mathbb{W}$ are the entries of the $n \times n$ symmetric weight matrix $\boldsymbol{W}$. If the symmetric maximal correlation r_1 is nonnegative, then with it, the estimation*

$$\frac{1 - r_1}{2} \le \min_{\substack{B \subset \mathbb{R} \, Borel\text{-}set \\ \psi, \psi' H \, i.d. \\ \mathbb{P}_{\mathbb{D}}(\psi \in B) \le 1/2}} \mathbb{P}_{\mathbb{W}}(\psi' \in \overline{B} | \psi \in B) \le \sqrt{1 - r_1^2}$$

holds, where we used the previous notation.

Consequently, the symmetric maximal correlation somehow regulates the minimum conditional probability that provided a categorical random variable takes values in a category set (with probability less than $1/2$) then another copy of it (their joint distribution is $\mathbb{W}$) will take values in the complementary category set. The larger r_1, the smaller this minimum conditional probability is. In particular, if r_1 is the largest absolute value eigenvalue of $\boldsymbol{I} - \boldsymbol{L}_D$ (apart from the trivial 1), then r_1 is the usual maximal correlation of Gebelein and Rényi.

We also remark that any or both of the starting random variables ξ, η can as well be random vectors (with real components). For example, if they have p- and q-dimensional Gaussian distribution respectively, than their maximum correlation is the largest canonical correlation between them, and it is realized by appropriate linear combinations of the components of ξ and η, respectively. Moreover, we can find canonical correlations one after the other with corresponding function pairs (under some orthogonality constraints), as many as the rank of the cross-covariance matrix of ξ and η. In fact, the whole decomposition relies on the SVD of a matrix calculated from this cross-covariance matrix and the individual covariance matrices of ξ and η. Note that this SVD-based treatment of the canonical correlation analysis was discussed in [Bol83] in details.

1.3.4 Treating nonlinearities via reproducing kernel Hilbert spaces

Reproducing kernel Hilbert spaces were introduced in the middle of the 20th century by Aronszajn and Parzen [Ar], but the theory itself is an elegant application of already known theorems of functional analysis, first of all the Riesz–Fréchet representation theorem and the theory of integral operators [Ri-SzN], tracing back to the beginning of the 20th century. Later on, in the last decades of the 20th century and even in our days, reproducing kernel Hilbert spaces live their renaissance and are applied in modern statistical methods and data mining, see, e.g., the independent component analysis (ICA) of Bach and Jordan [Bach]. Here I want to discuss their use in data clustering, and resolve the mystery around them.

A popular approach to data clustering (sometimes this is what called spectral clustering) is the following. Our data points $\mathbf{x}_1, \ldots, \mathbf{x}_n$ are already in a metric space, called *data space*, but cannot be well classified by the k-means algorithm for no integer $0 < k < n$, because they cannot be linearly separated (for example, there are obviously two clusters of points in $\mathbb{R}^2$, but they are separated by an annulus and the k-means algorithm with $k = 2$ is not able to recover them). Though, we can map the points into a usually higher (possibly, infinite) dimensional, or more abstract space with a non-linear mapping so that the images are already well clustered in the new, so-called *feature space*. However, in practical higher-dimensional problems, when we do not have the faintest idea about the clusters and no visualization is possible, unfortunately, we cannot give such mappings explicitly. Moreover, the feature space usually has much higher dimension than the original one, which fact is frequently referred to as the *curse of dimensionality*. The main point of the kernel method to be introduced is that it is not even necessary to perform the mapping, it suffices to select a kernel – based on the inner product of the original points – that is no longer a linear kernel, but a more complicated, still admissible kernel (the exact meaning is given in Definition 15), and defines a new inner product within the feature space. Then we process statistical algorithms that need only this kernel and nothing else.

Let $\mathcal{H}$ be a Hilbert space of functions $\mathcal{X} \to \mathbb{R}$ (where $\mathcal{X}$ is an arbitrary set, for the time being). The evaluation mapping $L_x : \mathcal{H} \to \mathbb{R}$ works on an $f \in \mathcal{H}$ so that

$$L_x(f) = f(x). \tag{1.30}$$

Definition 14 (Reproducing kernel Hilbert space) *A Hilbert space $\mathcal{H}$ of (real) functions on the set $\mathcal{X}$ is a reproducing kernel Hilbert space, briefly RKHS, if the point evaluation functional L_x of (1.30) exists and is continuous for all $x \in \mathcal{X}$.*

The name reproducing kernel comes from the fact that – by the Riesz–Fréchet representation theorem – the result of such a continuous mapping can be written as an inner product. This theorem states that a Hilbert space (in our case $\mathcal{H}$) and its dual (in our case the set of

$\mathcal{H} \to \mathbb{R}$ continuous linear functionals, e.g., L_x) are isometrically isomorphic. Therefore, to any L_x there uniquely corresponds a $K_x \in \mathcal{H}$ such that

$$L_x(f) = \langle f, K_x \rangle_{\mathcal{H}}, \quad \forall f \in \mathcal{H}. \tag{1.31}$$

Since K_x is itself an $\mathcal{X} \to \mathbb{R}$ function, it can be evaluated at any $y \in \mathcal{X}$. We define the bivariate function $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ as

$$K(x, y) := K_x(y) \tag{1.32}$$

and call it the *reproducing kernel* for the Hilbert space $\mathcal{H}$. Then using formulas (1.30), (1.31), and (1.32), we get that on the one hand,

$$K(x, y) = K_x(y) = L_y(K_x) = \langle K_x, K_y \rangle_{\mathcal{H}},$$

and on the other hand,

$$K(y, x) = K_y(x) = L_x(K_y) = \langle K_y, K_x \rangle_{\mathcal{H}}.$$

By the symmetry of the (real) inner product it follows that the reproducing kernel is symmetric and it is also reproduced as the inner product of special functions in the RKHS:

$$K(x, y) = \langle K_x, K_y \rangle_{\mathcal{H}} = \langle K(x, .), K(., y) \rangle_{\mathcal{H}}, \tag{1.33}$$

hence, K is positive definite (for the precise notion see the forthcoming Definition 15). This is the diabolic kernel trick.

Vice versa, if we are given a positive definite kernel function on $\mathcal{X} \times \mathcal{X}$ at the beginning, then there exists an RKHS such that with appropriate elements of it, the inner product relation (1.33) holds. (In fact, we are not given, we just select an appropriate kernel function.) The mystery of RKHS just lies in this converse statement.

For this purpose, let us first define the most important types of kernel functions and discuss how more and more complicated ones can be derived from the simplest ones.

Definition 15 *A symmetric two-variate function $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is called positive definite kernel (equivalently, admissible, valid, or Mercer kernel) if for any $n \in \mathbb{N}$ and $x_1, \ldots, x_n \in \mathcal{X}$, the symmetric matrix of entries $K(x_i, x_j) = K(x_j, x_i)$ $(i, j = 1, \ldots n)$ is positive semidefinite.*

We remark that a symmetric real matrix is positive semidefinite if and only if it is a Gram matrix, and hence, its entries become inner products, but usually not of the entries in its arguments. However, the simplest kernel function, the so-called *linear kernel*, does this job. It is defined as

$$K_{\text{lin}}(x, y) = \langle x, y \rangle_{\mathcal{X}},$$

if $\mathcal{X}$ is subset of a Euclidean space.

From a valid kernel, one can get other valid kernels with the following operations:

1. If $K_1, K_2 : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ are positive definite kernels, then the kernel K defined by $K(x, y) = K_1(x, y) + K_2(x, y)$ is also positive definite.

2. If $K_1, K_2 : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ are positive definite kernels, then the kernel K defined by $K(x, y) = K_1(x, y)K_2(x, y)$ (sometimes called Hadamard or Schur product) is also positive definite. Especially, if K is a positive definite kernel, then so does cK with any $c > 0$.

Consequently, if h is a polynomial with positive coefficients and $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ is a positive definite kernel, then the kernel $K_h : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ defined by

$$K_h(x, y) = h(K(x, y)) \tag{1.34}$$

is also positive definite. Since the exponential function can be approximated by polynomials with positive coefficients and the positive definiteness is closed under pointwise convergence, the same is true if h is the exponential function: $h(x) = e^x$, perhaps some transformation of it.

Putting these facts together, though it is known, here we give another short proof for the admissibility of the so-called *Gaussian kernel*.

Proposition 4 ([Bol13]) *The Gaussian kernel*

$$K_{\text{Gauss}}(x, y) = e^{-\frac{\|x-y\|^2}{2\sigma^2}}$$

is positive definite, with every parameter value $\sigma > 0$.

Proof. In view of

$$\|x - y\|^2 = \langle x, x\rangle + \langle y, y\rangle - 2\langle x, y\rangle,$$

this kernel can be written as the product of two positive definite kernels in the following way:

$$K_{\text{Gauss}}(x, y) = K_1(x, y)K_2(x, y),$$

where

$$K_1(x, y) = e^{-\frac{\langle x,x\rangle+\langle y,y\rangle}{2\sigma^2}} \quad \text{and} \quad K_2(x, y) = e^{\frac{\langle x,y\rangle}{\sigma^2}}.$$

Here K_2 is positive definite as it is the exponential function of the positive definite kernel $\frac{1}{\sigma^2}K_{\text{lin}}$. To show that K_1 is positive definite, by definition, we have to verify that for any $n \in \mathbb{N}$ and $x_1, \ldots, x_n \in \mathcal{X}$, the symmetric matrix of entries

$$K_1(x_i, x_j) = e^{-\frac{\langle x_i,x_i\rangle}{2\sigma^2}} \cdot e^{-\frac{\langle x_j,x_j\rangle}{2\sigma^2}}, \quad i, j = 1, \ldots n$$

is positive semidefinite. But it is a rank 1 matrix, its only non-zero eigenvalue being equal to its trace, which is positive. This finishes the proof. □

If $\mathcal{X} = \{x_1, \ldots, x_n\}$ and $\boldsymbol{S}$ is an $n \times n$ symmetric similarity matrix comprised of the pairwise similarities between the entries of $\mathcal{X}$, then the kernel K defined by the $n \times n$ symmetric, positive definite matrix $e^{\lambda \boldsymbol{S}}$ is called *diffusion kernel*, where $0 < \lambda < 1$ is parameter (sometimes called decay factor). Let us recapitulate that the eigenvalues of the $e^{\lambda \boldsymbol{S}}$ matrix are the numbers $e^{\lambda\lambda_i}$ $(i = 1, \ldots, n)$, where λ_i's are real eigenvalues of $\boldsymbol{S}$. Therefore the diffusion kernel is always strictly positive definite, even if $\boldsymbol{S}$ is not positive semidefinite.

The following converse theorem is due to Aronszajn [Ar] who attributed it to E. H. Moore.
Theorem (Aronszajn [Ar]) *For any positive definite kernel $K : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ there exists a unique, possibly infinite-dimensional Hilbert space $\mathcal{H}$ of functions on $\mathcal{X}$, for which K is a reproducing kernel.*

If we want to emphasize that the RKHS corresponds to the kernel K, we will denote it by $\mathcal{H}_K$. We may wish to realize the elements of an RKHS $\mathcal{H}_K$ in a more straightforward Hilbert space $\mathcal{F}$. Assume that there is a (usually not linear) map $\phi : \mathcal{X} \to \mathcal{F}$ such that when $x \in \mathcal{X}$ is mapped into $\phi(x) \in \mathcal{F}$, then

$$K(x, y) = \langle \phi(x), \phi(y)\rangle_{\mathcal{F}}$$

is the desired positive definite kernel. At the same time, in view of (1.33),

$$K(x,y) = \langle K_x, K_y \rangle_{\mathcal{H}_K}$$

where recall that $K_x = K(x,.)$ is an $\mathcal{X} \to \mathbb{R}$ function, hence, cannot be identical to $\phi(x)$, but they can be connected with the following transformation. Let T be a linear operator from $\mathcal{F}$ to the space of functions $\mathcal{X} \to \mathbb{R}$ defined by

$$(Tf)(y) = \langle f, \phi(y) \rangle_{\mathcal{F}}, \quad y \in \mathcal{X},\ f \in \mathcal{F}.$$

Then

$$(T\phi(x))(y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{F}} = K(x,y) = K_x(y),$$

therefore

$$T\phi(x) = K_x, \quad \forall x \in \mathcal{X} \tag{1.35}$$

and hence, $\mathcal{H}_K$ becomes the range of T. This was the informal proof of the more precise statements about this correspondence.

In practical applications, we usually have a finite sample $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n)$. Based on it, the *empirical feature map* $\hat{\phi} : \mathcal{X} \to \mathbb{R}^n$ can be constructed in the following way (see e.g. [Sch-Sm]):

$$\hat{\phi}(\mathbf{x}) = \boldsymbol{K}^{-1/2} \phi_n(\mathbf{x}) \tag{1.36}$$

with $\phi_n(\mathbf{x}) = (K(\mathbf{x}, \mathbf{x}_1), \dots, K(\mathbf{x}, \mathbf{x}_n))^T$, the counterpart of $K(\mathbf{x}, .)$ based on the n-element set $\mathcal{X}$, and the $n \times n$ symmetric real matrix $\boldsymbol{K} = (K_{ij})$ of entries $K_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)$, $i, j = 1, \dots, n$. Assume that $\boldsymbol{K}$ is positive definite, otherwise (if positive semidefinite with at least one zero eigenvalue) we will use generalized inverse when calculating $\boldsymbol{K}^{-1/2}$. Let us apply (1.36) for $\mathbf{x}_i$'s. Since

$$\phi_n(\mathbf{x}_i) = \boldsymbol{K}\mathbf{e}_i,$$

where $\mathbf{e}_i$ is the i-th unit vector in $\mathbb{R}^n$ (it has 0 coordinates, except the i-th one which is equal to 1), the relation

$$\hat{\phi}(\mathbf{x}_i) = \boldsymbol{K}^{-1/2} \phi_n(\mathbf{x}_i) = \boldsymbol{K}^{1/2} \mathbf{e}_i$$

holds. Further,

$$\langle \hat{\phi}(\mathbf{x}_i), \hat{\phi}(\mathbf{x}_j) \rangle = (\boldsymbol{K}^{1/2}\mathbf{e}_i)^T (\boldsymbol{K}^{1/2}\mathbf{e}_j) = \mathbf{e}_i^T \boldsymbol{K} \mathbf{e}_j = K_{ij}, \quad i, j = 1, \dots, n.$$

Howsoever we cannot navigate well in the artificially constructed spaces, this whole abstraction was not in vain. Observe that for the data points $\mathbf{x}_i$'s we need not even calculate $\hat{\phi}(\mathbf{x}_i)$'s, the spectral clustering of these images can be done based on their pairwise distances:

$$\begin{aligned} \|\hat{\phi}(\mathbf{x}_i) - \hat{\phi}(\mathbf{x}_j)\|^2 &= \langle \hat{\phi}(\mathbf{x}_i), \hat{\phi}(\mathbf{x}_j) \rangle + \langle \hat{\phi}(\mathbf{x}_i), \hat{\phi}(\mathbf{x}_i) \rangle - 2\langle \hat{\phi}(\mathbf{x}_j), \hat{\phi}(\mathbf{x}_j) \rangle \\ &= K(\mathbf{x}_i, \mathbf{x}_i) + K(\mathbf{x}_j, \mathbf{x}_j) - 2K(\mathbf{x}_i, \mathbf{x}_j) \end{aligned}$$

$(i, j = 1, \dots, n)$. Thus, to evaluate the pairwise distances between any pairs of the n features, merely the kernel values are needed. Sometimes the kernel is some transformation of a similarity matrix of n objects, even if we do not have them as finite-dimensional points. In other cases, we have finite dimensional measurements on the objects, but merely the $n \times n$ empirical covariance matrix is stored. If our data are multivariate Gaussian, this matrix suffices for further processing, in other cases, we can calculate a polynomial or Gaussian kernel based on it, with the understanding that it may be the true similarity between our

non-Gaussian data which are, in fact, in an abstract space. For example, the $n \times n$ similarity matrix of the features, applying a Gaussian kernel afterwards, gives matrix entries

$$K_{\mathrm{Gauss}}(\hat{\phi}(\mathbf{x}_i), \hat{\phi}(\mathbf{x}_j)) = e^{-\frac{\|\hat{\phi}(\mathbf{x}_i)-\hat{\phi}(\mathbf{x}_j)\|^2}{2\sigma^2}} = e^{-\frac{K(\mathbf{x}_i,\mathbf{x}_i)+K(\mathbf{x}_j,\mathbf{x}_j)-2K(\mathbf{x}_i,\mathbf{x}_j)}{2\sigma^2}}$$

which can be further processed through Laplacian based clustering, see Section 1.1.2.

In this way, linear methods are applicable in an implicitly constructed space, instead of having to use non-linear methods in the original one. Here we only use the kernel which is calculated from the inner products of the data points through several transformations. The philosophy behind the above techniques is that sometimes sophisticated, composite kernels are more capable to reveal the structure of our data or to classify them, especially if they are not from a Gaussian distribution or consist of different type of measurements (e.g. location, brightness, color, texture of pixels).

But what kind of a kernel should be used? This is the important question. Many authors, e.g. [Lux, Lux-Be-Bo], recommend the Gaussian kernel. For the data points $\mathbf{x}_1, \ldots, \mathbf{x}_n$ to be classified they construct the Gaussian kernel and the $n \times n$ symmetric, positive definite kernel matrix is considered as weight matrix $\boldsymbol{W}$ of a graph (in [Ng-Jo-We] the authors use zero diagonal). Then they perform spectral clustering based on the Laplacian or normalized Laplacian matrix corresponding to $\boldsymbol{W}$. In this way, applying the k-means algorithm for the so obtained k-dimensional (in fact, $(k-1)$-dimensional) representatives, they obtain nice clusters. This is because the data points of $\mathcal{X}$ are mapped into a feature space $\mathcal{F}$ such that the only implicitly known images $\phi(\mathbf{x}_i)$ $(i = 1, \ldots, n)$ define a graph similarity, starting with which, the usual representation based spectral clustering works well. Hence, the graph construction is just an intermediate step for the subsequent metric clustering. Even if we are given a graph in advance, we may calculate the k-dimensional representatives of the vertices (with a relatively small k, based on the k bottom Laplacian eigenvectors), and then we classify them using kernel methods (e.g., substitute them into the Gaussian kernel).

The advantage of the Gaussian kernel is that it is translation-invariant, and the infinite dimensional feature space can be described by Fourier theory. Since the Fourier transforms of functions in $\mathcal{H}_K$ decay rapidly, the induced RKHS consists of smooth functions. In this way, Gaussian kernels may be used as smoothing functions, for example, in the ACE (Alternating Conditional Expectation) algorithm elaborated for the generalized non-parametric regression problem, see [Bre-Fri] for details.

If the underlying space $\mathcal{X}$ is a probability space of random variables with finite variance, and the product space is endowed with a joint distribution (see Section 1.3), we may look for the so-called $\mathcal{F}$-correlation of two random variables which is the largest possible correlation between their ϕ-maps in the feature space. In [Bach] it is proved that if $\mathcal{F}$ is the feature space corresponding to the RKHS defined by the Gaussian kernel on $\mathbb{R}$ (with any positive parameter σ), then the $\mathcal{F}$-correlation of two random variables is zero if and only if they are independent, akin to the maximal correlation of Gebelein and Rényi. Therefore, by mapping our data into the feature space, usual linear methods – like principal component analysis or canonical correlation analysis – become non-linear ones, and able to find independent components instead of uncorrelated ones, which fact has significance if our data come from a non-Gaussian distribution. This is the base of the so-called independent component analysis (ICA), see [Bach, Sch-Sm-Mu] for details. Note that with a finite-dimensional feature space, the $\mathcal{F}$-correlation cannot characterize independence.

We remark that for *image segmentation* purposes, [Shi-Ma] uses two or more Gaussian kernels: one contains the Euclidean distances of the pixels, and the others those of their brightness, color, texture, etc. Eventually, they take the product of the two or more positive definite kernels. If we multiply kernels (with Hadamard product), it means that entries in the

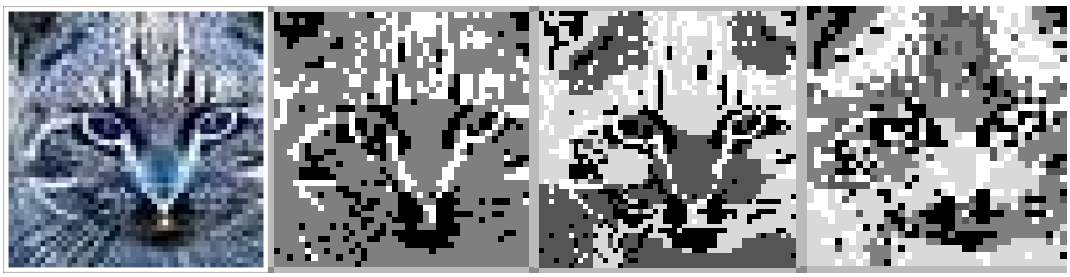

Figure 1.1: The original picture and the pixels colored with 3, 4, and 5 different colors according to their cluster memberships (made by László Nagy).

same position of the kernel matrices are multiplied together. During the whole calculation for n data points, we only use the $n \times n$ symmetric, positive definite, usually sparse kernel matrix.

As for the normalized modularity matrix, we used the eigenvectors corresponding to its $k-1$ largest eigenvalues to find k clusters of 2304 pixels based on a Gaussian kernel. More precisely, we assigned the points $\mathbf{x}_1, \ldots, \mathbf{x}_{2304} \in \mathbb{R}^3$ to the pixels, the coordinates of which characterize the spacial location, color, and brightness of the pixels. With the positive parameter σ, the similarity between pixels i and j was $w_{ij} = e^{-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|^2}{2\sigma^2}}$ for $i \neq j$. Figure 1.1 shows the original picture, and the picture when the pixels were colored according to their cluster memberships with number of clusters 3,4, and 5. Since the largest absolute value eigenvalues of the 2304×2304 normalized modularity matrix are

$$0.137259,\ 0.0142548,\ 0.000925228,\ -0.000670733,\ -0.000670674,\ \ldots$$

and the number of the positive eigenvalues is three, with a gap after the second one, the 3- or 4-cluster solution seems the most reasonable. It is an intriguing question – unsolved so far – whether the dimension k of the original data points can be detected in the spectrum of $\boldsymbol{M}_D$ when n is large.

Chapter 2

Treating randomness in large networks and clustering with small discrepancy

Here we apply the results of Chapter 1 for the spectral clustering of large networks. Networks are modeled either by edge-weighted graphs or contingency tables, and usually subject to random errors due to their evolving and flexible nature. Asymptotic properties of SD and SVD of the involved matrices are discussed when not only the number of the graph's vertices or that of the rows and columns of the contingency table tend to infinity, but the cluster sizes also grow proportionally with them. Mostly, perturbation results for the SD and SVD of blown-up matrices burdened with a Wigner-type error matrix are investigated.

The idea is that if we understand the spectral properties of the noisy blown-up matrices, it helps us to find similar structures, i.e., clusters with small within- and between-cluster discrepancies, in large deterministic networks. Theorems 24, 25, 26, 27, 28 give an answer to this. Since the ideas of many proofs are used later, we present them with unique notation.

2.1 Perturbation results for symmetric block structures

We will use the following very general kind of a random noise.

Definition 16 *Let w_{ij} $(1 \le i \le j \le n)$ be independent, real-valued random variables defined on the same probability space, $w_{ji} = w_{ij}$, $\mathbb{E}(w_{ij}) = 0$ $(\forall i, j)$, and the w_{ij}'s are uniformly bounded, i.e., there is a constant $K > 0$ – that does not depend of n – such that $|w_{ij}| \le K$, $\forall i, j$. Then the $n \times n$ real symmetric random matrix $\boldsymbol{W}_n = (w_{ij})_{1 \le i,j \le n}$ is called symmetric Wigner-noise.*

This random matrix is the generalization of that introduced by E. Wigner when formulated his famous semicircle law. Note that the condition of uniform boundedness of the entries could be relaxed: the entries may have Gaussian distribution or sub-Gaussian moments. However, most of the subsequent results and the frequently used Theorem of Füredi and Komlós [Fü-Ko] rely on Definition 16. This theorem implies that for the spectral norm of an $n \times n$ symmetric Wigner noise

$$\|\boldsymbol{W}_n\| = \max_{1 \le i \le n} |\lambda_i(\boldsymbol{W}_n)| \le 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3} \log n) \tag{2.1}$$

holds with probability tending to 1 as $n \to \infty$, where σ^2 is the uniform bound for the variances of the w_{ij}'s.

Definition 17 *The $n \times n$ matrix $\boldsymbol{B}_n$ is a symmetric blown-up matrix if there is a positive integer $k < n$, a $k \times k$ symmetric probability matrix $\boldsymbol{P}$ with entries p_{ij} $(0 < p_{ij} < 1)$, and there are positive integers $n_1, \ldots, n_k$, $\sum_{i=1}^{k} n_i = n$ such that – after rearranging its rows and columns with the same permutation – the matrix $\boldsymbol{B}_n$ can be divided into k^2 blocks, where the block (i, j) is an $n_i \times n_j$ matrix with entries all equal to p_{ij} $(1 \le i, j \le k)$.*

Let us fix $\boldsymbol{P}$, blow it up to an $n \times n$ matrix $\boldsymbol{B}_n$, and consider the noisy matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ as $n_1, \ldots, n_k \to \infty$, roughly speaking, at the same rate. For this purpose, an exact growth rate condition is formulated as follows.

Definition 18 *Under Growth Rate Condition 1, briefly, **GC1**, the following is understood: $n = \sum_{i=1}^{k} n_i \to \infty$ in such a way that $\frac{n_i}{n} \ge c$ with some constant $0 < c \le \frac{1}{k}$.*

While perturbing $\boldsymbol{B}_n$ by $\boldsymbol{W}_n$, assume that for the uniform bound of the entries of $\boldsymbol{W}_n$ the condition

$$K \le \min\{\min_{i,j \in \{1,\ldots,k\}} p_{ij}\,,\, 1 - \max_{i,j \in \{1,\ldots,k\}} p_{ij}\} \tag{2.2}$$

is satisfied. In this way, the entries of $\boldsymbol{A}_n$ are in the [0,1] interval, and $\boldsymbol{A}_n$ defines a random edge-weighted graph $G_n = (V, \boldsymbol{A}_n)$ on n vertices.

With an appropriate Wigner-noise we can achieve that the noisy matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ contains 1's in the (a, b)-th block with probability p_{ab}, and 0's otherwise. Indeed, for indices $1 \le a < b \le k$ and $i \in V_a$, $j \in V_b$ let

$$w_{ij} := \begin{cases} 1 - p_{ab} & \text{with probability} \quad p_{ab} \\ -p_{ab} & \text{with probability} \quad 1 - p_{ab} \end{cases} \tag{2.3}$$

be independent random variables; further, for $a = 1, \ldots, k$ and $i, j \in V_a$ $(i \le j)$ let

$$w_{ij} := \begin{cases} 1 - p_{aa} & \text{with probability} \quad p_{aa} \\ -p_{aa} & \text{with probability} \quad 1 - p_{aa} \end{cases} \tag{2.4}$$

be also independent, otherwise $\boldsymbol{W}$ is symmetric. This $\boldsymbol{W}$ satisfies the conditions of Definition 16 with uniformly bounded entries of zero expectation and variance bounded by

$$\sigma^2 = \max_{1 \le i \le j \le k} p_{ij}(1 - p_{ij}) \le \frac{1}{4}.$$

So, the noisy weighted graph $G_n = (V, \boldsymbol{A}_n)$ becomes a *generalized random graph* on the *planted partition* $V_1, \ldots, V_k$ of the vertices such that vertices of V_a and V_b are connected independently, with probability p_{ab}, $1 \le a \le b \le k$; see e.g., [Bi-Ch, Co-La, McSh] and the forthcoming Definition 21. This so-called *stochastic block-model* was first mentioned in [Hol-Las-Lei]. Note that this model is the generalization of the classical Erdős–Rényi random graph (the first random graph of the history, introduced in [Erd-Reny] and also discussed in [Bo]), which corresponds to the $k = 1$ case.

Definition 19 *Let $\boldsymbol{A}_n$ be an $n \times n$ symmetric matrix and $\mathcal{P}_n$ be a property which mostly depends on the SD of $\boldsymbol{A}_n$. Then $\boldsymbol{A}_n$ can have the property $\mathcal{P}_n$ in the two following senses. If $\boldsymbol{A}_n$ is the weight matrix of an edge-weighted graph G_n, then we equivalently say that G_n has the property $\mathcal{P}_n$, and denote this fact by $\boldsymbol{A}_n \in \mathcal{P}_n$ or $G_n \in \mathcal{P}_n$.*

WP1 *The property $\mathcal{P}_n$ holds for $\boldsymbol{A}_n$ with probability tending to 1 if*

$$\lim_{n\to\infty} \mathbb{P}\left(\boldsymbol{A}_n \in \mathcal{P}_n\right) = 1.$$

AS *The property $\mathcal{P}_n$ holds for $\boldsymbol{A}_n$ almost surely if*

$$\mathbb{P}\left(\exists\, n_0 \in \mathbb{N}\, s.t.\ for \quad n \geq n_0 \,:\, \boldsymbol{A}_n \in \mathcal{P}_n\right) = 1.$$

*Here we may assume **GC1** of Definition 18 for the growth of n together with the cluster sizes.*

(Note that in combinatorics literature, sometimes **WP1** is called **AS**.) Obviously, **AS** always implies **WP1**. Conversely, if in addition to **WP1**, $\sum_{n=1}^{\infty} \mathbb{P}\left(\boldsymbol{A}_n \notin \mathcal{P}_n\right) < \infty$ also holds, then, by the Borel–Cantelli lemma, $\boldsymbol{A}_n$ has $\mathcal{P}_n$ **AS**.

To established almost sure properties, the following sharp concentration result plays a crucial role.

Theorem (Alon, Krivelevich, Vu [Al-Kr-Vu]) *Let a_{ij} $(1 \leq i \leq j \leq n)$ be independent random variables with absolute value at most 1. Define a_{ij} for $1 \leq j < i \leq n$ by $a_{ij} = a_{ji}$. Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$ be the eigenvalues of the random symmetric matrix $\boldsymbol{A} = (a_{ij})_{i,j=1}^{n,n}$. Then for every positive integer $1 \leq i \leq \frac{n}{2}$, the probability that λ_i deviates from its median by more than t is at most $e^{-\frac{t^2}{32i^2}}$. The same estimate holds for the probability that λ_{n-i+1} deviates from its median by more than t.*

The statement – apart from a constant factor in the exponent – remains valid if a_{ij}'s are uniformly bounded with some constant K. The authors also prove that the eigenvalues are highly concentrated on their own expectations, since λ_i and its median are $\mathcal{O}(i)$ apart. Indeed, let m_i denote the median of λ_i. Then

$$|\mathbb{E}(\lambda_i) - m_i| \leq \mathbb{E}|\lambda_i - m_i| = \int_0^\infty \mathbb{P}(|\lambda_i - m_i| > t)\, dt$$
$$\leq \int_0^\infty e^{-\frac{t^2}{32i^2}}\, dt = 8\sqrt{2\pi} i.$$

Therefore, for all $t \gg i$ we have

$$\mathbb{P}\left(|\lambda_i - \mathbb{E}(\lambda_i)| > t\right) \leq e^{-\frac{(1-o(1))t^2}{32i^2}} \quad \text{if} \quad 1 \leq i \leq \frac{n}{2}, \tag{2.5}$$

and the same estimate holds for the probability $\mathbb{P}\left(|\lambda_{n-i+1} - \mathbb{E}(\lambda_{n-i+1})| > t\right)$. The authors also note that for the adjacency matrix of a random graph, when the entries are in the [0,1] interval, the estimate of their theorem can be improved to $4e^{-\frac{t^2}{8i^2}}$.

In view of (2.5), for the spectral norm, i.e., the largest absolute value eigenvalue, of an $n \times n$ Wigner-noise $\boldsymbol{W}_n$, the following relation holds with every positive real number t:

$$\mathbb{P}\left(|\|\boldsymbol{W}_n\| - \mathbb{E}(\|\boldsymbol{W}_n\|)| > t\right) \leq \exp\left(-\frac{(1-o(1))t^2}{32K^2}\right),$$

where K is the uniform bound for the entries of $\boldsymbol{W}_n$.

This inequality and the fact that, in view of (2.1), $\|\boldsymbol{W}_n\| = \mathcal{O}(\sqrt{n})$ together imply that $\mathbb{E}\|\boldsymbol{W}_n\| = \mathcal{O}(\sqrt{n})$. Therefore, there exist positive constants c_1 and c_2 such that they do not depend on n (they only depend on K) that

$$\mathbb{P}\left(\,\|\boldsymbol{W}\| > c_1\sqrt{n}\,\right) \leq e^{-c_2 n}. \tag{2.6}$$

As the right-hand side of (2.6) is the general term of a convergent series, by the Borel–Cantelli lemma it follows that the spectral norm of $\boldsymbol{W}_n$ is of order $\sqrt{n}$, **AS**. This observation will provide the base of the almost sure results of this chapter.

2.1.1 General blown-up structures

The spectrum of a symmetric blown-up matrix is characterized as follows.

Proposition 5 ([Bol05]) *Under the growth rate condition **GC1**, all the non-zero eigenvalues of the $n \times n$ blown-up matrix $\boldsymbol{B}_n$ of the $k \times k$ symmetric probability matrix $\boldsymbol{P}$ are of order n in absolute value.*

Proof. As there are at most k linearly independent rows of $\boldsymbol{B}_n$, $r = \operatorname{rank}(\boldsymbol{B}_n) \leq k$. Let $\beta_1, \ldots, \beta_r > 0$ be the non-zero eigenvalues of $\boldsymbol{B}_n$ with corresponding orthonormal eigenvectors $\mathbf{u}_1, \ldots, \mathbf{u}_r \in \mathbb{R}^n$. For notational convenience, we discard the subscripts: let $\beta \neq 0$ be an eigenvalue with corresponding eigenvector $\mathbf{u}$, $\|\mathbf{u}\| = 1$. It is easy to see that $\mathbf{u}$ is a step-vector: it has n_i coordinates equal to $u(i)$ $(i = 1, \ldots, k)$, where $n_1, \ldots, n_k$ are the blow-up sizes. Then, with these coordinates, the eigenvalue–eigenvector equation

$$\boldsymbol{B}\mathbf{u} = \beta\mathbf{u}$$

has the form

$$\sum_{j=1}^{k} n_j p_{ij} u(j) = \beta u(i), \quad i = 1, \ldots, k. \tag{2.7}$$

With the notation

$$\tilde{\mathbf{u}} = (u(1), \ldots, u(k))^T, \qquad \boldsymbol{N} = \operatorname{diag}(n_1, \ldots, n_k), \tag{2.8}$$

(2.7) can be rewritten in the form

$$\boldsymbol{P}\boldsymbol{N}\tilde{\mathbf{u}} = \beta\tilde{\mathbf{u}}. \tag{2.9}$$

Further, introducing the transformation

$$\mathbf{v} = \boldsymbol{N}^{1/2}\tilde{\mathbf{u}}, \tag{2.10}$$

Equation (2.9) is equivalent to

$$\boldsymbol{N}^{1/2}\boldsymbol{P}\boldsymbol{N}^{1/2}\mathbf{v} = \beta\mathbf{v}. \tag{2.11}$$

It is easy to see that the transformation (2.10) results in a unit-norm vector. Furthermore, applying the transformation (2.10) to the $\tilde{\mathbf{u}}_i$ vectors obtained from the $\mathbf{u}_i$ $(i = 1, \ldots, r)$, the orthogonality is also preserved. Consequently, $\mathbf{v}_i = \boldsymbol{N}^{1/2}\tilde{\mathbf{u}}_i$ is an eigenvector corresponding to the eigenvalue β_i of the $k \times k$ matrix $\boldsymbol{N}^{1/2}\boldsymbol{P}\boldsymbol{N}^{1/2}$, $i = 1, \ldots, r$. With the shrinking

$$\widetilde{\boldsymbol{N}} = \frac{1}{n}\boldsymbol{N}, \tag{2.12}$$

(2.11) is also equivalent to

$$\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{v} = \frac{\beta}{n}\mathbf{v},$$

that is the $k \times k$ matrix $\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}$ has nonzero eigenvalues $\frac{\beta_i}{n}$ with orthonormal eigenvectors $\mathbf{v}_i$ $(i = 1, \ldots, r)$.

Now we want to establish relations between the eigenvalues of $\boldsymbol{P}$ and $\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}$. Since we are interested in the absolute values of the nonzero eigenvalues, we will use singular values (recall that the singular values of a symmetric matrix are the absolute values of its real eigenvalues). Also, we are interested only in the first r eigenvalues, where $r = \operatorname{rank}(\boldsymbol{B}) = \operatorname{rank}(\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2})$, therefore, it suffices to consider vectors $\mathbf{x}$, for which $\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x} \neq \mathbf{0}$

and apply the Fischer–Courant–Weyl minimax principle to them. In view of this, for $i \in \{1,\dots,r\}$ and an arbitrary i-dimensional subspace $H \subset \mathbb{R}^n$:

$$\min_{\mathbf{x}\in H} \frac{\|\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\mathbf{x}\|} = \min_{\mathbf{x}\in H} \frac{\|\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|} \cdot \frac{\|\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|} \cdot \frac{\|\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\mathbf{x}\|}$$
$$\geq s_k(\widetilde{\boldsymbol{N}}^{1/2}) \cdot \min_{\mathbf{x}\in H} \frac{\|\boldsymbol{N}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|} \cdot s_k(\widetilde{\boldsymbol{N}}^{1/2}) \geq c \cdot \min_{\mathbf{x}\in H} \frac{\|\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|}{\|\widetilde{\boldsymbol{N}}^{1/2}\mathbf{x}\|},$$

with the constant c of the growth rate condition **GC1** (see Definition 18). Now taking the maximum for all possible i-dimensional subspace H, we obtain that $|\lambda_i(\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2})| \geq c|\lambda_i(\boldsymbol{P})| > 0$. On the other hand,

$$|\lambda_i(\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2})| \leq \|\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}\| \leq \|\widetilde{\boldsymbol{N}}^{1/2}\| \cdot \|\boldsymbol{P}\| \cdot \|\widetilde{\boldsymbol{N}}^{1/2}\| \leq \|\boldsymbol{P}\| \leq k.$$

These together imply that $\lambda_i(\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2})$ can be bounded from below and from above with a positive constant that does not depend on n and n_i's, it only depends on c and $\|\boldsymbol{P}\|$. Hence, because of $\lambda_i(\widetilde{\boldsymbol{N}}^{1/2}\boldsymbol{P}\widetilde{\boldsymbol{N}}^{1/2}) = \frac{\beta_i}{n}$, we obtain that $\beta_1,\dots,\beta_r = \Theta(n)$. This finishes the proof. □

For simplicity, in the sequel, we will assume that $\operatorname{rank}(\boldsymbol{P}) = k$, consequently, $\operatorname{rank}(\boldsymbol{B}_n) = k$ too.

Theorem 13 ([Bol05]) *Let $\boldsymbol{B}_n$ be an $n \times n$ blown-up matrix of the $k \times k$ symmetric probability matrix $\boldsymbol{P}$ with non-zero eigenvalues $\beta_1,\dots,\beta_k$, and $\boldsymbol{W}_n$ be an $n \times n$ Wigner-noise. Then there are k eigenvalues $\lambda_1,\dots,\lambda_k$ of the noisy random matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ such that*

$$|\lambda_i - \beta_i| \leq 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3}\log n), \quad i = 1,\dots,k \tag{2.13}$$

and for the other $n-k$ eigenvalues

$$|\lambda_j| \leq 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3}\log n), \quad j = k+1,\dots,n \tag{2.14}$$

*holds **AS** as $n \to \infty$ under **GC1**.*

Proof. The statement immediately follows by applying the Weyl's perturbation theorem for the spectrum of the symmetric matrix $\boldsymbol{B}_n$ characterized in Proposition 5, where the spectral norm of the perturbation $\boldsymbol{W}_n$ is estimated by (2.1). This proves the order of the eigenvalues **WP1**. In view of (2.6) and the Borel–Cantelli lemma, it implies that this is an **AS** property as well. □

Consequently, taking into consideration the order $\Theta(n)$ of the non-zero eigenvalues of $\boldsymbol{B}_n$, there is a spectral gap between the k largest absolute value and the other eigenvalues of $\boldsymbol{A}_n$, and this is of order $\Delta - 2\varepsilon$, where

$$\varepsilon = 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3}\log n) \quad \text{and} \quad \Delta = \min_{1\leq i\leq k} |\beta_i|. \tag{2.15}$$

In this way, Theorem 13 guarantees the existence of k protruding, so-called *structural* eigenvalues of $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$. With the help of this theorem we are also able to estimate the distances between the corresponding eigen-subspaces of the matrices $\boldsymbol{B}_n$ and $\boldsymbol{A}_n$.

Let us denote the unit-norm eigenvectors corresponding to the largest eigenvalues $\beta_1,\dots,\beta_k$ of $\boldsymbol{B}_n$ by $\mathbf{u}_1,\dots,\mathbf{u}_k$ and those corresponding to the largest eigenvalues $\lambda_1,\dots,\lambda_k$ of $\boldsymbol{A}_n$ by $\mathbf{x}_1,\dots,\mathbf{x}_k$. Let $F := \operatorname{Span}\{\mathbf{u}_1,\dots,\mathbf{u}_k\} \subset \mathbb{R}^n$ be the k-dimensional eigen-subspace, and let $\operatorname{dist}(\mathbf{x}, F)$ denote the Euclidean distance between the vector $\mathbf{x} \in \mathbb{R}^n$ and the subspace F.

Proposition 6 ([Bol05]) *With the above notation, the following estimate holds* ***AS*** *for the sum of the squared distances between* $\mathbf{x}_1, \ldots, \mathbf{x}_k$ *and* F*:*

$$\sum_{i=1}^{k} \operatorname{dist}^2(\mathbf{x}_i, F) \leq k \frac{\varepsilon^2}{(\Delta - \varepsilon)^2} = \mathcal{O}(\frac{1}{n}). \tag{2.16}$$

Proof. Let us choose one of the eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_k$ of $\boldsymbol{A}_n$ and denote it simply by $\mathbf{x}$ with corresponding eigenvalue λ. To estimate the distance between $\mathbf{x}$ and F, we expand $\mathbf{x}$ in the basis $\mathbf{u}_1, \ldots, \mathbf{u}_n$ with coefficients $t_1, \ldots, t_n \in \mathbb{R}$:

$$\mathbf{x} = \sum_{i=1}^{n} t_i \mathbf{u}_i.$$

The eigenvalues $\beta_1, \ldots, \beta_k$ of the matrix $\boldsymbol{B}_n$ corresponding to $\mathbf{u}_1, \ldots, \mathbf{u}_k$ are of order n (by Proposition 5), whereas the other eigenvalues are zeros.

Then, on the one hand

$$\boldsymbol{A}_n \mathbf{x} = (\boldsymbol{B}_n + \boldsymbol{W}_n)\mathbf{x} = \sum_{i=1}^{k} t_i \beta_i \mathbf{u}_i + \boldsymbol{W}_n \mathbf{x}, \tag{2.17}$$

and on the other hand

$$\boldsymbol{A}_n \mathbf{x} = \lambda \mathbf{x} = \sum_{i=1}^{n} t_i \lambda \mathbf{u}_i. \tag{2.18}$$

Equating the right-hand sides of (2.17) and (2.18) we get that

$$\sum_{i=1}^{k} t_i (\lambda - \beta_i)\mathbf{u}_i + \sum_{i=k+1}^{n} t_i \lambda \mathbf{u}_i = \boldsymbol{W}_n \mathbf{x}.$$

Then the Pythagorean theorem yields

$$\sum_{i=1}^{k} t_i^2 (\lambda - \beta_i)^2 + \sum_{i=k+1}^{n} t_i^2 \lambda^2 = \|\boldsymbol{W}_n \mathbf{x}\|^2 = \mathbf{x}^T \boldsymbol{W}_n^T \boldsymbol{W}_n \mathbf{x} \leq \varepsilon^2, \tag{2.19}$$

since $\|\mathbf{x}\| = 1$ and the largest eigenvalue of $\boldsymbol{W}_n^T \boldsymbol{W}_n$ is ε^2.

The squared distance between $\mathbf{x}$ and F is $\operatorname{dist}^2(\mathbf{x}, F) = \sum_{i=k+1}^{n} t_i^2$. In view of $|\lambda| \geq \Delta - \varepsilon$,

$$\begin{aligned}(\Delta - \varepsilon)^2 \operatorname{dist}^2(\mathbf{x}, F) = (\Delta - \varepsilon)^2 \sum_{i=k+1}^{n} t_i^2 &\leq \sum_{i=k+1}^{n} t_i^2 \lambda^2 \\ &\leq \sum_{i=1}^{k} t_i^2 (\lambda - \beta_i)^2 + \sum_{i=k+1}^{n} t_i^2 \lambda^2 \leq \varepsilon^2.\end{aligned}$$

Note that in the last inequality we used (2.19). From here,

$$\operatorname{dist}^2(\mathbf{x}, F) \leq \frac{\varepsilon^2}{(\Delta - \varepsilon)^2} = \mathcal{O}(\frac{1}{n}) \tag{2.20}$$

where the order of the estimate follows from the order of ε and Δ of (2.15).

Applying (2.20) for the eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_k$ of $\boldsymbol{A}_n$, and adding the k inequalities together, we obtain the same order of magnitude for the sum of the squared distances, which finishes the proof. □

Let $G_n = (V, \boldsymbol{A}_n)$ be the random edge-weighted graph on the n-element vertex set and edge-weight matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$, where for the uniform bound of the entries of $\boldsymbol{W}_n$ (2.2) is assumed. Denote by $V_1, \ldots, V_k$ the partition of V with respect to the blow-up of $\boldsymbol{B}_n$ (it defines a clustering of the vertices). Proposition 6 implies the well-clustering property of the representatives of the vertices of G_n in the following representation. Let $\boldsymbol{X}$ be the $n \times k$ matrix containing the eigenvectors $\mathbf{x}_1, \ldots, \mathbf{x}_k$ of $\boldsymbol{A}_n$ in its columns. Let the k-dimensional representatives of the vertices be the row vectors of $\boldsymbol{X}$ and $S_k^2(P_k; \boldsymbol{X})$ denote the k-variance – see (1.10) – of these representatives in the clustering $P_k = (V_1, \ldots, V_k)$.

Theorem 14 ([Bol05]) *Under the noise condition (2.2), for the k-variance of the above representation of the noisy weighted graph $G_n = (V, \boldsymbol{A}_n)$, the relation*

$$S_k^2(\boldsymbol{X}) = \mathcal{O}(\frac{1}{n})$$

*holds **AS** as $n \to \infty$ under **GC1**.*

Proof. Since F consists of step-vectors over the k-partition $P_k = (V_1, \ldots, V_k)$, by an analysis of variance argument (see [Bol92]), $S_k^2(P_k; \boldsymbol{X})$ is equal to the left-hand side of (2.16), therefore, it is $\mathcal{O}(1/n)$. This is also inherited to $S_k^2(\boldsymbol{X}) = \min_{P_k' \in \mathcal{P}_k} S_k^2(P_k'; \boldsymbol{X})$. □

Consequently, the addition of any kind of a Wigner-noise to a weight matrix that has a blown-up structure will not change the order of its structural eigenvalues, and the block structure of it can be concluded from the vertex representatives of the noisy matrix, where the representation is performed by means of the corresponding eigenvectors.

In [Bol08a] we showed that Laplacian spectra cannot be well treated under perturbations: the Laplacian eigenvalues of the above noisy graph G_n are all of order n, except the single 0. We disregard the cumbersome calculations, but this SD is included in Table 2.1. Also, Wigner-type perturbations cannot be treated in this case for the following reasons. Obviously, Laplacians of edge-disjoint simple graphs are added together; moreover, Laplacians of edge-weighted graphs are also added together. Unfortunately, no edge-weighted graph corresponds to a Wigner-noise, which usually has negative entries, and cannot be the weight-matrix of an edge-weighted graph. Nonetheless, perturbation results, analogous to those of Theorem 13 for the adjacency spectrum, can be proved for the normalized Laplacian spectrum of the noisy graph in the miniature world of the $[0, 2]$ interval.

Proposition 7 ([Bol08a]) *Let $\boldsymbol{B}_n$ be the blown-up matrix of the $k \times k$ symmetric probability matrix $\boldsymbol{P}$ of rank k. Under the growth rate condition **GC1**, there exists a constant $\delta \in (0, 1)$, independent of n, such that there are k eigenvalues of the normalized Laplacian of the edge-weighted graph $(V, \boldsymbol{B}_n)$ within the union of intervals $[0, 1-\delta]$ and $[1+\delta, 2]$; whereas, 1 is an eigenvalue with multiplicity $n-k$. Equivalently, there are $k-1$ eigenvalues of the normalized modularity matrix of $(V, \boldsymbol{B}_n)$ with absolute values at least δ; whereas, 0 is an eigenvalue with multiplicity $n-k+1$.*

This statement, as well as the following results are not proved here, as they are proved more generally in the next section for normalized contingency tables. In Proposition 12 we will prove that the normalized matrix $\boldsymbol{D}(\boldsymbol{B}_n)^{-1/2} \cdot \boldsymbol{B}_n \cdot \boldsymbol{D}(\boldsymbol{B}_n)^{-1/2}$ is also a blown-up matrix and it has k non-zero singular values within the interval $[\delta, 1]$. This implies that it has k non-zero eigenvalues within $[-1, -\delta] \cup [\delta, 1]$. Consequently, $\boldsymbol{I}_n - \boldsymbol{D}(\boldsymbol{B}_n)^{-1/2} \cdot \boldsymbol{B}_n \cdot \boldsymbol{D}(\boldsymbol{B}_n)^{-1/2}$ has k non-1 eigenvalues within $[0, 1-\delta]$ and $[1+\delta, 2]$.

Proposition 7 states that the normalized Laplacian eigenvalues of $\boldsymbol{B}_n$, which are not equal to 1, are bounded away from 1. Equivalently, the non-zero eigenvalues of the normalized modularity matrix belonging to the edge-weighted graph $(V, \boldsymbol{B}_n)$ are bounded away from zero (there are only $k-1$ such ones, as 1 is not an eigenvalue of this matrix if the underlying graph is connected, see Section 1.1.3). We claim that this property is inherited to the the normalized Laplacian or normalized modularity matrix of the noisy graph $G_n = (V, \boldsymbol{A}_n)$, where $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ with a Wigner-noise $\boldsymbol{W}_n$. More precisely, the following statement is formulated.

Theorem 15 ([Bol08a]) *Let $G_n = (V, \boldsymbol{A}_n)$ be random edge-weighted graph with $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$, where $\boldsymbol{B}_n$ is the blown-up matrix of the $k \times k$ probability matrix $\boldsymbol{P}$ of rank k, and $\boldsymbol{W}_n$ is a Wigner-noise that satisfies (2.2). Then there exists a positive constant $\delta \in (0,1)$, independent of n, such that for every $0 < \tau < 1/2$, the following statement holds **AS** as $n \to \infty$ under the growth rate condition **GC1**: there are exactly k eigenvalues of the normalized Laplacian of G_n that are located in the union of intervals $[0, 1-\delta+n^{-\tau}]$ and $[1+\delta-n^{-\tau}, 2]$, while all the others are in the interval $(1-n^{-\tau}, 1+n^{-\tau})$. Equivalently, there are exactly $k-1$ eigenvalues of the normalized modularity matrix of G_n that are at least $\delta - n^{-\tau}$, while all the others are at most $n^{-\tau}$, in absolute value.*

This statement also follows from the the analogous one stated for rectangular matrices. Here $m = n$, and hence, the so-called **GC2**, required there, is automatically satisfied here. Note that the uniform bound of the entries of $\boldsymbol{W}_n$ guarantees that the random matrix $\boldsymbol{A}_n$ has nonnegative entries and its normalized Laplacian spectrum is in the $[0, 2]$ interval.

Now, let $\mathbf{u}'_0, \ldots, \mathbf{u}'_{k-1}$ be unit-norm, pairwise orthogonal eigenvectors corresponding to the non-one eigenvalues (including the 0) of the normalized Laplacian of $\boldsymbol{B}_n$. The n-dimensional vectors obtained by the transformations

$$\mathbf{x}_i = \boldsymbol{D}(\boldsymbol{B}_n)^{-1/2}\mathbf{u}'_i \quad (i = 0, \ldots, k-1)$$

($\mathbf{x}_0 = \mathbf{1}$) are vector components of the optimal k-dimensional representation of the weighted graph $(V, \boldsymbol{B}_n)$, see Theorem 3. The $n \times k$ matrix $\boldsymbol{X}^* = (\mathbf{x}_0, \ldots, \mathbf{x}_{k-1})$ contains the optimal vertex representatives in its rows.

Let $0 < \tau < 1/2$ be arbitrary and $\epsilon := n^{-\tau}$. Let us also denote the unit-norm, pairwise orthogonal eigenvectors corresponding to the k eigenvalues of the normalized Laplacian of $G_n = (V, \boldsymbol{A}_n)$, separated from 1, by $\mathbf{v}_0, \ldots, \mathbf{v}_{k-1} \in \mathbb{R}^n$ (their existence is guaranteed by Theorem 15). Note that $\mathbf{v}_1, \ldots, \mathbf{v}_{k-1}$ also correspond to the $k-1$ structural (largest absolute value) eigenvalues of the normalized modularity matrix of G_n. Further, set

$$F' := \mathrm{Span}\{\mathbf{u}'_0, \ldots, \mathbf{u}'_{k-1}\}.$$

Proposition 8 ([Bol08a]) *With the above notation, for the distance between $\mathbf{v}_i$ and F', the following estimate holds **AS** as $n \to \infty$ under **GC1**:*

$$\mathrm{dist}(\mathbf{v}_i, F) \leq \frac{\epsilon}{(\delta - \epsilon)} = \frac{1}{(\frac{\delta}{\epsilon} - 1)}, \quad i = 0, \ldots, k-1. \tag{2.21}$$

Observe that the statement is similar to that of Proposition 6 with δ instead of Δ and ϵ instead of ε. The right-hand side of (2.21) is of order $n^{-\tau}$ that tends to zero, as $n \to \infty$. For the proof see the upcoming Section 2.2.

Proposition 8 implies the well-clustering property of the vertex representatives by means of the transformed eigenvectors

$$\mathbf{y}_i = \boldsymbol{D}(\boldsymbol{A}_n)^{-1/2}\mathbf{v}_i, \quad i = 0, \ldots, k-1.$$

The optimal k-dimensional representatives of the random edge-weighted graph $G_n = (V, \boldsymbol{A}_n)$ are row vectors of the $n \times k$ matrix $\boldsymbol{Y}^* = (\mathbf{y}_0, \ldots, \mathbf{y}_{k-1})$. The weighted k-variance of the representatives, defined by (1.14) of Chapter 1, is $\tilde{S}_k^2(P_k; \boldsymbol{Y}^*)$ with respect to the k-partition $P_k = (V_1, \ldots, V_k)$ of the vertices corresponding to the blow-up. This is the same as the weighted k-variance obtained by the $(k-1)$-dimensional representatives disregarding $\mathbf{y}_0$ and keeping only $\mathbf{y}_1, \ldots, \mathbf{y}_{k-1}$ in the normalized modularity setup.

Theorem 16 ([Bol05]) *With the above notation,*

$$\tilde{S}_k^2(\boldsymbol{Y}^*) \leq \frac{k}{(\frac{\delta}{\epsilon} - 1)^2}$$

holds ***AS*** *as* $n \to \infty$ *under* ***GC1****.*

Proof. An easy analysis of variance argument (see [Bol11c]) shows that

$$\tilde{S}_k^2(P_k; \boldsymbol{Y}^*) = \sum_{i=0}^{k-1} \mathrm{dist}^2(\mathbf{v}_i, F')$$

and hence, $\tilde{S}_k^2(\boldsymbol{Y}^*) \leq \tilde{S}_k^2(P_k; \boldsymbol{Y}^*)$ that finishes the proof. □

Spectra and spectral clusters of some generalized random graphs, artificially generated based on different types of probability matrices are shown in Section 3.1.1 of [Bol13].

We also investigated the following 'weak link' structure, which is not a blown up structure, though has some structural eigenvalues that obey a multivariate Gaussian law.

Let $k < n$ be a fixed positive integer. Now the underlying structure is the following: our edge-weighted graph consists of k disjoint components on $n_1, \ldots, n_k$ vertices, respectively. With $n = \sum_{i=1}^{k} n_i$, let $\boldsymbol{B}$ denote the $n \times n$ symmetric weight matrix, which is block-diagonal: $\boldsymbol{B} = \boldsymbol{B}^{(1)} \oplus \cdots \oplus \boldsymbol{B}^{(k)}$, where $\boldsymbol{B}^{(i)}$ is an $n_i \times n_i$ symmetric matrix with non-diagonal entries μ_i's and diagonal ones ν_i's ($\mu_i > 0$ and ν_i are real numbers), $i = 1, \ldots, k$. This means that within the connected components of the edge-weighted graph $(V, \boldsymbol{B})$ each pair of vertices is connected with an edge of the same weight, and loops are also allowed (when $\nu_i \neq 0$). The spectrum of $\boldsymbol{B}$ is the union of the spectra of $\boldsymbol{B}^{(i)}$'s. It is easy to verify that the eigenvalues of $\boldsymbol{B}^{(i)}$ are $(n_i - 1)\mu_i + \nu_i$ with eigen-direction $\mathbf{1}_{n_i}$ and $\nu_i - \mu_i$ with multiplicity $n_i - 1$ and corresponding eigen-subspace $\mathbf{1}_{n_i}^{\perp} \in \mathbb{R}^{n_i}$ $(i = 1, \ldots, k)$.

Now $\boldsymbol{B}$ is not a blown-up matrix, unless $\nu_i = \mu_i$ $(i = 1, \ldots, k)$. However, keeping μ_i's and ν_i's fixed, we can increase the size of $\boldsymbol{B}$ in such a way that $n_1, \ldots, n_k \to \infty$ under the growth rate condition **GC1**. In the sequel, we use the notation $\boldsymbol{B}_n$ for the expanding $\boldsymbol{B}$. We put a Wigner-noise $\boldsymbol{W}_n$ on $\boldsymbol{B}_n$. About the spectral properties of the weight matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ of the random edge-weighted graph $G_n = (V, \boldsymbol{A}_n)$, the following result can be stated.

Theorem 17 ([Bol04]) *Let* $\boldsymbol{W}_n$ *be an* $n \times n$ *Wigner-noise (with uniform bound and variance of the entries* K *and* σ^2*) and the matrix* $\boldsymbol{B}_n$ *be defined as above. The numbers* k*,* K*,* σ*,* μ_i*, and* ν_i $(i = 1, \ldots, k)$ *are kept fixed as* $n_1, \ldots, n_k$ *tend to infinity under* ***GC1****. Then, for the eigenvalues* $\lambda_1, \ldots, \lambda_n$ *of* $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ *the following inequalities hold* ***AS****. There is an ordering of the* k *largest absolute value eigenvalues* $\lambda_1, \ldots, \lambda_k$ *such that*

$$|\lambda_i - [(n_i - 1)\mu_i + \nu_i]| \leq 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3} \log n), \quad i = 1, \ldots, k;$$

among the other eigenvalues, for $i = 1, \ldots, k$ *there are* $n_i - 1$ λ_j*'s with*

$$|\lambda_j - [\nu_i - \mu_i]| \leq 2\sigma\sqrt{n} + \mathcal{O}(n^{1/3} \log n).$$

The proof is similar to that of Theorem 13 if we take into consideration the spectrum of $\boldsymbol{B}_n$, the Weyl's perturbation theory and the bound (2.1) for the spectral norm of a Wigner-noise. The complete proof is found in [Bol04], where *asymptotic k-variate normality* for the random vector $(\lambda_1, \ldots, \lambda_k)$ was also proved with covariance matrix $2\sigma^2\boldsymbol{I}_k$, and it was shown that the k-variance of the vertices of $G_n = (V, \boldsymbol{A}_n)$ – in the Euclidean representation defined by the corresponding eigenvectors – is $\mathcal{O}(1/n)$.

Theorem 17 implies that the k largest absolute value eigenvalues of this type of a random matrix $\boldsymbol{A}_n$ are of order $\Theta(n)$, and there must be a spectral gap between them and the remaining eigenvalues **AS** as $n \to \infty$ under **GC1**. In view of the asymptotic normality, the k largest eigenvalues are highly concentrated on their expectation of order n, independently, with finite variance. For instance, such data structures occur, when the n objects come from k loosely connected strata ($k < n$). Note that in [Gran], the importance of so-called weak links between social strata is emphasized. In our model, the weak links correspond to the entries of the Wigner-noise, see [Bol08b].

The Laplacian spectrum of the graph $(V, \boldsymbol{B})$ is characterized in [Bol08a], but carries no important information. The normalized Laplacian spectrum is again the union of those of the blocks. The normalized Laplacian matrix belonging to the block i is

$$\{[(n_i-1)\mu_i+\nu_i]\boldsymbol{I}_{n_i}\}^{-1/2}\boldsymbol{L}(\boldsymbol{B}^{(i)})\{[(n_i-1)\mu_i+\nu_i]\boldsymbol{I}_{n_i}\}^{-1/2}$$
$$= \frac{1}{(n_i-1)\mu_i+\nu_i}\boldsymbol{L}(\boldsymbol{B}^{(i)}), \quad i = 1, \ldots, k.$$

Hence, the normalized Laplacian spectrum of $(V, \boldsymbol{B})$ is as follows: the zero with multiplicity k and the numbers $\frac{n_i\mu_i}{(n_i-1)\mu_i+\nu_i}$ with multiplicity $n_i - 1$, $i = 1, \ldots, k$. Note that the letter ones tend to 1 as $n_i \to \infty$ ($i = 1, \ldots, k$). Here the loops do contribute to the spectrum.

In Table 2.1 we summarize the adjacency, Laplacian, normalized Laplacian and modularity spectra and spectral subspaces of the three main types of block- and blown-up matrices based on the previous results. Through this table we want to demonstrate that whenever the rank of the $k \times k$ probability matrix $\boldsymbol{P}$ is k, the blown-up matrix (under the usual conditions for the blow-up sizes) will asymptotically have k (or, in the modularity case, $k-1$) structural eigenvalues (separated from the others) with corresponding eigen-subspace such that the derived vertex-representatives will reveal the k underlying clusters. Latter fact is based on the piecewise constant structure of the (not necessarily) unique eigenvectors. What is only important that the eigen-subspace corresponding to the structural eigenvalues has dimension k ($k-1$ in the modularity case) and is separated from the eigen-subspace corresponding to the eigenvalues in the remainder of the spectrum. If the probability matrix has rank k, whatever small the difference between its entries is (they can even be equal), the differences between the structural and the other eigenvalues, akin to the spectral subspaces, are magnified, which results in the separation of the clusters. Of course, the speed of this separation depends on the relative values of the entries of the probability matrix and that of the blown-up sizes.

Graph $G = (V, \boldsymbol{B})$	Adjacency matrix $\boldsymbol{B}$	Laplacian matrix $\boldsymbol{D} - \boldsymbol{B}$	Normalized Laplacian $\boldsymbol{I} - \boldsymbol{D}^{-1/2}\boldsymbol{B}\boldsymbol{D}^{-1/2}$	Normalized modularity $\boldsymbol{D}^{-1/2}\boldsymbol{B}\boldsymbol{D}^{-1/2} - \sqrt{\mathbf{d}}\sqrt{\mathbf{d}}^T$
$\boldsymbol{B} = \oplus_{i=1}^{k} \boldsymbol{B}_i$, where the $n_i \times n_i$ $\boldsymbol{B}_i$ has diagonal ν_i and off-diagonal μ_i $V = (V_1, \ldots, V_k)$ $\lvert V_i \rvert = n_i$ $(i = 1, \ldots, k)$	$\lambda_i = (n_i - 1)\mu_i + \nu_i$ $(i = 1, \ldots, k)$ with piecewise constant eigenvectors over V_i's; $\nu_i - \mu_i$ with multiplicity $n_i - 1$, and eigenvectors with 0-sum coordinates over V_i $(i = 1, \ldots, k)$	0 with multiplicity k and piecewise constant eigenvectors over V_i's; $n_i\mu_i$ with multiplicity $n_i - 1$, and eigenvectors with 0-sum coordinates over V_i $(i = 1, \ldots, k)$	0 with multiplicity k and stepwise constant eigenvectors over V_i's; $\frac{n_i\mu_i}{\nu_i + (n_i - 1)\mu_i}$ with multiplicity $n_i - 1$, and eigenvectors with 0-sum coordinates over V_i $(i = 1, \ldots, k)$	1 with multiplicity $k - 1$ and stepwise constant eigenvectors over V_i's; $1 - \frac{n_i\mu_i}{\nu_i + (n_i - 1)\mu_i}$ with multiplicity $n_i - 1$, and eigenvectors with 0-sum coordinates over V_i $(i = 1, \ldots, k)$
$G = K_{n_1, \ldots, n_k}$ with independent sets V_i's, $\lvert V_i \rvert = n_i$ $(i = 1, \ldots, k)$. w.l.g. assume that $n_1 \leq \cdots \leq n_k$ $(n = \sum_{i=1}^{k} n_i)$	0 with multiplicity $n - k$ with eigenvectors of 0-sum coordinates over V_i's; the other k eigenvalues are in $[-n_k, -n_1] \cup [n - n_k, n - n_1]$ with piecewise constant eigenvectors	0 single; n with multiplicity $k - 1$ and piecewise constant eigenvectors over V_i's; $n - n_i$ with multiplicity $n_i - 1$ and eigenvectors with 0-sum coordinates over V_i's	0 single; 1 with multiplicity $n - k$; $k - 1$ eigenvalues in $[1 + \delta, 2]$, where δ does not depend on n under $\frac{n_i}{n} \geq c$ $(i = 1, \ldots, k)$	0 with multiplicity $n - k$; $k - 1$ eigenvalues in $[-1, -\delta]$, where δ does not depend on n under $\frac{n_i}{n} \geq c$ $(i = 1, \ldots, k)$
$\boldsymbol{B}$ is the blown-up matrix of $\boldsymbol{P} = (p_{ij})$ $(i, j = 1, \ldots, k)$, with blow-up sizes $n_1, \ldots, n_k$ and clusters $V_1, \ldots, V_k$; $\lvert V_i \rvert = n_i$ $(n = \sum_{i=1}^{k} n_i)$ $\operatorname{rank}(\boldsymbol{P}) = k$, $\frac{n_i}{n} \geq c$	0 with multiplicity $n - k$ with eigenvectors of 0-sum coordinates over V_i's; k non-zero eigenvalues $\lambda_1, \ldots, \lambda_k = \Theta(n)$ with piecewise constant eigenvectors over $V_1, \ldots, V_k$	0 single; $\lambda_1, \ldots, \lambda_{k-1} = \Theta(n)$ with piecewise constant eigenvectors over V_i's; $\gamma_i = \sum_{j \neq i} n_j p_{ij}$ with multiplicity $n_i - 1$ and zero-sum coordinates over V_i $(i = 1, \ldots, k)$; $\sum_{i=1}^{k-1} \lambda_i = \sum_{i=1}^{k} \gamma_i$	$\exists 0 < \delta < 1$ s.t. there are k eigenvalues (including the 0) in $[0, 1 - \delta] \cup [1 + \delta, 2]$ with piecewise constant eigenvectors over V_i's, and the 1 with multiplicity $n - k$	$\exists 0 < \delta < 1$ s.t. there are $k - 1$ eigenvalues (excluding the 1) in $[-1, -\delta] \cup [\delta, 1)$ with piecewise constant eigenvectors over V_i's, and the 0 with multiplicity $n - k$

Table 2.1: Spectra and spectral subspaces of some special block- and blown-up matrices

2.1.2 Recognizing the structure

In the previous section we investigated how the addition of a completely random Wigner-noise influences the behavior of the structural eigenvalues of the matrix having a deterministic structure. Wigner-type matrices became important in quantum mechanics, whereas in case of real-life matrices they are merely benchmarks of a random noise added to the underlying linear structure of the edge-weight matrix of communication, social, or biological networks. Whatever hard is to recognize the structure concealed by the noise, in a number of models it is possible by means of spectral techniques and large deviation principles. When we perform spectral clustering, it is a crucial question how many structural eigenvalues – with corresponding eigenvectors – to retain for the vertex-representation.

Note that numerical algorithms for the SD or SVD of a matrix with size of millions are not immediately applicable, and some newly developed randomized algorithms are to be used instead, e.g., [Ac-Mc]. These algorithms exploit the randomness of the underlying matrix, and rely on the fact that a random noise will not change the order of magnitude of the structural eigenvalues. Sometimes, instead of depriving the matrix of the noise, rather a noise is added (by digitalizing the entries of or sparsifying the underlying matrix with an appropriate randomization) to make the matrix more easily decomposable by means of the classical methods (e.g., the Lánczos method). These algorithms are only capable to find a low-rank approximation of the noisy matrix, where this rank cannot exceed the number of the structural eigenvalues of the original matrix. If there are no such eigenvalues, this property is also inherited to the randomized matrix, so the worst that can happen: we get known of the fact that there is no linear structure in our matrix at all, but it is a noise itself. In all the other cases we obtain a good approximation for the part of the spectrum needed, exploiting the randomness in our original data.

Both the number of eigenvalues to be retained and algorithmic questions can be analyzed by means of the results of this section. Also note that Wigner-type noises over a remarkable structure are not only numerically tractable, but have significance in real-life networks too. For example, sociologist Granovetter in his paper [Gran] "The strength of weak ties" shows that sometimes weak ties better help people to find job than strong (historical or family) relations in which they are stuck.

We investigate the converse question too: what kind of random matrices have a blown-up matrix as a skeleton, except of a 'small' perturbation? The following proposition states that under very general conditions an $n \times n$ random symmetric matrix with nonnegative, uniformly bounded entries (so that it can be the weight matrix of an edge-weighted graph) has at least one eigenvalue greater than of order $\sqrt{n}$.

Proposition 9 ([Bol05]) *Let $\boldsymbol{A}$ be an $n \times n$ random symmetric matrix such that $0 \le a_{ij} \le 1$ and the entries are independent for $i \le j$. Further, let us assume that there are positive constants c_1 and c_2 and $0 < \delta \le \Delta \le 1/2$ such that, with the notation $X_i = \sum_{j=1}^{n} a_{ij}$,*

$$\mathbb{E}(X_i) \ge c_1 n^{\frac{1}{2}+\delta} \quad \text{and} \quad \mathrm{Var}(X_i) \le c_2 n^{\frac{1}{2}+\Delta}, \quad i = 1, \dots, n.$$

Then for every $0 < \varepsilon < \delta$:

$$\lim_{n\to\infty} \mathbb{P}\left(\lambda_{\max}(\boldsymbol{A}) \ge c_1 n^{\frac{1}{2}+\varepsilon}\right) = 1.$$

Note that the above conditions automatically hold true if there is a constant $0 < \mu_0 < 1$ such that $\mathbb{E}(a_{ij}) \ge \mu_0$ for all i, j pairs. This is the case in the theorems of Füredi–Komlós [Fü-Ko] and Juhász [Juh81]. In our case there can be a lot of zero entries, we require only that in each row there are at least $c_1 n^{1/2+\delta}$ entries with expectation greater than or equal to any

small fixed positive constant μ_0. As the matrix is symmetric, this also holds for its columns. Therefore, among the n^2 entries there must be at least $\Theta(n^{1+2\delta})$ ones (but not anyhow) with expectation at least a fixed $0 < \mu_0 < 1$, all the others can be zeros. Also note that by the Perron–Frobenius theory, the largest eigenvalue of $\boldsymbol{A}$ is always positive.

To prove Proposition 9, the following lemma is needed.

Lemma 2 (Bernstein inequality) *Let $X_1, \ldots, X_n$ be independent random variables, $|X_i| \leq K$, $X := \sum_{i=1}^n X_i$. Then for every $a > 0$:*

$$\mathbb{P}\left(|X - \mathbb{E}(X)| > a\right) \leq e^{-\frac{a^2}{2(\mathrm{Var}(X) + Ka/3)}}.$$

Proof of Proposition 9. As a consequence of the Frobenius theorem, $\lambda_{\max}(\boldsymbol{A}) \geq \min_i X_i$, hence

$$\mathbb{P}\left(\lambda_{\max}(\boldsymbol{A}) \geq c_1 n^{\frac{1}{2}+\varepsilon}\right) \geq \mathbb{P}\left(\min_i X_i \geq c_1 n^{\frac{1}{2}+\varepsilon}\right),$$

and it is enough to prove that the latter probability tends to 1 as $n \to \infty$. We will prove that the probability of the complementary event tends to 0:

$$\mathbb{P}\left(\text{for at least one } i:\ X_i < c_1 n^{\frac{1}{2}+\varepsilon}\right) \leq n\mathbb{P}\left(\text{for a general } i:\ X_i < c_1 n^{\frac{1}{2}+\varepsilon}\right). \tag{2.22}$$

From now on, we will drop the suffix i and X denotes the sum of the entries in an arbitrary row of $\boldsymbol{A}$. As X is the sum of n independent random variables satisfying the conditions of Lemma 2 with $K = 1$,

$$\begin{aligned}
\mathbb{P}\left(X < c_1 n^{\frac{1}{2}+\varepsilon}\right) &= \mathbb{P}\left(\mathbb{E}(X) - X > \mathbb{E}(X) - c_1 n^{\frac{1}{2}+\varepsilon}\right) \\
&\leq \mathbb{P}\left(|X - \mathbb{E}(X)| > \mathbb{E}(X) - c_1 n^{\frac{1}{2}+\varepsilon}\right) \\
&\leq \mathbb{P}\left(|X - \mathbb{E}(X| > c_1 n^{\frac{1}{2}}(n^\delta - n^\varepsilon)\right) \\
&\leq e^{-\frac{c_1^2 n (n^\delta - n^\varepsilon)^2}{2(c_2 n^{\frac{1}{2}+\Delta} + n^{\frac{1}{2}}(n^\delta - n^\varepsilon)/3)}} \\
&\leq e^{-c_3 n^{\frac{1}{2}} \frac{(n^\delta - n^\varepsilon)^2}{n^\Delta}} \\
&= e^{-c_3 n^{\frac{1}{2}-\Delta}(n^\delta - n^\varepsilon)^2}
\end{aligned}$$

with some positive constant c_3, in view of the inequalities $0 < \varepsilon < \delta \leq \Delta \leq 1/2$. Thus, the right-hand side of (2.22) can be estimated from above by

$$\frac{n}{e^{c_3 n^{\frac{1}{2}-\Delta}(n^\delta - n^\varepsilon)^2}} \leq \frac{n}{e^{c_4 n^\gamma}}$$

with some constants $c_4 > 0$ and $\gamma > 0$, because of $0 < \varepsilon < \delta \leq \Delta \leq 1/2$. The last term above tends to 0 as $n \to \infty$, that finishes the proof. □

Note that the constants δ and Δ were only responsible for the speed of the convergence. Now we will use Proposition 9 to deprive a random symmetric matrix of the noise.

Theorem 18 ([Bol05]) *Let $(\boldsymbol{A}_n)$ be a sequence of $n \times n$ symmetric matrices with uniformly bounded, nonnegative entries, where n tends to infinity. Assume that $\boldsymbol{A}_n$ has exactly k eigenvalues of order greater than $\sqrt{n}$ (k is fixed), and there is a k-partition of the vertices of $G_n = (V, \boldsymbol{A}_n)$ such that the k-variance of the representatives – in the representation with the corresponding eigenvectors – is $\mathcal{O}(1/n)$. Then there is an explicit construction for a symmetric blown-up matrix $\boldsymbol{B}_n$ (on $k \times k$ blocks) such that $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{E}_n$, with $\|\boldsymbol{E}_n\| = \mathcal{O}(\sqrt{n})$.*

Instead of the complete proof, we only describe the construction, since the estimations are special cases of those performed more generally, for rectangular matrices, when we prove Theorem 23. First we prove the following lemma which will frequently be used in the sequel.

Lemma 3 ([Bol-Tus94]) *Let* $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^n$ *and* $\mathbf{y}_1, \ldots, \mathbf{y}_k \in \mathbb{R}^n$ *be orthonormal sets* $(k \leq n)$*. Then another orthonormal set of vectors* $\mathbf{v}_1, \ldots, \mathbf{v}_k$ *within* $F := \mathrm{Span}\{\mathbf{y}_1, \ldots, \mathbf{y}_k\}$ *can be found such that*

$$\sum_{i=1}^{k} \|\mathbf{x}_i - \mathbf{v}_i\|^2 \leq 2 \sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{x}_i, F).$$

Proof of the Lemma. We give a construction for such $\mathbf{v}_i$'s which are 'close' to the individual $\mathbf{x}_i$'s, respectively. Note that $\boldsymbol{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$ and $\boldsymbol{Y} = (\mathbf{y}_1, \ldots, \mathbf{y}_k)$ are $n \times k$ suborthogonal matrices. Since the vectors $\mathbf{v}_i$'s also form an orthonormal set within F, they can be obtained by applying a rotation (within F) for the $\mathbf{y}_i$'s. That is, we are looking for a $k \times k$ orthogonal matrix $\boldsymbol{R}$ such that $(\mathbf{v}_1, \ldots, \mathbf{v}_k) = \boldsymbol{Y}\boldsymbol{R}$ and, with it,

$$\sum_{i=1}^{k} \|\mathbf{x}_i - \mathbf{v}_i\|^2 = \mathrm{tr}[(\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{R})^T (\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{R})] \leq 2 \sum_{i=1}^{k} d^2(\mathbf{x}_i, F) \tag{2.23}$$

holds. By the properties of the trace operator,

$$\begin{aligned} \mathrm{tr}[(\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{R})^T (\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{R})] &= \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{X}) + \mathrm{tr}(\boldsymbol{R}^T\boldsymbol{Y}^T\boldsymbol{Y}\boldsymbol{R}) - 2\mathrm{tr}(\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{R}) \\ &= \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{X}) + \mathrm{tr}[(\boldsymbol{Y}^T\boldsymbol{Y})(\boldsymbol{R}\boldsymbol{R}^T)] - 2\mathrm{tr}(\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{R}) \\ &= 2[k - \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{R})] \end{aligned} \tag{2.24}$$

is obtained, where we used that $\boldsymbol{X}^T\boldsymbol{X} = \boldsymbol{Y}^T\boldsymbol{Y} = \boldsymbol{R}\boldsymbol{R}^T = \boldsymbol{I}_k$. The expression in (2.24) is minimal if and only if $\mathrm{tr}(\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{R})$ is maximal as a function of $\boldsymbol{R}$. At this point, we use a Proposition of [Bol83], according to which $\mathrm{tr}[(\boldsymbol{X}^T\boldsymbol{Y})\boldsymbol{R}]$ is maximal if $(\boldsymbol{X}^T\boldsymbol{Y})\boldsymbol{R}$ is symmetric, and the maximum is $\sum_{i=1}^{k} s_i$, with s_i's being the singular values of $\boldsymbol{X}^T\boldsymbol{Y}$.

Now, let $\boldsymbol{V}\boldsymbol{S}\boldsymbol{U}^T$ be the SVD of $\boldsymbol{X}^T\boldsymbol{Y}$, where $\boldsymbol{V}$ and $\boldsymbol{U}$ are $k \times k$ orthogonal matrices and $\boldsymbol{S}$ is $k \times k$ diagonal matrix with the singular values in its main diagonal. Note that the singular values s_i's are the cosines of the *principal (canonical) angles* between the subspaces $\mathrm{Span}\{\mathbf{x}_1, \ldots, \mathbf{x}_k\}$ and F. By the above SVD, $(\boldsymbol{X}^T\boldsymbol{Y})\boldsymbol{R} = \boldsymbol{V}\boldsymbol{S}\boldsymbol{U}^T\boldsymbol{R}$ is symmetric if $\boldsymbol{U}^T\boldsymbol{R} = \boldsymbol{V}^T$, i.e., $\boldsymbol{R} = \boldsymbol{U}\boldsymbol{V}^T$. Consequently, the minimum that can be attained in (2.24) is equal to

$$2 \sum_{i=1}^{k} (1 - s_i). \tag{2.25}$$

Eventually, the sum of the squared distances in (2.23) can also be written in terms of the singular values $s_1, \ldots, s_k$. Since $\boldsymbol{Y}\boldsymbol{Y}^T$ is the matrix of the orthogonal projection onto F,

$$\begin{aligned} \sum_{i=1}^{k} d^2(\mathbf{x}_i, F) &= \mathrm{tr}[(\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{X})^T (\boldsymbol{X} - \boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{X})] = \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{X}) - \mathrm{tr}(\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{X}) \\ &= k - \sum_{i=1}^{k} s_i^2 = \sum_{i=1}^{k} (1 - s_i^2), \end{aligned} \tag{2.26}$$

where we also used that the $k \times k$ symmetric matrix $\boldsymbol{X}^T\boldsymbol{Y}\boldsymbol{Y}^T\boldsymbol{X}$ has eigenvalues $s_1^2, \ldots, s_k^2$.

Comparing (2.25) and (2.26), it remains to show that $\sum_{i=1}^{k}(1-s_i) \leq \sum_{i=1}^{k}(1-s_i^2)$. But s_i's are the singular values of the matrix $\boldsymbol{X}^T\boldsymbol{Y}$, therefore denoting by $s_{\max}(.)$ the maximum singular value of the matrix in the argument, we have

$$s_i \leq s_{\max}(\boldsymbol{X}^T\boldsymbol{Y}) \leq s_{\max}(\boldsymbol{X}) \cdot s_{\max}(\boldsymbol{Y}) = 1,$$

since all positive singular values of the suborthogonal matrices $\boldsymbol{X}$ and $\boldsymbol{Y}$ are equal to 1. Hence, $s_i \geq s_i^2$ implies the desired relation (2.23). □

Construction of Theorem 18. Let $\mathbf{x}_1, \ldots, \mathbf{x}_k$ denote the eigenvectors corresponding to $\lambda_1, \ldots, \lambda_k$, the k largest absolute value (of order larger than $\sqrt{n}$) eigenvalues of $\boldsymbol{A}_n$. The representatives – that are row vectors of the $n \times k$ matrix $\boldsymbol{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$ – by the assumption of the theorem, form k clusters in $\mathbb{R}^k$ with k-variance less than c/n with some constant c. Let $V_1, \ldots, V_k$ denote the clusters; properly reordering the rows of $\boldsymbol{X}$, together they give the index set $V = \{1, \ldots, n\}$. Let $\mathbf{x}^1, \ldots, \mathbf{x}^n \in \mathbb{R}^k$ be the Euclidean representatives of the vertices (the rows of $\boldsymbol{X}$), and let $\bar{\mathbf{x}}^1, \ldots, \bar{\mathbf{x}}^k$ denote the cluster centers. Now let us choose the following representation of the vertices. The representatives are row vectors of the $n \times k$ matrix $\tilde{\boldsymbol{X}}$ such that the first n_1 rows of $\tilde{\boldsymbol{X}}$ are equal to $\bar{\mathbf{x}}^1$, the next n_2 rows of $\tilde{\boldsymbol{X}}$ are equal to $\bar{\mathbf{x}}^2$, ..., and so on; the last n_k rows of $\tilde{\boldsymbol{X}}$ are equal to $\bar{\mathbf{x}}^k$. Finally, let $\mathbf{y}_1, \ldots, \mathbf{y}_k \in \mathbb{R}^n$ be the column vectors of $\tilde{\boldsymbol{X}}$. By an easy (already applied) analysis of variance algorithm,

$$S_k^2(\boldsymbol{X}) = \sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{x}_i, F),$$

where the k-dimensional subspace F is spanned by the vectors $\mathbf{y}_1, \ldots, \mathbf{y}_k$; further, by the assumption of the theorem, $S_k^2(\boldsymbol{X}) < \frac{c}{n}$.

Then, in view of Lemma 3, a set $\mathbf{v}_1, \ldots, \mathbf{v}_k$ of orthonormal vectors within F can be found such that

$$\sum_{i=1}^{k} \|\mathbf{x}_i - \mathbf{v}_i\|^2 \leq 2\frac{c}{n}$$

holds. It is important that $\mathbf{v}_i$'s are also step-vectors over the vertex-clusters $V_1, \ldots, V_k$.

Finally, for the matrix $\boldsymbol{A}_n = \sum_{i=1}^{n} \lambda_i \mathbf{x}_i \mathbf{x}_i^T$, the blown-up matrix $\boldsymbol{B}_n = \sum_{i=1}^{k} \lambda_i \mathbf{v}_i \mathbf{v}_i^T$ is constructed. Then the spectral norm of the error matrix $\boldsymbol{E}_n = \boldsymbol{A}_n - \boldsymbol{B}_n$ is $\mathcal{O}(\sqrt{n})$, as it will be proved more generally, in the rectangular case. □

Theoretically, for any graph on n vertices, the Szemerédi Regularity Lemma (see [Kometal, Sim-Sos, Szem]) guarantees the existence of an equitable regular partition of the vertices with 'small' pairwise discrepancies such that the upper limit for the maximum number of clusters does not depend on n (it only depends on the discrepancy to be attained).

Definition 20 *We say that the disjoint pair $V_i, V_j \subset V$ $(i \neq j)$ is ϵ-regular, if for any $A \subset V_i$, $B \subset V_j$, with $|A| > \epsilon|V_i|$, $|B| > \epsilon|V_j|$,*

$$|\rho(A,B) - \rho(V_i, V_j)| < \epsilon$$

holds, where $\rho(A,B)$ denotes the edge-density between the disjoint vertex-subsets A and B. More precisely, denoting by $e(A,B)$ the number of cut-edges between A and B,

$$\rho(A,B) = \frac{e(A,B)}{|A| \cdot |B|}.$$

Theorem (Szemerédi Regularity Lemma [Szem]) *For every positive real ϵ and positive integer m there are positive integers $N = N(\epsilon, m)$ and $M = M(\epsilon, m)$ with the following property: for every simple graph $G = (V, E)$ with $n \geq N$ vertices there is a partition of V into $k + 1$ classes $V_0, V_1, \ldots, V_k$ such that*

- $m \leq k \leq M$,
- $|V_0| \leq \epsilon n$,
- $|V_1| = |V_2| = \cdots = |V_k|$,
- *all but at most ϵk^2 of the pairs (V_i, V_j) are ϵ-regular.*

If the graph is sparse – the number of edges is $e = o(n^2)$ – then $k = 1$, otherwise k can be arbitrarily large (but the upper limit for it does not depend on n, it merely depends on ϵ).

Definition 21 *Let n be a natural number and $k \leq n$ be a positive integer. The graph $G_n(\boldsymbol{P}, \mathcal{P}_k)$ is a generalized random graph with probability matrix $\boldsymbol{P}$ and proper k-partition $\mathcal{P}_k = (V_1, \ldots, V_k)$ of the vertices if it satisfies the following. The vertex set is V, $|V| = n$; the $k \times k$ symmetric matrix $\boldsymbol{P}$ is such that its entries satisfy $0 \leq p_{ij} \leq 1$ $(1 \leq i \leq j \leq k)$. Then vertices of V_i and V_j are connected independently, with probability p_{ij}, $1 \leq i \leq j \leq k$.*

If our random graph is a generalized random graph, then $e(V_i, V_j)$ is the sum of $|V_i| \cdot |V_j|$ independent, identically distributed Bernoulli variables with parameter p_{ij} $(1 \leq i, j \leq k)$, where p_{ij}'s are entries of the probability matrix $\boldsymbol{P}$. Hence, $e(A, B)$ is a binomially distributed random variable with expectation $|A| \cdot |B| \cdot p_{ij}$ and variance $|A| \cdot |B| \cdot p_{ij}(1 - p_{ij})$. Therefore, by the Bernstein inequality (Lemma 2 with the choice of $K = 1$) and with $A \subset V_i$, $B \subset V_j$, $|A| > \epsilon |V_i|$, $|B| > \epsilon |V_j|$ we have that

$$\begin{aligned}
\mathbb{P}\left(|\rho(A, B) - p_{ij}| > \epsilon\right) &= \mathbb{P}\left(|e(A, B) - |A| \cdot |B| \cdot p_{ij}| > \epsilon \cdot |A| \cdot |B|\right) \\
&\leq e^{-\frac{\epsilon^2 |A|^2 |B|^2}{2[|A||B| p_{ij}(1-p_{ij}) + \epsilon |A||B|/3]}} \\
&= e^{-\frac{\epsilon^2 |A||B|}{2[p_{ij}(1-p_{ij}) + \epsilon/3]}} \\
&\leq e^{-\frac{\epsilon^4 |V_i||V_j|}{2[p_{ij}(1-p_{ij}) + \epsilon/3]}},
\end{aligned}$$

that tends to 0 as $|V_i| = n_i \to \infty$ and $|V_j| = n_j \to \infty$. Hence, any pair V_i, V_j is ϵ-regular, **AS**. Note that a matrix approximation theorem to derive a constructive version of the Szemerédi Regularity Lemma is introduced in [Fr-Kan]. Further aspects will be discussed in Section 2.3.

2.2 Noisy contingency tables

In this section, the results of the previous section will be extended to the stability of the SVD of large noisy contingency tables.

Definition 22 *The $m \times n$ real matrix $\boldsymbol{W}_{m \times n}$ is a Wigner-noise if its entries w_{ij} $(1 \leq i \leq m,\ 1 \leq j \leq n)$ are independent random variables, $\mathbb{E}(w_{ij}) = 0$, and the w_{ij}'s are uniformly bounded (i.e., there is a constant $K > 0$, independently of m and n, such that $|w_{ij}| \leq K$, $\forall i, j$).*

Achlioptas and McSherry [Ac-Mc] extended the Füredi–Komlós theorem [Fü-Ko] to rectangular matrices. In view of this, the spectral norm of an $m \times n$ Wigner-noise $\boldsymbol{W}_{m\times n}$ is at most of order $\sqrt{m+n}$, with probability tending to 1 as $n, m \to \infty$.

Definition 23 *The $m\times n$ real matrix $\boldsymbol{B}_{m\times n}$ is a blown-up matrix if there is an $a\times b$ so-called probability matrix $\boldsymbol{P}$ with entries $0 < p_{ij} < 1$, and there are positive integers $m_1, \ldots, m_a$ with $\sum_{i=1}^{a} m_i = m$ and $n_1, \ldots, n_b$ with $\sum_{i=1}^{b} n_i = n$ such that – after rearranging its rows and columns – the matrix $\boldsymbol{B}_{m\times n}$ can be divided into $a \times b$ blocks, where the block (i, j) is an $m_i \times n_j$ matrix with entries all equal to p_{ij} $(1 \le i \le a,\, 1 \le j \le b)$.*

Let us fix $\boldsymbol{P}$, blow it up to an $m \times n$ matrix $\boldsymbol{B}_{m\times n}$, and consider the noisy matrix $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$ as $m, n \to \infty$ under one or both of the subsequent growth rate conditions.

Definition 24 *The following growth rate conditions for the growth of the sizes and that of the cluster sizes of an $m \times n$ rectangular array are introduced.*

GC1 *There exists a constant $0 < c \le \frac{1}{a}$ such that $\frac{m_i}{m} \ge c$ $(i = 1, \ldots, a)$ and a constant $0 < d \le \frac{1}{b}$ such that $\frac{n_i}{n} \ge d$ $(i = 1, \ldots, b)$.*

GC2 *There exist constants $C \ge 1$, $D \ge 1$, and $C_0 > 0$, $D_0 > 0$ such that $m \le C_0 n^C$ and $n \le D_0 m^D$ for sufficiently large m and n.*

We remark the following:

- **GC1** implies that
$$c \le \frac{m_k}{m_i} \le \frac{1}{c} \quad \text{and} \quad d \le \frac{n_\ell}{n_j} \le \frac{1}{d} \tag{2.27}$$
hold for any pair of indices $k, i \in \{1, \ldots, a\}$ and $\ell, j \in \{1, \ldots, b\}$.

- **GC2** implies that for sufficiently large m and n,
$$(\frac{1}{D_0})^{\frac{1}{D}} n^{\frac{1}{D}} \le m \le C_0 n^C \quad \text{and} \quad (\frac{1}{C_0})^{\frac{1}{C}} m^{\frac{1}{C}} \le n \le D_0 m^D.$$
Therefore, **GC2** mildly regulates the relation between m and n, but they need not tend to infinity at the same rate.

While perturbing $\boldsymbol{B}_{m\times n}$ by $\boldsymbol{W}_{m\times n}$, assume that for the uniform bound of the entries of $\boldsymbol{W}_{m\times n}$ the condition
$$K \le \min\{\min_{\substack{i\in\{1,\ldots,a\}\\ j\in\{1,\ldots,b\}}} p_{ij}\,,\, 1 - \max_{\substack{i\in\{1,\ldots,a\}\\ j\in\{1,\ldots,b\}}} p_{ij}\} \tag{2.28}$$
is satisfied. In this way, the entries of $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$ are in the [0,1] interval, and hence, $\boldsymbol{A}_{m\times n}$ defines a standardized contingency table. We are interested in the asymptotic properties of the SVD of this expanding contingency table sequence as $n, m \to \infty$ under the growth rate conditions.

With an appropriate Wigner-noise we can again guarantee that the noisy binary table $\boldsymbol{A}_{m\times n}$ contains 1's in the (a, b)-th block with probability p_{ab}, and 0's otherwise.

Let $R_1, \ldots, R_a$ and $C_1, \ldots, C_b$ denote the row- and column clusters induced by the blow-up. In the random 0-1 contingency table $\boldsymbol{A}_{m\times n}$, the row and column categories of R_i and C_j are in interaction with probability p_{ij}. Such schemes are sought for in microarray analysis and they called *checker-board patterns*, see [Klugetal] for details. In terms of microarrays,

the above property means that genes of the same cluster R_i equally influence conditions of the cluster C_j.

Analogously to the quadratic case, we define **AS** and **WP1** properties $\mathcal{P}_{m,n}$ for the rectangular array $\boldsymbol{A}_{m\times n}$. We may assume **GC1** and/or **GC2** for the simultaneous growth of m and n. In fact, **GC2** will only be used for normalized noisy tables.

To find **AS** estimate for the spectral norm, i.e., the largest singular value of $\boldsymbol{W}_{m\times n}$, the theorem of Alon, Krivelevich, and Vu (see Section 2.1) can be adopted to rectangular matrices in the following manner. Let $\boldsymbol{W}_{m\times n}$ be a Wigner-noise with entries uniformly bounded by K. The $(m+n)\times(m+n)$ symmetric matrix

$$\widetilde{\boldsymbol{W}} = \frac{1}{K}\begin{pmatrix} \boldsymbol{0} & \boldsymbol{W}_{m\times n} \\ \boldsymbol{W}^T_{m\times n} & \boldsymbol{0} \end{pmatrix}$$

satisfies the conditions of this theorem, and its largest and smallest eigenvalues are

$$\lambda_i(\widetilde{\boldsymbol{W}}) = -\lambda_{n+m-i+1}(\widetilde{\boldsymbol{W}}) = \frac{1}{K}\cdot s_i(\boldsymbol{W}_{m\times n}), \quad i = 1,\dots,\min\{m,n\},$$

while the others are zeros, where $\lambda_i(.)$ and $s_i(.)$ denote the i-th largest eigenvalue and singular value of the matrix in the argument, respectively. Therefore,

$$\boldsymbol{P}\left(|s_1(\boldsymbol{W}_{m\times n}) - \mathbb{E}(s_1(\boldsymbol{W}_{m\times n}))| > t\right) \le \exp\left(-\frac{(1-o(1))t^2}{32K^2}\right). \tag{2.29}$$

The fact that $s_1(\boldsymbol{W}_{m\times n}) = \|\boldsymbol{W}_{m\times n}\| = \mathcal{O}(\sqrt{m+n})$ **WP1** and inequality (2.29) together ensure that $\mathbb{E}(\|\boldsymbol{W}\|) = \mathcal{O}(\sqrt{m+n})$. Hence, no matter how $\mathbb{E}\|\boldsymbol{W}_{m\times n}\|$ behaves when $m,n\to\infty$, the following rough estimate holds: there exist positive constants c_1 and c_2, depending merely on the common bound K of the entries of $\boldsymbol{W}_{m\times n}$, such that

$$\mathbb{P}\left(\|\boldsymbol{W}_{m\times n}\| > c_1\sqrt{m+n}\right) \le e^{-c_2(m+n)}. \tag{2.30}$$

Since the right-hand side of (2.30) forms a convergent series, the spectral norm of a Wigner-noise $\boldsymbol{W}_{m\times n}$ is of order $\sqrt{m+n}$ **AS**. This observation will provide the base of the subsequent **AS** results, which are also **WP1** ones.

2.2.1 Singular values of a noisy contingency table

Analogously to the quadratic case, the following is proved for the noisy table.

Proposition 10 ([Bol-Fr-Kr10]) *Under **GC1**, all the non-zero singular values of the blown-up contingency table $\boldsymbol{B}_{m\times n}$ are of order $\sqrt{mn}$.*

Theorem 19 ([Bol-Fr-Kr10]) *Let $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$ be an $m\times n$ random matrix, where $\boldsymbol{B}_{m\times n}$ is a blown-up matrix with positive singular values $s_1,\dots,s_r$ ($r = \operatorname{rank}(\boldsymbol{B}_{m\times n})$) and $\boldsymbol{W}_{m\times n}$ is a Wigner-noise. Then, under **GC1**, the matrix $\boldsymbol{A}_{m\times n}$ has r singular values $z_1,\dots,z_r$, such that*

$$|z_i - s_i| = \mathcal{O}(\sqrt{m+n}), \quad i = 1,\dots,r$$

and for the other singular values

$$z_j = \mathcal{O}(\sqrt{m+n}), \quad j = r+1,\dots,\min\{m,n\}$$

*holds, **AS**.*

Summarizing, with the notation

$$\varepsilon := \|\boldsymbol{W}_{m\times n}\| = \mathcal{O}(\sqrt{m+n}) \quad \text{and} \quad \Delta := \min_{1\le i\le r} s_i(\boldsymbol{B}_{m\times n}) = \Theta(\sqrt{mn}) \tag{2.31}$$

there is a spectral gap of size $\Delta - 2\varepsilon$ between the r largest and the other singular values of the perturbed matrix $\boldsymbol{A}_{m\times n}$, and this gap is significantly larger than ε.

2.2.2 Clustering the rows and columns via singular vector pairs

Here perturbation results for the corresponding singular vector pairs are established. To this end, with the help of Theorem 19, we estimate the distances between the corresponding right- and left-hand side eigenspaces of the matrices $\boldsymbol{B}_{m\times n}$ and $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$.

Let $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \mathbb{R}^m$ and $\mathbf{u}_1, \ldots, \mathbf{u}_n \in \mathbb{R}^n$ be the orthonormal left- and right-hand side singular vectors of $\boldsymbol{B}_{m\times n}$. They are step-vectors over the row and column clusters determined by the blow-up; further, they satisfy

$$\boldsymbol{B}_{m\times n}\mathbf{u}_i = s_i\mathbf{v}_i, \quad i = 1, \ldots, r \quad \text{and} \quad \boldsymbol{B}_{m\times n}\mathbf{u}_j = 0, \quad j = r+1, \ldots, n.$$

Let us also denote the unit-norm, pairwise orthogonal left- and right-hand side singular vectors corresponding to the r structural singular values $z_1, \ldots, z_r$ of $\boldsymbol{A}_{m\times n}$ by $\mathbf{y}_1, \ldots, \mathbf{y}_r \in \mathbb{R}^m$ and $\mathbf{x}_1, \ldots, \mathbf{x}_r \in \mathbb{R}^n$, respectively. Then, for them $\boldsymbol{A}_{m\times n}\mathbf{x}_i = z_i\mathbf{y}_i$ holds $(i = 1, \ldots, r)$. Let

$$F := \mathrm{Span}\{\mathbf{v}_1, \ldots, \mathbf{v}_r\} \quad \text{and} \quad G := \mathrm{Span}\{\mathbf{u}_1, \ldots, \mathbf{u}_r\}$$

denote the spanned linear subspaces of step-vectors in $\mathbb{R}^m$ and $\mathbb{R}^n$, respectively.

Proposition 11 ([Bol-Fr-Kr10]) *With the above notation, under **GC1**, the following estimate holds **AS**:*

$$\sum_{i=1}^{r} \mathrm{dist}^2(\mathbf{y}_i, F) \le r\frac{\varepsilon^2}{(\Delta-\varepsilon)^2} = \mathcal{O}\left(\frac{m+n}{mn}\right), \tag{2.32}$$

and analogously,

$$\sum_{i=1}^{r} \mathrm{dist}^2(\mathbf{x}_i, G) \le r\frac{\varepsilon^2}{(\Delta-\varepsilon)^2} = \mathcal{O}\left(\frac{m+n}{mn}\right). \tag{2.33}$$

By Proposition 11, the individual distances between the original and the perturbed subspaces and also the sum of these distances tend to zero **AS** as $m, n \to \infty$ under **GC1**.

Now let $\boldsymbol{A}_{m\times n}$ be a microarray on m genes and n conditions, with a_{ij} denoting the expression level of gene i under condition j. We assume that $\boldsymbol{A}_{m\times n}$ is a noisy random matrix obtained by adding a Wigner-noise $\boldsymbol{W}_{m\times n}$ to the blown-up matrix $\boldsymbol{B}_{m\times n}$. Let us denote by $R_1, \ldots, R_a$ the partition of the genes and by $C_1, \ldots, C_b$ the partition of the conditions with respect to the blow-up (they can also be thought of as clusters of the genes and conditions).

Proposition 11 also implies the well-clustering property of the representatives of the genes and conditions in the following representation. Let $\boldsymbol{Y}$ be the $m \times r$ matrix containing the left-hand side singular vectors $\mathbf{y}_1, \ldots, \mathbf{y}_r$ of $\boldsymbol{A}_{m\times n}$ in its columns. Likewise, let $\boldsymbol{X}$ be the $n \times r$ matrix containing the right-hand side singular vectors $\mathbf{x}_1, \ldots, \mathbf{x}_r$ of $\boldsymbol{A}_{m\times n}$ in its columns. Let the r-dimensional representatives of the genes be the row vectors of $\boldsymbol{Y}$: $\mathbf{y}^1, \ldots, \mathbf{y}^m \in \mathbb{R}^r$, while the r-dimensional representatives of the conditions be the row vectors of $\boldsymbol{X}$: $\mathbf{x}^1, \ldots, \mathbf{x}^n \in \mathbb{R}^r$. Let $S_a^2(\boldsymbol{Y})$ denote the a-variance, introduced in (1.10) of the genes' representatives:

$$S_a^2(\boldsymbol{Y}) = \min_{\{R_1', \ldots, R_a'\}} \sum_{i=1}^{a} \sum_{j\in R_i'} \|\mathbf{y}^j - \bar{\mathbf{y}}^i\|^2, \quad \text{where} \quad \bar{\mathbf{y}}^i = \frac{1}{m_i}\sum_{j\in R_i'} \mathbf{y}^j,$$

while $S_b^2(\boldsymbol{X})$ denotes the b-variance of the conditions' representatives:

$$S_b^2(\boldsymbol{X}) = \min_{\{C'_1,\ldots,C'_b\}} \sum_{i=1}^{b} \sum_{j\in C'_i} \|\mathbf{x}^j - \bar{\mathbf{x}}^i\|^2, \quad \text{where} \quad \bar{\mathbf{x}}^i = \frac{1}{n_i} \sum_{j\in C'_i} \mathbf{x}^j,$$

when the partitions $\{R'_1,\ldots,R'_a\}$ and $\{C'_1,\ldots,C'_b\}$ vary over all a- and b-partitions of the genes and conditions, respectively.

Theorem 20 ([Bol-Fr-Kr10]) *With the above notation, under **GC1**, for the a- and b-variances of the representation of the microarray $\boldsymbol{A}_{m\times n}$ the relations*

$$S_a^2(\boldsymbol{Y}) = \mathcal{O}\left(\frac{m+n}{mn}\right) \quad \textit{and} \quad S_b^2(\boldsymbol{X}) = \mathcal{O}\left(\frac{m+n}{mn}\right)$$

*hold **AS**.*

The proofs are analogous to those of the quadratic case. The message of the above theorem is that the addition of any kind of a Wigner-noise to a rectangular matrix that has a blown-up structure $\boldsymbol{B}_{m\times n}$, will not change the order of the structural singular values, and the block structure of $\boldsymbol{B}_{m\times n}$ can be reconstructed from the representatives of the row and column items of the noisy matrix $\boldsymbol{A}_{m\times n}$. So far, we have only used **GC1**, and no restriction for the relation between m and n has been made. For noisy normalized contingency tables, **GC2** is also needed.

2.2.3 Perturbation results for noisy normalized contingency tables

Let $\boldsymbol{B}_{m\times n}$ be the blown-up matrix of the $a\times b$ probability matrix $\boldsymbol{P}$ of positive entries and rank r, with blow-up sizes $m_1,\ldots,m_a$ and $n_1,\ldots,n_b$. We perform the normalization, described in Section 1.2, on $\boldsymbol{B}_{m\times n}$. We are interested in the order of the singular values of $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$ when the normalization is applied to it. To this end, we introduce the following notation

$$\begin{aligned}
\boldsymbol{D}_{Brow} &= \operatorname{diag}(d_{Brow,1},\ldots,d_{Brow,m}) = \operatorname{diag}\left(\sum_{j=1}^{n} b_{1j},\ldots,\sum_{j=1}^{n} b_{mj}\right)\\
\boldsymbol{D}_{Bcol} &= \operatorname{diag}(d_{Bcol,1},\ldots,d_{Bcol,n}) = \operatorname{diag}\left(\sum_{i=1}^{m} b_{i1},\ldots,\sum_{i=1}^{m} b_{in}\right)\\
\boldsymbol{D}_{Arow} &= \operatorname{diag}(d_{Arow,1},\ldots,d_{Arow,m}) = \operatorname{diag}\left(\sum_{j=1}^{n} a_{1j},\ldots,\sum_{j=1}^{n} a_{mj}\right)\\
\boldsymbol{D}_{Acol} &= \operatorname{diag}(d_{Acol,1},\ldots,d_{Acol,n}) = \operatorname{diag}\left(\sum_{i=1}^{m} a_{i1},\ldots,\sum_{i=1}^{m} a_{in}\right)
\end{aligned}$$

for the diagonal matrices containing the row- and column-sums of $\boldsymbol{B}_{m\times n}$ and $\boldsymbol{A}_{m\times n}$ in their main diagonals, respectively. For notational convenience, we discard the subscript $m\times n$ and set

$$\boldsymbol{B}_D = \boldsymbol{D}_{Brow}^{-1/2}\boldsymbol{B}\boldsymbol{D}_{Bcol}^{-1/2} \quad \text{and} \quad \boldsymbol{A}_D = \boldsymbol{D}_{Arow}^{-1/2}\boldsymbol{A}\boldsymbol{D}_{Acol}^{-1/2} \tag{2.34}$$

for the transformed matrices while carrying out the normalization on $\boldsymbol{B}$ and $\boldsymbol{A}$, respectively. It is known that the leading singular value of $\boldsymbol{B}_D$ is equal to 1, and it is a single singular

value if $\boldsymbol{B}$ is non-degenerate, which will be assumed from now on. Let $1 > s_0 = s_1 \geq \cdots \geq s_{r-1}$ denote a non-zero singular values of $\boldsymbol{B}_D$ with unit-norm singular vector pairs $\mathbf{v}_i$, $\mathbf{u}_i$ $(i = 0, \ldots, r-1)$, where $r = \mathrm{rank}(\boldsymbol{B}_D) = \mathrm{rank}(\boldsymbol{B}) = \mathrm{rank}(\boldsymbol{P})$. With the transformations

$$\mathbf{v}_{corr,i} = \boldsymbol{D}_{Brow}^{-1/2}\mathbf{v}_i \quad \text{and} \quad \mathbf{u}_{corr,i} = \boldsymbol{D}_{Bcol}^{-1/2}\mathbf{u}_i \quad (i = 0, 1, \ldots, r-1)$$

the so-called correspondence vector pairs are obtained. To $\mathbf{v}_0, \mathbf{u}_0$, the trivial (constantly 1) vector pair corresponds. If the coordinates $v_{corr,i}(j)$, $u_{corr,i}(\ell)$ of such a pair are considered as possible values of two discrete random variables ψ_i and ϕ_i (often called the i-th correspondence factor pair) taking on values with the marginal distributions, then – akin to the canonical analysis – their correlation is s_i, and this is the largest possible correlation under the condition that they are uncorrelated to the first $i-1$ correspondence factors within their own sets, as discussed in Section 1.3.

In Section 1.2, we proved that the row vectors of the matrices $\mathbf{v}_{corr,1}, \ldots, \mathbf{v}_{corr,k-1}$ and $\mathbf{u}_{corr,1}, \ldots, \mathbf{u}_{corr,k-1}$ are optimal $(k-1)$-dimensional representatives of the rows and columns of the contingency table (where the trivial 0-index coordinates can as well be included to obtain a k-dimensional representation). Here we will establish **AS** properties for the k-variances of these representatives when also make use of the growth rate condition **GC2**.

Proposition 12 ([Bol-Fr-Kr10]) *Given the $m \times n$ blown-up matrix $\boldsymbol{B}$ of the $a \times b$ probability matrix $\boldsymbol{P}$, under **GC1** there exists a constant $\delta \in (0, 1)$, independent of m and n, such that all the non-zero singular values of $\boldsymbol{B}_D$ are in the interval $[\delta, 1]$.*

Proof. It is easy to see that $\boldsymbol{B}_D$ is the blown-up matrix of the $a \times b$ probability matrix $\tilde{\boldsymbol{P}}$ with entries

$$\tilde{p}_{ij} = \frac{p_{ij}}{\sqrt{(\sum_{\ell=1}^{b} p_{i\ell}n_\ell)(\sum_{k=1}^{a} p_{kj}m_k)}}.$$

Likewise as stated in Proposition 10, the blown-up matrix $\boldsymbol{B}_D$ has exactly $r = \mathrm{rank}(\boldsymbol{P}) = \mathrm{rank}(\tilde{\boldsymbol{P}})$ non-zero singular values that are the singular values of the $a \times b$ matrix $\boldsymbol{P}' = \boldsymbol{D}_a^{1/2}\tilde{\boldsymbol{P}}\boldsymbol{D}_b^{1/2}$ (where $\boldsymbol{D}_a = \mathrm{diag}(m_1, \ldots, m_a)$ and $\boldsymbol{D}_b = \mathrm{diag}(n_1, \ldots, n_b)$) with entries

$$p'_{ij} = \frac{p_{ij}\sqrt{m_i}\sqrt{n_j}}{\sqrt{(\sum_{\ell=1}^{b} p_{i\ell}n_\ell)(\sum_{k=1}^{a} p_{kj}m_k)}} = \frac{p_{ij}}{\sqrt{(\sum_{\ell=1}^{b} p_{i\ell}\frac{n_\ell}{n_j})(\sum_{k=1}^{a} p_{kj}\frac{m_k}{m_i})}}.$$

Since the matrix $\boldsymbol{P}$ contains no zero entries, under **GC1**, the matrix $\boldsymbol{P}'$ varies on a compact set of $a \times b$ matrices determined by the inequalities (2.27). The range of the non-zero singular values depends continuously on the matrix that does not depend on m and n. Therefore, the minimum non-zero singular value does not depend on m and n. Because the largest singular value is 1, this finishes the proof. □

Theorem 21 ([Bol-Fr-Kr10]) *Consider the normalized matrix obtained from the noisy matrix $\boldsymbol{A}_{m\times n} = \boldsymbol{B}_{m\times n} + \boldsymbol{W}_{m\times n}$, see (2.34); $r = \mathrm{rank}\boldsymbol{B}_{m\times n}$. Under **GC1** and **GC2**, there exists a positive constant δ, that does not depend on m and n, such that for every $0 < \tau < 1/2$ the following statement holds **AS**: the r largest singular values of the normalized table are in the interval $[\delta - \max\{n^{-\tau}, m^{-\tau}\}, 1 + \max\{n^{-\tau}, m^{-\tau}\}]$, while all the others are at most $\max\{n^{-\tau}, m^{-\tau}\}$.*

Note that under the condition (2.28), the noisy matrix $\boldsymbol{A}_{m\times n}$ has nonnegative entries (it is a contingency table), and so, the noisy contingency table has largest singular value 1. Therefore, its r largest singular values are in the interval $[\delta - \max\{n^{-\tau}, m^{-\tau}\}, 1]$.

Proof. We will drop the subscript $m \times n$. First notice that

$$\boldsymbol{A}_D = \boldsymbol{D}_{Arow}^{-1/2} \boldsymbol{A} \boldsymbol{D}_{Acol}^{-1/2} = \boldsymbol{D}_{Arow}^{-1/2} \boldsymbol{B} \boldsymbol{D}_{Acol}^{-1/2} + \boldsymbol{D}_{Arow}^{-1/2} \boldsymbol{W} \boldsymbol{D}_{Acol}^{-1/2}. \tag{2.35}$$

The entries of $\boldsymbol{D}_{Brow}$ and those of $\boldsymbol{D}_{Bcol}$ are of order $\Theta(n)$ and $\Theta(m)$, respectively. Now we prove that for every $i = 1, \dots, m$ and $j = 1, \dots, n$, $|d_{Arow,i} - d_{Brow,i}| < n \cdot n^{-\tau}$ and $|d_{Acol,j} - d_{Bcol,j}| < m \cdot m^{-\tau}$ hold, **AS**. To this end, we use the Bernstein inequality for large deviations (see Lemma 2):

$$\begin{aligned}
&\mathbb{P}\left(|d_{Arow\,i} - d_{Brow\,i}| > n \cdot n^{-\tau}\right) = \mathbb{P}\left(\left|\sum_{j=1}^{n} w_{ij}\right| > n^{1-\tau}\right) \\
&\quad < \exp\left\{-\frac{n^{2-2\tau}}{2(\mathrm{Var}(\sum_{j=1}^{n} w_{ij}) + Kn^{1-\tau}/3)}\right\} \le \exp\left\{-\frac{n^{2-2\tau}}{2(n\sigma^2 + Kn^{1-\tau}/3)}\right\} \\
&\quad = \exp\left\{-\frac{n^{1-2\tau}}{2(\sigma^2 + Kn^{-\tau}/3)}\right\}, \quad i = 1, \dots, m,
\end{aligned}$$

where the constant K is the uniform bound for $|w_{ij}|$'s and σ^2 is the bound for their variances. In view of **GC2** the following estimate holds with some $C_0 > 0$ and $C \ge 1$ (constants of **GC2**) and large enough n:

$$\begin{aligned}
&\mathbb{P}\left(|d_{Arow,i} - d_{Brow,i}| > n^{1-\tau} \quad \forall i \in \{1, \dots, m\}\right) \\
&\quad \le m \cdot \exp\left\{-\frac{n^{1-2\tau}}{2(\sigma^2 + Kn^{-\tau}/3)}\right\} \le C_0 \cdot n^C \cdot \exp\left\{-\frac{n^{1-2\tau}}{2(\sigma^2 + Kn^{-\tau}/3)}\right\} \\
&\quad = \exp\left\{\ln C_0 + C \ln n - \frac{n^{1-2\tau}}{2(\sigma^2 + Kn^{-\tau}/3)}\right\}.
\end{aligned} \tag{2.36}$$

The estimation of the probability

$$\mathbb{P}\left(|d_{Acol,j} - d_{Bcol,j}| > m^{1-\tau} \quad \forall j \in \{1, \dots, n\}\right)$$

can be treated analogously (with $D_0 > 0$ and $D \ge 1$ of **GC2**). The right-hand side of (2.36) forms a convergent series; therefore, by the Borel–Cantelli lemma,

$$\min_{i \in \{1,\dots,m\}} |d_{Arow,i}| = \Theta(n), \qquad \min_{j \in \{1,\dots,n\}} |d_{Acol,j}| = \Theta(m) \tag{2.37}$$

hold **AS**.

Now it is straightforward to bound the norm of the second term of (2.35) by

$$\|\boldsymbol{D}_{Arow}^{-1/2}\| \cdot \|\boldsymbol{W}\| \cdot \|\boldsymbol{D}_{Acol}^{-1/2}\|. \tag{2.38}$$

As by (2.30), $\|\boldsymbol{W}\| = \mathcal{O}(\sqrt{m+n})$ holds **AS**, the quantity (2.38) is at most of order $\sqrt{\frac{m+n}{mn}}$, **AS**. Hence, it is less than $\max\{n^{-\tau}, m^{-\tau}\}$, **AS**.

In order to estimate the norm of the first term of (2.35) let us write it in the form

$$\begin{aligned}
\boldsymbol{D}_{Arow}^{-1/2} \boldsymbol{B} \boldsymbol{D}_{Acol}^{-1/2} = \boldsymbol{D}_{Brow}^{-1/2} \boldsymbol{B} \boldsymbol{D}_{Bcol}^{-1/2} + \left[\boldsymbol{D}_{Arow}^{-1/2} - \boldsymbol{D}_{Brow}^{-1/2}\right] \boldsymbol{B} \boldsymbol{D}_{Bcol}^{-1/2} \\
+ \boldsymbol{D}_{Arow}^{-1/2} \boldsymbol{B} \left[\boldsymbol{D}_{Acol}^{-1/2} - \boldsymbol{D}_{Bcol}^{-1/2}\right].
\end{aligned} \tag{2.39}$$

The first term is just $\boldsymbol{B}_D$, so due to Proposition 12, we should prove only that the norms of both remainder terms are less than $\max\{n^{-\tau}, m^{-\tau}\}$, **AS**. These two terms have a similar appearance, therefore it is enough to estimate one of them. For example, the second term can be bounded by

$$\|\boldsymbol{D}_{Arow}^{-1/2} - \boldsymbol{D}_{Brow}^{-1/2}\| \cdot \|\boldsymbol{B}\| \cdot \|\boldsymbol{D}_{Bcol}^{-1/2}\|. \tag{2.40}$$

The estimation of the first factor in (2.40) is as follows:

$$\begin{aligned} \|\boldsymbol{D}_{Arow}^{-1/2} - \boldsymbol{D}_{Brow}^{-1/2}\| &= \max_{i\in\{1,\dots,m\}} \left(\frac{1}{\sqrt{d_{Arow,i}}} - \frac{1}{\sqrt{d_{Brow,i}}} \right) \\ &= \max_{i\in\{1,\dots,m\}} \frac{|d_{Arow,i} - d_{Brow,i}|}{\sqrt{d_{Arow,i}\cdot d_{Brow,i}}(\sqrt{d_{Arow,i}} + \sqrt{d_{Brow,i}})} \\ &\le \max_{i\in\{1,\dots,m\}} \frac{|d_{Arow,i} - d_{Brow,i}|}{\sqrt{d_{Arow,i}\cdot d_{Brow,i}}} \cdot \max_{i\in\{1,\dots,m\}} \frac{1}{(\sqrt{d_{Arow,i}} + \sqrt{d_{Brow,i}})}. \end{aligned} \tag{2.41}$$

By relations (2.37), $\sqrt{d_{Arow,i}\cdot d_{Brow,i}} = \Theta(n)$ for any $i = 1, \dots, m$, and hence,

$$\frac{|d_{Arow,i} - d_{Brow,i}|}{\sqrt{d_{Arow,i}\cdot d_{Brow,i}}} \le n^{-\tau},$$

AS; further, $\max_{i\in\{1,\dots,m\}} \frac{1}{\sqrt{d_{Arow,i}}+\sqrt{d_{Brow,i}}} = \Theta(\frac{1}{\sqrt{n}})$, **AS**.

Therefore, the left-hand side of (2.41) can be estimated by $n^{-\tau-1/2}$ from above, **AS**. For the further factors in (2.40) we obtain $\|\boldsymbol{B}\| = \Theta(\sqrt{mn})$ (see Proposition 10), while $\|\boldsymbol{D}_{Bcol}^{-1/2}\| = \Theta(\frac{1}{\sqrt{m}})$, **AS**. These together imply that

$$n^{-\tau-1/2} \cdot n^{1/2} m^{1/2} \cdot m^{-1/2} \le n^{-\tau} \le \max\{n^{-\tau}, m^{-\tau}\}.$$

This finishes the estimation of the first term in (2.35), and by the Weyl's perturbation theorem, the proof too. □

We remark the following. If in the definition of the Wigner-noise we used Gaussian distributed entries, the large deviation principle could be replaced by the simple estimation of the Gaussian probabilities with any $\kappa > 0$:

$$\mathbb{P}\left(\left| \frac{1}{n} \sum_{j=1}^{n} w_{ij} \right| > \kappa \right) < \min\left(1, \frac{4\sigma}{\kappa\sqrt{2\pi n}} \exp\left\{ -\frac{n}{2\sigma^2}\kappa^2 \right\} \right).$$

Setting $\kappa = n^{-\tau}$ we get an estimate, analogous to (2.36).

Now, we are ready to estimate the k-variances. Recall that the normalized table of a blown-up matrix is also a blown-up matrix. Therefore, the non-zero singular values of $\boldsymbol{B}_D$ are the numbers $1 = s_0 > s_1 \ge \cdots \ge s_{r-1} > 0$ with unit-norm singular vector pairs $\mathbf{v}_i$, $\mathbf{u}_i$ having piecewise constant structure ($i = 0, \dots, r-1$). Set

$$F := \mathrm{Span}\{\mathbf{v}_0, \dots, \mathbf{v}_{r-1}\} \quad \text{and} \quad G := \mathrm{Span}\{\mathbf{u}_0, \dots, \mathbf{u}_{r-1}\}.$$

Let $0 < \tau < 1/2$ be arbitrary and $\epsilon := \max\{n^{-\tau}, m^{-\tau}\}$. Let us also denote the unit-norm, pairwise orthogonal left- and right-hand side singular vectors corresponding to the singular values $z_0, \dots, z_{r-1} \in [\delta - \epsilon, 1 + \epsilon]$ of $\boldsymbol{A}_D$ – guaranteed by Theorem 21 under **GC2** – by $\mathbf{y}_0, \dots, \mathbf{y}_{r-1} \in \mathbb{R}^m$ and $\mathbf{x}_0, \dots, \mathbf{x}_{r-1} \in \mathbb{R}^n$, respectively.

Proposition 13 ([Bol-Fr-Kr10]) *With the above notation, under **GC1** and **GC2**, the following estimate holds **AS** for the distance between $\mathbf{y}_i$ and F:*

$$\operatorname{dist}(\mathbf{y}_i, F) \leq \frac{\epsilon}{(\delta - \epsilon)} = \frac{1}{(\frac{\delta}{\epsilon} - 1)} \quad i = 0, \ldots, r-1 \tag{2.42}$$

and analogously, for the distance between $\mathbf{x}_i$ and G:

$$\operatorname{dist}(\mathbf{x}_i, G) \leq \frac{\epsilon}{(\delta - \epsilon)} = \frac{1}{(\frac{\delta}{\epsilon} - 1)} \quad i = 0, \ldots, r-1. \tag{2.43}$$

Proof. Follow the method of proving Proposition 11 – under **GC1** – with δ instead of Δ and ϵ instead of ε. Here **GC2** is necessary only for $\boldsymbol{A}_D$ to have r structural singular values. □

Note that the left-hand sides of (2.42) and (2.43) are **AS** of order $\max\{n^{-\tau}, m^{-\tau}\}$, tending to zero as $m, n \to \infty$ under **GC1** and **GC2**.

Proposition 13 implies the well-clustering property of the representatives of the two discrete variables by means of the noisy correspondence vector pairs

$$\mathbf{y}_{corr,i} := \boldsymbol{D}_{Arow}^{-1/2}\mathbf{y}_i, \quad \mathbf{x}_{corr,i} := \boldsymbol{D}_{Acol}^{-1/2}\mathbf{x}_i \quad i = 0, \ldots, r-1.$$

Let $\boldsymbol{Y}_{corr} = (\mathbf{y}_{corr,0}, \ldots, \mathbf{y}_{corr,r-1})$ and $\boldsymbol{X}_{corr} = (\mathbf{x}_{corr,0}, \ldots, \mathbf{x}_{corr,r-1})$ be the matrices containing the optimal representatives of the rows and columns of the noisy contingency table in their columns. Let $\mathbf{y}^1_{corr}, \ldots \mathbf{y}^m_{corr} \in \mathbb{R}^r$ and $\mathbf{x}^1_{corr}, \ldots \mathbf{x}^n_{corr} \in \mathbb{R}^r$ denote the row vectors of $\boldsymbol{Y}_{corr}$ and $\boldsymbol{X}_{corr}$, i.e., the representatives of the genes and conditions, respectively. With respect to the marginal distributions, the weighted a- and b-variances of these representatives are defined (see 1.14) by

$$\tilde{S}^2_a(\boldsymbol{Y}_{corr}) = \min_{\{R'_1, \ldots, R'_a\}} \sum_{i=1}^{a} \sum_{j \in R'_i} d_{Arow,j} \|\mathbf{y}^j_{corr} - \bar{\mathbf{y}}^i_{corr}\|^2$$

and

$$\tilde{S}^2_b(\boldsymbol{X}_{corr}) = \min_{\{C'_1, \ldots, C'_b\}} \sum_{i=1}^{b} \sum_{j \in C'_i} d_{Acol,j} \|\mathbf{x}^j_{corr} - \bar{\mathbf{x}}^i_{corr}\|^2,$$

where $\{R'_1, \ldots, R'_a\}$ and $\{C'_1, \ldots, C'_b\}$ are a- and b-partitions of the genes and conditions, respectively; further,

$$\bar{\mathbf{y}}^i_{corr} = \sum_{j \in A'_i} d_{Arow,j}\mathbf{y}^j_{corr} \quad \text{and} \quad \bar{\mathbf{x}}^i_{corr} = \sum_{j \in B'_i} d_{Acol,j}\mathbf{x}^j_{corr}.$$

Theorem 22 ([Bol-Fr-Kr10]) *With the above notation, under **GC1** and **GC2**,*

$$\tilde{S}^2_a(\boldsymbol{Y}_{corr}) \leq \frac{r}{(\frac{\delta}{\epsilon} - 1)^2} \quad \textit{and} \quad \tilde{S}^2_b(\boldsymbol{X}_{corr}) \leq \frac{r}{(\frac{\delta}{\epsilon} - 1)^2}$$

*hold **AS**, where $\epsilon = \max\{n^{-\tau}, m^{-\tau}\}$ with every $0 < \tau < 1/2$.*

Proof. With the considerations of the Proof of Theorem 20,

$$\tilde{S}^2_a(\boldsymbol{Y}_{corr}) \leq \sum_{i=1}^{a} \sum_{j \in A_i} d_{Arow,j} \|\mathbf{y}^j_{corr} - \bar{\mathbf{y}}^i_{corr}\|^2 = \sum_{i=1}^{r} \operatorname{dist}^2(\mathbf{y}_i, F),$$

and

$$\tilde{S}_b^2(\boldsymbol{X}_{corr}) \leq \sum_{i=1}^{b} \sum_{j \in B_i} d_{Acol,j} \|\mathbf{x}_{corr}^{(j)} - \bar{\mathbf{x}}_{corr}^{i}\|^2 = \sum_{i=1}^{r} \mathrm{dist}^2(\mathbf{x}_i, G),$$

hence the result of Proposition 13 can be used. □

Under **GC1** and **GC2**, with m, n large enough, Theorem 22 implies that after performing correspondence analysis on the $m \times n$ noisy matrix $\boldsymbol{A}$, the representation through the correspondence vectors corresponding to $\boldsymbol{A}_D$ will also reveal the block structure behind $\boldsymbol{A}$.

2.2.4 Finding the blown-up skeleton

One might wonder where the singular values of an $m \times n$ matrix $\boldsymbol{A} = (a_{ij})$ are located if $a := \max_{i,j} |a_{ij}|$ is independent of m and n. On the one hand, the maximum singular value cannot exceed $\mathcal{O}(\sqrt{mn})$, as it is at most $\sqrt{\sum_{i=1}^{m} \sum_{j=1}^{n} a_{ij}^2}$. On the other hand, let $\boldsymbol{Q}$ be an $m \times n$ random matrix with entries a or $-a$ (independently of each other). Consider the spectral norm of all such matrices and take the minimum of them: $\min_{\boldsymbol{Q} \in \{-a,+a\}^{m \times n}} \|\boldsymbol{Q}\|$. This quantity measures the minimum linear structure that a matrix of the same size and magnitude as $\boldsymbol{A}$ can possess. As the Frobenius norm of $\boldsymbol{Q}$ is $a\sqrt{mn}$, in view of the inequalities between the spectral and Frobenius norms, the above minimum is at least $\frac{a}{\sqrt{2}}\sqrt{m+n}$, which is exactly the order of the spectral norm of a Wigner-noise. So an $m \times n$ random matrix with independent and uniformly bounded entries under very general circumstances has at least one singular value of order greater than $\sqrt{m+n}$. Assume there are k such singular values and the representatives by means of the corresponding singular vector pairs can be well classified into k clusters in terms of the k-variances. Under these conditions we can reconstruct a blown-up structure behind our matrix.

Theorem 23 ([Bol-Fr-Kr10]) *Let $\boldsymbol{A}_{m \times n}$ be a sequence of $m \times n$ matrices of uniformly bounded, nonnegative entries, where m and n tend to infinity. Assume that $\boldsymbol{A}_{m \times n}$ has exactly k singular values of order greater than $\sqrt{m+n}$ (k is fixed). If there are integers $a \geq k$ and $b \geq k$ such that the a- and b-variances of the optimal row- and column-representatives are $\mathcal{O}(\frac{m+n}{mn})$, then there is an explicit construction for a blown-up matrix $\boldsymbol{B}_{m \times n}$ (on $a \times b$ blocks) such that $\boldsymbol{A}_{m \times n} = \boldsymbol{B}_{m \times n} + \boldsymbol{E}_{m \times n}$, with $\|\boldsymbol{E}_{m \times n}\| = \mathcal{O}(\sqrt{m+n})$.*

Proof. In the sequel the subscripts m and n will be dropped, for notational convenience. We will speak in terms of microarrays (genes and conditions). Let $\mathbf{y}_1, \ldots, \mathbf{y}_k \in \mathbb{R}^m$ and $\mathbf{x}_1, \ldots, \mathbf{x}_k \in \mathbb{R}^n$ denote the left- and right-hand side unit-norm singular vectors corresponding to $z_1, \ldots, z_k$, the singular values of $\boldsymbol{A}$ of order larger than $\sqrt{m+n}$. The k-dimensional representatives of the genes and conditions – that are row vectors of the $m \times k$ matrix $\boldsymbol{Y} = (\mathbf{y}_1, \ldots, \mathbf{y}_k)$ and those of the $n \times k$ matrix $\boldsymbol{X} = (\mathbf{x}_1, \ldots, \mathbf{x}_k)$, respectively – by the assumption of the theorem form a and b clusters in $\mathbb{R}^k$, respectively, with sum of inner variances $\mathcal{O}(\frac{m+n}{mn})$. Reorder the rows and columns of $\boldsymbol{A}$ according to their respective cluster memberships. Denote by $\mathbf{y}^1, \ldots, \mathbf{y}^m \in \mathbb{R}^k$ and $\mathbf{x}^1, \ldots, \mathbf{x}^n \in \mathbb{R}^k$ the Euclidean representatives of the genes and conditions (the rows of the reordered $\boldsymbol{Y}$ and $\boldsymbol{X}$), and let $\bar{\mathbf{y}}^1, \ldots, \bar{\mathbf{y}}^a \in \mathbb{R}^k$ and $\bar{\mathbf{x}}^1, \ldots, \bar{\mathbf{x}}^b \in \mathbb{R}^k$ denote the cluster centers, respectively. Now let us choose the following new representation of the genes and conditions. The genes' representatives are row vectors of the $m \times k$ matrix $\widetilde{\boldsymbol{Y}}$ such that the first m_1 rows of $\widetilde{\boldsymbol{Y}}$ are equal to $\bar{\mathbf{y}}^1$, the next m_2 rows to $\bar{\mathbf{y}}^2$, and so on ... the last m_a rows of $\widetilde{\boldsymbol{Y}}$ are equal to $\bar{\mathbf{y}}^a$. Then likewise, the conditions' representatives are row vectors of the $n \times k$ matrix $\widetilde{\boldsymbol{X}}$ such that the first n_1 rows of $\widetilde{\boldsymbol{X}}$ are equal to $\bar{\mathbf{x}}^1$, the next n_2 rows to $\bar{\mathbf{x}}^2$, and so on ... the last n_b rows of $\widetilde{\boldsymbol{X}}$ are equal to $\bar{\mathbf{x}}^b$.

By the considerations of Theorem 20 and the assumption for the clusters,

$$\sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{y}_i, F) = S_a^2(\boldsymbol{Y}) = \mathcal{O}(\frac{m+n}{mn}) \tag{2.44}$$

and

$$\sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{x}_i, G) = S_b^2(\boldsymbol{X}) = \mathcal{O}(\frac{m+n}{mn}) \tag{2.45}$$

hold respectively, where the k-dimensional subspace $F \subset \mathbb{R}^m$ is spanned by the column vectors of $\widetilde{\boldsymbol{Y}}$, and the k-dimensional subspace $G \subset \mathbb{R}^n$ is spanned by the column vectors of $\widetilde{\boldsymbol{X}}$. We follow the construction given in Lemma 3 of a set $\mathbf{v}_1, \ldots, \mathbf{v}_k$ of orthonormal vectors within F and another set $\mathbf{u}_1, \ldots, \mathbf{u}_k$ of orthonormal vectors within G such that

$$\sum_{i=1}^{k} \|\mathbf{y}_i - \mathbf{v}_i\|^2 = \min_{\mathbf{v}'_1,\ldots,\mathbf{v}'_k} \sum_{i=1}^{k} \|\mathbf{y}_i - \mathbf{v}'_i\|^2 \leq 2 \sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{y}_i, F) \tag{2.46}$$

and

$$\sum_{i=1}^{k} \|\mathbf{x}_i - \mathbf{u}_i\|^2 = \min_{\mathbf{u}'_1,\ldots,\mathbf{u}'_k} \sum_{i=1}^{k} \|\mathbf{x}_i - \mathbf{u}'_i\|^2 \leq 2 \sum_{i=1}^{k} \mathrm{dist}^2(\mathbf{x}_i, G) \tag{2.47}$$

hold, where the minimum is taken over orthonormal sets of vectors $\mathbf{v}'_1, \ldots, \mathbf{v}'_k \in F$ and $\mathbf{u}'_1, \ldots, \mathbf{u}'_k \in G$, respectively. The construction of the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ is as follows ($\mathbf{u}_1, \ldots, \mathbf{u}_k$ can be constructed in the same way). Let $\mathbf{v}'_1, \ldots, \mathbf{v}'_k \in F$ be an arbitrary orthonormal system (obtained, e.g., by the Schmidt orthogonalization method; note that in the Lemma 3 they were given at the beginning). Let $\boldsymbol{V}' = (\mathbf{v}'_1, \ldots, \mathbf{v}'_k)$ be an $m \times k$ matrix and

$$\boldsymbol{Y}^T \boldsymbol{V}' = \boldsymbol{Q} \boldsymbol{S} \mathbf{Z}^T$$

be SVD, where the matrix $\boldsymbol{S}$ contains the singular values of the $k \times k$ matrix $\boldsymbol{Y}^T \boldsymbol{V}'$ in its main diagonal and zeros otherwise, while $\boldsymbol{Q}$ and $\mathbf{Z}$ are $k \times k$ orthogonal matrices (containing the corresponding unit-norm singular vector pairs in their columns). The orthogonal matrix $\boldsymbol{R} = \mathbf{Z}\boldsymbol{Q}^T$ will give the convenient orthogonal rotation of the vectors $\mathbf{v}'_1, \ldots, \mathbf{v}'_k$. That is, the column vectors of the matrix $\boldsymbol{V} = \boldsymbol{V}'\boldsymbol{R}$ form also an orthonormal set that is the desired set $\mathbf{v}_1, \ldots, \mathbf{v}_k$. Define the error terms $\mathbf{r}_i$ and $\mathbf{q}_i$, respectively:

$$\mathbf{r}_i = \mathbf{y}_i - \mathbf{v}_i \quad \text{and} \quad \mathbf{q}_i = \mathbf{x}_i - \mathbf{u}_i \qquad (i = 1, \ldots, k).$$

In view of (2.44) – (2.47),

$$\sum_{i=1}^{k} \|\mathbf{r}_i\|^2 = \mathcal{O}(\frac{m+n}{mn}) \quad \text{and} \quad \sum_{i=1}^{k} \|\mathbf{q}_i\|^2 = \mathcal{O}(\frac{m+n}{mn}). \tag{2.48}$$

Consider the following decomposition:

$$\boldsymbol{A} = \sum_{i=1}^{k} z_i \mathbf{y}_i \mathbf{x}_i^T + \sum_{i=k+1}^{\min\{m,n\}} z_i \mathbf{y}_i \mathbf{x}_i^T .$$

The spectral norm of the second term is at most of order $\sqrt{m+n}$. Now consider the first term,

$$\begin{aligned}\sum_{i=1}^{k} z_i \mathbf{y}_i \mathbf{x}_i^T &= \sum_{i=1}^{k} z_i (\mathbf{v}_i + \mathbf{r}_i)(\mathbf{u}_i^T + \mathbf{q}_i^T) = \\ &= \sum_{i=1}^{k} z_i \mathbf{v}_i \mathbf{u}_i^T + \sum_{i=1}^{k} z_i \mathbf{v}_i \mathbf{q}_i^T + \sum_{i=1}^{k} z_i \mathbf{r}_i \mathbf{u}_i^T + \sum_{i=1}^{k} z_i \mathbf{r}_i \mathbf{q}_i^T .\end{aligned} \tag{2.49}$$

Since $\mathbf{v}_1, \ldots, \mathbf{v}_k$ and $\mathbf{u}_1, \ldots, \mathbf{u}_k$ are unit vectors, the last three terms in (2.49) can be estimated by means of the relations

$$\begin{aligned}\|\mathbf{v}_i \mathbf{u}_i^T\| &= \sqrt{\|\mathbf{v}_i \mathbf{u}_i^T \mathbf{u}_i \mathbf{v}_i^T\|} = 1 \qquad (i = 1, \ldots, k), \\ \|\mathbf{v}_i \mathbf{q}_i^T\| &= \sqrt{\|\mathbf{q}_i \mathbf{v}_i^T \mathbf{v}_i \mathbf{q}_i^T\|} = \|\mathbf{q}_i\| \qquad (i = 1, \ldots, k), \\ \|\mathbf{r}_i \mathbf{u}_i^T\| &= \sqrt{\|\mathbf{r}_i \mathbf{u}_i^T \mathbf{u}_i \mathbf{r}_i^T\|} = \|\mathbf{r}_i\| \qquad (i = 1, \ldots, k), \\ \|\mathbf{r}_i \mathbf{q}_i^T\| &= \sqrt{\|\mathbf{r}_i \mathbf{q}_i^T \mathbf{q}_i \mathbf{r}_i^T\|} = \|\mathbf{q}_i\| \cdot \|\mathbf{r}_i\| \qquad (i = 1, \ldots, k).\end{aligned}$$

Taking into consideration that z_i cannot exceed $\Theta(\sqrt{mn})$, while k is fixed and, due to (2.48), we get that the spectral norms of the last three terms in (2.49) – for their finitely many subterms the triangle inequality is applicable – are at most of order $\sqrt{m+n}$. Let $\boldsymbol{B}$ be the first term, i.e.,

$$\boldsymbol{B} = \sum_{i=1}^{k} z_i \mathbf{v}_i \; u_i^T .$$

Then $\|\boldsymbol{A} - \boldsymbol{B}\| = \mathcal{O}(\sqrt{m+n})$.

By their definition, the vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ and the vectors $\mathbf{u}_1, \ldots, \mathbf{u}_k$ are in the subspaces F and G, respectively. Both spaces consist of step-vectors; thus the matrix $\boldsymbol{B}$ is a blown-up matrix containing $a \times b$ blocks. The noise matrix is

$$\boldsymbol{E} = \sum_{i=1}^{k} z_i \mathbf{v}_i \mathbf{q}_i^T + \sum_{i=1}^{k} z_i \mathbf{r}_i \mathbf{u}_i^T + \sum_{i=1}^{k} z_i \mathbf{r}_i \mathbf{q}_j^T + \sum_{i=k+1}^{\min\{m,n\}} z_i \mathbf{y}_i \mathbf{x}_i^T$$

which finishes the proof. □

Then, provided the conditions of Theorem 23 hold, by the construction given in the proof above, an algorithm can be written that uses several SVD's and produces the blown-up matrix $\boldsymbol{B}$. This $\boldsymbol{B}$ can be considered as the best blown-up approximation of the microarray $\boldsymbol{A}$. At the same time, clusters of the genes and conditions are also obtained. More precisely, first we conclude the clusters from the SVD of $\boldsymbol{A}$, rearrange the rows and columns of $\boldsymbol{A}$ accordingly, and afterwards we use the above construction. If we decide to perform correspondence analysis on $\boldsymbol{A}$, then by (2.35) and (2.39), $\boldsymbol{B}_D$ will give a good approximation to $\boldsymbol{A}_D$ and likewise, the correspondence vectors obtained by the SVD of $\boldsymbol{A}_D$ will give representatives of the genes and conditions.

Clustering microarray data via the k-means algorithm is also discussed in [Dhil], but with an other objective function. To find the SVD for large rectangular matrices, randomized algorithms are favored, e.g., [Ac-Mc]. In case of random matrices with an underlying linear structure (outstanding singular values), the random noise of the algorithm is just added to the noise in our data, but their sum is also a Wigner-noise, so it does not change the effect of our algorithm in finding the clusters. Under the conditions of Theorem 23, the

separated error matrix is comparable with the noise matrix, and this fact guarantees that the underlying block structure can be extracted.

2.3 Discrepancy and spectra

Here under clustering we understand partition of the vertex-set into subsets of similar vertices, i.e., members of a cluster behave similarity toward members of each (other or own) cluster. We will generalize the Laplacian and modularity based spectral clustering methods to recover so-called regular cluster pairs such that the information flow between the pairs and within the clusters is as homogeneous as possible. The notion of volume-regularity is also extended to contingency tables. For this purpose, we take into account both ends of the normalized Laplacian spectrum, i.e., large absolute value, so-called structural eigenvalues of the normalized modularity matrix, or the largest singular values of the normalized contingency table. First we introduce the notion of multiway discrepancy for rectangular arrays of nonnegative entries, of which the quadratic edge-weight matrices are special cases.

2.3.1 Estimating the singular values of normalized contingency tables by the multiway discrepancy

Definition 25 *The multiway discrepancy of the rectangular array $\boldsymbol{C}$ of nonnegative entries in the proper k-partition $R_1, \ldots, R_k$ of its rows and $C_1, \ldots, C_k$ of its columns is*

$$\mathrm{md}(\boldsymbol{C}; R_1, \ldots, R_k, C_1, \ldots, C_k) = \max_{\substack{1 \le a, b \le k \\ X \subset R_a,\, Y \subset C_b}} \frac{|c(X,Y) - \rho(R_a, C_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)|}{\sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)}}, \quad (2.50)$$

where $c(X,Y)$, $\mathrm{Vol}(X)$, and $\mathrm{Vol}(Y)$ are defined in Section 1.2.2, whereas $\rho(R_a, C_b) = \frac{c(R_a, C_b)}{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}$ denotes the relative density between R_a and C_b. The minimum k-way discrepancy of $\boldsymbol{C}$ itself is

$$\mathrm{md}_k(\boldsymbol{C}) = \min_{\substack{(R_1, \ldots, R_k) \\ (C_1, \ldots, C_k)}} \mathrm{md}(\boldsymbol{C}; R_1, \ldots, R_k, C_1, \ldots, C_k).$$

We will also extend this notion to an edge-weighted graph G and denote it by $\mathrm{md}_k(G)$. In that setup, $\boldsymbol{C}$ plays the role of the edge-weight matrix: symmetric in the undirected; quadratic, but usually not symmetric in the directed case; and it is the adjacency matrix if G is a simple graph.

Note that the division by $\sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)}$ ensures that the multiway discrepancy is not affected by the scaling of the entries of $\boldsymbol{C}$, akin to the normalized table $\boldsymbol{C}_D$, introduced in Section 1.2. Therefore, without loss of generality, $\sum_{i=1}^{n}\sum_{j=1}^{m} c_{ij} = 1$ will be assumed.

Observe that $\mathrm{md}(\boldsymbol{C}; R_1, \ldots, R_k, C_1, \ldots, C_k)$ is the smallest α such that for every R_a, C_b pair and for every $X \subset R_a$, $Y \subset C_b$,

$$|c(X,Y) - \rho(R_a, C_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)| \le \alpha\sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)} \quad (2.51)$$

holds. Therefore, in the k-partitions of the rows and columns, giving the minimum k-way discrepancy (say, α^*) of $\boldsymbol{C}$, every R_a, C_b pair is α^*-regular in terms of the volumes, and α^* is the smallest possible discrepancy that can be attained with proper k-partitions. It resembles the notion of ϵ-regular pairs in the Szemerédi Regularity Lemma [Szem], albeit with given number of vertex-clusters, which are usually not equitable; further, with volumes, instead of cardinalities.

Historically, the notion of discrepancy together with the expander mixing lemma was introduced for simple, regular graphs, see e.g., Alon, Spencer, Hoory, Linial, Widgerson [Al-Sp, Ho-Lin-Wid], and extended to Hermitian matrices by Bollobás, Nikiforov [Bo-Nik]. In Chung, Graham, Wilson [Chu-G-W], the authors use the term quasirandom for simple graphs that satisfy any of some equivalent properties, some of them closely related to discrepancy and eigenvalue separation. Chung and Graham [Chu-G] prove that for simple graphs 'small' discrepancy $\mathrm{disc}(G)$ (with our notation, $\mathrm{md}_1(G)$) is caused by eigenvalue 'separation': the second largest singular value (which is also the second largest absolute value eigenvalue), s_1, of the normalized adjacency matrix is 'small', i.e., separated from the trivial singular value $s_0 = 1$, which is the edge of the spectrum. More exactly, they prove $\mathrm{disc}(G) \leq s_1$, hence giving some kind of generalization of the *expander mixing lemma for irregular graphs*.

In the other direction, for Hermitian matrices, Bollobás and Nikiforov [Bo-Nik] estimate the second largest singular value of an $n \times n$ Hermitian matrix $\boldsymbol{A}$ by $C\mathrm{disc}(\boldsymbol{A}) \log n$ (where C is an absolute constant), and show that this is best possible up to a multiplicative constant. Bilu and Linial [Bil-Lin] prove the converse of the expander mixing lemma for simple regular graphs, but their key Lemma 3.3, producing this statement, goes beyond regular graphs. In Alon et al. [Aletal], the authors relax the notion of eigenvalue separation to essential eigenvalue separation (by introducing a parameter for it, and requiring the separation only for the eigenvalues of a relatively large part of the graph). Then they prove relations between the constants of this kind of eigenvalue separation and the discrepancy.

For a general rectangular array $\boldsymbol{C}$ of nonnegative entries, Butler [But] proves the following forward and backward statement in the $k = 1$ case:

$$\mathrm{disc}(\boldsymbol{C}) \leq s_1 \leq 150\mathrm{disc}(\boldsymbol{C})(1 - 8 \ln \mathrm{disc}(\boldsymbol{C})), \tag{2.52}$$

where his $\mathrm{disc}(\boldsymbol{C})$ is our $\mathrm{md}_1(\boldsymbol{C})$ and, with our notation, s_1 is the largest nontrivial singular value of $\boldsymbol{C}_D$ (he denotes is with σ_2). Since $s_1 < 1$, the upper estimate makes sense for very small discrepancy, in particular, for $\mathrm{disc}(\boldsymbol{C}) \leq 8.868 \times 10^{-5}$. The lower estimate of (2.52) further generalizes the expander mixing lemma to rectangular matrices, but it can be proved with the same tools as in the quadratic case (see the forthcoming Proposition 14 in Section 2.3.3).

The above papers consider the overall discrepancy in the sense that $\mathrm{disc}(\boldsymbol{C})$ or $\mathrm{disc}(G)$ measure the largest possible deviation between the actual and expected connectedness of arbitrary (sometimes disjoint) subsets X, Y, where under expected the hypothesis of independence is understood (which corresponds to the rank 1 approximation of the normalized matrix). Our purpose is, in the multicluster scenario, to find similar relations between the minimum k-way discrepancy and the SVD of the normalized matrix, for given k. In one direction, we are able to prove the following.

Theorem 24 ([Bol16]) *For every non-degenerate real matrix $\boldsymbol{C}$ of nonnegative entries and integer $1 \leq k \leq \mathrm{rank}(\boldsymbol{C})$,*

$$s_k \leq 9\mathrm{md}_k(\boldsymbol{C})(k + 2 - 9k \ln \mathrm{md}_k(\boldsymbol{C})) \tag{2.53}$$

holds, provided $0 < \mathrm{md}_k(\boldsymbol{C}) < 1$, where s_k is the k-th largest nontrivial singular value of the normalized matrix $\boldsymbol{C}_D$ of $\boldsymbol{C}$, defined in (1.23).

Note that $\mathrm{md}_k(\boldsymbol{C}) = 0$ if $\boldsymbol{C}$ has a block structure with k row- and column-blocks, in which case $s_k = 0$ also holds. Likewise, $\mathrm{md}_k(\boldsymbol{C}) < 1$ is not a peculiar requirement, since in view of $s_k < 1$, the upper bound of the theorem has relevance only for $\mathrm{md}_k(\boldsymbol{C})$ much smaller than 1; for example, for $\mathrm{md}_1(\boldsymbol{C}) \leq 1.866 \times 10^{-3}$, $\mathrm{md}_2(\boldsymbol{C}) \leq 8.459 \times 10^{-4}$, $\mathrm{md}_3(\boldsymbol{C}) \leq 5.329 \times 10^{-4}$, etc.

Before proving the theorem, we encounter some lemmas of other authors that will be used, possibly with some modifications.

Lemma 3 of Bollobás and Nikiforov [Bo-Nik] is the key to prove their main result. This lemma states that to every $0 < \varepsilon < 1$ and vector $\mathbf{x} \in \mathbb{C}^n$, $\|\mathbf{x}\| = 1$, there exists a vector $\mathbf{y} \in \mathbb{C}^n$ such that its coordinates take no more than $\lceil \frac{8\pi}{\varepsilon} \rceil \lceil \frac{4}{\varepsilon} \log \frac{2n}{\varepsilon} \rceil$ distinct values and $\|\mathbf{x} - \mathbf{y}\| \le \varepsilon$. We will rather use the construction of the following lemma, which is indeed a consequence of Lemma 3 of [Bo-Nik].

Lemma 4 (Lemma 3 of Butler [But]) *To any vector $\mathbf{x} \in \mathbb{C}^n$, $\|\mathbf{x}\| = 1$ and diagonal matrix $\boldsymbol{D}$ of positive real diagonal entries, one can construct a step-vector $\mathbf{y} \in \mathbb{C}^n$ such that $\|\mathbf{x} - \boldsymbol{D}\mathbf{y}\| \le \frac{1}{3}$, $\|\boldsymbol{D}\mathbf{y}\| \le 1$, and the nonzero entries of $\mathbf{y}$ are of the form $\left(\frac{4}{5}\right)^j e^{\frac{\ell}{29} 2\pi i}$ with appropriate integers j (taking on $\mathcal{O}(\log n)$ distinct values) and ℓ ($0 \le \ell \le 28$).*

Note that starting with an $\mathbf{x}$ of real coordinates, we do not need all the 29 values of ℓ, only two of them will show up, as it follows from a better understanding of the construction of [But]. In fact, by the idea of [Bo-Nik], j's come from dividing the coordinates of $\boldsymbol{D}^{-1}\mathbf{x}/\|\boldsymbol{D}^{-1}\mathbf{x}\|$ in decreasing absolute values into groups, where the cut-points are powers of $\frac{4}{5}$. With the notation $\mathbf{x} = (x_s))_{s=1}^n$, if x_s is in the j-th group, then the corresponding coordinate of the approximating complex vector $\mathbf{y} = (y_s)_{s=1}^n$ is as follows. If $x_s = 0$, then $y_s = 0$, otherwise $y_s = \left(\frac{4}{5}\right)^j e^{\left(\lfloor \frac{29\theta}{2\pi} \rfloor / 29\right) 2\pi i}$, where θ is the argument of x_s, $0 \le \theta < 2\pi$, and therefore, $\ell = \lfloor \frac{29\theta}{2\pi} \rfloor$ is an integer between 0 and 28. However, when the coordinates of $\mathbf{x}$ are real numbers, then only the values 0 and 14 of ℓ can occur, since θ can take only one of the values 0 or π, depending on whether x_s is positive or negative. We will intensively use this observation in our proof.

Lemma 5 (Lemma 4 of Butler [But]) *Let $\boldsymbol{M}$ be a matrix with largest singular value σ and corresponding unit-norm singular vector pair $\mathbf{v}, \mathbf{u}$. If $\mathbf{x}$ and $\mathbf{y}$ are vectors such that $\|\mathbf{x}\| \le 1$, $\|\mathbf{y}\| \le 1$, $\|\mathbf{v} - \mathbf{x}\| \le \frac{1}{3}$, $\|\mathbf{u} - \mathbf{y}\| \le \frac{1}{3}$, then $\sigma \le \frac{9}{2}\langle \mathbf{x}, \boldsymbol{M}\mathbf{y} \rangle$.*

Lemma 6 (Theorem 3 of Thompson [Thomp]) *Let the $n \times n$ matrix $\boldsymbol{A}$ have singular values $\alpha_1 \ge \cdots \ge \alpha_n$ and $1 \le k \le n$ be a fixed integer. Then an $n \times n$ matrix $\mathbf{X}$ exists with $\mathrm{rank}(\mathbf{X}) \le k$ such that $\boldsymbol{B} = \boldsymbol{A} + \mathbf{X}$ has singular values $\beta_1 \ge \cdots \ge \beta_n$ if and only if*

$$\alpha_{i+k} \le \beta_i \le \alpha_{i-k}, \qquad i = 1, \ldots, n$$

with the understanding that $\alpha_j = +\infty$ if $j \le 0$ and $\alpha_j = 0$ if $j \ge n$.

Proof of Theorem 24. Assume that $\alpha^* = \mathrm{md}_k(\boldsymbol{C}) \in (0, 1)$ and it is attained with the proper k-partition $R_1, \ldots, R_k$ of the rows and $C_1, \ldots, C_k$ of the columns of $\boldsymbol{C}$; i.e., for every R_a, C_b pair and $X \subset R_a$, $Y \subset C_b$ we have

$$|c(X, Y) - \rho(R_a, C_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)| \le \alpha^* \sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)}. \tag{2.54}$$

Our purpose is to put Inequality (2.54) in a matrix form by using indicator vectors and introducing the $m \times n$ auxiliary matrix

$$\boldsymbol{F} = \boldsymbol{C} - \boldsymbol{D}_{row}\boldsymbol{R}\boldsymbol{D}_{col}, \tag{2.55}$$

where $\boldsymbol{R} = (\rho(R_a, C_b))$ is the $m \times n$ block-matrix of $k \times k$ blocks with entries equal to $\rho(R_a, C_b)$ over the block $R_a \times C_b$. With the indicator vectors $\mathbf{1}_X$ and $\mathbf{1}_Y$ of $X \subset R_a$ and $Y \subset C_b$, Inequality (2.54) has the following equivalent form:

$$|\langle \mathbf{1}_X, \boldsymbol{F}\mathbf{1}_Y \rangle| \le \alpha^* \sqrt{\langle \mathbf{1}_X, \boldsymbol{C}\mathbf{1}_n \rangle \langle \mathbf{1}_m, \boldsymbol{C}\mathbf{1}_Y \rangle}, \tag{2.56}$$

where $\mathbf{1}_n$ denotes the all 1's vector of size n. At the same time, Equation (2.55) yields

$$\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2} = \boldsymbol{D}_{row}^{-1/2}\boldsymbol{C}\boldsymbol{D}_{col}^{-1/2} - \boldsymbol{D}_{row}^{1/2}\boldsymbol{R}\boldsymbol{D}_{col}^{1/2} = \boldsymbol{C}_D - \boldsymbol{D}_{row}^{1/2}\boldsymbol{R}\boldsymbol{D}_{col}^{1/2}.$$

Since the rank of the matrix $\boldsymbol{D}_{row}^{1/2}\boldsymbol{R}\boldsymbol{D}_{col}^{1/2}$ is at most k, by the upper estimate of Lemma 6 (with the rolecast $\boldsymbol{A} = \boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}$, $\boldsymbol{B} = \boldsymbol{C}_D$, $\mathbf{X} = \boldsymbol{D}_{row}^{1/2}\boldsymbol{R}\boldsymbol{D}_{col}^{1/2}$, and $i = k+1$) [1] we obtain the following upper estimate for s_k, that is the $(k+1)$-th largest (including the trivial 1) singular value of $\boldsymbol{C}_D$:

$$s_k \le s_{max}(\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}) = \|\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}\|,$$

where $\|.\|$ denotes the spectral norm.

Let $\mathbf{v} \in \mathbb{R}^m$ be the left and $\mathbf{u} \in \mathbb{R}^n$ be the right unit-norm singular vector corresponding to the maximal singular value of $\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}$, i.e.,

$$|\langle \mathbf{v}, (\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2})\mathbf{u}\rangle| = \|\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}\|.$$

In view of Lemma 4, there are step-vectors $\mathbf{x} \in \mathbb{C}^m$ and $\mathbf{y} \in \mathbb{C}^n$ such that $\|\mathbf{v} - \boldsymbol{D}_{row}^{1/2}\mathbf{x}\| \le \frac{1}{3}$ and $\|\mathbf{u} - \boldsymbol{D}_{col}^{1/2}\mathbf{y}\| \le \frac{1}{3}$; further, $\|\boldsymbol{D}_{row}^{1/2}\mathbf{x}\| \le 1$ and $\|\boldsymbol{D}_{col}^{1/2}\mathbf{y}\| \le 1$. Then Lemma 5 yields

$$\|\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2}\| \le \frac{9}{2}\left|\langle(\boldsymbol{D}_{row}^{1/2}\mathbf{x}), (\boldsymbol{D}_{row}^{-1/2}\boldsymbol{F}\boldsymbol{D}_{col}^{-1/2})(\boldsymbol{D}_{col}^{1/2}\mathbf{y})\rangle\right| = \frac{9}{2}|\langle \mathbf{x}, \boldsymbol{F}\mathbf{y}\rangle|.$$

Now we will use the construction of the proof of the Lemma 4 in the special case when the vectors $\mathbf{v} = (v_s))_{s=1}^m$ and $\mathbf{u} = (u_s))_{s=1}^n$, to be approximated, have real coordinates. Therefore, only the following three types of coordinates of the approximating complex vectors $\mathbf{x} = (x_s))_{s=1}^m$ and $\mathbf{y} = (y_s)_{s=1}^n$ will appear. If $v_s = 0$, then $x_s = 0$; if $v_s > 0$, then $x_s = (\frac{4}{5})^j$ with some integer j; if $v_s < 0$, then $x_s = (\frac{4}{5})^j e^{\frac{28}{29}\pi i}$ with some integer j. Likewise, if $u_s = 0$, then $y_s = 0$; if $u_s > 0$, then $y_s = (\frac{4}{5})^\ell$ with some integer ℓ; if $u_s < 0$, then $y_s = (\frac{4}{5})^\ell e^{\frac{28}{29}\pi i}$ with some integer ℓ. With these observations, the step-vectors $\mathbf{x}$ and $\mathbf{y}$ can be written as the following finite sums with respect to the integers j and ℓ:

$$\mathbf{x} = \sum_j (\frac{4}{5})^j \mathbf{x}^{(j)}, \quad \mathbf{x}^{(j)} = \sum_{a=1}^{k} (\mathbf{1}_{\mathcal{X}_{ja1}} + e^{\frac{28}{29}\pi i}\mathbf{1}_{\mathcal{X}_{ja2}}), \quad \text{where}$$

$$\mathcal{X}_{ja1} = \{s:\ v_s > 0,\ s \in R_a\} \quad \text{and} \quad \mathcal{X}_{ja2} = \{s:\ v_s < 0,\ s \in R_a\};$$

likewise,

$$\mathbf{y} = \sum_\ell (\frac{4}{5})^\ell \mathbf{y}^{(\ell)}, \quad \mathbf{y}^{(\ell)} = \sum_{b=1}^{k} (\mathbf{1}_{\mathcal{Y}_{\ell b1}} + e^{\frac{28}{29}\pi i}\mathbf{1}_{\mathcal{Y}_{\ell b2}}), \quad \text{where}$$

$$\mathcal{Y}_{\ell b1} = \{s:\ u_s > 0,\ s \in C_b\} \quad \text{and} \quad \mathcal{Y}_{\ell b2} = \{s:\ u_s < 0,\ s \in C_b\}.$$

It is important that the $2k$ indicator vectors appearing in the decomposition of any $\mathbf{x}^{(j)}$ or $\mathbf{y}^{(\ell)}$ are disjointly supported, and so, all the coordinates of these vectors are of absolute value

[1] Actually, Lemma 6 is about square matrices, but in the possession of a rectangular one, we can supplement it with zero rows or columns to make it quadratic; further, the nonzero singular values of the so obtained square matrix are the same as those of the rectangular one, supplemented with additional zero singular values that will not alter the shifted interlacing facts.

1. These considerations give rise to the following estimation.

$$\begin{aligned}
|\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| &\le \sum_{a=1}^{k}\sum_{p=1}^{2}\sum_{b=1}^{k}\sum_{q=1}^{2} \left|\langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{F}\mathbf{1}_{\mathcal{Y}_{\ell bq}}\rangle\right| \\
&\overset{(2.56)}{\le} \sum_{a=1}^{k}\sum_{p=1}^{2}\sum_{b=1}^{k}\sum_{q=1}^{2} \alpha^* \sqrt{\langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{C}\mathbf{1}_n\rangle\langle \mathbf{1}_m, \boldsymbol{C}\mathbf{1}_{\mathcal{Y}_{\ell bq}}\rangle} \\
&\le \alpha^* 2k \sqrt{\sum_{a=1}^{k}\sum_{p=1}^{2}\sum_{b=1}^{k}\sum_{q=1}^{2} \langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{C}\mathbf{1}_n\rangle\langle \mathbf{1}_m, \boldsymbol{C}\mathbf{1}_{\mathcal{Y}_{\ell bq}}\rangle} \\
&= 2k\alpha^* \sqrt{\langle \sum_{a=1}^{k}\sum_{p=1}^{2} \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{C}\mathbf{1}_n\rangle\langle \mathbf{1}_m, \boldsymbol{C}\sum_{b=1}^{k}\sum_{q=1}^{2} \mathbf{1}_{\mathcal{Y}_{\ell bq}}\rangle} \\
&= 2k\alpha^* \sqrt{\langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle\langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle},
\end{aligned} \tag{2.57}$$

where in the first inequality we used the triangle inequality and $|e^{\frac{28}{29}\pi i}| = 1$, in the second one we used (2.56), while in the third one, the Cauchy–Schwarz inequality with $4k^2$ terms.

In the last step we exploited that the indicator vectors composing $\mathbf{x}^{(j)}$ and $\mathbf{y}^{(\ell)}$ are disjointly supported. We also introduced the notation $|\mathbf{z}| = (|z_s|)_{s=1}^n$ for the real vector, the coordinates of which are the absolute values of the corresponding coordinates of the (possibly complex) vector $\mathbf{z}$. (Note that the so introduced $|\mathbf{z}|$ is a vector, unlike $\|\mathbf{z}\| = (\sum_{s=1}^n |z_s|^2)^{1/2}$.) In the same spirit, let $|\boldsymbol{M}|$ denote the matrix whose entries are the absolute values of the corresponding entries of $\boldsymbol{M}$ (we will use this only for real matrices). With this formalism, this is the right moment to prove the following inequalities that will be used soon to finish the proof:

$$\sum_{\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| \le 2\langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle, \quad \sum_{j} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| \le 2\langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle. \tag{2.58}$$

Since the two inequalities are of the same flavor, it suffices to prove only the first one. Note that it is here, where we use the exact definition of $\boldsymbol{F}$ as follows.

$$\begin{aligned}
\sum_{\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| &\le \langle |\mathbf{x}^{(j)}|, |\boldsymbol{F}| \sum_{\ell} |\mathbf{y}^{(\ell)}|\rangle \\
&\le \langle |\mathbf{x}^{(j)}|, (\boldsymbol{C} + \boldsymbol{D}_{row}\boldsymbol{R}\boldsymbol{D}_{col})\mathbf{1}_n\rangle| = 2\langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle
\end{aligned}$$

because $|\mathbf{y}^{(\ell)}|$ is a 0-1 vector and $\boldsymbol{C} + \boldsymbol{D}_{row}\boldsymbol{R}\boldsymbol{D}_{col}$ is a (real) matrix of nonnegative entries. We also used that the i-th coordinate of the vector $(\boldsymbol{C} + \boldsymbol{D}_{row}\boldsymbol{R}\boldsymbol{D}_{col})\mathbf{1}_n$ for $i \in R_a$ is

$$d_{row,i}\left(1 + \sum_{b=1}^{k} \rho(R_a, C_b)\mathrm{Vol}(C_b)\right) = 2d_{row,i}$$

(here we utilized that the sum of the entries of $\boldsymbol{C}$ is 1), and therefore,

$$(\boldsymbol{C} + \boldsymbol{D}_{row}\boldsymbol{R}\boldsymbol{D}_{col})\mathbf{1}_n = 2\boldsymbol{C}\mathbf{1}_n.$$

Finally, we will finish the proof with similar considerations as in [But]. Let us further estimate

$$\langle \mathbf{x}, \boldsymbol{F}\mathbf{y}\rangle = \sum_{j}\sum_{\ell} \langle (\tfrac{4}{5})^j \mathbf{x}^{(j)}, \boldsymbol{F}(\tfrac{4}{5})^{\ell}\mathbf{y}^{(\ell)}\rangle.$$

Put $\gamma := \log_{4/5} \alpha^*$; in view of $\alpha^* < 1$, $\gamma > 0$ holds. Then we divide the above summation into three parts as follows.

$$|\langle \mathbf{x}, \boldsymbol{F}\mathbf{y}\rangle| \le \sum_j \sum_\ell (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle|$$

$$= \underbrace{\sum_{|j-\ell|\le\gamma} (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle|}_{\text{(a)}} + \underbrace{\sum_{j-\ell>\gamma} (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle|}_{\text{(b)}} + \underbrace{\sum_{j-\ell<-\gamma} (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle|}_{\text{(c)}}.$$

The three terms are estimated separately. Term (a) can be bounded from above as follows:

$$\begin{aligned}
\sum_{|j-\ell|\le\gamma} (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| &\overset{(2.57)}{\le} 2k\alpha^* \sum_{|j-\ell|\le\gamma} \sqrt{(\frac{4}{5})^{2j} \langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle (\frac{4}{5})^{2\ell} \langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle} \\
&\overset{(*)}{\le} k\alpha^* \sum_{|j-\ell|\le\gamma} \left[(\frac{4}{5})^{2j} \langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle + (\frac{4}{5})^{2\ell} \langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle \right] \\
&\overset{(**)}{\le} k\alpha^* (2\gamma+1) \left[\sum_j (\frac{4}{5})^{2j} \langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle + \sum_\ell (\frac{4}{5})^{2\ell} \langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle \right], \\
&\overset{(***)}{\le} 2k\alpha^* (2\gamma+1),
\end{aligned}$$

where in the first inequality, the estimate of (2.57), and in (*), the geometric-arithmetic mean inequality were used; (**) comes from the fact that in the second line, the first term depends merely on j, while the second one merely on ℓ, and so, for fixed j or ℓ, any term can show up at most $2\gamma + 1$ times; (***) is due to the easy observation that

$$\sum_j (\frac{4}{5})^{2j} \langle |\mathbf{x}^{(j)}|, \boldsymbol{C}\mathbf{1}_n\rangle = \|\boldsymbol{D}_{row}^{1/2}\mathbf{x}\|^2 \le 1, \quad \sum_\ell (\frac{4}{5})^{2\ell} \langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle = \|\boldsymbol{D}_{col}^{1/2}\mathbf{y}\|^2 \le 1. \quad (2.59)$$

Terms (b) and (c) are of similar appearance (the role of j and ℓ is symmetric in them), therefore, we will estimate only (b). Here $j - \ell > \gamma$, yielding $j + \ell > 2\ell + \gamma$. Therefore,

$$\begin{aligned}
\sum_{j-\ell>\gamma} (\frac{4}{5})^{j+\ell} |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| &\le \sum_\ell (\frac{4}{5})^{2\ell+\gamma} \sum_j |\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell)}\rangle| \overset{(2.58)}{\le} \sum_\ell (\frac{4}{5})^{2\ell+\gamma} 2\langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle \\
&= 2(\frac{4}{5})^\gamma \sum_\ell (\frac{4}{5})^{2\ell} \langle \mathbf{1}_m, \boldsymbol{C}|\mathbf{y}^{(\ell)}|\rangle \overset{(2.59)}{\le} 2(\frac{4}{5})^\gamma
\end{aligned}$$

where, in the second and third inequalities, (2.58) and (2.59) were used. Consequently, (c) can also be estimated from above with $2(\frac{4}{5})^\gamma$.

Collecting the so obtained estimates together, we get

$$\begin{aligned}
s_k \le \frac{9}{2} |\langle \mathbf{x}, \boldsymbol{F}\mathbf{y}\rangle| &\le \frac{9}{2} \left[2k\alpha^*(2\gamma+1) + 4(\frac{4}{5})^\gamma \right] = 9\alpha^* \left[2k \frac{\ln \alpha^*}{\ln \frac{4}{5}} + k + 2 \right] \\
&\le 9\alpha^* [2k(-4.5) \ln \alpha^* + k + 2] = 9\alpha^* (k + 2 - 9k \ln \alpha^*),
\end{aligned}$$

that was to be proved. □

Note that for $k = 1$, our upper bound is tighter than that of (2.52), see Theorem 2 of [But].

Observe that for small discrepancies, the right-hand side of (2.53) is a strictly increasing function of $\mathrm{md}_k(\boldsymbol{C})$ when it is 'small'. Actually, the same function of $\mathrm{md}(\boldsymbol{C}; R_1, \ldots, R_k, C_1, \ldots, C_k)$ is also a valid upper estimate for s_k whenever the row-partitions $R_1, \ldots, R_k$ and the column-partitions $C_1, \ldots, C_k$ are such that $\mathrm{md}(\boldsymbol{C}; R_1, \ldots, R_k, C_1, \ldots, C_k) < 1$ holds. Since the function $f(x) = 9x(k + 2 - 9k \ln x)$ is strictly increasing near zero, $\mathrm{md}_k(\boldsymbol{C})$ is the best upper estimate.

2.3.2 Estimating the multiway discrepancy of contingency tables by the singular values and subspace deviations of the normalized table

In the forward direction, we did not manage to estimate the k-way discrepancy from above merely by means of the k-th largest non-trivial singular value of the normalized table, but had to use the k-variances of the optimal $(k-1)$-dimensional row- and column-representatives too. In the proof, we applied a bit different notion of the multiway discrepancy, but at the end, we will discuss the relation between it and that of Definition 25. Actually, we used a notion similar to that of the volume regularity introduced in Alon and coauthors [Aletal], where the authors also give an algorithm that computes a regular partition of a given (possibly sparse) graph in polynomial time giving some kind of construction for the Szemerédi Regularity Lemma.

Definition 26 *The row–column cluster pair $R \subset Row$, $C \subset Col$ of the contingency table $\boldsymbol{C}$ of total volume 1 is α-volume regular if for every $X \subset R$ and $Y \subset C$ the relation*

$$|c(X, Y) - \rho(R, C)\mathrm{Vol}(X)\mathrm{Vol}(Y)| \leq \alpha\sqrt{\mathrm{Vol}(R)\mathrm{Vol}(C)} \tag{2.60}$$

holds, where $\rho(R, C)$ is the relative inter-cluster density of the row–column pair R, C, introduced in Definition 25.

Theorem 25 ([Bol14b]) *Let $\boldsymbol{C}$ be a non-degenerate contingency table of m rows and n columns, with row- and column sums $d_{row,1}, \ldots, d_{row,m}$ and $d_{col,1}, \ldots, d_{col,n}$, respectively. Assume that $\sum_{i=1}^{n}\sum_{j=1}^{m} c_{ij} = 1$ and there are no dominant rows and columns: $d_{row,i} = \Theta(\frac{1}{m})$, $i = 1, \ldots, m$ and $d_{col,j} = \Theta(\frac{1}{n})$, $j = 1, \ldots, n$ as $m, n \to \infty$. Let the singular values of $\boldsymbol{C}_D$ be*

$$1 = s_0 > s_1 \geq \cdots \geq s_{k-1} > \varepsilon \geq s_i, \quad i \geq k.$$

The partition $(R_1, \ldots, R_k)$ of Row and $(C_1, \ldots, C_k)$ of Col are defined so that they minimize the weighted k-variances $\tilde{S}_k^2(\boldsymbol{X})$ and $\tilde{S}_k^2(\boldsymbol{Y})$ of the optimal row and column representatives collected in $\boldsymbol{X}$ and $\boldsymbol{Y}$ (defined in Section 1.2). Assume that there are constants $0 < K_1, K_2 \leq \frac{1}{k}$ such that $|R_i| \geq K_1 n$ and $|C_i| \geq K_2 m$ $(i = 1, \ldots, k)$, respectively. Then the R_i, C_j pairs are $\mathcal{O}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})\tilde{S}_k(\boldsymbol{Y})) + \varepsilon)$-volume regular $(i, j = 1, \ldots, k)$.

For the proof, we need the definition of the cut-norm and the relation between it and the spectral norm (see also [Fr-Kan, Gh-Trev]).

Definition 27 *The cut-norm of the real matrix $\boldsymbol{A}$ with row-set Row and column-set Col is*

$$\|\boldsymbol{A}\|_{\square} = \max_{R \subset Row,\, C \subset Col} \left| \sum_{i \in R} \sum_{j \in C} a_{ij} \right|.$$

Lemma 7 *For the $m \times n$ real matrix $\boldsymbol{A}$,*

$$\|\boldsymbol{A}\|_\square \le \sqrt{mn}\|\boldsymbol{A}\|,$$

where the right hand side contains the spectral norm, i.e., the largest singular value of $\boldsymbol{A}$.

Proof of Lemma 7.

$$\begin{aligned}\|\boldsymbol{A}\|_\square &= \max_{\mathbf{x}\in\{0,1\}^m,\,\mathbf{y}\in\{0,1\}^n} |\mathbf{x}^T\boldsymbol{A}\mathbf{y}| = \max_{\mathbf{x}\in\{0,1\}^m,\,\mathbf{y}\in\{0,1\}^n} \left|(\frac{\mathbf{x}}{\|\mathbf{x}\|})^T\boldsymbol{A}(\frac{\mathbf{y}}{\|\mathbf{y}\|})\right| \cdot \|\mathbf{x}\| \cdot \|\mathbf{y}\|| \\ &\le \sqrt{mn} \max_{\|\mathbf{x}\|=1,\,\|\mathbf{y}\|=1} |\mathbf{x}^T\boldsymbol{A}\mathbf{y}| = \sqrt{mn}\|\boldsymbol{A}\|,\end{aligned}$$

since for $\mathbf{x} \in \{0,1\}^m$, $\|\mathbf{x}\| \le \sqrt{m}$, and for $\mathbf{y} \in \{0,1\}^n$, $\|\mathbf{y}\| \le \sqrt{n}$, that finishes the proof. □

The definition of the cut-norm and the result of the above lemma naturally extends to symmetric matrices with $m = n$. Note that in [Szeg], B. Szegedy estimates the cut-norm of a graphon from above by the spectral norm of the corresponding compact operator. Since our normalization is for matrices and not for graphons, the estimate of Lemma 7 does contain the size of the matrix.

Proof of Theorem 25. Let $\boldsymbol{C}_D = \sum_{i=0}^{r-1} s_i \mathbf{v}_i \mathbf{u}_i^T$ be SVD, where $r = \operatorname{rank}(\boldsymbol{C}) = \operatorname{rank}(\boldsymbol{C}_D)$. Recall that provided $\boldsymbol{C}$ is non-degenerate, the largest singular value $s_0 = 1$ of $\boldsymbol{C}_D$ is single with corresponding singular vector pair $\mathbf{v}_0 = \boldsymbol{D}_{row}^{1/2}\mathbf{1}_m$ and $\mathbf{u}_0 = \boldsymbol{D}_{col}^{1/2}\mathbf{1}_n$, respectively. The optimal k-dimensional representatives of the rows and columns are row vectors of the matrices $\boldsymbol{X} = (\mathbf{x}_0, \dots, \mathbf{x}_{k-1})$ and $\boldsymbol{Y} = (\mathbf{y}_0, \dots, \mathbf{y}_{k-1})$, where $\mathbf{x}_i = \boldsymbol{D}_{row}^{-1/2}\mathbf{v}_i$ and $\mathbf{y}_i = \boldsymbol{D}_{col}^{-1/2}\mathbf{u}_i$, respectively ($i = 0, \dots, k-1$), in view of Theorem 9. (Note that the first columns of equal coordinates can as well be omitted.) Assume that the minimum weighted k-variance is attained at the k-partition $(R_1, \dots, R_k)$ of the rows and $(C_1, \dots, C_k)$ of the columns, respectively. By the usual analysis of variance argument, it follows that

$$\tilde{S}_k^2(\boldsymbol{X}) = \sum_{i=0}^{k-1} \operatorname{dist}^2(\mathbf{v}_i, F), \quad \tilde{S}_k^2(\boldsymbol{Y}) = \sum_{i=0}^{k-1} \operatorname{dist}^2(\mathbf{u}_i, G),$$

where $F = \operatorname{Span}\{\boldsymbol{D}_{row}^{1/2}\mathbf{w}_1, \dots, \boldsymbol{D}_{row}^{1/2}\mathbf{w}_k\}$ and $G = \operatorname{Span}\{\boldsymbol{D}_{col}^{1/2}\mathbf{z}_1, \dots, \boldsymbol{D}_{col}^{1/2}\mathbf{z}_k\}$ with the so-called normalized row partition vectors $\mathbf{w}_1, \dots, \mathbf{w}_k$ of coordinates $w_{ji} = \frac{1}{\sqrt{\operatorname{Vol}(R_i)}}$ if $j \in R_i$ and 0, otherwise; and column partition vectors $\mathbf{z}_1, \dots, \mathbf{z}_k$ of coordinates $z_{ji} = \frac{1}{\sqrt{\operatorname{Vol}(C_i)}}$ if $j \in C_i$ and 0, otherwise ($i = 1, \dots, k$). Note that the vectors $\boldsymbol{D}_{row}^{1/2}\mathbf{w}_1, \dots, \boldsymbol{D}_{row}^{1/2}\mathbf{w}_k$ and $\boldsymbol{D}_{col}^{1/2}\mathbf{z}_1, \dots, \boldsymbol{D}_{col}^{1/2}\mathbf{z}_k$ form orthonormal systems in $\mathbb{R}^n$ and $\mathbb{R}^m$, respectively (but they are, usually, not complete). By Lemma 3, we can find orthonormal systems $\tilde{\mathbf{v}}_0, \dots, \tilde{\mathbf{v}}_{k-1} \in F$ and $\tilde{\mathbf{u}}_0, \dots, \tilde{\mathbf{u}}_{k-1} \in G$ such that

$$\tilde{S}_k^2(\boldsymbol{X}) \le \sum_{i=0}^{k-1} \|\mathbf{v}_i - \tilde{\mathbf{v}}_i\|^2 \le 2\tilde{S}_k^2(\boldsymbol{X}), \quad \tilde{S}_k^2(\boldsymbol{Y}) \le \sum_{i=0}^{k-1} \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|^2 \le 2\tilde{S}_k^2(\boldsymbol{Y}). \tag{2.61}$$

We approximate $\boldsymbol{C}_D$ by the rank k matrix $\sum_{i=0}^{k-1} s_i \tilde{\mathbf{v}}_i \tilde{\mathbf{u}}_i^T$ with the following accuracy (in spectral norm):

$$\left\| \sum_{i=0}^{r-1} s_i \mathbf{v}_i \mathbf{u}_i^T - \sum_{i=0}^{k-1} s_i \tilde{\mathbf{v}}_i \tilde{\mathbf{u}}_i^T \right\| \le \sum_{i=0}^{k-1} s_i \left\| \mathbf{v}_i \mathbf{u}_i^T - \tilde{\mathbf{v}}_i \tilde{\mathbf{u}}_i^T \right\| + \left\| \sum_{i=k}^{r-1} s_i \mathbf{v}_i \mathbf{u}_i^T \right\|, \tag{2.62}$$

where the spectral norm of the last term is at most ε, and the individual terms of the first one are estimated from above in the following way.

$$\begin{aligned} s_i\|\mathbf{v}_i\mathbf{u}_i^T - \tilde{\mathbf{v}}_i\tilde{\mathbf{u}}_i^T\| &\le \|(\mathbf{v}_i\mathbf{u}_i^T - \tilde{\mathbf{v}}_i\mathbf{u}_i^T) + (\tilde{\mathbf{v}}_i\mathbf{u}_i^T - \tilde{\mathbf{v}}_i\tilde{\mathbf{u}}_i^T)\| \\ &\le \|(\mathbf{v}_i - \tilde{\mathbf{v}}_i)\mathbf{u}_i^T\| + \|\tilde{\mathbf{v}}_i(\mathbf{u}_i - \tilde{\mathbf{u}}_i)^T\| \\ &= \sqrt{\|(\mathbf{v}_i - \tilde{\mathbf{v}}_i)\mathbf{u}_i^T\mathbf{u}_i(\mathbf{v}_i - \tilde{\mathbf{v}}_i)^T\|} + \sqrt{\|(\mathbf{u}_i - \tilde{\mathbf{u}}_i)\tilde{\mathbf{v}}_i^T\tilde{\mathbf{v}}_i(\mathbf{u}_i - \tilde{\mathbf{u}}_i)^T\|} \\ &= \sqrt{(\mathbf{v}_i - \tilde{\mathbf{v}}_i)^T(\mathbf{v}_i - \tilde{\mathbf{v}}_i)} + \sqrt{(\mathbf{u}_i - \tilde{\mathbf{u}}_i)^T(\mathbf{u}_i - \tilde{\mathbf{u}}_i)} \\ &= \|\mathbf{v}_i - \tilde{\mathbf{v}}_i\| + \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\| \end{aligned}$$

where we exploited that the spectral norm (i.e., the largest singular value) of an $m \times n$ matrix $\boldsymbol{A}$ is equal to either the squareroot of the largest eigenvalue of the matrix $\boldsymbol{A}\boldsymbol{A}^T$ or equivalently, that of $\boldsymbol{A}^T\boldsymbol{A}$. In the above calculations all of these matrices are of rank 1, hence, the largest eigenvalue of the symmetric, positive semidefinite matrix under the squareroot is the only non-zero eigenvalue of it, therefore, it is equal to its trace; finally, we used the commutativity of the trace, and in the last line we have the usual vector norm.

Therefore, the first term in (2.62) can be estimated from above with

$$\begin{aligned} \sum_{i=0}^{k-1}\|\mathbf{v}_i\mathbf{u}_i^T - \tilde{\mathbf{v}}_i\tilde{\mathbf{u}}_i^T\| &\le \sqrt{k}\sqrt{\sum_{i=0}^{k-1}\|\mathbf{v}_i - \tilde{\mathbf{v}}_i\|^2} + \sqrt{k}\sqrt{\sum_{i=0}^{k-1}\|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|^2} \\ &\le \sqrt{k}(\sqrt{2\tilde{S}_k^2(\boldsymbol{X})} + \sqrt{2\tilde{S}_k^2(\boldsymbol{Y})}) = \sqrt{2k}(\tilde{S}_k(\boldsymbol{X}) + \tilde{S}_k(\boldsymbol{Y})), \end{aligned}$$

where we also used the upper estimate of (2.61).

Based on these considerations and relation between the cut-norm and the spectral norm, the densities to be estimated in the defining formula (2.60) of volume regularity can be written in terms of step-vectors in the following way. The vectors $\hat{\mathbf{v}}_i := \boldsymbol{D}_{row}^{-1/2}\tilde{\mathbf{v}}_i$ are stepwise constant on the partition $(R_1, \dots, R_k)$ of the rows, whereas the vectors $\hat{\mathbf{u}}_i := \boldsymbol{D}_{col}^{-1/2}\tilde{\mathbf{u}}_i$ are stepwise constant on the partition $(C_1, \dots, C_k)$ of the columns, $i = 0, \dots, k-1$. The matrix

$$\sum_{i=0}^{k-1} s_i\hat{\mathbf{v}}_i\hat{\mathbf{u}}_i^T$$

is therefore an $n \times m$ block-matrix on $k \times k$ blocks corresponding to the above partition of the rows and columns. Let $\hat{c}_{ab}$ denote its entries in the ab block $(a, b = 1, \dots, k)$. Using (2.62), the rank k approximation of the matrix $\boldsymbol{C}$ is performed with the following accuracy of the perturbation $\boldsymbol{E}$ in spectral norm:

$$\|\boldsymbol{E}\| = \left\|\boldsymbol{C} - \boldsymbol{D}_{row}(\sum_{i=0}^{k-1} s_i\hat{\mathbf{v}}_i\hat{\mathbf{u}}_i^T)\boldsymbol{D}_{col}\right\| = \left\|\boldsymbol{D}_{row}^{1/2}(\boldsymbol{C}_D - \sum_{i=0}^{k-1} s_i\tilde{\mathbf{v}}_i\tilde{\mathbf{u}}_i^T)\boldsymbol{D}_{col}^{1/2}\right\|.$$

Therefore, the entries of $\boldsymbol{C}$ can be decomposed as

$$c_{ij} = d_{row,i}d_{col,j}\hat{c}_{ab} + \eta_{ij} \qquad (i \in R_a, \quad j \in C_b)$$

where the cut-norm of the $n \times m$ error matrix $\boldsymbol{E} = (\eta_{ij})$ restricted to $R_a \times C_b$ (otherwise it contains entries all zeros) and denoted by $\boldsymbol{E}_{ab}$, is estimated as follows. Making use of

Lemma 7,

$$\begin{aligned}\|\boldsymbol{E}_{ab}\|_\square &\le \sqrt{mn}\|\boldsymbol{E}_{ab}\| \le \sqrt{nm}\cdot\|\boldsymbol{D}_{row,a}^{1/2}\|\cdot(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon)\cdot\|\boldsymbol{D}_{col,b}^{1/2}\| \\ &\le \sqrt{nm}\sqrt{c_1\frac{\mathrm{Vol}(R_a)}{|R_a|}}\cdot\sqrt{c_2\frac{\mathrm{Vol}(C_b)}{|C_b|}}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon) \\ &= \sqrt{c_1c_2}\cdot\sqrt{\frac{n}{|R_a|}}\cdot\sqrt{\frac{m}{|C_b|}}\cdot\sqrt{\mathrm{Vol}(R_a)}\sqrt{\mathrm{Vol}(C_b)}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon) \\ &\le \sqrt{\frac{c_1c_2}{K_1K_2}}\sqrt{\mathrm{Vol}(R_a)}\sqrt{\mathrm{Vol}(C_b)}(\sqrt{2k}s+\varepsilon) \\ &= c\sqrt{\mathrm{Vol}(R_a)}\sqrt{\mathrm{Vol}(C_b)}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon)\end{aligned}$$

where the $n\times n$ diagonal matrix $\boldsymbol{D}_{row,a}$ inherits $\boldsymbol{D}_{row}$'s diagonal entries over R_a, whereas the $m\times m$ diagonal matrix $\boldsymbol{D}_{col,b}$ inherits $\boldsymbol{D}_{col}$'s diagonal entries over C_b, otherwise they are zeros. Further, the constants c_1, c_2 are due to the fact that there are no dominant rows and columns, while K_1, K_2 are from the cluster size balancing conditions. Hence,

$$\|\boldsymbol{E}_{ab}\|_\square \le c\sqrt{\mathrm{Vol}(R_a)}\sqrt{\mathrm{Vol}(C_b)}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon)$$

where the constant c does not depend on n and m. Consequently, for $a,b=1,\dots,k$ and $X\subset R_a$, $Y\subset C_b$,

$$\begin{aligned}&|c(X,Y)-\rho(R_a,C_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)| = \\ &\left|\sum_{i\in X}\sum_{j\in Y}(d_{row,i}d_{col,j}\hat{c}_{ab}+\eta_{ij})-\frac{\mathrm{Vol}(X)\mathrm{Vol}(Y)}{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}\sum_{i\in R_a}\sum_{j\in C_b}(d_{row,i}d_{col,j}\hat{c}_{ab}+\eta_{ij})\right| = \\ &\left|\sum_{i\in X}\sum_{j\in Y}\eta_{ij}-\frac{\mathrm{Vol}(X)\mathrm{Vol}(Y)}{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}\sum_{i\in R_a}\sum_{j\in C_b}\eta_{ij}\right| \le 2\|\boldsymbol{E}_{ab}\|_\square \\ &\le 2c(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon)\sqrt{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}\end{aligned}$$

that gives the required statement for $a,b=1,\dots,k$, and finishes the proof. □

So we managed to prove the following. Given the $m\times n$ contingency table $\boldsymbol{C}$, consider the spectral clusters $R_1,\dots,R_k$ of its rows and $C_1,\dots,C_k$ of its columns, obtained by applying the weighted k-means algorithm to the $(k-1)$-dimensional row- and column representatives, defined as the row vectors of the matrices $(\boldsymbol{D}_{row}^{-1/2}\mathbf{v}_1,\dots,\boldsymbol{D}_{row}^{-1/2}\mathbf{v}_{k-1})$ and $(\boldsymbol{D}_{col}^{-1/2}\mathbf{u}_1,\dots,\boldsymbol{D}_{col}^{-1/2}\mathbf{u}_{k-1})$, respectively, where $\mathbf{v}_i,\mathbf{u}_i$ is the unit norm singular vector pair corresponding to s_i $(i=1,\dots,k-1)$. In fact, these partitions minimize the weighted k-variances $\tilde{S}_k^2(\boldsymbol{X})$ and $\tilde{S}_k^2(\boldsymbol{Y})$ of these row- and column-representatives. Then, under some balancing conditions for $d_{row,i}$'s and $d_{col,j}$'s (there are no dominant rows and columns) and for the cluster sizes, we proved that $\mathrm{md}_k'(\boldsymbol{C})=\mathcal{O}(\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+s_k)$, where $\mathrm{md}_k'(\boldsymbol{C})$ is a somewhat modified version of the k-way discrepancy; the only difference is that in the definition of $\mathrm{md}_k'(\boldsymbol{C})$ we substitute $\sqrt{\mathrm{Vol}(R_a)\mathrm{Vol}(C_b)}$ for $\sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)}$ in the denominator of (2.50). In accordance with the original definition of the discrepancy in the Szemerédi Regularity Lemma [Szem] for simple graphs, in (2.50), we may take the maximum over subsets $X\subset V_a$, $Y\subset V_b$ such that $\mathrm{Vol}(X)\ge\epsilon\mathrm{Vol}(V_a)$ and $\mathrm{Vol}(Y)\ge\epsilon\mathrm{Vol}(V_b)$ with some fixed $\epsilon>0$. If we impose similar conditions on the row- and column-subsets, our result also implies that $\mathrm{md}_k(\boldsymbol{C})$ is of order $\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+s_k$.

The message of Theorems 24 and 25 is that the k-way discrepancy, when it is 'small' enough, suppresses s_k. Conversely, s_k together with 'small' enough $\tilde{S}_k(\boldsymbol{X})$ and $\tilde{S}_k(\boldsymbol{Y})$ also suppresses the k-way discrepancy. By using perturbation theory of spectral subspaces, in [Bol14a] (in the framework of edge-weighted graphs), we also discuss that a 'large' gap between s_{k-1} and s_k suppresses $\tilde{S}_k(\boldsymbol{X})$ and $\tilde{S}_k(\boldsymbol{Y})$. Therefore, if we want to find row–column cluster pairs of small discrepancy, we must select a k such that there is a remarkable gap between s_{k-1} and s_k; further s_k is small enough. Moreover, by using this k and the construction in the proof of the forward statement of Theorem 25, we are able to find these clusters with spectral clustering tools. It makes sense, for example, when we want to find clusters of genes and conditions simultaneously in microarrays so that genes of the same row-cluster would 'equally' influence conditions of the same column-cluster.

We also remark the following. When we perform correspondence analysis on a large $m \times n$ contingency table and consider the rank k approximation of it, the entries of this matrix will not necessarily be positive at all. Nonetheless, the entries $\hat{c}_{ij}$'s of the block-matrix constructed in the proof of Theorem 25 will already be positive provided the weighted k-variances $\tilde{S}_k(\boldsymbol{X})$ and $\tilde{S}_k(\boldsymbol{Y})$ are 'small' enough. Let us discuss this issue more precisely.

In accord with the notation used in the proof, denote by ab in the lower index if the matrix is restricted to the $R_a \times C_b$ block (otherwise it has zero entries). Then for the squared Frobenius norm of the rank k approximation of $\boldsymbol{D}_{row}^{-1}\boldsymbol{C}\boldsymbol{D}_{col}^{-1}$, restricted to the ab block, we have that

$$\begin{aligned}\left\|\boldsymbol{D}_{row,a}^{-1}\boldsymbol{C}_{ab}\boldsymbol{D}_{col,b}^{-1}-(\sum_{i=0}^{k-1}s_i\hat{\mathbf{v}}_i\hat{\mathbf{u}}_i^T)_{ab}\right\|_2^2 &= \sum_{i\in R_a}\sum_{j\in C_b}(\frac{c_{ij}}{d_{row,i}d_{col,j}}-\hat{c}_{ab})^2\\ =\sum_{i\in R_a}\sum_{j\in C_b}(\frac{c_{ij}}{d_{row,i}d_{col,j}}-\bar{c}_{ab})^2 &+ |R_a||C_b|(\bar{c}_{ab}-\hat{c}_{ab})^2\end{aligned} \tag{2.63}$$

where we used the Steiner equality with the average $\bar{c}_{ab}$ of the entries of $\boldsymbol{D}_{row}^{-1}\boldsymbol{C}\boldsymbol{D}_{col}^{-1}$ in the ab block. Now we estimate the above Frobenius norm by a constant multiple of the spectral norm, where for the spectral norm

$$\begin{aligned}&\left\|\boldsymbol{D}_{row,a}^{-1}\boldsymbol{C}_{ab}\boldsymbol{D}_{col,b}^{-1}-(\sum_{i=0}^{k-1}s_i\hat{\mathbf{v}}_i\hat{\mathbf{u}}_i^T)_{ab}\right\| = \left\|\boldsymbol{D}_{row,a}^{-1/2}(\boldsymbol{C}_{corr}-\sum_{i=0}^{k-1}s_i\tilde{\mathbf{v}}_i\tilde{\mathbf{u}}_i^T)_{ab}\boldsymbol{D}_{col,b}^{-1/2}\right\|\\ &\le \max_{i\in R_a}\frac{1}{\sqrt{d_{row,i}}}\cdot\max_{j\in C_b}\frac{1}{\sqrt{d_{col,j}}}\cdot[\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon]\end{aligned}$$

holds. Therefore,

$$\begin{aligned}&\left\|\boldsymbol{D}_{row,a}^{-1}\boldsymbol{C}_{ab}\boldsymbol{D}_{col,b}^{-1}-(\sum_{i=0}^{k-1}s_i\hat{\mathbf{v}}_i\hat{\mathbf{u}}_i^T)_{ab}\right\|_2^2\\ &\le \min\{|R_a|,|C_b|\}\cdot\max_{i\in R_a}\frac{1}{d_{row,i}}\cdot\max_{j\in C_b}\frac{1}{d_{col,j}}\cdot[\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon]^2.\end{aligned}$$

Consequently, in view of (2.63),

$$(\bar{c}_{ab}-\hat{c}_{ab})^2 \le \frac{1}{\max\{|R_a|,|C_b|\}}\cdot\max_{i\in R_a}\frac{1}{d_{row,i}}\cdot\max_{j\in C_b}\frac{1}{d_{col,j}}\cdot[\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon]^2.$$

But using the conditions on the block sizes and the row- and column-sums of Theorem 25, provided

$$\sqrt{2k}(\tilde{S}_k(\boldsymbol{X})+\tilde{S}_k(\boldsymbol{Y}))+\varepsilon) = \mathcal{O}\left(\frac{1}{(\min\{m,n\})^{\frac{1}{2}+\tau}}\right)$$

holds with some 'small' $\tau > 0$, the relation $\bar{c}_{ab} - \hat{c}_{ab} \to 0$ also holds as $n, m \to \infty$. Therefore, both $\hat{c}_{ab}$ and $\hat{c}_{ab} d_{row,i} d_{col,j}$ are positive over such blocks that are not constantly zero in the original table if m and n are large enough.

2.3.3 Multiway discrepancy of undirected graphs

The notion of multiway discrepancy naturally extends to edge-weighted graphs.

Definition 28 *The multiway discrepancy of the undirected, edge-weighted graph $G = (V, \boldsymbol{W})$ in the proper k-partition $(V_1, \ldots, V_k)$ of its vertices is*

$$\mathrm{md}(G; V_1, \ldots, V_k) = \max_{\substack{1 \le a \le b \le k \\ X \subset V_a,\, Y \subset V_b}} \frac{|w(X,Y) - \rho(V_a, V_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)|}{\sqrt{\mathrm{Vol}(X)\mathrm{Vol}(Y)}}.$$

The minimum k-way discrepancy of the undirected edge-weighted graph $G = (V, \boldsymbol{W})$ is

$$\mathrm{md}_k(G) = \min_{(V_1,\ldots,V_k)} \mathrm{md}(G; V_1, \ldots, V_k).$$

A result, analogous to that of Theorem 24 can now be proved in terms of the normalized modularity matrix of G.

Theorem 26 ([Bol16]) *Let $G = (V, \boldsymbol{W})$ be an edge-weighted, undirected graph, $\boldsymbol{W}$ is irreducible. Then for any integer $1 \le k < \mathrm{rank}(\boldsymbol{W})$,*

$$|\mu_k| \le 9\mathrm{md}_k(G)(k + 2 - 9k \ln \mathrm{md}_k(G)) \tag{2.64}$$

holds, provided $0 < \mathrm{md}_k(G) < 1$, where μ_k is the k-th largest absolute value eigenvalue of the normalized modularity matrix $\boldsymbol{M}_D$ of G.

Proof (of Theorem 26). The proof follows the same considerations as the proof of Theorem 24 with the difference that here we use symmetric matrices. In particular, $\boldsymbol{R} = (\rho(V_a, V_b))$ is an $n \times n$ symmetric block-matrix of $k \times k$ blocks corresponding to the partition $V_1, \ldots, V_k$ of the vertices for which $\alpha^* = \mathrm{md}_k(G) = \mathrm{md}(G; V_1, \ldots, V_k)$; consequently, the matrix $\boldsymbol{F} = \boldsymbol{W} - \boldsymbol{DRD}$ is also symmetric. Therefore, in accord with (2.56) and Definition 28: for every V_a, V_b pair and $X \subset V_a$, $Y \subset V_b$ $(1 \le a \le b \le k)$ we have

$$|\langle \mathbf{1}_X, \boldsymbol{F}\mathbf{1}_Y \rangle| \le \alpha^* \sqrt{\langle \mathbf{1}_X, \boldsymbol{W}\mathbf{1}_n \rangle \langle \mathbf{1}_n, \boldsymbol{W}\mathbf{1}_Y \rangle}. \tag{2.65}$$

The left and right singular vectors $(\mathbf{v}, \mathbf{u} \in \mathbb{R}^n)$ corresponding to the maximal singular value of the real symmetric matrix $\boldsymbol{D}^{-1/2}\boldsymbol{F}\boldsymbol{D}^{-1/2}$ satisfy $\mathbf{u} = \pm\mathbf{v}$ (the sign is the same as the sign of the eigenvalue of the maximal absolute value). If $\mathbf{u} = \mathbf{v}$, then $\mathcal{Y}_{\ell bq} = \mathcal{X}_{\ell bq}$ for every ℓ, $b = 1, \ldots, k$, and $q = 1, 2$. If $\mathbf{u} = -\mathbf{v}$, then $\mathcal{Y}_{\ell b1} = \mathcal{X}_{\ell b2}$ and $\mathcal{Y}_{\ell b2} = \mathcal{X}_{\ell b1}$ for every ℓ and $b = 1, \ldots, k$. Consequently, in the estimates of (2.57), when we use the absolute values of the coordinates of the vectors $\mathbf{x}^{(j)}$ and $\mathbf{y}^{(\ell)}$, constructed in the rectangular case, the inequalities remain valid. Namely,

$$|\langle \mathbf{x}^{(j)}, \boldsymbol{F}\mathbf{y}^{(\ell))} \rangle| \le \sum_{a=1}^{k} \sum_{p=1}^{2} \sum_{b=1}^{k} \sum_{q=1}^{2} \left| \langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{F}\mathbf{1}_{\mathcal{X}_{\ell bq}} \rangle \right|.$$

Here the summation is for every $1 \le a, b \le k$ pair. However, if $a \le b$, then by (2.65) we get

$$\left| \langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{F}\mathbf{1}_{\mathcal{X}_{\ell bq}} \rangle \right| \le \alpha^* \sqrt{\langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{W}\mathbf{1}_n \rangle \langle \mathbf{1}_n, \boldsymbol{W}\mathbf{1}_{\mathcal{X}_{\ell bq}} \rangle};$$

whereas, if $a > b$, then by the symmetry of $\boldsymbol{F}$:

$$\left|\langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{F}\mathbf{1}_{\mathcal{X}_{\ell bq}}\rangle\right| = \left|\langle \mathbf{1}_{\mathcal{X}_{\ell bq}}, \boldsymbol{F}\mathbf{1}_{\mathcal{X}_{jap}}\rangle\right| \leq \alpha^* \sqrt{\langle \mathbf{1}_{\mathcal{X}_{lbq}}, \boldsymbol{W}\mathbf{1}_n\rangle\langle \mathbf{1}_n, \boldsymbol{W}\mathbf{1}_{\mathcal{X}_{jap}}\rangle}$$
$$= \sqrt{\langle \mathbf{1}_{\mathcal{X}_{jap}}, \boldsymbol{W}\mathbf{1}_n\rangle\langle \mathbf{1}_n, \boldsymbol{W}\mathbf{1}_{\mathcal{X}_{\ell bq}}\rangle}.$$

Therefore, for $a \neq b$, the same term appears twice, and all the subsequent estimates remain valid by substituting $\boldsymbol{W}$ for $\boldsymbol{C}$ and $\boldsymbol{D}$ for both $\boldsymbol{D}_{row}$ and $\boldsymbol{D}_{col}$. This completes the proof. □

Recall that Bilu and Linial [Bil-Lin] proved the following converse of the expander mixing lemma for simple d-regular graphs on n vertices. Assume that for any disjoint vertex-subsets S, T: $|e(S,T) - \frac{|S||T|d}{n}| \leq \alpha\sqrt{|S||T|}$. Then all but the largest adjacency eigenvalue of G are bounded (in absolute value) by $O(\alpha(1+\log\frac{d}{\alpha}))$. Note that for a d-regular graph the adjacency eigenvalues are d times larger than the normalized adjacency ones, and the deviation between $e(S,T)$ and the one what is expected in a random d-regular graph, is also proportional to our (1-way) discrepancy in terms of the volumes (note that $\mathrm{Vol}(S)$ is also proportional to $|S|$). Though they use disjoint subsets S, T, their upper estimate for the absolute value of the second largest (in absolute value) eigenvalue with the (1-way) discrepancy α is $C\alpha(1-A\log\alpha)$ with some absolute constants A, C. Hence, the upper estimate of (2.52) or that of (2.64) in the $k = 1$ case are reminiscent of this.

In the other direction, for the $k = 1$ case, a straightforward generalization of the *expander mixing lemma for irregular graphs* is the following.

Proposition 14 ([Bol11a])

$$\mathrm{disc}(G) = \mathrm{md}_1(G) \leq \|\boldsymbol{M}_D\| = s_1 = |\mu_1|,$$

where $\|\boldsymbol{M}_D\|$ is the spectral norm of the normalized modularity matrix of G.

Though, with different notation sometimes even a stronger version of this statement is proved in [Bol11a, But, Chu-G], we give another short proof here.

Proof. Via separation theorems for singular values, $s_1 = |\mu_1|$ is the maximum of the bilinear form $\mathbf{v}^T\boldsymbol{M}_D\mathbf{u}$ over the unit sphere. Let $X, Y \subset V$ be arbitrary, and denote by $\mathbf{1}_X, \mathbf{1}_Y \in \mathbb{R}^n$ the indicator vectors of them. Then

$$\|\boldsymbol{M}_D\| = \max_{\|\mathbf{u}\|=\|\mathbf{v}\|=1} |\mathbf{v}^T\boldsymbol{M}_D\mathbf{u}| \geq \left|\left(\frac{\boldsymbol{D}^{1/2}\mathbf{1}_X}{\|\boldsymbol{D}^{1/2}\mathbf{1}_X\|}\right)^T \boldsymbol{M}_D \left(\frac{\boldsymbol{D}^{1/2}\mathbf{1}_Y}{\|\boldsymbol{D}^{1/2}\mathbf{1}_Y\|}\right)\right|$$
$$= \frac{|\mathbf{1}_X^T\boldsymbol{M}\mathbf{1}_Y|}{\|\boldsymbol{D}^{1/2}\mathbf{1}_X\| \cdot \|\boldsymbol{D}^{1/2}\mathbf{1}_Y\|} = \frac{|w(X,Y) - \mathrm{Vol}(X)\mathrm{Vol}(Y)|}{\sqrt{\mathrm{Vol}(X)}\sqrt{\mathrm{Vol}(Y)}}.$$

Taking the maxima on the right-hand side over subsets $X, Y \subset V$, the desired statement follows. Note that the estimate is also valid if we take maxima over disjoint X, Y pairs only. □

For an arbitrary integer k, in the range $1 \leq k < \mathrm{rank}(\boldsymbol{W})$, the following analogue of Theorem 25 for undirected, edge-weighted graphs was proved in [Bol11b, Bol14a].

Theorem 27 ([Bol14a]) *Let $G = (V, \boldsymbol{W})$ be an edge-weighted graph on n vertices, with generalized degrees $d_1, \ldots, d_n$ and degree-matrix $\boldsymbol{D}$. Assume that G is connected, $\mathrm{Vol}(V) = 1$, and there are no dominant vertices: $d_i = \Theta(1/n)$, $i = 1, \ldots, n$ as $n \to \infty$. Let the eigenvalues of the normalized modularity matrix $\boldsymbol{M}_D$ of G, enumerated in decreasing absolute values, be*

$$|\mu_1| \geq \cdots \geq |\mu_{k-1}| > \varepsilon \geq |\mu_k| \geq \cdots \geq |\mu_n| = 0.$$

The partition $(V_1, \ldots, V_k)$ of V is defined so that it minimizes the weighted k-variance $\tilde{S}_k^2(\boldsymbol{X}^*)$ of the optimal vertex representatives obtained as row vectors of the $n \times (k-1)$ matrix $\boldsymbol{X}^*$ of column vectors $\boldsymbol{D}^{-1/2}\mathbf{u}_i$, where $\mathbf{u}_i$ is the unit-norm eigenvector corresponding to μ_i $(i = 1, \ldots, k-1)$. Assume that there is a constant $0 < K \leq \frac{1}{k}$ such that $|V_i| \geq Kn$, $i = 1, \ldots, k$. With the notation $\sigma_k = \sqrt{\tilde{S}_k^2(\boldsymbol{X}^*)}$, the (V_i, V_j) pairs are $\mathcal{O}(\sqrt{2k}\sigma_k + \varepsilon)$-volume regular $(i \neq j)$ and for the clusters V_i $(i = 1, \ldots, k)$ the following holds: for all $X, Y \subset V_i$,

$$|w(X,Y) - \rho(V_i)\mathrm{Vol}(X)\mathrm{Vol}(Y)| = \mathcal{O}(\sqrt{2k}\sigma_k + \varepsilon)\mathrm{Vol}(V_i),$$

where $\rho(V_i) = \frac{w(V_i,V_i)}{\mathrm{Vol}^2(V_i)}$ is the relative intra-cluster density of V_i.

In fact, inspired by Alon et al. [Aletal], in [Bol14a] we used a bit different notation and concept of α-volume regular pairs, namely, for every $X \subseteq V_a$, $Y \subseteq V_b$ we required

$$|w(X,Y) - \rho(V_a, V_b)\mathrm{Vol}(X)\mathrm{Vol}(Y)| \leq \alpha\sqrt{\mathrm{Vol}(V_a)\mathrm{Vol}(V_b)}.$$

In the above formula, the right had side contains the squareroots of the volumes of the clusters, unlike (2.51), which contains the squareroots of the volumes of X and Y. However, with the same argument as in the rectangular case: in the spirit of the Szemerédi Regularity Lemma [Szem], if we require (2.51) to hold only for X, Y's satisfying $\mathrm{Vol}(X) \geq \epsilon\mathrm{Vol}(V_i)$, $\mathrm{Vol}(Y) \geq \epsilon\mathrm{Vol}(V_j)$ with some fixed ϵ, then the so modified k-way discrepancy is $\mathcal{O}(\sqrt{2k}\sigma_k + |\mu_k|)$, and so does $\mathrm{md}_k(G)$.

In [Bol14a] we also discuss that, in view of subspace perturbation theorems, the larger the gap between $|\mu_{k-1}|$ and $|\mu_k|$, the smaller σ_k is. So the message is, that here the eigenvectors corresponding to the largest absolute value eigenvalues have to be used, unlike usual spectral clustering methods which automatically use the bottom eigenvalues of the Laplacian or normalized Laplacian matrix (latter one is just $\boldsymbol{I} - \boldsymbol{W}_D$). The clusters or cluster-pairs of small discrepancy behave like expanders or bipartite expanders. In another context, they resemble the generalized random or quasirandom graphs of Lovász, Sós, Simonovits [Lov-Sos, Sim-Sos], see Section 3 for further details.

Note that in the special two-cluster case, due to Theorem 5, the 2-variance of the optimal one-dimensional representatives can be directly estimated from above by the gap between the two largest absolute value eigenvalues of $\boldsymbol{M}_D$, and hence, the statement of Theorem 27 simplifies as follows. The optimal pair (V_1, V_2) based on minimizing the weighted 2-variance of the coordinates of $\mathbf{u}_1$ is $\mathcal{O}(\sqrt{\frac{1-\delta}{1-\varepsilon}})$-volume regular, where $\delta = |\mu_1|$ and $\varepsilon = |\mu_2|$, provided $\boldsymbol{W}$ is non-degenerate (the underlying graph is connected, but not bipartite). With other methods, the same estimate is obtained in [Bol11a], where we treat the case when the two largest absolute value eigenvalues of the normalized modularity matrix are positive, though it can be adopted to the other situations too; see, for example, the dual Cheeger inequality of Luca Trevisan [Trev].

For a general k, we can make the following considerations, by using the forthcoming well-known fact for the perturbation of spectral subspaces, see also [Bol13].

Lemma 8 (Theorem VII.3.2 of [Bhat]) *Let $\boldsymbol{A}$ and $\boldsymbol{B}$ be symmetric matrices; S_1 and S_2 are subsets of $\mathbb{R}$ or $\mathbb{C}$ such that $\mathrm{dist}(S_1, S_2) = \delta > 0$. Let $\boldsymbol{P}_{\boldsymbol{A}}(S_1)$ and $\boldsymbol{P}_{\boldsymbol{B}}(S_2)$ be orthogonal projections onto the subspace spanned by the eigenvectors of the matrix in the lower index, corresponding to the eigenvalues within the subset in the argument. Then with any unitary invariant norm:*

$$\|\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)\|_{\mathrm{un}} \leq \frac{c_1}{\delta}\|\boldsymbol{P}_A(S_1)(\boldsymbol{A} - \boldsymbol{B})\boldsymbol{P}_B(S_2)\|_{\mathrm{un}} \leq \frac{c_1}{\delta}\|\boldsymbol{A} - \boldsymbol{B}\|_{\mathrm{un}}$$

where c_1 is a constant.

Note that, in another context, Szőkefalvy proved that $c_1 = \pi/2$. When S_1 and S_2 are separated by an annulus, then the constant improves to $c_1 = 1$; further, with the Frobenius norm, $c_1 = 1$ will always do. If $\boldsymbol{P}_A(S_1)$ and $\boldsymbol{P}_B^{\perp}(S_2)$ project onto subspaces of the same dimension, then either the spectral or the Frobenius norm of $\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)$ can be expressed in terms of the sines of the so-called canonical (principal) angles between these subspaces and $\|\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)\|_2$ is considered as the distance between them. This is why the special case of Lemma 8, when Frobenius norm is used, is called Davis–Kahan $\sin(\boldsymbol{\theta})$ theorem.

Now assume that the normalized modularity spectrum (with decreasing absolute values) of $G_n = (V, \boldsymbol{W})$ satisfies

$$|\mu_1| \geq \cdots \geq |\mu_{k-1}| \geq \delta > \varepsilon \geq |\mu_k| \geq \cdots \geq |\mu_n| = 0.$$

Our purpose is to estimate the k-variance of the optimal representatives with the gap $\theta := \delta - \varepsilon$. We will use the notation of the Proof of Theorem 27 (actually, it is proved in B [Bol13], but the proof is analogous to that of Theorem 25) and apply Lemma 8 for the perturbation of spectral subspaces of the symmetric matrices

$$\boldsymbol{A} = \sum_{i=0}^{n-1} \mu_i \mathbf{u}_i \mathbf{u}_i^T \quad \text{and} \quad \boldsymbol{B} = \sum_{i=0}^{k-1} \mu_i \mathbf{v}_i \mathbf{v}_i^T$$

in the following situation. The subsets $S_1 = \{\mu_k, \ldots, \mu_{n-1}\}$ and $S_2 = \{\mu_0, \ldots, \mu_{k-1}\}$ of the eigenvalues of $\boldsymbol{D}^{-1/2}\boldsymbol{W}\boldsymbol{D}^{-1/2}$ are separated by an annulus, where $\operatorname{dist}(S_1, S_2) = \theta > 0$. Denote by $\boldsymbol{P}_A(S_1)$ and $\boldsymbol{P}_B(S_2)$ the projections onto the spectral subspaces of $\boldsymbol{A}$ and $\boldsymbol{B}$ spanned by the eigenvectors corresponding to the eigenvalues in S_1 and S_2, respectively:

$$\boldsymbol{P}_A(S_1) = \sum_{j=k}^{n-1} \mathbf{u}_j \mathbf{u}_j^T, \quad \boldsymbol{P}_B(S_2) = \sum_{i=0}^{k-1} \mathbf{v}_i \mathbf{v}_i^T.$$

Then Lemma 8 implies that

$$\|\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)\|_2 \leq \frac{1}{\theta}\|\boldsymbol{P}_A(S_1)(\boldsymbol{A} - \boldsymbol{B})\boldsymbol{P}_B(S_2)\|_2, \tag{2.66}$$

where $\|.\|_2$ denotes the Frobenius norm.

Since $\boldsymbol{P}_A^{\perp}(S_1)$ and $\boldsymbol{P}_B(S_2)$ project onto subspaces of the same dimension, the non-zero singular values of $\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)$ are the sines of the the canonical (principal) angles α_i's between the subspaces $\operatorname{Span}\{\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{k-1}\}$ and $\operatorname{Span}\{\mathbf{v}_0, \mathbf{v}_1, \ldots, \mathbf{v}_{k-1}\}$. Further, $\|\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)\|_2$ is the distance between the two above subspaces, where $\sin\alpha_0 = 0$.

On the left hand side, $\|\boldsymbol{P}_A(S_1)\boldsymbol{P}_B(S_2)\|_2 = \sqrt{\sum_{i=0}^{k-1} \sin^2\alpha_i}$, and in view of $\|\mathbf{u}_i - \mathbf{v}_i\| = 2\sin\frac{\alpha_i}{2}$ and the quadratic analogue of (2.61), this is between $\frac{\sqrt{3}}{2}\sigma_k$ and σ_k. On the right hand side,

$$\begin{aligned}
&\boldsymbol{P}_A(S_1)\boldsymbol{A}\boldsymbol{P}_B(S_2) - \boldsymbol{P}_A(S_1)\boldsymbol{B}\boldsymbol{P}_B(S_2) \\
=&(\boldsymbol{P}_A(S_1)\boldsymbol{A})\boldsymbol{P}_B(S_2) - \boldsymbol{P}_A(S_1)(\boldsymbol{P}_B(S_2)\boldsymbol{B}) \\
=&\sum_{i=0}^{k-1}\sum_{j=k}^{n-1}(\mu_j - \mu_i)(\mathbf{u}_j^T\mathbf{v}_i)\mathbf{u}_j\mathbf{v}_i^T
\end{aligned}$$

where the Frobenius norm of the rank one matrices $\mathbf{u}_j\mathbf{v}_i^T$ is 1, and the inner product $\mathbf{u}_j^T\mathbf{v}_i$ is the 'smaller', the $\mathbf{v}_i$ is the 'closer' to $\mathbf{u}_i$. Therefore, by Inequality (2.66), σ_k is the 'smaller', θ is the 'larger' and the $|\mu_j - \mu_i|$ differences for $i = 0, \ldots, k-1$; $j = k, \ldots, n-1$ are 'closer'

to δ. If $|\mu_k| = \varepsilon$ is 'small', then $|\mu_1|, \ldots, |\mu_{k-1}|$ should be 'close' to each other ($\mu_0 = 1$ does not play an important role because of $\mathbf{u}_0 = \mathbf{v}_0$).

In some special cases, $\sigma_k = 0$, and then, $\mathrm{md}_k(G) \leq B|\mu_k| = Bs_k$ follows from the above results. In particular, $\sigma_k = 0$ whenever the vectors $\boldsymbol{D}^{-1/2}\mathbf{u}_1, \ldots, \boldsymbol{D}^{-1/2}\mathbf{u}_{k-1}$ are step-vectors over the same proper k-partition of the vertices. Some examples:

- If $k = 1$, then the unit-norm eigenvector corresponding to $\mu_0 = 1$ is $\mathbf{u}_0 = \sqrt{\mathbf{d}}$, and $\boldsymbol{D}^{-1/2}\mathbf{u}_0 = \mathbf{1}$ is the all 1's vector. Consequently, the variance of its coordinates is $\sigma_1 = 0$. But in this case, by Proposition 14, we already know that $\mathrm{disc}(G)$ can be estimated from above merely by $|\mu_1| = s_1$.

- If $k = 2$ and G is bipartite, then $\mu_1 = -1$, $s_1 = 1$, and σ_2^2, i.e., the 2-variance of the coordinates of the transformed eigenvector corresponding to μ_1 can be small if $|\mu_2|$ is separated from $|\mu_1| = 1$ (like the bipartite expanders of Alon [Alon]).

- Let $k = 2$ and G be *bipartite, biregular* on the independent vertex-subsets V_1, V_2. That is, all the edge-weights within V_1 or V_2 are zeros, and the 0-1 weights between vertices of V_1 and V_2 are such that $d_i = k_1$ if $i \in V_1$ and $d_i = k_2$ if $i \in V_2$ with the understanding that $|V_1|k_1 = |V_2|k_2$ (both are the total number of edges in G). It is easy to see that the unit-norm eigenvector corresponding to the eigenvalue $\mu_1 = -1$ is $\mathbf{u}_1 = \boldsymbol{D}^{1/2}\mathbf{1}_{V_1} - \boldsymbol{D}^{1/2}\mathbf{1}_{V_2}$, and $\boldsymbol{D}^{-1/2}\mathbf{u}_1 = \mathbf{1}_{V_1} - \mathbf{1}_{V_2}$. Therefore, the representatives of vertices of V_1 are all 1's, and those of V_2 are -1's, so $\sigma_2^2 = 0$. Consequently, $\mathrm{md}_2(G) \leq B|\mu_2|$, with some absolute constant B. Up to the constant, this was another proof of Lemma 3.2 of Evra et al. [Ev-Go-Lu]. They call their result expander mixing lemma for bipartite graphs, and use cardinalities instead of volumes, but in this special case, these cardinalities are proportional to the volumes both within V_1 and V_2.

- Let G_n be a generalized random graph (see Definition 21) over the symmetric $k \times k$ probability matrix $\boldsymbol{P} = (p_{ab})$, i.e., there is a proper k-partition, $V_1, \ldots, V_k$, of its vertices such that $|V_a| = n_a$ $(a = 1, \ldots, k)$, $\sum_{a=1}^{k} n_a = n$, and for any $1 \leq a \leq b \leq k$, vertices $i \in V_a$ and $j \in V_b$ are connected independently, with the same probability p_{ab}. This is the k-cluster generalization of the classical Erdős–Rényi random graph, see also [Lov-Sos] for their generalized quasirandom counterparts. In [Bol08a] (see also Chapter 3) we characterize the adjacency and normalized Laplacian spectra of such graphs, that extends to their normalized modularity spectra as follows: both $|\mu_k| = s_k$ and σ_k^2 tend to zero almost surely when $n \to \infty$, under some balancing conditions for the cluster sizes ($\frac{n_a}{n} \geq c$ with some constant c, for $a = 1, \ldots, k$). By Theorem 27, it also holds for the k-way discrepancy in the clustering $V_1, \ldots, V_k$. However, this is not surprising, since this almost sure limit for the k-way discrepancy is easily obtained with large deviation principles too, see [Bol05] and the considerations made at the end of Section 2.1.2.

Summarizing, in the $k = 1$ case: when the second singular value $|\mu_1| = s_1$ is small (much smaller than $s_0 = 1$), then the overall discrepancy is small. However, for $k > 1$, a small s_k is necessary but not sufficient for a small k-way discrepancy. In addition, the weighted k-variance σ_k^2 should be small too. With subspace perturbation theorems, it is small if s_k is much smaller than s_{k-1}. Hence, a gap in the normalized modularity spectrum may be an indication for the number of clusters. The two directions together may give a hint about the optimal choice of k if a practitioner wants to find a k-clustering of the rows and columns (or just of the vertices of a graph) with small pairwise discrepancies. If there does not exist a fairly 'small' k with this property, then in the worst case scenario, the

Szemerédi Regularity Lemma [Szem] with an enormously large number of clusters (which number only depends on the maximum pairwise discrepancy to be attained, and does not depend on n) comes into existence. Weak versions of this lemma (where $V_1, \dots, V_k$ are not necessarily equitable) are also available, see e.g., [Borgsetal1, Gh-Trev, Lov-Szeg07]. Eventually, note that B. Szegedy [Szeg] and T. Tao (in his blog https://terrytao.wordpress.com/2012/12/03/the-spectral-proof-of-the-szemeredi-regularity-lemma/) also use the normalized adjacency eigenvalues in their decreasing absolute values to give a matrix proof of the Szemerédi Regularity Lemma.

2.3.4 Multiway discrepancy of directed graphs

A directed edge-weighted graph $G = (V, \boldsymbol{W})$ is given by its quadratic, but usually not symmetric edge-weight matrix $\boldsymbol{W} = (w_{ij})$ of zero diagonal, where w_{ij} is the nonnegative weight of the $i \to j$ edge ($i \neq j$). The row-sums $d_{out,i} = \sum_{j=1}^{n} w_{ij}$ and column-sums $d_{in,j} = \sum_{i=1}^{n} w_{ij}$ of $\boldsymbol{W}$ are the *out- and in-degrees*, while $\boldsymbol{D}_{out} = \mathrm{diag}(d_{out,1}, \dots, d_{out,n})$ and $\boldsymbol{D}_{in} = \mathrm{diag}(d_{in,1}, \dots, d_{in,n})$ are the diagonal *out- and in-degree matrices*, respectively. Now Definition 25 can be formulated as follows.

Definition 29 *The multiway discrepancy of the directed, weighted graph $G = (V, \boldsymbol{W})$ in the in-clustering $(V_{in,1}, \dots, V_{in,k})$ and out-clustering $(V_{out,1}, \dots, V_{out,k})$ of its vertices is*

$$\begin{aligned}&\mathrm{md}(G; V_{in,1}, \dots, V_{in,k}, V_{out,1}, \dots, V_{out,k})\\ &= \max_{\substack{1 \le a,b \le k\\ X \subset V_{out,a},\, Y \subset V_{in,b}}} \frac{|w(X,Y) - \rho(V_{out,a}, V_{in,b})\mathrm{Vol}_{out}(X)\mathrm{Vol}_{in}(Y)|}{\sqrt{\mathrm{Vol}_{out}(X)\mathrm{Vol}_{in}(Y)}},\end{aligned}$$

where $w(X,Y)$ is the sum of the weights of the $X \to Y$ edges, whereas $\mathrm{Vol}_{out}(X) = \sum_{i \in X} d_{out,i}$ and $\mathrm{Vol}_{in}(Y) = \sum_{j \in Y} d_{in,j}$ are the out- and in-volumes, respectively. The minimum k-way discrepancy of the directed edge-weighted graph $G = (V, \boldsymbol{W})$ is

$$\mathrm{md}_k(G) = \min_{\substack{(V_{in,1}, \dots, V_{in,k})\\ (V_{out,1}, \dots, V_{out,k})}} \mathrm{md}(G; V_{in,1}, \dots, V_{in,k}, V_{out,1}, \dots, V_{out,k}).$$

Butler [But1] and Chung and Kenter [Chu-K] treat the $k = 1$ case, and for a general k, Theorem 24 implies the following.

Theorem 28 ([Bol16]) *Let $G = (V, \boldsymbol{W})$ be directed edge-weighted graph, $\boldsymbol{W}$ is non-degenerate. Then for any integer $1 \le k \le \mathrm{rank}(\boldsymbol{W})$,*

$$s_k \le 9\mathrm{md}_k(G)(k + 2 - 9k \ln \mathrm{md}_k(G))$$

holds, provided $0 < \mathrm{md}_k(G) < 1$, where s_k is the k-th largest nontrivial singular value of the normalized edge-weight matrix $\boldsymbol{W}_D = \boldsymbol{D}_{out}^{-1/2} \boldsymbol{W} \boldsymbol{D}_{in}^{-1/2}$ of G.

Together with BSM students, in 2012 we applied this spectral method to find migration patterns in the set of 75 countries, and found 3 underlying immigration and emigration trait clusters. The SVD-based algorithm is the same as the one introduced in the construction of the proof of Theorem 23 for rectangular matrices of nonnegative entries.

Chapter 3

Theoretical applications and further perspectives

We applied the theory of convergent graph sequences (as described in Borgs and coauthors [Borgsetal1, Borgsetal2]) for vertex- and edge-weighted graphs and contingency tables. Roughly speaking, graphs and contingency tables of a convergent sequence become more and more similar to each other in small details, which fact is made exact in terms of the convergence of the homomorphism densities when 'small' simple graphs or binary tables are mapped into the 'large' graphs or tables constituting the sequence. The convergence can as well be formulated with the help of the cut-distance, and limit objects are defined. This cut-distance also makes it possible to classify graphs and contingency tables, or to assign them to given prototypes, providing classical problems of supervised or unsupervised learning, see e.g., Hastie et al. [Ha-Ti-Fri].

Testable graph and contingency table parameters are, in fact, nonparametric statistics of the sample (given in the form of graphs and contingency tables, briefly, networks) that can be consistently estimated based on a smaller sample, selected with an appropriate randomization from the underlying huge network. Indeed, real-world graphs or rectangular arrays of nonnegative entries are sometimes considered as samples from a large (partly unobservable and steadily changing) network, and we want to conclude the parameters of the network from the same parameters of its smaller parts. The theory guarantees that this can be done if the investigated parameter is testable. In [Bol10] we extended the notion of testable parameters to contingency tables. In [Bol-Ko-Kr12] we proved that certain balanced minimum multiway cut densities are testable and that the noisy graph sequences of Chapter 2 converge in the sense of [Borgsetal1] too. Since these results are routine adaptations of the equivalent convergence facts of [Borgsetal2] for weighted graphs, here we only present the results of [Bol14a] about the testability of the normalized modularity spectra and spectral subspaces.

Then we will consider generalized random and quasirandom graphs from the point of view of discrepancy and spectra, and – in the multicluster scenario – define so-called *generalized quasirandom properties* that large graphs may posses at the same time, irrespective of stochastic models.

Eventually, we discuss how the EM (Expectation–Maximization) algorithm can be used to estimate the parameters of the *homogeneous and heterogeneous block models*. These mixture models are semiparametric in the sense that we are looking for a convenient clustering of the vertices (starting with spectral clustering) and, at the same time, we want to estimate the

parameters of Erdős–Rényi or logit type models corresponding to the subgraphs or bipartite subgraphs induced by the clustering.

3.1 Convergent weighted graph sequences

3.1.1 Testability of the normalized modularity spectra and spectral subspaces

Borgs and his coauthors [Borgsetal1] defined the testability of simple graph parameters and proved equivalent notions of this testability. They also anticipated that their results remain valid if they consider weighted graph sequences (G_n) with edge-weights β_{ij}'s in the [0,1] interval and *no dominant vertex-weights* $\alpha_i(G_n) > 0$ $(i = 1, \dots, n)$ such that $\max_i \frac{\alpha_i(G_n)}{\alpha_{G_n}} \to 0$ as $n \to \infty$, where $\alpha_{G_n} = \sum_{i=1}^{n} \alpha_i(G_n)$. To this end, in [Bol-Ko-Kr12], we slightly modified the definition of a testable graph parameter for weighted graphs in the following way.

Definition 30 *A weighted graph parameter f is testable if for every $\varepsilon > 0$ there is a positive integer $m < n$ such that if G_n satisfies $\max_i \frac{\alpha_i(G_n)}{\alpha_{G_n}} \leq \frac{1}{m}$, then*

$$\mathbb{P}(|f(G_n) - f(\eta(m, G_n))| > \varepsilon) \leq \varepsilon,$$

where $\eta(m, G_n)$ is a simple graph on m vertices selected randomly from G_n in the following manner: m vertices of G_n are chosen with replacement, with respective probabilities proportional to the vertex-weights; given the m-element vertex-set, the edges (in this small graph) come into existence conditionally independently, with probabilities of the corresponding edge-weights (in the large graph).

By the above definition, a testable weighted graph parameter can consistently be estimated based on a fairly large sample. Analogously to the results of [Borgsetal1] for simple graphs, in [Bol-Ko-Kr12] we established equivalent statements of this testability, from among which here we use the following.

Proposition 15 ([Bol-Ko-Kr12]) *Let f be a testable weighted graph parameter. Then for every convergent weighted graph sequence (G_n), with no dominant vertex-weights, $f(G_n)$ is also convergent as $n \to \infty$.*

The notion of the convergence of a weighted graph sequence is defined in [Borgsetal1], where the authors also describe the limit object as a symmetric, measurable function $W : [0,1] \times [0,1] \to [0,1]$, called *graphon*. The so-called *cut distance* between the graphons W and U is $\delta_\square(W, U) = \inf_\nu \|W - U^\nu\|_\square$, where the cut norm of the graphon W is defined by

$$\|W\|_\square = \sup_{S,T \subset [0,1]} \left| \int_{S \times T} W(x, y)\, dx\, dy \right|,$$

and the above infimum is taken over all measure preserving bijections $\nu : [0,1] \to [0,1]$, while U^ν denotes the transformed U after performing the same measure preserving bijection ν on both sides of the unit square. Graphons are considered modulo measure preserving maps, and under graphon the whole equivalence class is understood. In this way, to a convergent weighted graph sequence (G_n), there is a unique limit graphon W such that $\delta_\square(G_n, W) \to 0$ as $n \to \infty$, where $\delta_\square(G_n, W)$ is defined as $\delta_\square(W_{G_n}, W)$ with the step-function graphon W_{G_n} assigned to G_n in the following way. The sides of the unit square are divided into

intervals $I_1, \dots, I_n$ of lengths $\alpha_1(G_n)/\alpha_{G_n}, \dots, \alpha_n(G_n)/\alpha_{G_n}$, and over the rectangle $I_i \times I_j$ the stepfunction takes on the value $\beta_{ij}(G_n)$.

In [Bol-Ko-Kr12], we proved the testability of some normalized and unnormalized balanced multiway cut densities such that we imposed balancing conditions on the cluster volumes. Under similar conditions, for fixed number of clusters k, the unnormalized and normalized multiway cuts and modularities are also testable, provided our edge-weighted graph has no dominant vertices. The proofs rely on statistical physics notions of [Borgsetal2], utilizing the fact that the graph convergence implies the convergence of the ground state energy (minimum of the energy function over the set of k-partitions of vertices). In [Re-Bo], the authors showed that the Newman–Girvan modularity is an energy function (Hamiltonian), and hence, testability of the maximum or minimum normalized modularities, under appropriate balancing conditions, can be shown analogously. Here we rather discuss the testability of spectra and k-variances, because, when using spectral clustering methods, these provide us with polynomial time algorithms, though only approximate solutions are obtained with this spectral relaxation technique.

In Theorem 6.6 of [Borgsetal2], the authors prove that the normalized spectrum of a convergent graph sequence converges in the following sense. Let W be a graphon and (G_n) be a sequence of weighted graphs with uniformly bounded edge-weights, tending to W. For simplicity, we assume that $|V(G_n)| = n$. Let $|\lambda_{n,1}| \geq |\lambda_{n,2}| \geq \cdots \geq |\lambda_{n,n}|$ be the adjacency eigenvalues of G_n indexed by their decreasing absolute values, and let $\mu_{n,i} = \lambda_{n,i}/n$ $(i = 1, \dots, n)$ be the normalized eigenvalues. Further, let T_W be the $L^2[0,1] \to L^2[0,1]$ integral operator corresponding to W:

$$T_W f(x) = \int_0^1 W(x,y) f(y)\, dy.$$

It is well-known that this operator is self-adjoint and compact, and hence, it has a discrete real spectrum, whose only possible point of accumulation is the 0. Let $\mu_i(W)$ denote the i-th largest absolute value eigenvalue of T_W. Then for every $i \geq 1$, $\mu_{n,i} \to \mu_i(W)$ as $n \to \infty$. In fact, the authors prove a bit more (see Theorem 6.7 of [Borgsetal2]): if a sequence W_n of uniformly bounded graphons converges to a graphon W, then for every $i \geq 1$, $\mu_i(W_n) \to \mu_i(W)$ as $n \to \infty$. Note that the spectrum of W_G is the normalized spectrum of G, together with countably infinitely many 0's. Therefore, the convergence of the spectrum of (G_n) is the consequence of that of (W_{G_n}).

We will prove that in the absence of dominant vertices, the normalized modularity spectrum is testable. To this end, both the modularity matrix and the graphon are related to kernels of special integral operators, described in Section 1.3, and briefly revisited as follows. Let (ξ, ξ') be a pair of identically distributed real-valued random variables defined over the product space $\mathcal{X} \times \mathcal{X}$ having a symmetric joint distribution $\mathbb{W}$ with the same marginals $\mathbb{P}$. Assume that the dependence between ξ and ξ' is regular, i.e., their joint distribution $\mathbb{W}$ is absolutely continuous with respect to the product measure $\mathbb{P} \times \mathbb{P}$, and let w denote its Radon–Nikodym derivative, see Rényi [Reny59b]. Let $H = L^2(\xi)$ and $H' = L^2(\xi')$ be the Hilbert spaces of random variables which are functions of ξ and ξ' and have zero expectation and finite variance with respect to $\mathbb{P}$. Recall that H and H' are isomorphic Hilbert spaces with the covariance as inner product; further, they are embedded as subspaces into the L^2-space defined similarly over the product space. (Note that here H and H' are also isomorphic in the sense that for any $\psi \in H$ there exists a $\psi' \in H'$ and vice versa, such that ψ and ψ' are identically distributed.)

Consider the linear operator taking conditional expectation between H' and H with respect to the joint distribution. It is an integral operator and will be denoted by $P_{\mathbb{W}} : H' \to$

H as it is a projection restricted to H' and projects onto H. To $\psi' \in H'$, the operator $P_{\mathbb{W}}$ assigns $\psi \in H$ such that $\psi = \mathbb{E}_{\mathbb{W}}(\psi' \,|\, \xi)$, i.e.,

$$\psi(x) = \int_{\mathcal{Y}} w(x,y)\psi'(y)\,\mathbb{P}(dy), \quad x \in \mathcal{X},$$

so $P_{\mathbb{W}}$ is an integral operator with kernel w. If

$$\int_{\mathcal{X}}\int_{\mathcal{X}} w^2(x,y)\mathbb{P}(dx)\mathbb{P}(dy) < \infty,$$

then $P_{\mathbb{W}}$ is a Hilbert–Schmidt operator, therefore compact and has spectral decomposition

$$P_{\mathbb{W}} = \sum_{i=1}^{\infty} \lambda_i \langle ., \psi_i' \rangle_{H'} \psi_i,$$

where for the eigenvalues $|\lambda_i| \leq 1$ holds. The eigenvalue–eigenfunction equation looks like

$$P_{\mathbb{W}}\psi_i' = \lambda_i \psi_i \quad (i = 1, 2, \dots),$$

where ψ_i and ψ_i' are identically distributed, whereas their joint distribution is $\mathbb{W}$. In Section 1.3 we saw that $P_{\mathbb{W}}$ is self-adjoint and it takes the constantly 1 random variable of H' into the constantly 1 random variable of H; however, the $\psi_0 = 1, \psi_0' = 1$ pair is not considered as a function pair with eigenvalue $\lambda_0 = 1$, since they have no zero expectation.

In the next theorem, $\boldsymbol{W}_n$ denotes the edge-weight matrix of an edge-weighted graph, and not a Wigner-noise.

Theorem 29 ([Bol14a]) *Let $G_n = (V_n, \boldsymbol{W}_n)$ be the general entry of a convergent sequence of connected edge-weighted graphs whose edge-weights are in [0,1] and the vertex-weights are the generalized degrees. Assume that there are no dominant vertices. Let W denote the limit graphon of the sequence (G_n), and let*

$$|\mu_{n,1}| \geq |\mu_{n,2}| \geq \cdots \geq |\mu_{n,n}| = 0$$

be the normalized modularity spectrum of G_n (the eigenvalues are indexed by their decreasing absolute values). Further, let $\mu_i(P_{\mathbb{W}})$ be the i-th largest absolute value eigenvalue of the integral operator $P_{\mathbb{W}} : L^2(\xi') \to L^2(\xi)$ taking conditional expectation with respect to the joint measure $\mathbb{W}$ embodied by the normalized limit graphon W, and ξ, ξ' are identically distributed random variables with the marginal distribution of their symmetric joint distribution $\mathbb{W}$. Then for every $i \geq 1$,

$$\mu_{n,i} \to \mu_i(P_{\mathbb{W}}) \quad \textit{as} \quad n \to \infty.$$

Proof. In case of a finite $\mathcal{X}$ (vertex set) we have an edge-weighted graph, and we will show that the spectrum of the operator taking conditional expectation with respect to the joint distribution, determined by the edge-weights, corresponds to the spectrum of its normalized modularity matrix.

Indeed, let $\mathcal{X} = V$, $|V| = n$, and $G_n = (V, \boldsymbol{W})$ be an edge-weighted graph on the $n \times n$ weight matrix of the edges $\boldsymbol{W}$ with entries W_{ij}'s; now, they do not necessarily sum to 1. (For the time being, n is kept fixed, so – for the sake of simplicity – we do not denote the dependence of $\boldsymbol{W}$ on n). Let the vertices be also weighted with special weights $\alpha_i(G_n) := \sum_{j=1}^n W_{ij}$, $i = 1, \dots, n$. Then the step-function graphon W_{G_n} is such that $W_{G_n}(x,y) = W_{ij}$ whenever $x \in I_i$ and $y \in I_j$, where the (not necessarily contiguous) intervals $I_1, \dots, I_n$ form a partition of [0,1] such that the length of I_i is $\alpha_i(G_n)/\alpha_{G_n}$ $(i = 1, \dots, n)$.

Let us transform $\boldsymbol{W}$ into a symmetric joint distribution $\mathbb{W}_n$ over $V \times V$. The entries $w_{ij} = W_{ij}/\alpha_{G_n}$ $(i, j = 1, \ldots, n)$ embody this discrete joint distribution of random variables ξ and ξ' which are identically distributed with marginal distribution $d_1, \ldots, d_n$, where $d_i = \alpha_i(G_n)/\alpha_{G_n}$ $(i = 1, \ldots, n)$. With the previous notation, $H = L^2(\xi)$, $H' = L^2(\xi')$, and the operator $P_{\mathbb{W}_n} : H' \to H$ taking conditional expectation is an integral operator with now discrete kernel $K_{ij} = \frac{w_{ij}}{d_i d_j}$, see also Section 1.3. The fact that ψ, ψ' is an eigenfunction pair of $P_{\mathbb{W}_n}$ with eigenvalue λ means that

$$\frac{1}{d_i}\sum_{j=1}^{n} w_{ij}\psi'(j) = \sum_{j=1}^{n} \frac{w_{ij}}{d_i d_j}\psi'(j)d_j = \lambda\psi(i), \tag{3.1}$$

where $\psi(j) = \psi'(j)$ denotes the value of ψ or ψ' taken on with probability d_j (recall that ψ and ψ' are identically distributed with joint distribution $\mathbb{W}$). The above equation is equivalent to

$$\sum_{j=1}^{n} \frac{w_{ij}}{\sqrt{d_i}\sqrt{d_j}}\sqrt{d_j}\psi(j) = \lambda\sqrt{d_i}\psi(i),$$

therefore the vector of coordinates $\sqrt{d_i}\psi(i)$ $(i = 1, \ldots, n)$ is a unit-norm eigenvector of the normalized modularity matrix with eigenvalue λ (note that the normalized modularity spectrum does not depend on the scale of the edge-weights, it is the same whether we use W_{ij}'s or w_{ij}'s as edge-weights). Consequently, the eigenvalues of the conditional expectation operator are the same as the eigenvalues of the normalized modularity matrix, and the possible values taken on by the eigenfunctions of the conditional expectation operator are the same as the coordinates of the transformed eigenvectors of the normalized modularity matrix forming the column vectors of the matrix $\mathbf{X}^*$ of the optimal $(k-1)$-dimensional representatives, see Sections 1.1.3 and 1.1.5.

Let f be a stepwise constant function on [0,1], taking on value $\psi(i)$ on I_i. Then $\mathrm{Var}(\psi) = 1$ is equivalent to $\int_0^1 f^2(x)\,dx = 1$. Let K_{G_n} be the stepwise constant graphon defined as $K_{G_n}(x, y) = K_{ij}$ for $x \in I_i$ and $y \in I_j$. With this, the eigenvalue–eigenvector equation (3.1) looks like

$$\lambda f(x) = \int_0^1 K_{G_n}(x, y)f(y)\,dy.$$

The spectrum of K_{G_n} is the normalized modularity spectrum of G_n together with countably infinitely many 0's (it is of finite rank, and therefore, trivially compact), and because of the convergence of the weighted graph sequence G_n, in lack of dominant vertices, the sequence of graphons K_{G_n} also converges. Indeed, the $W_{G_n} \to W$ convergence in the cut metric means the weak convergence of the induced discrete measures $\mathbb{W}_n$'s to the continuous $\mathbb{W}$. Since K_{G_n} and K are so-called copula transformations of these distributions, in lack of dominant vertices (this causes the convergence of the marginals) the convergence in distribution also holds, which in turn implies the $K_{G_n} \to K$ convergence in the cut metric.

The limit graphon K of (K_{G_n}) is the kernel of the integral operator taking conditional expectation with respect to the joint distribution $\mathbb{W}$. It is easy to see that this operator is also a Hilbert–Schmidt operator, and therefore, compact. With these considerations, the remainder of the proof is analogous to the proof of Theorem 6.7 of [Borgsetal2], where the authors prove that if the sequence (W_{G_n}) of graphons converges to the limit graphon W, then both ends of the spectra of the integral operators, induced by W_{G_n}'s as kernels, converge to the ends of the spectrum of the integral operator induced by W as kernel. We apply this argument for the spectra of the integral operators induced by the kernels K_{G_n}'s and K. This finishes the proof. □

Note that in Lovász–Szegedy [Lov-Szeg11], kernel operators are also discussed, but not with this normalization.

By Proposition 15, provided there are no dominant vertices, Theorem 29 implies that for any fixed positive integer k, the $(k-1)$-tuple of the largest absolute value eigenvalues of the normalized modularity matrix is testable.

Theorem 30 ([Bol14a]) *Assume that there are constants $0 < \varepsilon < \delta \leq 1$ such that the normalized modularity spectrum (with decreasing absolute values) of any G_n satisfies*

$$|\mu_{n,1}| \geq \cdots \geq |\mu_{n,k-1}| \geq \delta > \varepsilon \geq |\mu_{n,k}| \geq \cdots \geq |\mu_{n,n}| = 0.$$

With the notions of Theorem 29, and assuming that there are no dominant vertices of G_n's, the subspace spanned by the transformed eigenvectors $\boldsymbol{D}_n^{-1/2}\mathbf{u}_{n,1}, \ldots, \boldsymbol{D}_n^{-1/2}\mathbf{u}_{n,k-1}$ belonging to the $k-1$ largest absolute value eigenvalues of the normalized modularity matrix of G_n also converges to the corresponding $(k-1)$-dimensional subspace of $P_{\mathbb{W}}$. More precisely, if $\boldsymbol{P}_{n,k-1}$ denotes the projection onto the subspace spanned by the transformed eigenvectors belonging to $k-1$ largest absolute value eigenvalues of the normalized modularity matrix of G_n, and $\boldsymbol{P}_{k-1}$ denotes the projection onto the analogous eigen-subspace of $P_{\mathbb{W}}$, then $\|\boldsymbol{P}_{n,k-1} - \boldsymbol{P}_{k-1}\| \to 0$ as $n \to \infty$ (in spectral norm).

Proof. If we apply the convergence fact $\mu_{n,i} \to \mu_i(P_{\mathbb{W}})$ for indices $i = k-1$ and k, we get that there will be a gap of order $\delta - \varepsilon - o(1)$ between $|\mu_{k-1}(P_{\mathbb{W}})|$ and $|\mu_k(P_{\mathbb{W}})|$ too.

Let $P_{\mathbb{W},n}$ denote the n-rank approximation of $P_{\mathbb{W}}$ (keeping its n largest absolute value eigenvalues, together with the corresponding eigenfunctions) in spectral norm. The projection $\boldsymbol{P}_{k-1}$ $(k < n)$ operates on the eigen-subspace spanned by the eigenfunctions belonging to the $k-1$ largest absolute value eigenvalues of $P_{\mathbb{W},n}$ in the same way as on the corresponding $(k-1)$-dimensional subspace determined by $P_{\mathbb{W}}$. With these considerations, we apply the perturbation theory of eigen-subspaces with the following unitary invariant norm: the trace- or Schatten-norm of the Hilbert–Schmidt operator A is $\|A\|_{\mathrm{tr}} = (\sum_{i=1}^{\infty} \lambda_i^4(A))^{1/4}$. Our argument with the finite $(k-1)$-rank projections is the following. Denoting by $P_{\mathbb{W}_n}$ the integral operator belonging to the normalized modularity matrix of G_n (with kernel K_{G_n} introduced in the proof of Theorem 29),

$$\begin{aligned}\|\boldsymbol{P}_{n,k-1} - \boldsymbol{P}_{k-1}\| = \|\boldsymbol{P}_{n,k-1}^{\perp}\boldsymbol{P}_{k-1}\| \leq \|\boldsymbol{P}_{n,k-1}^{\perp}\boldsymbol{P}_{k-1}\|_{\mathrm{tr}} \\ \leq \frac{c}{\delta - \varepsilon - o(1)}\|P_{\mathbb{W}_n} - P_{\mathbb{W},n}\|_{\mathrm{tr}}\end{aligned}$$

with constant c that is at most $\pi/2$ (see Lemma 8). But

$$\|P_{\mathbb{W}_n} - P_{\mathbb{W},n}\|_{\mathrm{tr}} \leq \|P_{\mathbb{W}_n} - P_{\mathbb{W}}\|_{\mathrm{tr}} + \|P_{\mathbb{W}} - P_{\mathbb{W},n}\|_{\mathrm{tr}}$$

where the last term tends to 0 as $n \to \infty$, since the tail of the spectrum (taking the fourth power of the eigenvalues) of a Hilbert–Schmidt operator converges. For the convergence of the first term we use Lemma 7.1 of [Borgsetal1], which states that the trace-norm of an integral operator can be estimated from above by four times the cut norm of the corresponding kernel. But the convergence in the cut distance of the corresponding kernels to zero follows from the considerations made in the proof of Theorem 29. This finishes the proof. □

As the weighted k-variance depends continuously on the above subspaces (see the the proof of Theorem 27), Theorem 30 implies the *testability of the weighted k-variance* as well.

The above results suggest that in the absence of dominant vertices, even the normalized modularity matrix of a smaller part of the underlying weighted graph, selected at random

with an appropriate procedure, is able to reveal its cluster structure. Hence, the gain regarding the computational time of this spectral clustering algorithm is twofold: we only use a smaller part of the graph and the spectral decomposition of its normalized modularity matrix runs in polynomial time in the reduced number of the vertices. Under the vertex- and cluster-balance conditions, this method can give quite good approximations for the multiway cuts and helps us to find the number of clusters and identify the cluster structure. In addition, taking into account both the positive and negative, large absolute value eigenvalues together with eigenvectors, regular cuts can also be detected, as the investigated spectral characteristics give good estimates for the volume regularity's constant of the cluster pairs by Theorem 27. Such regular cuts are looked for in social or biological networks, e.g., if we want to find equally functioning synopses of the brain.

3.1.2 Noisy graph sequences

Now, we use the above theory for perturbations, showing that special noisy weighted graph sequences of Chapter 2 converge in the sense of the homomorphism densities too. We consider edge- and vertex-weighted graphs, where the edge-weights are nonnegative, and normalized so that they are in the [0,1] interval. Recall that if every edge-weight is positive, then our graph is called soft-core. The vertex-weights are always positive. Let $\mathcal{G}$ denote the set of all such edge- and vertex-weighted graphs. If the vertex-weights of $G \in \mathcal{G}$ on n vertices are the real numbers $\alpha_1, \dots, \alpha_n > 0$, then the *volume* of G is defined by $\alpha_G = \sum_{i=1}^{n} \alpha_i$, while that of the vertex-subset S by $\alpha_S = \sum_{i \in S} \alpha_i$. Note that this notion of volume coincides with that of Chapter 1 only if the vertex-weights are the generalized degrees.

From now on, the vertex-weights are equal (say, equal to 1), and a weighted graph G on n vertices is determined by its $n \times n$ symmetric weighted adjacency matrix $\boldsymbol{A}$. Let $G_{\boldsymbol{A}}$ denote the weighted graph with unit vertex-weights and edge-weights that are entries of $\boldsymbol{A}$. We will use the notion of a Wigner-noise (Definition 16) and blown-up matrix (Definition 17).

Let us fix the $k \times k$ symmetric probability matrix $\boldsymbol{P}$ of entries $0 < p_{ij} < 1$, and blow it up to an $n \times n$ blown-up matrix $\boldsymbol{B}_n$ of blow-up sizes $n_1, \dots, n_k$ (note that $G_{\boldsymbol{B}_n}$ is a soft-core graph). Consider the noisy matrix $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{W}_n$ as $n_1, \dots, n_k \to \infty$ at the same rate, where $\boldsymbol{W}_n$ is an $n \times n$ Wigner-noise. While perturbing $\boldsymbol{B}_n$ by $\boldsymbol{W}_n$, assume that for the uniform bound of the entries of $\boldsymbol{W}_n$ the condition (2.2) is satisfied. In this way, the entries of $\boldsymbol{A}_n$ are in the [0,1] interval, and hence, $G_{\boldsymbol{A}_n} \in \mathcal{G}$. We remark that $G_{\boldsymbol{W}_n} \notin \mathcal{G}$, but $W_{G_{\boldsymbol{W}_n}}$ is a bounded graphon, and the theory applies to it. In Section 2.1 we showed that by adding an appropriate Wigner-noise to $\boldsymbol{B}_n$, we can achieve that $\boldsymbol{A}_n$ becomes a 0-1 matrix: its entries are equal to 1 with probability p_{ij} and 0 otherwise within the block of size $n_i \times n_j$ (after rearranging its rows and columns). In this case, the corresponding noisy graph $G_{\boldsymbol{A}_n}$ is a generalized random graph of Definition 21.

By routine large deviation techniques, in [Bol-Ko-Kr12] we were able to prove that the cut-norm of the stepfunction graphon, assigned to a Wigner-noise, tends to zero with probability 1 as $n \to \infty$.

Proposition 16 ([Bol-Ko-Kr12]) *For any sequence $\boldsymbol{W}_n$ of Wigner-noises*

$$\lim_{n \to \infty} \|W_{G_{\boldsymbol{W}_n}}\|_{\square} = 0$$

almost surely.

The main idea of the proof is that the definition of the cut-norm of a stepfunction graphon

and formulas (7.2), (7.3) of [Borgsetal1] yield

$$\|W_{G_{\boldsymbol{W}_n}}\|_\square = \frac{1}{n^2}\max_{U,T\subset[n]}\left|\sum_{i\in U}\sum_{j\in T} w_{ij}\right| \le 6\max_{U\subset[n]}\frac{1}{n^2}\left|\sum_{i\in U}\sum_{j\in[n]\setminus U} w_{ij}\right|,$$

where the entries behind the latter double summation are independent random variables. Hence, the Azuma's inequality is applicable, and the statement follows by the Borel–Cantelli lemma. Then Proposition 16 implies the following.

Proposition 17 ([Bol-Ko-Kr12]) *Let $\boldsymbol{A}_n := \boldsymbol{B}_n + \boldsymbol{W}_n$ be the edge-weight matrix of the general term of the noisy graph sequence (G_{A_n}), where $\boldsymbol{B}_n$ is the blown-up matrix of the $k\times k$ symmetric probability matrix $\boldsymbol{P} = (p_{ij})$ with blow-up sizes $n_1,\dots,n_k \to \infty$ such that that $\lim_{n\to\infty}\frac{n_i}{n} = r_i$ $(i = 1,\dots,k)$, $n = \sum_{i=1}^k n_i$; further, for the uniform bound K of the entries of the Wigner-noise $\boldsymbol{W}_n$ the condition (2.2) is assumed. Under these conditions, $(G_{\boldsymbol{A}_n}) \subset \mathcal{G}$ converges almost surely in the $\delta_\square$ metric. The almost sure limit is the stepfunction graphon W_H, where the vertex- and edge-weights of the weighted graph H are*

$$\alpha_i(H) = r_i \quad (i = 1,\dots,k), \qquad \beta_{ij}(H) = p_{ij} \quad (i,j = 1,\dots,k).$$

3.2 Generalized quasirandom properties of expanding graph sequences

With an appropriate Wigner-noise, the noisy graph sequence $(G_{\boldsymbol{A}_n})$ of Proposition 17 becomes a generalized graph sequence on the model graph H, defined in Definition 21 and characterized from the point of view of discrepancy and spectra. Then we will establish so-called generalized quasirandom properties of expanding graph sequences, which are closely related to the properties of the generalized random graphs, and state implications between them, irrespective of the stochastic model.

3.2.1 Generalized random and quasirandom graphs

For the generalied random graphs, defined in Definition 21, the following can be proved (with subspace perturbation theorems and large deviations), see [Bol13], Chapter 2, and [Bol-El16].

Proposition 18 *Let $G_n(\boldsymbol{P},\mathcal{P}_k)$ be a generalized random graph on n vertices with vertex-classes $\mathcal{P}_k = (C_1,\dots,C_k)$ of sizes $n_1,\dots n_k$ and $k\times k$ symmetric probability matrix $\boldsymbol{P}$. Let k be a fixed positive integer and $n\to\infty$ in such a way that $\frac{n_u}{n} \ge c$ $(u = 1,\dots,k)$ with some constant $0 < c \le \frac{1}{k}$ (called balancing condition). Then the following hold almost surely for the adjacency matrix $\boldsymbol{A}_n = (a_{ij}^{(n)})$ and the normalized modularity matrix $\boldsymbol{M}_{D,n}$ of $G_n(\boldsymbol{P},\mathcal{P}_k)$.*

1. *$\boldsymbol{A}_n$ has k structural eigenvalues that are $\Theta(n)$ in absolute value, while the remaining eigenvalues are $O(\sqrt{n})$. Further, the k-variance $S^2_{k,n}$ of the k-dimensional vertex representatives, based on the eigenvectors corresponding to the structural eigenvalues of $\boldsymbol{A}_n$ (see (1.10)), is $\mathcal{O}(\frac{1}{n})$.*
2. *There exists a positive constant $0 < \delta < 1$ independent of n (it only depends on k) such that $\boldsymbol{M}_{D,n}$ has exactly $k-1$ structural eigenvalues of absolute value greater than δ, while all the other eigenvalues are less than $n^{-\tau}$ in absolute value, for every $0 < \tau < \frac{1}{2}$. Further, the weighted k-variance $\tilde{S}^2_{k,n}$ of the $(k-1)$-dimensional vertex representatives, based on the transformed eigenvectors corresponding to the structural eigenvalues of $\boldsymbol{M}_{D,n}$ (see (1.14)), is $\mathcal{O}(n^{-2\tau})$, for every $0 < \tau < \frac{1}{2}$.*

3. *There exists a constant* $0 < \theta < 1$ *independent of* n *(it only depends on* k*) such that* $\mathrm{disc}_1(G_n(\boldsymbol{P}, \mathcal{P}_k)) > \theta$, $\ldots$, $\mathrm{disc}_{k-1}(G_n(\boldsymbol{P}, \mathcal{P}_k)) > \theta$, *and the* k*-way discrepancy* $\mathrm{disc}_k(G_n(\boldsymbol{P}, \mathcal{P}_k); C_1, \ldots, C_k)$ *is* $\mathcal{O}(n^{-\tau})$, *for every* $0 < \tau < \frac{1}{2}$.

4. *For every* $1 \le u \le v \le k$ *and* $i \in C_u$*:*

$$\sum_{j \in C_v} a_{ij}^{(n)} = p_{uv} n_v + o(n).$$

For every $1 \le u \le v \le k$ *and* $i, j \in C_u$*:*

$$\sum_{t \in C_v} a_{it}^{(n)} a_{jt}^{(n)} = p_{uv}^2 n_v + o(n).$$

Proof. Property 1 follows from Theorems 13 and 14, while Property 2 from Theorems 15 and 16. Property 3 is the consequence of Theorems 26 and 27.

The proof of Property 4 is as follows. Consider the generalized random graph sequence $G_n(\boldsymbol{P}, \mathcal{P}_k)$, the subgraphs and the bipartite subgraphs of which have the following expected degrees. We will drop the index n, and use the notation $\boldsymbol{A} = (a_{ij})$ for the entries of its adjacency matrix. As for the C_u, C_v pair $(1 \le u \le v \le k)$, for any $i \in C_u$, the average degree of i with regard to C_v is

$$\mathbb{E}(\sum_{j \in C_v} a_{ij}) = n_v p_{uv},$$

each vertex in C_u has the same expected number of neighbors in C_v.

Observe that for $i \in C_u$, the sum $\sum_{j \in C_v} a_{ij}$ has binomial distribution with the above expectation and variance $n_v p_{uv}(1 - p_{uv})$. Therefore, by Lemma 2, the within- and between-cluster average degrees are highly concentrated on their expectations as $n \to \infty$ under the balancing conditions $\frac{n_u}{n} \ge c$ $(u = 1, \ldots, k)$ for the cluster sizes. Indeed, for any $0 < \varepsilon < 1$:

$$\mathbb{P}(|\frac{1}{n_v} \sum_{j \in C_v} a_{ij} - p_{uv}| > \varepsilon) = \mathbb{P}(|\sum_{j \in C_v} a_{ij} - n_v p_{uv}| > n_v \varepsilon) \le e^{-\frac{n_v^2 \varepsilon^2}{2(n_v p_{uv}(1-p_{uv}) + \varepsilon/3)}}$$

that tends to 0 even with the choice $\varepsilon = n^{-\tau}$, $0 < \tau < \frac{1}{2}$. Therefore, it holds almost surely that

$$|\sum_{j \in C_v} a_{ij} - n_v p_{uv}| \le n_v n^{-\tau} = \frac{n_v}{n} n^{1-\tau} = o(n).$$

This finishes the proof of the first part of 4.

As for every $1 \le u \le v \le k$, the number of common neighbors in C_v of any $i, j \in C_u$ $(i \ne j)$ pair has binomial distribution with expectation $n_v p_{uv}^2$ and variance $n_v p_{uv}^2(1 - p_{uv}^2)$, with the same calculations as above, we obtain that

$$|\sum_{t \in C_v} a_{it} a_{jt} - p_{uv}^2 n_v| = o(n)$$

holds almost surely. This finishes the proof of the second part of 4. □

Recall that Proposition 17 implies that $G_n(\boldsymbol{P}, \mathcal{P}_k) \to W_H$ almost surely when $n \to \infty$, under the strict balancing condition $\frac{n_i}{n} \to r_i$ $(i = 1, \ldots, k)$. In Lovász–Sós [Lov-Sos] the following definition of a generalized quasirandom graph sequence was given.

Definition 31 *Given a model graph H on k vertices with vertex-weights $r_1, \dots, r_k$ and edge-weights $p_{uv} = p_{vu}$, $1 \le u \le v \le k$ (entries of $\boldsymbol{P}$), the sequence (G_n) is H-quasirandom if $G_n \to W_H$ as $n \to \infty$ in term of the homomorphism densities.*

The authors of [Lov-Sos] also proved that the vertex set V of a generalized quasirandom graph G_n can be partitioned into classes $C_1, \dots, C_k$ in such a way that $\frac{|C_u|}{n} \to r_u$ $(u = 1, \dots, k)$ as $n \to \infty$, and the subgraph of G_n induced by C_u is the general term of a quasirandom graph sequence with edge-density tending to p_{uu} $(u = 1, \dots, k)$, whereas the bipartite subgraph between C_u and C_v is the general term of a quasirandom bipartite graph sequence with edge-density tending to p_{uv} $(u \neq v)$ as $n \to \infty$.

Because of the limit relation in the definition of the generalized quasirandom graphs, the properties, discussed in Proposition 18, are – with some modification – valid for them. Actually, the authors in [Borgsetal1] proved that for any k, the k largest absolute value normalized adjacency eigenvalues of a convergent graph sequence converge (to the corresponding eigenvalues of the limiting graphon). In Section 3.1.1, we proved the same for the normalized modularity spectra of convergent graph sequences, see Theorem 29. As for the multiway discrepancies and spectra, we can use Theorems 26 and 27.

However, the order $\sqrt{n}$ and $n^{-\tau}$ of the non-structural eigenvalues in the adjacency and normalized modularity spectrum, respectively, is not necessarily valid for the generalized quasirandom graphs; instead, $o(n)$ and $o(1)$ can be stated for their order. Indeed, in the case of a generalized random graph we can separate a Wigner-noise, the corresponding graphon to which tends to zero very quickly. The slower separation in the spectrum is supported by simulations and the following construction.

Vera T. Sós suggested the following construction of a generalized quasirandom graph with given k, $\boldsymbol{P}$, and vertex-weights $r_1, \dots, r_k$ of the model graph H. Consider the instance when there are k clusters $C_1, \dots, C_k$ of the vertices of sizes $n_1, \dots, n_k$ such that $\frac{n_u}{n} = r_u$ $(u = 1, \dots, k)$. Let us choose the independent irrational numbers α_{uv} $(1 \le u \le v \le k)$. Then the subgraph on the vertex-set C_u is constructed as follows: $i, j \in C_u$, $i < j$ are connected if and only if

$$\{(i-j)^2 \alpha_{uu}\} < p_{uu} \quad (u = 1, \dots, k),$$

where $\{.\}$ denotes the fractional part of a real number. The bipartite subgraph between C_u and C_v is constructed as follows: $i \in C_u$ and $j \in C_v$ are connected if and only if

$$\{(i-j)^2 \alpha_{uv}\} < p_{uv} \quad (1 \le u < v \le k).$$

In the $k = 1$ case, Bollobás and Erdős [Bo-Erd] recommended this construction, and Pinch [Pinch] in terms of the codegrees proved that it indeed produces a quasirandom graph. The analytical number theoretical considerations of [Bo, Kup-Nied] and particularly of [Pinch] imply that, for any $1 \le u \le v \le k$, the sequence

$$y_t := (\{(t-i)^2 \alpha_{uv}\}, \{(t-j)^2 \alpha_{uv}\})$$

is *well-distributed symmetrically* in $[0,1]^2$, uniformly in $i - j$. This means that the sequences (y_{t+h}) are uniformly distributed symmetrically in $[0,1]^2$ for $h \in \mathbb{Z}$. With $h = \sum_{\ell=1}^{v-1} n_\ell$ and the considerations of [Pinch], we get that

$$\begin{aligned} &\left|\{t \in C_v \,:\, \{(t-i)^2 \alpha_{uv}\} < p_{uv} \quad \text{and} \quad \{(t-j)^2 \alpha_{uv}\} < p_{uv}\}\right| \\ &= p_{uv}^2 n_v + o(n_v) = p_{uv}^2 n_v + o(n) \end{aligned}$$

for any $i, j \in C_u$ $(i \neq j)$ pair, when $n \to \infty$ and $\frac{n_u}{n} \to r_u$ $(u = 1, \dots, k)$. It is important that the role of i, j is symmetric here: both are in C_u and connected to any $t \in C_v$ with the

same rule. Possibly, it suffices to assume the weaker balancing condition to guarantee that $n_1, \dots, n_k$ tend to infinity at the same rate: $n \to \infty$ in such a way that $\frac{n_u}{n} \geq c$ $(u = 1, \dots, k)$, with some constant $0 < c \leq \frac{1}{k}$. This ensures that $o(n_1) = \cdots = o(n_k) = o(n)$.

For more examples of quasirandom graphs in the $k = 1$ case see [Bo, Thom87]. For an illustration of generalized random and quasirandom graphs see Figures 3.1, 3.2, 3.3, where we used the probability matrix

$$\boldsymbol{P} = \begin{pmatrix} 0.7 & 0.1 & 0.15 & 0.2 & 0.25 \\ 0.1 & 0.75 & 0.3 & 0.35 & 0.4 \\ 0.15 & 0.3 & 0.8 & 0.45 & 0.5 \\ 0.2 & 0.35 & 0.45 & 0.85 & 0.55 \\ 0.25 & 0.4 & 0.5 & 0.55 & 0.9 \end{pmatrix}.$$

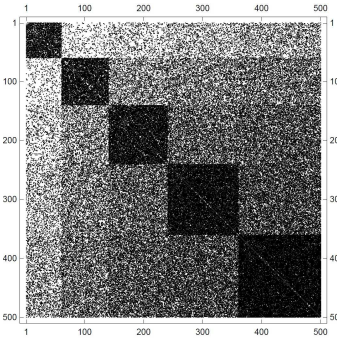

Figure 3.1: Generalized random graph generated with $k = 5$, cluster sizes $60, 80, 100, 120, 140$ and probability matrix $\boldsymbol{P}$. The first non-trivial eigenvalues of $\boldsymbol{M}_D$ are $0.304, 0.214, 0.17, 0.153,$ $-0.097, -0.094, -0.093,$ $-0.092, -0.091, \dots,$ with a gap after the 4th one.

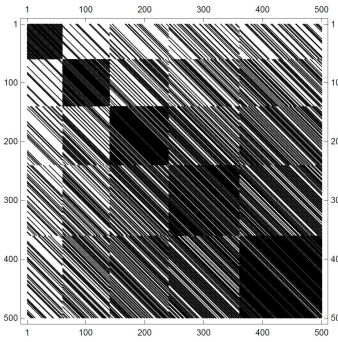

Figure 3.2: Generalized quasirandom graph constructed with $k = 5$, cluster sizes $60, 80, 100, 120, 140$ and probability matrix $\boldsymbol{P}$. The first non-trivial eigenvalues of $\boldsymbol{M}_D$ are $0.318, 0.207, 0.154, 0.115,$ $-0.100, -0.099, -0.091,$ $-0.090, 0.084, \dots,$ exhibiting decreasing eigenvalues up to the 4th one.

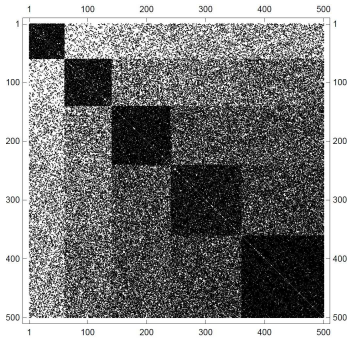

Figure 3.3: The former generalized quasirandom graph after appropriately permuting the vertices within the blocks (made by Ahmed Elbanna).

3.2.2 Generalized quasirandom properties

Some properties similar to those of Proposition 18 are now formulated for expanding deterministic graph sequences.

Conjecture 1 *Consider the sequence of graphs G_n with vertex-set V_n, adjacency matrix $\boldsymbol{A}_n = (a_{ij}^{(n)})$, and normalized modularity matrix $\boldsymbol{M}_{D,n}$. Let k be a fixed positive integer and $|V_n| = n \to \infty$ in such a way that there are no dominant vertices. Then the following properties are equivalent:*

P0. There exists a vertex- and edge-weighted graph H on k vertices such that $G_n \to W_H$ as $n \to \infty$ in terms of the homomorphism densities.

PI. $\boldsymbol{A}_n$ *has* k *structural eigenvalues* $\lambda_{1,n}, \dots, \lambda_{k,n}$ *such that the normalized eigenvalues converge:* $\frac{1}{n}|\lambda_{i,n}| \to q_i$ *as* $n \to \infty$ $(i = 1, \dots, k)$ *with some positive reals* $q_1, \dots, q_k$*, and the remaining eigenvalues are* $o(n)$ *in absolute value. The* k*-variance* $S^2_{k,n}$ *of the* k*-dimensional vertex representatives, based on the eigenvectors corresponding to the structural eigenvalues of* $\boldsymbol{A}_n$*, is* $o(1)$*.*

PII. There exists a constant $0 < \delta < 1$ *(independent of* n*) such that* $\boldsymbol{M}_{D,n}$ *has* $k-1$ *structural eigenvalues that are greater than* δ *in absolute value, while the remaining eigenvalues are* $o(1)$*. Further, the weighted* k*-variance* $\tilde{S}^2_{k,n}$ *of the* $(k-1)$*-dimensional vertex representatives, based on the transformed eigenvectors corresponding to the structural eigenvalues of* $\boldsymbol{M}_{D,n}$*, is* $o(1)$*.*

PIII. There are vertex-classes $\mathcal{P}_k = (C_1, \dots, C_k)$ *and a constant* $0 < \theta < 1$ *(independent of* n*) such that* $\mathrm{md}_1(G_n) > \theta, \dots, \mathrm{md}_{k-1}(G_n) > \theta$*, and* $\mathrm{md}_k(G_n; C_1, \dots, C_k) = o(1)$*.*

PIV. There are vertex-classes $\mathcal{P}_k = (C_1, \dots, C_k)$ *and a* $k \times k$ *symmetric probability matrix* $\boldsymbol{P} = (p_{uv})$ *(independent of* n*), such that every vertex of* C_u *has asymptotically* $n_v p_{uv}$ *neighbors in* C_v *for any* $1 \le u \le v \le k$ *pair. Further, for the codegrees (number of common neighbors) the following holds: every two different vertices* $i, j \in C_u$ *have asymptotically* $p^2_{uv} n_v$ *common neighbors in* C_v *for any* $1 \le u \le v \le k$ *pair. More exactly, for every* $1 \le u \le v \le k$ *and* $i, j \in C_u$

$$\sum_{t \in C_v} a^{(n)}_{it} = p_{uv} n_v + o(n)$$

and

$$\sum_{t \in C_v} a^{(n)}_{it} a^{(n)}_{jt} = p^2_{uv} n_v + o(n)$$

hold, where $n_v = |C_v|$*,* $v = 1, \dots, k$*.*

We will not prove the implications here, but have some remarks about the way how some implications follow from former results of others and theorems proved in the present dissertation; in particular, from Theorems 13, 14, 15, 16, 18, 26, and 27 of Chapter 2, and Theorems 29 and 30 of Chapter 3.

P0 is equivalent to PIV, due to the results of [Chu-G-W, Lov-Sos, Sim-Sos, Thom87, Thom89]. We can use that by [Lov-Sos], the vertex set of the generalized quasirandom graph G_n can be partitioned into classes $C_1, \dots, C_k$ in such a way that $\frac{|C_u|}{|V|} \to r_u$ $(u = 1, \dots, k)$ and the subgraph of G_n induced by C_u is the general term of a quasirandom graph sequence with edge-density tending to p_{uu} $(u = 1, \dots, k)$, whereas the bipartite subgraph between C_u and C_v is the general term of a quasirandom bipartite graph sequence with edge-density tending to p_{uv} $(u \ne v)$ as $n \to \infty$. The converse is trivial. Then, with some modification, theorems of Thomason [Thom87, Thom89] about (p, α)-jumbled graphs are applicable for the subgraphs and bipartite subgraphs, where p is some p_{uv} and α is related to the k-way discrepancy. Lovász [Lov08] also discusses that quasirandom graphs (as our subgraphs) are asymptotically regular, while, in view of Thomason [Thom89], the bipartite quasirandom graphs (as our bipartite subgraphs) are asymptotically biregular. According to Chung–Graham–Wilson [Chu-G-W], these properties are weaker than the other properties of quasirandomness, therefore asymptotic behavior of the codegrees should be characterized too, see [Bo-Erd, Lov-Sos, Thom87, Thom89].

As for the equivalence between P0 and PI, we can use that the convergence of a graph sequence implies the convergence of its normalized spectrum, see [Borgsetal1]. By the considerations of the proof of Proposition 5, we can relate the eigenvalues of the limiting graphon

W_H to the eigenvalues of a $k \times k$ matrix. Since the other eigenvalues of W_H are 0, it follows that $|\lambda_{i,n}| = o(n)$ $(i > k)$. As the spectral subspace corresponding to $\lambda_{1,n}, \dots, \lambda_{k,n}$ also converges to that of the step-vectors, $S^2_{k,n} = o(n)$ follows, see Theorem 14. In the backward direction, the convergence of the spectra usually does not imply the convergence of the graph or graphon sequence, but in the case of the quasirandom graphs it does as noted in [Borgsetal1]. Here we use both the separation in the spectrum and the convergence of the spectral subspaces, i.e., of the k-variances. By Theorem 18 and 23, we are able to find a blown-up matrix $\boldsymbol{B}_n$ of rank k and an error-matrix $\boldsymbol{E}_n$ with $\|\boldsymbol{E}_n\| = o(n)$ such that $\boldsymbol{A}_n = \boldsymbol{B}_n + \boldsymbol{E}_n$. It is important that, provided $S^2_{k,n}$ is 'small' enough, the so constructed $\boldsymbol{B}_n$ can have positive entries (see the considerations after the proof of Theorem 25), so that it can be the blown-up matrix of a $k \times k$ probability matrix.

In the equivalence between PI and PII we use that there are no dominant vertices, and the ideas of the proof of Theorem 15 and 16 extend to this case.

For the PII–PIII equivalence we plan to use the back and forth statements of Theorems 26 and 27. This equivalence suggest that low discrepancy clusters and cluster pairs can be obtained by spectral clustering tools, and so, justify the discrepancy minimizing spectral clustering.

Note that in the $k = 1$ case, the P0, PI, PIV properties are in accord with some of the properties of Chung–Graham–Wilson [Chu-G-W] and Simonovts–Sós [Sim-Sos], whereas PII, PIII rather harmonize with the properties of Chung–Graham [Chu-G] that do not contain a statement about the leading eigenvalue. As in the case of $k = 1$ there is only one leading eigenvalue, this does not make too much difference, but in the $k > 1$ case the statements should deal with their asymptotic behavior too, see [Bol15] for some details.

Summarizing, we believe that the P0–>PIV–>PIII–>PII–>PI–>P0 implications can be proved and so, they close the circle. In view of the above, PIV–>PIII is the only missing chain. However, we hope that the results of [Thom87, Thom89] can be adopted for the subgraphs and bipartite subgraphs to estimate the number of common neighbors by means of discrepancy. In fact, our discrepancy is a bit different of α and our $k \times k$ probability matrix $\boldsymbol{P}$ is a bit different of p ($k = 1$ case) of the notion of (p, α)-jumbledness.

3.3 Parameter estimation in probabilistic mixture models

Here we discuss semiarametric models such that the background distribution of the random graph, given the clusters of vertices, depends on some parameters, and we will use the EM algorithm to estimate them. We consider two basic types of block models: a homogeneous, and a heterogeneous one. In the literature, sometimes the multiclass model, consisting of Erdős–Rényi type homogeneous blocks, is also called heterogeneous. Under heterogeneity, we understand that – even within the building blocks – the entries heve different expectations, therefore our random graph is not a noisy one of Chapter 2.

3.3.1 EM algorithm for estimating the parameters of the homogeneous block model

The so-called stochastic block model, first introduced in Holland and coauthors [Hol-Las-Lei], later investigated by Bickel et al.[Bi-Ch] among others, was already discussed in Chapter 2 and Section 3.2.2. In fact, this is a generalized random graph model, but now it is formulated in terms of probabilistic mixture models. The assumptions of the model are the following.

Given a simple graph $G = (V, \boldsymbol{A})$ ($|V| = n$, with adjacency matrix $\boldsymbol{A}$) and integer k ($1 < k < n$), we are looking for the hidden k-partition $(C_1, \dots, C_k)$ of the vertices such that

- vertices are independently assigned to cluster C_a with probability π_a, $a = 1, \dots, k$; $\sum_{a=1}^{k} \pi_a = 1$;
- given the cluster memberships, vertices of C_a and C_b are connected independently, with probability

$$\mathbb{P}(i \sim j \,|\, i \in C_a, j \in C_b) = p_{ab}, \quad 1 \le a, b \le k.$$

The parameters are collected in the vector $\boldsymbol{\pi} = (\pi_1, \dots, \pi_k)$ and the $k \times k$ symmetric probability matrix $\boldsymbol{P}$ of p_{ab}'s. We will call this model homogeneous because of the constant within- and between-cluster edge probabilities. The classical Erdős–Rényi random graph is the special case of it with $k = 1$.

Our statistical sample is the $n \times n$ symmetric, 0-1 adjacency matrix $\boldsymbol{A} = (a_{ij})$ of G. There are no loops, so the diagonal entries are zeros. Based on $\boldsymbol{A}$, we want to estimate the parameters of the above block model. (It may seem that we have a one-element sample here, however, there are $\binom{n}{2}$ independent random variables, the adjacencies, in the background.) Using the theorem of mutually exclusive and exhaustive events, the likelihood function is constructed by the (within the blocks) i.i.d. Bernoulli distributed entries:

$$\begin{aligned}&\frac{1}{2} \sum_{1 \le a,b \le k} \pi_a \pi_b \prod_{i \in V_a, j \in V_b, i \ne j} p_{ab}^{a_{ij}} (1 - p_{ab})^{(1 - a_{ij})} \\ &= \frac{1}{2} \sum_{1 \le a,b \le k} \pi_a \pi_b \cdot p_{ab}^{e_{ab}} (1 - p_{ab})^{(n_{ab} - e_{ab})}.\end{aligned}$$

This is the mixture of binomial distributions, where e_{ab} is the number of edges connecting vertices of C_a and C_b ($a \ne b$), while e_{aa} is twice the number of edges with both endpoints in V_a; further,

$$n_{ab} = |C_a| \cdot |C_b| \quad (a \ne b) \quad \text{and} \quad n_{aa} = |C_a| \cdot (|C_a| - 1) \quad (a = 1, \dots, k)$$

are the numbers of possible edges between C_a, C_b and within C_a, respectively.

Here $\boldsymbol{A}$ is the incomplete data specification as the cluster memberships are missing. Therefore, it is straightforward to use the *Expectation-Maximization*, briefly EM algorithm, proposed by Dempster–Laird–Rubin [De-La-Ru]. This special application for mixtures is also related to the *collaborative filtering*, see e.g., [Ung].

First we complete our data matrix $\boldsymbol{A}$ with latent membership vectors $\boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_n$ of the vertices that are k-dimensional i.i.d. $Poly(1, \boldsymbol{\pi})$ (polynomially distributed) random vectors. More precisely, $\boldsymbol{\Delta}_i = (\Delta_{1i}, \dots, \Delta_{ki})$, where $\Delta_{ai} = 1$ if $i \in V_a$ and zero otherwise. Thus, the sum of the coordinates of any $\boldsymbol{\Delta}_i$ is 1, and $\mathbb{P}(\Delta_{ai} = 1) = \pi_a$.

Based on these, the likelihood function above is

$$\frac{1}{2} \sum_{1 \le a,b \le k} \pi_a \pi_b \cdot p_{ab}^{\sum_{i \ne j} \Delta_{ai} \Delta_{bj} a_{ij}} \cdot (1 - p_{ab})^{\sum_{i \ne j} \Delta_{ai} \Delta_{bj} (1 - a_{ij})}$$

that is maximized in the alternating **E** and **M** steps of the EM algorithm.

Note that that the complete likelihood would be the squareroot of

$$\prod_{a=1}^{k} \prod_{i=1}^{n} \prod_{b=1}^{k} [p_{ab}^{\sum_{j:\, j \ne i} \Delta_{bj} a_{ij}} \cdot (1 - p_{ab})^{\sum_{j:\, j \ne i} \Delta_{bj} (1 - a_{ij})}]^{\Delta_{ai}} \tag{3.2}$$

which is valid only in case of known cluster memberships.

Starting with initial parameter values $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{P}^{(0)}$ and membership vectors $\boldsymbol{\Delta}_1^{(0)}, \dots, \boldsymbol{\Delta}_n^{(0)}$, the t-th step of the iteration is the following ($t = 1, 2, \dots$).

- **E**-step: we calculate the conditional expectation of each Δ_i conditioned on the model parameters and on the other cluster assignments obtained in step $t-1$ and collectively denoted by $M^{(t-1)}$. By the Bayes theorem, the responsibility of vertex i for cluster a is
$$\pi_{ai}^{(t)} = \mathbb{E}(\Delta_{ai} \mid M^{(t-1)}) = \frac{\mathbb{P}(M^{(t-1)}|\Delta_{ai} = 1) \cdot \pi_a^{(t-1)}}{\sum_{b=1}^{k} \mathbb{P}(M^{(t-1)}|\Delta_{bi} = 1) \cdot \pi_b^{(t-1)}}$$
($a = 1, \dots, k$; $i = 1, \dots, n$). For each i, $\pi_{ai}^{(t)}$ is proportional to the numerator, where
$$\mathbb{P}(M^{(t-1)}|\Delta_{ai} = 1) = \prod_{b=1}^{k} (p_{ab}^{(t-1)})^{\sum_{j \neq i} \Delta_{bj}^{(t-1)} a_{ij}} \cdot (1 - p_{ab}^{(t-1)})^{\sum_{j \neq i} \Delta_{bj}^{(t-1)} (1 - a_{ij})}$$
is the part of the likelihood (3.2) effecting vertex i under the condition $\Delta_{ai} = 1$.

- **M**-step: we maximize the truncated binomial likelihood
$$p_{ab}^{\sum_{i \neq j} \pi_{ai}^{(t)} \pi_{bj}^{(t)} a_{ij}} \cdot (1 - p_{ab})^{\sum_{i \neq j} \pi_{ai}^{(t)} \pi_{bj}^{(t)} (1 - a_{ij})}$$
with respect to the parameter p_{ab}, for all a, b pairs separately. Obviously, the maximum is attained by the following estimators of p_{ab}'s comprising the symmetric matrix $\boldsymbol{P}^{(t)}$: $p_{ab}^{(t)} = \frac{\sum_{i,j:\, i \neq j} \pi_{ai}^{(t)} \pi_{bj}^{(t)} a_{ij}}{\sum_{i,j:\, i \neq j} \pi_{ai}^{(t)} \pi_{bj}^{(t)}}$ ($1 \leq a \leq b \leq k$), where edges connecting vertices of clusters a and b are counted fractionally, multiplied by the membership probabilities of their endpoints.

The maximum likelihood estimator of $\boldsymbol{\pi}$ in the t-th step is $\boldsymbol{\pi}^{(t)}$ of coordinates $\pi_a^{(t)} = \frac{1}{n} \sum_{i=1}^{n} \pi_{ai}^{(t)}$ ($a = 1, \dots, k$), while that of the membership vector $\boldsymbol{\Delta}_i$ is obtained by discrete maximization: $\Delta_{ai}^{(t)} = 1$ if $\pi_{ai}^{(t)} = \max_{b \in \{1,\dots,k\}} \pi_{bi}^{(t)}$ and 0, otherwise. (In case of ambiguity, the cluster with the smallest index is selected.) This choice of $\boldsymbol{\pi}$ will increase (better to say, not decrease) the likelihood function. Note that it is not necessary to assign vertices uniquely to the clusters, the responsibility π_{ai} of a vertex i can as well be considered as the intensity of vertex i belonging to cluster a, and it defines a fuzzy clustering.

According to the general theory of the EM algorithm [De-La-Ru], in exponential families (as in the present case), convergence to a local maximum can be guaranteed (depending on the starting values), but it runs in polynomial time in the number of vertices n. However, the speed and limit of the convergence depends on the starting clustering, which can be chosen by means of preliminary application of some spectral clustering methods of the previous chapters.

3.3.2 EM algorithm for estimating the parameters of the inhomogeneous block model

Loglinear type models to describe contingency tables were proposed, e.g., by [Hol-Lei, Laur] and widely used in statistics. Together with the Rasch model [Rasch], they give the foundation of our unweighted graph and bipartite graph models, the building blocks of our EM iteration.

With different parameterization, Chatterjee et al. [Ch-Dia-Sl] and V. Csiszár et al. [Csetal1] introduced the following random graph model, where the degree sequence is a sufficient statistic. We have an unweighted, undirected random graph on n vertices without loops, such that edges between distinct vertices come into existence independently, but not with the same probability as in the classical Erdős–Rényi model [Erd-Reny]. This random graph can uniquely be characterized by its $n \times n$ symmetric adjacency matrix $\boldsymbol{A} = (A_{ij})$ which has zero diagonal and the entries above the main diagonal are independent Bernoulli random variables whose parameters $p_{ij} = \mathbb{P}(A_{ij} = 1)$ obey the following rule. Actually, we formulate this rule for the $\frac{p_{ij}}{1-p_{ij}}$ ratios, the so-called *odds*:

$$\frac{p_{ij}}{1-p_{ij}} = \alpha_i \alpha_j \quad (1 \le i < j \le n), \tag{3.3}$$

where the parameters $\alpha_1, \dots, \alpha_n$ are positive reals. This model is called α *model* in [Csetal1]. With the parameter transformation $\beta_i = \ln \alpha_i$ $(i = 1, \dots n)$, it is equivalent to the β *model* of [Ch-Dia-Sl] which applies to the *logits*:

$$\ln \frac{p_{ij}}{1-p_{ij}} = \beta_i + \beta_j \quad (1 \le i < j \le n)$$

with real parameters $\beta_1, \dots, \beta_n$.

Conversely, the probabilities p_{ij} and $1 - p_{ij}$ can be expressed in terms of the parameters, like

$$p_{ij} = \frac{\alpha_i \alpha_j}{1 + \alpha_i \alpha_j} \quad \text{and} \quad 1 - p_{ij} = \frac{1}{1 + \alpha_i \alpha_j}.$$

We are looking for the ML estimate of the parameter vector $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)$ or $\boldsymbol{\beta} = (\beta_1, \dots, \beta_n)$ based on the observed unweighted, undirected graph as a statistical sample.

Let $\mathbf{D} = (D_1, \dots, D_n)$ denote the degree-vector of the above random graph, where $D_i = \sum_{j=1}^{n} A_{ij}$ $(i = 1, \dots n)$. The random vector $\mathbf{D}$, as a function of the sample entries A_{ij}'s, is a *sufficient statistic* for the parameter $\boldsymbol{\alpha}$, or equivalently, for $\boldsymbol{\beta}$, see [Ch-Dia-Sl, Csetal1]. Let (a_{ij}) be the matrix of the sample realizations (the adjacency entries of the observed graph), $d_i = \sum_{j=1}^{n} a_{ij}$ be the actual degree of vertex i $(i = 1, \dots, n)$ and $\mathbf{d} = (d_1, \dots, d_n)$ be the observed degree-vector. Since the joint distribution of the entries belongs to the exponential family, with canonical parameterization [De-La-Ru], the maximum likelihood estimate $\hat{\boldsymbol{\alpha}}$ (or equivalently, $\hat{\boldsymbol{\beta}}$) is derived from the fact that, with it, the observed degree d_i equals the expected one, that is $\mathbb{E}(D_i) = \sum_{i=1}^{n} p_{ij}$. Therefore, $\hat{\boldsymbol{\alpha}}$ is the solution of the following *maximum likelihood equation*:

$$d_i = \sum_{j \ne i}^{n} \frac{\alpha_i \alpha_j}{1 + \alpha_i \alpha_j} \quad (i = 1, \dots, n). \tag{3.4}$$

The ML estimate $\hat{\boldsymbol{\beta}}$ is easily obtained from $\hat{\boldsymbol{\alpha}}$ via taking the logarithms of its coordinates.

Before discussing the solution of the system of equations (3.4), let us see, what conditions a sequence of nonnegative integers should satisfy so that it could be realized as the degree sequence of a graph. The sequence $d_1, \dots, d_n$ of nonnegative integers is called *graphic* if there is an unweighted, undirected graph on n vertices such that its vertex-degrees are the numbers $d_1, \dots, d_n$ in some order. Without loss of generality, d_i's can be enumerated in non-increasing order. The Erdős–Gallai theorem [Erd-Gal] gives the following necessary and sufficient condition for a sequence to be graphic. The sequence $d_1 \ge \dots \ge d_n \ge 0$ of integers is graphic if and only if it satisfies the following two conditions: $\sum_{i=1}^{n} d_i$ is even and

$$\sum_{i=1}^{k} d_i \le k(k-1) + \sum_{i=k+1}^{n} \min\{k, d_i\}, \quad k = 1, \dots, n-1. \tag{3.5}$$

Note that for nonnegative (not necessarily integer) real sequences a continuous analogue of (3.5) is derived in [Ch-Dia-Sl]. For given n, the convex hull of the possible graphic degree sequences is a polytope, to be denoted by $\mathcal{D}_n$. Its extreme points are the so-called *threshold graphs* [Mah-Pel]. It is interesting that for $n = 3$ all undirected graphs are threshold, since there are 8 possible graphs on 3 nodes, and there are also 8 vertices of $\mathcal{D}_3$; the $n = 2$ case is also not of much interest, therefore we will treat the $n > 3$ cases only.

The authors of [Ch-Dia-Sl, Csetal1] prove that $\mathcal{D}_n$ is the topological closure of the set of expected degree sequences, and for given $n > 3$, if $\mathbf{d} \in \mathrm{int}(\mathcal{D}_n)$ is an interior point, then the maximum likelihood equation (3.4) has a unique solution. Later, it turned out that the converse is also true: in [Rin-Pe-Fi] the authors prove that the ML estimate exists if and only if the observed degree vector is an inner point of $\mathcal{D}_n$. On the contrary, when the observed degree vector is a boundary point of $\mathcal{D}_n$, there is at least one 0 or 1 probability p_{ij} which can be obtained only by a parameter vector such that at least one of the β_i's is not finite. In this case, the likelihood function cannot be maximized with a finite parameter vector, its supremum is approached with a parameter vector $\boldsymbol{\beta}$ with at least one coordinate tending to $+\infty$ or $-\infty$.

V. Csiszár et al. [Csetal1] recommended the following algorithm and proved its convergence to the unique solution of the system (3.4), provided $\mathbf{d} \in \mathrm{int}(\mathcal{D}_n)$. To motivate the iteration, we rewrite (3.4) as

$$d_i = \alpha_i \sum_{j \neq i} \frac{1}{\frac{1}{\alpha_j} + \alpha_i} \quad (i = 1, \ldots, n).$$

Then starting with initial parameter values $\alpha_1^{(0)}, \ldots, \alpha_n^{(0)}$ and using the observed degree sequence $d_1, \ldots, d_n$, which is an inner point of $\mathcal{D}_n$, the iteration is as follows:

$$\alpha_i^{(t)} = \frac{d_i}{\sum_{j \neq i} \frac{1}{\frac{1}{\alpha_j^{(t-1)}} + \alpha_i^{(t-1)}}} \quad (i = 1, \ldots, n)$$

for $t = 1, 2, \ldots$, until convergence.

Now we will discuss the bipartite graph model, which traces back to Haberman [Hab], Lauritzen [Laur], and Rasch [Rasch] who applied it for psychological and educational measurements. The frequently cited Rasch model involves categorical data, mainly binary variables, therefore the underlying random object can be thought of as a contingency table. According to the Rasch model, the entries of an $m \times n$ binary table $\boldsymbol{A}$ are independent Bernoulli random variables, where for the parameter p_{ij} of the entry A_{ij} the following holds:

$$\ln \frac{p_{ij}}{1 - p_{ij}} = \beta_i - \delta_j \quad (i = 1, \ldots m;\ j = 1, \ldots, n) \tag{3.6}$$

with real parameters $\beta_1, \ldots, \beta_m$ and $\delta_1, \ldots, \delta_n$. As an example, Rasch [Rasch] investigated binary tables where the rows corresponded to patients and the columns to items of some psychological test, whereas the j-th entry of the i-th row was 1 if person i answered test item j correctly and 0, otherwise. He also gave a description of the parameters: β_i was the ability of person i, while δ_j the difficulty of test item j. Therefore, in view of the model equation (3.6), the more intelligent the person and the less difficult the test, the larger the success to failure ratio was on a logarithmic scale.

Motivated by the Rasch model, given an $m \times n$ random binary table $\boldsymbol{A} = (A_{ij})$, or equivalently, a bipartite graph, our model is

$$\ln \frac{p_{ij}}{1 - p_{ij}} = \beta_i + \gamma_j \quad (i = 1, \ldots, m,\ j = 1, \ldots, n) \tag{3.7}$$

with real parameters $\beta_1, \dots, \beta_m$ and $\gamma_1, \dots, \gamma_n$; further, $p_{ij} = \mathbb{P}(A_{ij} = 1)$. In terms of the transformed parameters $b_i = e^{\beta_i}$ and $g_j = e^{\gamma_j}$, the model (3.7) is equivalent to

$$\frac{p_{ij}}{1 - p_{ij}} = b_i g_j \quad (i = 1, \dots, m,\ j = 1, \dots, n) \tag{3.8}$$

where $b_1, \dots, b_m$ and $g_1, \dots, g_n$ are positive reals. Conversely, the probabilities can be expressed in terms of the parameters:

$$p_{ij} = \frac{b_i g_j}{1 + b_i g_j} \quad \text{and} \quad 1 - p_{ij} = \frac{1}{1 + b_i g_j}. \tag{3.9}$$

Observe that if (3.7) holds with the parameters β_i's and γ_j's, then it also holds with the transformed parameters $\beta_i' = \beta_i + c$ $(i = 1, \dots, m)$ and $\gamma_j' = \gamma_j - c$ $(j = 1, \dots, n)$ with some $c \in \mathbb{R}$. Equivalently, if (3.8) holds with the positive parameters b_i's and g_j's, then it also holds with the transformed parameters

$$b_i' = b_i \kappa \quad \text{and} \quad g_j' = \frac{g_j}{\kappa} \tag{3.10}$$

with some $\kappa > 0$. Therefore, the parameters b_i and g_j are arbitrary to within a multiplicative constant.

Here the row-sums $R_i = \sum_{j=1}^n A_{ij}$ and the column-sums $C_j = \sum_{i=1}^m A_{ij}$ are the sufficient statistics for the parameters collected in $\mathbf{b} = (b_1, \dots, b_m)$ and $\mathbf{g} = (g_1, \dots, g_n)$. Indeed, the likelihood function is factorized as

$$\begin{aligned}
L_{\mathbf{b},\mathbf{g}}(\boldsymbol{A}) &= \prod_{i=1}^m \prod_{j=1}^n p_{ij}^{A_{ij}} (1 - p_{ij})^{1 - A_{ij}} \\
&= \left\{ \prod_{i=1}^m \prod_{j=1}^n \left(\frac{p_{ij}}{1 - p_{ij}} \right)^{A_{ij}} \right\} \prod_{i=1}^m \prod_{j=1}^n (1 - p_{ij}) \\
&= \left\{ \prod_{i=1}^m b_i^{\sum_{j=1}^n A_{ij}} \right\} \left\{ \prod_{j=1}^n g_j^{\sum_{i=1}^m A_{ij}} \right\} \prod_{i=1}^m \prod_{j=1}^n (1 - p_{ij}) \\
&= \left\{ \prod_{i=1}^m \prod_{j=1}^n \frac{1}{1 + b_i g_j} \right\} \left\{ \prod_{i=1}^m b_i^{R_i} \right\} \left\{ \prod_{j=1}^n g_j^{C_j} \right\}.
\end{aligned}$$

Since the likelihood function depends on $\boldsymbol{A}$ only through its row- and column-sums, by the Neyman–Fisher factorization theorem, $R_1, \dots, R_m, C_1, \dots, C_n$ is a sufficient statistic for the parameters. The first factor (the whole above expression) depends only on the parameters and the row- and column-sums, whereas the seemingly not present factor – which would depend merely on $\boldsymbol{A}$ – is constantly 1, indicating that the conditional joint distribution of the entries, given the row- and column-sums, is uniform in this model. Note that in [Bar1], the author characterizes random tables sampled uniformly from the set of 0-1 matrices with fixed margins. Given the margins, the contingency tables coming from the above model are uniformly distributed, and a typical table of this distribution is produced by the β-γ model with parameters estimated via the row- and column sums as sufficient statistics. In this way, here we obtain another view of the typical table of [Bar1].

Based on an observed $m \times n$ binary table (a_{ij}), since we are in exponential family, the likelihood equation is obtained by making the expectation of the sufficient statistic equal to its

sample value. Therefore, with the notation $r_i = \sum_{j=1}^{n} a_{ij}$ $(i = 1, \dots, m)$ and $c_j = \sum_{i=1}^{m} a_{ij}$ $(j = 1, \dots, n)$, the following *system of likelihood equations* is yielded:

$$\begin{aligned} r_i &= \sum_{j=1}^{n} \frac{b_i g_j}{1 + b_i g_j} = b_i \sum_{j=1}^{n} \frac{1}{\frac{1}{g_j} + b_i}, \quad i = 1, \dots m; \\ c_j &= \sum_{i=1}^{m} \frac{b_i g_j}{1 + b_i g_j} = g_j \sum_{i=1}^{m} \frac{1}{\frac{1}{b_i} + g_j}, \quad j = 1, \dots n. \end{aligned} \tag{3.11}$$

Note that for any sample realization of $\boldsymbol{A}$,

$$\sum_{i=1}^{m} r_i = \sum_{j=1}^{n} c_j \tag{3.12}$$

holds automatically. Therefore, there is a dependence between the equations of the system (3.11), indicating that the solution is not unique, in accord with our previous remark about the arbitrary scaling factor $\kappa > 0$ of (3.10). Based on the proofs of [Rin-Pe-Fi], apart from this scaling, the solution is unique if it exists at all. For our convenience, let $(\tilde{\mathbf{b}}, \tilde{\mathbf{g}})$ denote the equivalence class of the parameter vector $(\mathbf{b}, \mathbf{g})$, which consists of parameter vectors $(\mathbf{b}', \mathbf{g}')$ satisfying (3.10) with some $\kappa > 0$. So that to avoid this indeterminacy, we may impose conditions on the parameters, for example,

$$\sum_{i=1}^{m} \beta_i + \sum_{j=1}^{n} \gamma_j = 0. \tag{3.13}$$

Like the graphic sequences, here the following sufficient conditions can be given for the sequences $r_1 \geq \dots \geq r_m > 0$ and $c_1 \geq \dots \geq c_n > 0$ of integers to be row- and column-sums of an $m \times n$ matrix of 0-1 entries (see, e.g., [Bar2]):

$$\begin{aligned} \sum_{i=1}^{k} r_i &\leq \sum_{j=1}^{n} \min\{c_j, k\}, \quad k = 1, \dots, m; \\ \sum_{j=1}^{k} c_j &\leq \sum_{i=1}^{m} \min\{r_i, k\}, \quad k = 1, \dots, n. \end{aligned} \tag{3.14}$$

Observe that the $k = 1$ cases imply $r_1 \leq n$ and $c_1 \leq m$; whereas the $k = m$ and $k = n$ cases together imply $\sum_{i=1}^{m} r_i = \sum_{j=1}^{n} c_j$. This statement is the counterpart of the Erdős-Gallai conditions for bipartite graphs, where – due to (3.12) – the sum of the degrees is automatically even. In fact, the conditions in (3.14) are redundant: one of the conditions – either the one for the rows, or the one for the columns – suffices together with (3.12) and $c_1 \leq m$ or $r_1 \leq n$. The so obtained necessary and sufficient conditions define *bipartite realizable sequences* with the wording of [Ham-Pe-Su].

The convex hull of the bipartite realizable sequences $\mathbf{r} = (r_1, \dots, r_m)$ and $\mathbf{c} = (c_1, \dots, c_n)$ form a polytope in $\mathbb{R}^{m+n}$, actually, because of (3.12), in an $(m+n-1)$-dimensional hyperplane of it. It is called *polytope of bipartite degree sequences* and denoted by $\mathcal{P}_{m,n}$ in Hammer et al. [Ham-Pe-Su]. Analogously to the considerations of the α-β models, and applying the thoughts of the proofs in [Ch-Dia-Sl, Csetal1, Rin-Pe-Fi], $\mathcal{P}_{m,n}$ is the closure of the set of the expected row- and column-sum sequences in the above model. In [Ham-Pe-Su] it is proved that an $m \times n$ binary table, or equivalently a bipartite graph on the independent sets of m and n vertices, is on the boundary of $\mathcal{P}_{m,n}$ if it does not contain two vertex-disjoint

edges. In this case, the likelihood function cannot be maximized with a finite parameter set, its supremum is approached with a parameter vector with at least one coordinate β_i or γ_j tending to $+\infty$ or $-\infty$, or equivalently, with at least one coordinate b_i or g_j tending to $+\infty$ or 0. Stated as Theorem 6.3 in the supplementary material of [Rin-Pe-Fi], the maximum likelihood estimate of the parameters of model (3.8) exists if and only if the observed row- and column-sum sequence $(\mathbf{r}, \mathbf{c}) \in \mathrm{ri}\,(\mathcal{P}_{m,n})$, the relative interior of $\mathcal{P}_{m,n}$, satisfying (3.12). In this case for the probabilities, calculated by the formula (3.9) through the estimated positive parameter values $\hat{b}_i$'s and $\hat{g}_j$'s (solutions of(3.11)), $0 < p_{ij} < 1$ holds $\forall i, j$.

Under these conditions, we define an algorithm that converges to the unique (up to the above equivalence) solution of the maximum likelihood equation (3.11). More precisely, in [Bol-El15] we proved that if $(\mathbf{r}, \mathbf{c}) \in \mathrm{ri}\,(\mathcal{P}_{m,n})$, then our algorithm gives a unique equivalence class of the parameter vectors as the fixed point of the iteration, which therefore provides the ML estimate of the parameters.

Starting with positive parameter values $b_i^{(0)}$ $(i = 1, \ldots, m)$ and $g_j^{(0)}$ $(j = 1, \ldots, n)$ and using the observed row- and column-sums, the iteration is as follows:

$$I. \quad b_i^{(t)} = \frac{r_i}{\sum_{j=1}^{n} \frac{1}{\frac{1}{g_j^{(t-1)}} + b_i^{(t-1)}}}, \quad i = 1, \ldots m$$

$$II. \quad g_j^{(t)} = \frac{c_j}{\sum_{i=1}^{m} \frac{1}{\frac{1}{b_i^{(t)}} + g_j^{(t-1)}}}, \quad j = 1, \ldots n$$

for $t = 1, 2, \ldots$, until convergence.

In the several clusters case, we are putting the bricks together, and use the so-called $k-\beta$ model, introduced in V. Csiszár et al. [Csetal2]. The above discussed α-β and β-γ models will be the building blocks of a heterogeneous block model. Here the degree sequences are not any more sufficient for the whole graph, only for the building blocks of the subgraphs.

Given $1 \le k \le n$, we are looking for k-partition, in other words, clusters $C_1, \ldots, C_k$ of the vertices such that

- different vertices are independently assigned to a cluster C_u with probability π_u $(u = 1, \ldots, k)$, where $\sum_{u=1}^{k} \pi_u = 1$;
- given the cluster memberships, vertices $i \in C_u$ and $j \in C_v$ are connected independently, with probability p_{ij} such that

$$\ln \frac{p_{ij}}{1 - p_{ij}} = \beta_{iv} + \beta_{ju}$$

for any $1 \le u, v \le k$ pair. Equivalently,

$$\frac{p_{ij}}{1 - p_{ij}} = b_{ic_j} b_{jc_i}$$

where c_i is the cluster membership of vertex i and $b_{iv} = e^{\beta_{iv}}$.

To estimate the parameters, we again use the EM algorithm. The parameters are collected in the vector $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_k)$ and the $n \times k$ matrix $\boldsymbol{B}$ of b_{iv}'s $(i \in C_u,\ u, v = 1, \ldots, k)$. The likelihood function is the following mixture:

$$\sum_{1 \le u, v \le k} \pi_u \pi_v \prod_{i \in C_u, j \in C_v} p_{ij}^{a_{ij}} (1 - p_{ij})^{(1 - a_{ij})}.$$

First we complete our data matrix $\boldsymbol{A}$ with latent membership vectors $\boldsymbol{\Delta}_1, \dots, \boldsymbol{\Delta}_n$ of the vertices that are k-dimensional i.i.d. $Poly(1, \boldsymbol{\pi})$ (polynomially distributed) random vectors, as in Section 3.3.2.

Note that, if the cluster memberships where known, then the complete likelihood would be

$$\prod_{u=1}^{k}\prod_{i=1}^{n}\prod_{v=1}^{k}\prod_{j=1}^{n}[p_{ij}^{\Delta_{jv}a_{ij}}\cdot(1-p_{ij})^{\Delta_{jv}(1-a_{ij})}]^{\Delta_{iu}} \tag{3.15}$$

that is valid only in case of known cluster memberships.

Starting with initial parameter values $\boldsymbol{\pi}^{(0)}$, $\boldsymbol{B}^{(0)}$ and membership vectors $\boldsymbol{\Delta}_1^{(0)}, \dots, \boldsymbol{\Delta}_n^{(0)}$, the t-th step of the iteration is as follows ($t = 1, 2, \dots$).

- **E**-step: we calculate the conditional expectation of each $\boldsymbol{\Delta}_i$ conditioned on the model parameters and on the other cluster assignments obtained in step $t-1$, and collectively denoted by $M^{(t-1)}$.

  The responsibility of vertex i for cluster u in the t-th step is defined as the conditional expectation $\pi_{iu}^{(t)} = \mathbb{E}(\Delta_{iu} \,|\, M^{(t-1)})$, and by the Bayes theorem, it is

  $$\pi_{iu}^{(t)} = \frac{\mathbb{P}(M^{(t-1)}|\Delta_{iu}=1)\cdot\pi_u^{(t-1)}}{\sum_{v=1}^{k}\mathbb{P}(M^{(t-1)}|\Delta_{iv}=1)\cdot\pi_v^{(t-1)}}$$

  ($u = 1, \dots, k$; $i = 1, \dots, n$). For each i, $\pi_{iu}^{(t)}$ is proportional to the numerator, therefore the conditional probabilities $\mathbb{P}(M^{(t-1)}|\Delta_{iu}=1)$ should be calculated for $u = 1, \dots, k$. But this is just the part of the likelihood (3.15) effecting vertex i under the condition $\Delta_{iu} = 1$. Therefore, if $i \in C_u$, then

  $$\mathbb{P}(M^{(t-1)}|\Delta_{iu}=1) = \prod_{v=1}^{k}\prod_{j\in C_v,\, j\sim i}\frac{b_{iv}^{(t-1)}b_{ju}^{(t-1)}}{1+b_{iv}^{(t-1)}b_{ju}^{(t-1)}}\prod_{j\in C_v,\, j\not\sim i}\frac{1}{1+b_{iv}^{(t-1)}b_{ju}^{(t-1)}}.$$

- **M**-step: We update $\boldsymbol{\pi}^{(t)}$ and $\boldsymbol{\Delta}^{(t)}$: $\pi_u^{(t)} := \frac{1}{n}\sum_{i=1}^{n}\pi_{iu}^{(t)}$ and $\Delta_{iu}^{(t)} = 1$ if $\pi_{iu}^{(t)} = \max_v \pi_{iv}^{(t)}$ and 0, otherwise (in case of ambiguity, we select the smallest index for the cluster membership of vertex i).

  Then we estimate the parameters in the actual clustering of the vertices. In the within-cluster scenario, we use the parameter estimation of model (3.3), obtaining estimates of b_{iu}'s ($i \in C_u$) in each cluster separately ($u = 1, \dots, k$); here b_{iu} corresponds to α_i and the number of vertices is $|C_u|$. In the between-cluster scenario, we use the bipartite graph model (3.8) in the following way. For $u \neq v$, edges connecting vertices of C_u and C_v form a bipartite graph, based on which the parameters b_{iv} ($i \in C_u$) and b_{ju} ($j \in C_v$) are estimated with the above algorithm; here b_{iv}'s correspond to b_i's, b_{ju}'s correspond to g_j's, and the number of rows and columns of the rectangular array corresponding to this bipartite subgraph of $\boldsymbol{A}$ is $|C_u|$ and $|C_v|$, respectively. With the estimated parameters, collected in the $n \times k$ matrix $\boldsymbol{B}^{(t)}$, we go back to the E-step, etc.

As in the M-step we increase the likelihood in all parts, and in the E-step we relocate each vertex into the cluster where its likelihood is maximized, the nonnegative likelihood function is increased in each iteration. Since the likelihood function is bounded from above (unless in some inner cycle we start from the boundary of a polytope of bipartite realizable sequences), it must converge to a local maximum.

Note that here the parameter β_{iv} with $c_i = u$ embodies the affinity of vertex i of cluster C_u towards vertices of cluster C_v; and likewise, β_{ju} with $c_j = v$ embodies the affinity of vertex j of cluster C_v towards vertices of cluster C_u. By the model, this affinities are added together on the level of the logits. This model is applicable to social networks, where attitudes of individuals in the same social group (say, u) are the same toward members of another social group (say, v), though, this attitude also depends on the individual in group u. The model may also be applied to biological networks, where the clusters consist, for example, of different functioning synopses or other units of the brain.

After normalizing the β_{iv} ($i \in C_u$) and β_{ju} ($j \in C_v$) to meet the requirement of (3.13) for any $u \neq v$ pair, the sum of the parameters will be zero:

$$\sum_{i \in C_u} \beta_{iv} + \sum_{j \in C_v} \beta_{ju} = 0,$$

and the sign and magnitude of them indicates the affinity of nodes of C_u to make ties with the nodes of C_v, and vice versa. This becomes important when we want to compare the parameters corresponding to different cluster pairs. For the initial clustering, spectral clustering tools are to be used.

We applied the algorithm for randomly generated and real-world data, see Figure 3.4 for some simulation results. We remark that in the case of real-world graphs, while processing the iteration, we sometimes run into threshold subgraphs or bipartite subgraphs on the boundary of the polytope of bipartite degree sequences. Even in this case our iteration converged for most coordinates of the parameter vectors, while some b_{iv} coordinates tended to $+\infty$ or 0 (numerically, when stopping the iteration, they took on a very 'large' or 'small' value). This means that the affinity of node i towards nodes of the cluster j is infinitely 'large' or 'small', i.e., this node is liable to always or never make ties with nodes of cluster j, see [Bol-El15] for details.

When applied to real-world data, our final clusters showed a good agreement with the spectral clusters; therefore, the algorithm can be considered as a fine-tuning of the spectral clustering in that it gives estimates of the parameters which provide a local maximum of the overall likelihood with clusters near to the spectral ones. Unfortunately, without a good starting, the EM iteration can run into a local maximum with clusters carrying not exact meaning; however, spectral clustering itself is not capable of parameter estimation. In this way, spectral clustering provides the initial clusters for our EM iteration that estimates parameters in the within- and between-cluster scenario, giving a 'happy marriage' of these approaches with practice (see the citation from Ravi Kannan's talk in the Introduction).

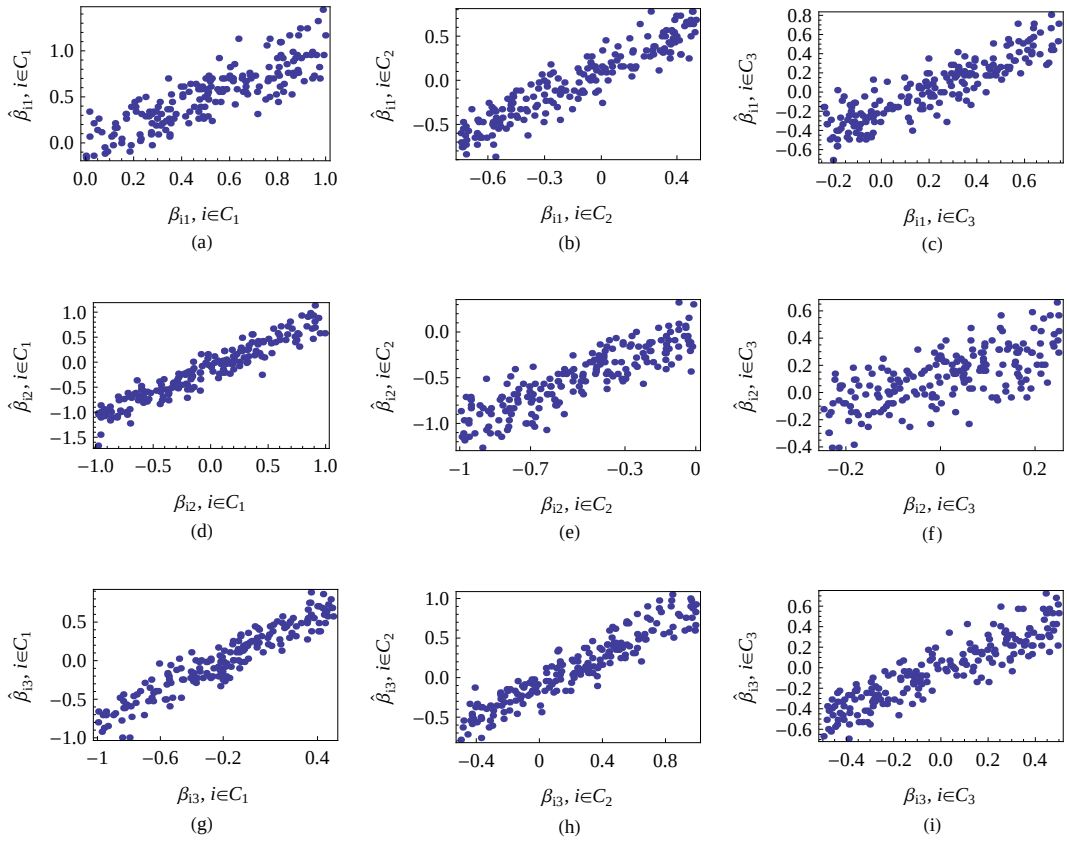

Figure 3.4: Data were generated based on parameters β_{iv}'s chosen uniformly in different intervals, $k = 3$, $|C_1| = 190$, $|C_2| = 193$, $|C_3| = 197$. The estimated versus the original parameters β_{iv}'s are shown for $i \in C_u$ $(u, v = 1, \ldots, k)$, where $\beta_{i1} \sim \mathcal{U}[0, 1]$ $(i \in C_1)$, $\beta_{i1} \sim \mathcal{U}[-0.75, 0.5]$ $(i \in C_2)$, $\beta_{i1} \sim \mathcal{U}[-0.25, 0.75]$ $(i \in C_3)$, $\beta_{i2} \sim \mathcal{U}[-1, 1]$ $(i \in C_1)$, $\beta_{i2} \sim \mathcal{U}[-1, 0]$ $(i \in C_2)$, $\beta_{i2} \sim \mathcal{U}[-0.25, 0.25]$ $(i \in C_3)$, $\beta_{i3} \sim \mathcal{U}[-1, 0.5]$ $(i \in C_1)$, $\beta_{i3} \sim \mathcal{U}[-0.5, 1]$ $(i \in C_2)$, and $\beta_{i3} \sim \mathcal{U}[-0.5, 0.5]$ $(i \in C_3)$, respectively. MSE=1.14634 (made by Ahmed Elbanna).

Bibliography

[Ac-Mc] Achlioptas, D. and McSherry, F., Fast computation of low-rank matrix approximations, *J. ACM* **54** (2007), Article 9.

[Alon] Alon, N., Eigenvalues and expanders. *Combinatorica* **6** (1986), 83–96.

[Al-Sp] Alon, N. and Spencer, J. H., *The Probabilistic Method*, Wiley (2000).

[Al-Kr-Vu] Alon, N., Krivelevich, M. and Vu, V. H., On the concentration of eigenvalues of random symmetric matrices, *Isr. J. Math.* **131** (2002), 259–267.

[Aletal] Alon, N., Coja-Oghlan, A., Han, H., Kang, M., Rödl, V. and Schacht, M., Quasi-randomness and algorithmic regularity for graphs with general degree distributions, *Siam J. Comput.* **39** (2010), 2336–2362.

[Ar] Aronszajn, N., Theory of Reproducing Kernels, *Trans. Am. Math. Soc.* **68** (1950), 337–404.

[Az-Gh] Azran, A. and Ghahramani, Z., Spectral methods for automatic multiscale data clustering. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006), New York NY* (Fitzgibbon A, Taylor CJ and Lecun Y eds), IEEE Computer Society, Los Alamitos, California (2006), pp. 190–197.

[Bach] Bach, F. R. and Jordan, M. I., Kernel Independent Component Analysis, *J. Mach. Learn. Res.* **3** (2002), 1–48.

[Bar1] Barvinok, A., What does a random contingency table look like? preprint (2009), arXiv:0806.3910 [math.CO].

[Bar2] Barvinok, A., Matrices with prescribed row and column sums, preprint (2010), arXiv:1010.5706 [math.CO].

[Benz] Benzécri, J. P. et al., *Pratique de l'analyse des données, 2. L'Analyse des correspondances*. Dunod, Paris (1980).

[Bhat] Bhatia, R., Matrix Analysis. Springer (1997).

[Bi-Ch] Bickel, P. J. and Chen, A., A nonparametric view of network models and Newman-Girvan and other modularities, *Proc. Natl. Acad. Sci. USA* **106** (2009), 21068–21073.

[Biggs] Biggs NL 1974 *Algebraic Graph Theory*. Cambridge Univ. Press, Cambridge.

[Bil-Lin] Bilu, Y. and Linial, N., Lifts, discrepancy and nearly optimal spectral gap, *Combinatorica* **26** (2006), 495–519.

[Bol81] Bolla M., Decomposition of Matrices in a Genetic Problem, *Biometrics* **37** (4) (1981), p. 845.

[Bol83] Bolla M., Mátrixok szinguláris felbontásának módszerei és statisztikai alkalmazásai. Egyetemi kisdoktori értekezés, ELTE, Budapest (1983).

[Bol-Tus85] Bolla, M. and Tusnády, G., The QRPS algorithm: a generalization of the QR algorithm for the singular values decomposition of rectangular matrices, *Periodica Math. Hung.* **16** (1985), 201–207.

[Bol87a] Bolla, M., Hilbert-terek lineáris operátorainak szinguláris felbontása: optimumtulajdonságok statisztikai alkalmazásai és numerikus módszerek, *Alk. Mat. Lapok* **13** (1987-88), 189–206.

[Bol87b] Bolla, M., Korrespondanciaanalizis, *Alk. Mat. Lapok* **13** (1987-88), 207–230.

[Bol91] Bolla, M., Relations between spectral and classification properties of multigraphs, DIMACS Technical Report (1991), 1991-27.

[Bol92] Bolla, M., Relations between spectral and classification properties of multigraphs. Kandidátusi disszertáció (1992), MTA Könyvtára.

[Bol93] Bolla, M., Spectra and Euclidean representation of hypergraphs, *Discret. Math.* **117** (1993), 19–39.

[Bol-Tus94] Bolla, M. and Tusnády, G., Spectra and Optimal Partitions of Weighted Graphs, *Discret. Math.* **128** (1994), 1–20.

[Boletal98] Bolla, M., Michaletzky, Gy., Tusnády, G., Ziermann, M., Extrema of sums of heterogeneous quadratic forms, *Linear Algebra and its Applications* **269** (1998), 331-365.

[Bol-Tus00] Bolla, M. and Tusnády, G., Hipergráfok összefüggőségének vizsgálata a spektrumon keresztül (Investigating connectivity of hypergraphs by spectra), *Mat. Lapok* 95/1-2 (2000), 1–27.

[Bol-Mol02] Bolla, M. and Molnár-Sáska, G., Isoperimetric Properties of Weighted Graphs Related to the Laplacian Spectrum and Canonical Correlations, *Studia Sci. Math. Hung.* **39** (2002), 425–441.

[Bol-Mol04] Bolla, M. and M.-Sáska, G., Optimization problems for weighted graphs and related correlation estimates, *Discret. Math.* **282** (2004), 23–33.

[Bol04] Bolla, M., Distribution of the Eigenvalues of Random Block-Matrices *Linear Algebra Appl.* **377** (2004), 219–240.

[Bol05] Bolla, M., Recognizing linear structure in noisy matrices, *Linear Algebra Appl.* **402** (2005), 228-244.

[Bol08a] Bolla, M., Noisy random graphs and their Laplacians, *Discret. Math.* **308** (2008), 4221–4230.

[Bol08b] Bolla, M., On the Spectra of Weighted Random Graphs Related to Social Networks. In *Social Networksd: Development, Evaluation and Influence*. Hannah L. Schneider and Lilli M. Huber eds, Nova Science Publishers (2008), New York, pp. 131–158.

[Bol-Fr-Kr10] Bolla, M., Friedl, K. and Krámli, A., Singular value decomposition of large random matrices (for two-way classification of microarrays), *J. Multivariate Anal.* **101** (2010), 434–446.

[Bol10] Bolla, M., Statistical inference on large contingency tables: convergence, testability, stability. In: Proc. of the COMPSTAT'2010: 19th International Conference on Computational Statistics, Paris, Physica-Verlag, Springer (2010), pp. 817-824.

[Bol11a] Bolla, M., Beyond the expanders, *Int. J. Comb.* (2011), Paper 787596.

[Bol11b] Bolla, M., Spectra and structure of weighted graphs, *Electronic Notes in Discrete Mathematics* **38** (2011), 149-154.

[Bol11c] Bolla, M., Penalized versions of the Newman–Girvan modularity and their relation to normalized cuts and k-means clustering, *Phys. Rev. E* **84** (2011),016108.

[Bol-Ko-Kr12] Bolla, M., Kói, T. and Krámli, A., Testability of minimum balanced multiway cut densities, *Discret. Appl. Math.* **160** (2012), 1019–1027.

[Bol13] Bolla, M., *Spectral Clustering and Biclustering. Learning Large Graphs and Contingency Tables*, Wiley (2013).

[Bol14a] Bolla, M., Modularity spectra, eigen-subspaces and structure of weighted graphs, *European J. Combin.* **35** (2014), 105–116.

[Bol14b] Bolla, M., SVD, discrepancy, and regular structure of contingency tables, *Discret. Appl. Math.* **176** (2014), 3–11.

[Boletal15] Bolla, M., Bullins, B., Chaturapruek, S., Chen, S., Friedl, K., Spectral properties of modularity matrices, *Linear Algebra and Its Applications* **73** (2015), 359-376.

[Bol-El15] Bolla, M., Elbanna, A., Estimating parameters of a probabilistic heterogeneous block model via the EM algorithm, *Journal of Probability and Statistics* (2015), Article ID 657965.

[Bol15] Bolla, M., Generalized quasirandom properties of graphs, preprint (2015), arXiv:1508.04369v3 [math.CO].

[Bol-El16] Bolla, M., Elbanna, A., Matrix and discrepancy view of generalized random and quasirandom graphs, *Special Matrices* **4**, Issue 1 (2016), 31-45.

[Bol16] Bolla, M., Relating multiway discrepancy and singular values of nonnegative rectangular matrices, *Discrete Applied Mathematics* **203** (2016), 26-34.

[Bo-Erd] Bollobás, B. and Erdős, P., An extremal problem of graphs with diameter 2, *Math. Mag.* **48** (1975), 419-427.

[Bo] Bollobás, B., *Random graphs*. Academic Press, New York (1987).

[Bo-Nik] Bollobás, B. and Nikiforov, V.,Hermitian matrices and graphs: singular values and discrepancy, *Discret. Math.* **285** (2004), 17–32.

[Borgsetal1] Borgs, C, Chayes, J. T., Lovász, L., T.-Sós, V. and Vesztergombi, K., Convergent graph sequences I: Subgraph Frequencies, metric properties, and testing, *Advances in Math.* **219** (2008), 1801–1851.

[Borgsetal2] Borgs, C, Chayes, J. T., Lovász, L., T.-Sós, V. and Vesztergombi, K., Convergent sequences of dense graphs II: Multiway cuts and statistical physics, *Ann. Math.* **176** (2012), 151–219.

[Bre-Fri] Breiman, L. and Friedman, J. H., Estimating optimal transformations for multiple regression and correlation, *J. Am. Stat. Assoc.* **80** (1985), 580–619.

[But] Butler, S., Using discrepancy to control singular values for nonnegative matrices, *Linear Algebra Appl.* **419** (2006), 486–493.

[But1] Butler, S., Relating singular values and discrepancy of weighted directed graphs. In *Proc. of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, Miami, FL, 2006*, SIAM, Philadelphia, PA (2006), pp. 1112–1116.

[Ch-Dia-Sl] Chatterjee, S., Diaconis, P. and Sly, A., Random graphs with a given degree sequence, *Ann. Stat.* **21** (2010), 1400–1435.

[Che] Cheeger, J., A lower bound for the smallest eigenvalue of the Laplacian, In *Problems in Analysis* (ed. R. C. Gunning), Princeton Univ. Press, Princeton NJ (1970), pp. 195–199.

[Chu] Chung, F., *Spectral Graph Theory*, CBMS Regional Conference Series in Mathematics **92**. American Mathematical Society, Providence RI (1997).

[Chu-G-W] Chung, F., Graham, R. and Wilson, R. K., Quasi-random graphs, *Combinatorica* **9** (1989), 345–362.

[Chu-G] Chung, F. and Graham, R., Quasi-random graphs with given degree sequences, *Random Struct. Algorithms* **12** (2008), 1–19.

[Chu-K] Chung, F. and Kenter, F., Discrepancy inequalities for directed graphs, *Discret. Appl. Math.* **176** (2014), 30–42.

[Co-La] Coja-Oghlan, A. and Lanka, A., Finding planted partitions in random graphs with general degree distributions, *J. Discret. Math.* **23** (2009), 1682–1714.

[Csetal1] Csiszár, V., Hussami, P., Komlós, J., Móri, T. F., Rejtő, L. and Tusnády, G., When the degree sequence is a sufficient statistic, *Acta Math. Hung.* **134** (2011), 45–53.

[Csetal2] Csiszár, V., Hussami, P., Komlós, J., Móri, T. F., Rejto, L. and Tusnady, G., Testing goodness of fit of random graph models, *Algorithms* **5** (2012), 629–635.

[Cv] Cvetković, D. M., Doob, M. and Sachs, H., *Spectra of Graphs.* Academic Press, New York (1979).

[Cv-Ro-Si] Cvetković, D. M., Rowlinson, P., and Simić, S., *Eigenspaces of Graphs*, Encyclopedia of Mathematics and Its Applications Vol. 66. Lavoisier Libraire, S.A.S. France (1997).

[De-La-Ru] Dempster, A. P., Laird, N. M. and Rubin, D. B., Maximum likelihood from incomplete data via the EM algorithm, *J. R. Statist. Soc. B* **39** (1977), 1–38.

[Dhil] Dhillon, I. S., Co-clustering documents and words using bipartite spectral graph partitioning. In *Proc. ACM Int. Conf. Knowledge Disc. Data Mining (KDD 2001)* (Provost FJ and Srikant R eds), Association for Computer Machinery, New York (2001), pp. 269–274.

[Dia-Str] Diaconis, P. and Stroock, D., Geometric bounds for eigenvalues of Markov chains. *Ann. Appl. Probab.* **1** (1991), 36–62.

[Erd-Reny] Erdős, P. Rényi, A., On the evolution of random graphs, *Publ. Math. Inst. Hung. Acad. Sci.* **5** (1960), 17–61.

[Erd-Gal] Erdős, P. and Gallai, T., Graphs with given degree of vertices (in Hungarian), *Mat. Lapok* **11** (1960), 264–274.

[Ev-Go-Lu] Evra, S., Golubev, K., Lubotzky, A., Mixing properties and the chromatic number of Ramanujan complexes, preprint (2014), arXiv:1407.7700 [math.CO].

[Fid72] Fiedler, M., Bounds for eigenvalues of doubly stochastic matrices, *Linear Algebra Appl.* **5** (1972), 299–310.

[Fid73] Fiedler, M., 1973 Algebraic connectivity of graphs, *Czech. Math. J.* **23** (1973), 298–305.

[Fr-Kan] Frieze, A. and Kannan, R., Quick approximation to matrices and applications, *Combinatorica* **19** (1999), 175–220.

[Fü-Ko] Füredi, Z. and Komlós, J., The eigenvalues of random symmetric matrices, *Combinatorica* **1** (1981), 233–241.

[Geb] Gebelein, H., Das statistische Problem der Korrelation als Variations und Eigenwertproblem und sein Zusammenhang mit der Ausgleichsrechnung. *Z. Angew. Math. Mech.* **21** (1941), 364–379.

[Gh-Trev] Gharan, S. H., Trevisan, L., A new regularity lemma and faster approximation algorithms for low threshold rank graphs. In *Proc. APPROX-RANDOM'13* (2013), pp. 303–316.

[Gran] Granovetter, M. S., The strength of weak ties, *Amer. J. Sociology* **78** (1973), 1360–1380.

[Green] Greenacre, M. J., *Theory and Applications of Correspondence Analysis*. Academic Press, New York (1984).

[Hab] Haberman, S. J., Log-linear models and frequency tables with small expected counts, *Ann. Statist.* **5** (1977), 1148–1169.

[Hag-Kah] Hagen, L. and Kahng, A. B., New spectral methods for ratio cut partitioning and clustering, *IEEE Trans. Comput. Aided Des.* **11** (1992), 1074–1085.

[Ham-Pe-Su] Hammer, P. L., Peled, U. N. and Sun, X., Difference graphs, *Discret. Appl. Math.* **28** (1990), 35–44.

[Ha-Ti-Fri] Hastie, T., Tibshirani, R. and Friedman, J., *The Elements of Statistical Learning. Data Mining, Inference, and Prediction*. Springer (2001).

[Hof69] Hoffman, A. L., The change in the least eigenvalue of the adjacency matrix of a graph under imbedding, *SIAM J. Appl. Math.* **17** (1969), 664–671.

[Hof70] Hoffman, A. J., On eigenvalues and colorings of graphs. In *Graph Theory and Its Applications* (ed. Harris B), Academic Press, New York (1970), pp. 79–91.

[Hof72] Hoffman, A. J., Eigenvalues and partitionings of the edges of a graph, *Linear Algebra Appl.* **5** (1972), 137–146.

[Hol-Lei] Holland, P. W. and Leinhardt, S., An exponential family of probability distributions for directed graphs, *J. Amer. Statist. Assoc.* **76** (1981), 33–50.

[Hol-Las-Lei] Holland, P., Laskey, K. B. and Leinhardt, S., Stochastic blockmodels: some first steps. *Social Networks* **5** (1983), 109–137.

[Ho-Lin-Wid] Hoory, S., Linial, N. and Widgerson, A., Expander graphs and their applications, *Bull. Amer. Math. Soc. (N. S.)* **43** (2006), 439–561.

[Juh-Mály] Juhász, F. and Mályusz, K., Problems of cluster analysis from the viewpoint of numerical analysis. In *Numerical Methods, Coll. Math. Soc. J. Bolyai* (ed. Rózsa P), Vol 22, pp. 405–415. North-Holland, Amsterdam (1980).

[Juh81] Juhász, F., On the spectrum of random graphs. In *Algebraic methods in graph theory, Coll. Math. Soc. J. Bolyai* (ed. Lovász L), Vol. 25, pp. 313–316. North–Holland, Amsterdam (1981).

[Klugetal] Kluger, Y., Basri, R., Chang, J. T. and Gerstein, M., Spectral biclustering of microarray data: coclustering genes and conditions, *Genome Res.* **13** (2003), 703–716.

[Kometal] Komlós, J., Shokoufanden, A., Simonovits, M. and Szemerédi, E., Szemerédi's Regularity Lemma and Its Applications in Graph Theory. In *Lecture Notes in Computer Science*, Vol 2292, pp. 84–112. Springer, Berlin (2002).

[Kup-Nied] Kuipers, L., Niederreiter, H., *Uniform distribution of sequences*, Wiley (1974).

[Laur] Lauritzen, S. L., *Extremal families and systems of sufficient statistics*. Lecture Notes in Statistics **49**, Springer (1988).

[Le-Gh-Trev] Lee, J. R., Gharan, S. O. and Trevisan, L., Multi-way spectral partitioning and higher-order Cheeger inequalities. In *Proc. 44th Annual ACM Symposium on the Theory of Computing (STOC 2012)*, pp. 1117–1130. New York NY (2012).

[Lov93] Lovász, L., Random walks on graphs: a survey. In *Combinatorics, Paul Erdős is Eighty. János Bolyai Society, Mathematical Studies* Vol. 2, pp. 1–46. Keszthely, Hungary (1993).

[Lov-Szeg07] Lovász, L. and Szegedy, B., Szemerédi's Lemma for the analyst. *Geom. Func. Anal.* **17** (2007), 252–270.

[Lov08] Lovász, L., Very large graphs. In In: *Current Developments in Mathematics* (Jerison D, Mazur B, Mrowka T, Schmid W, Stanley R and Yan ST eds), pp. 67-128. International Press, Somerville, MA (2008).

[Lov-Sos] Lovász, L. and T.-Sós V., Generalized quasirandom graphs, *J. Comb. Theory B* **98** (2008), 146–163.

[Lov-Szeg11] Lovász, L. and Szegedy, B., Finitely forcible graphons. *J. Comb. Theory B.* **101** (2011), 269–301.

[Lux] von Luxburg, U., A tutorial on spectral clustering. *Stat. Comput.* **17** (2007), 395–416.

[Lux-Be-Bo] von Luxburg, U., Belkin, M. and Bousquet, O., Consistency of spectral clustering. *Ann. Stat.* **36** (2008), 555–586.

[Mah-Pel] Mahadev, N. V. R. and Peled, U. N., Threshold graphs and related topics, *Ann. Discrete. Math.* **56**, North-Holland, Amsterdam (1995).

[McSh] McSherry, F., Spectral partitioning of random graphs. In *Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS 2001)*, pp. 529–537. Las Vegas, Nevada (2001).

[Mei-Shi] Meila, M. and Shi, J., Learning segmentation by random walks. In *Proc. 13th Neural Information Processing Systems Conference (NIPS 2001)* (Leen TK, Dietterich TG and Tresp V eds), pp. 873–879. MIT Press, Cambridge, USA (2001).

[Moh88] Mohar, B., Isoperimetric inequalities, growth and the spectrum of graphs. *Linear Algebra Appl.* **103** (1988), 119–131.

[Moh89] Mohar, B., Isoperimetric numbers of graphs. *J. Comb. Theor, Ser. B* **47** (1989), 274–291.

[Nash] Nash-Williams, C. St J. A., 1959 Random walks and electronic currents in networks, *Math. Proc. Cambridge Phil. Soc.* **55** (1959), 181–194.

[New-Gir] Newman, M. E. J. and Girvan, M., Finding and evaluating community structure in networks, *Phys. Rev. E* **69** (2004), 026113.

[New] Newman, M. E. J., *Networks, An Introduction*. Oxford University Press (2010).

[Ng-Jo-We] Ng, A. Y., Jordan, M. I. and Weiss, Y., On spectral clustering: analysis and an algorithm. In *Proc. 14th Neural Information Processing Systems Conference (NIPS 2001)* (Dietterich TG, Becker S and Ghahramani Z eds), pp. 849–856. MIT Press, Cambridge, USA (2001).

[Pinch] Pinch, R. G. E., Sequences well distributed in the square, *Math. Proc. Cambridge Phil. Soc.* 99 (1986), 19-22.

[Rao] Rao, C. R., Separation theorems for singular values of matrices and their applications in multivariate analysis, *J. Multivariate Anal.* **9** (1979), 362–377.

[Rasch] Rasch, G., On general laws and the meaning of measurement in psychology. In *Proc. of the Fourth Berkeley Symp. on Math. Statist. and Probab.*, pp. 321–333, University of California Press (1961).

[Re-Bo] Reichardt, J. and Bornholdt, S., Partitioning and modularity of graphs with arbitrary degree distribution, *Phys. Rev. E* **76** (2007) 015102(R).

[Reny59a] Rényi, A., On measures of dependence. *Acta Math. Acad. Sci. Hung.* **10** (1959), 441–451.

[Reny59b] Rényi, A., New version of the probabilistic generalization of the large sieve. *Acta Math. Acad. Sci. Hung.* **10** (1959), 218–226.

[Ri-SzN] Riesz, F. and Sz.-Nagy, B., *Leçons d'analyse fonctionnelle*. Academic Publishing House, Budapest (1952).

[Rin-Pe-Fi] Rinaldo, A., Petrovic, S. and Fienberg, S. E., Maximum likelihood estimation in the β-model, *Ann. Statist.* **41**, 1085–1110 (2013).

[Sch-Sm-Mu] Schölkopf, B., Smola, A. J. and Müller, K-R., Nonlinear component analysis as a kernel eigenvalue problem, *Neural Comput.* **10** (1998), 1299–1319.

[Sch-Sm] Schölkopf, B. and Smola, A. J. *Learning with Kernels*. MIT Press, Cambridge, USA (2002).

[Shi-Ma] Shi, J. and Malik, J., Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **22** (2000), 888–905.

[Shie-Pea] Shields, R. W. and Pearson, J. B., Structural controllability of multiinput linear systems, *IEEE Transactions on Automatic Control* **21** (1976), 203–212.

[Sim-Sos] Simonovits, M. and T.-Sós, V., Szemerédi's partition and quasi-randomness, *Random Struct. Algorithms* **2** (1991), 1–10.

[Sin-Jer] Sinclair, A. and Jerrum, M., Approximate counting, uniform generation and rapidly mixing Markov chains. *Inf. Comput.* **82** (1989), 93–133.

[Szeg] Szegedy, B., Limits of kernel operators and the spectral regularity lemma, *European J. Combin.* **32** (2011), 1156-1167.

[Szem] Szemerédi, E., Regular partitions of graphs. In *Colloque Inter. CNRS. No. 260, Problémes Combinatoires et Théorie Graphes* (Bermond J-C, Fournier J-C, Las Vergnas M and Sotteau D eds), 1976, pp. 399–401.

[Thom87] Thomason, A., Pseudo-random graphs, *Ann. Discret. Math.* **33**, 307–331 (1987).

[Thom89] Thomason, A., Dense expanders and pseudo-random bipartite graphs, *Discret. Math.* **75**, 381–386 (1989).

[Thomp] Thompson, R. C., The behavior of eigenvalues and singular values under perturbations of restricted rank, *Linear Algebra Appl.* **13** (1976), 69–78.

[Trev] Trevisan, L., Max cut and the smallest eigenvalue. In *Proc. 41th Annual ACM Symposium on the Theory of Computing (STOC 2009)*, pp. 1117–1130. Bethesda, Maryland USA (2009).

[Ung] Ungar, L. H. and Foster, D. P., A Formal Statistical Approach to Collaborative Filtering. In *Proc. Conference on Automatical Learning and Discovery (CONALD 98)*, pp. 1–6 (1998).