
Spectral Clustering of Signed Graphs via Matrix Power Means

Pedro Mercado^{1,2} Francesco Tudisco³ Matthias Hein²

Abstract

Signed graphs encode positive (attractive) and negative (repulsive) relations between nodes. We extend spectral clustering to signed graphs via the one-parameter family of *Signed Power Mean Laplacians*, defined as the matrix power mean of normalized standard and signless Laplacians of positive and negative edges. We provide a thorough analysis of the proposed approach in the setting of a general Stochastic Block Model that includes models such as the Labeled Stochastic Block Model and the Censored Block Model. We show that in expectation the signed power mean Laplacian captures the ground truth clusters under reasonable settings where state-of-the-art approaches fail. Moreover, we prove that the eigenvalues and eigenvector of the signed power mean Laplacian concentrate around their expectation under reasonable conditions in the general Stochastic Block Model. Extensive experiments on random graphs and real world datasets confirm the theoretically predicted behaviour of the signed power mean Laplacian and show that it compares favourably with state-of-the-art methods.

1. Introduction

The analysis of graphs has received a significant amount of attention due to their capability to encode interactions that naturally arise in social networks. Yet, the vast majority of graph methods has been focused on the case where interactions are of the same type, leaving aside the case where different kinds of interactions are available (Leskovec et al., 2010b). Graphs and networks with both positive and negative edge weights arise naturally in a number of social, biological and economic contexts. Social dynamics and relationships are intrinsically positive and negative: users of online social networks such as Slashdot and Epinions, for

example, can express positive interactions, like friendship and trust, and negative ones, like enmity and distrust. Other important application settings are the analysis of gene expressions in biology (Fujita et al., 2012) or the analysis of financial and economic time sequences (Ziegler et al., 2010; Pavlidis et al., 2006), where similarity and variable dependence measures commonly used may attain both positive and negative values (e.g. the Pearson correlation coefficient).

Although the majority of the literature has focused on graphs that encode only positive interactions, the analysis of signed graphs can be traced back to social balance theory (Cartwright & Harary, 1956; Harary, 1953; Davis, 1967), where the concept of a k -balance signed graph is introduced. The analysis of signed networks has been then pushed forward through the study of a variety of tasks in signed graphs, as for example edge prediction (Kumar et al., 2016; Leskovec et al., 2010a; Falher et al., 2017), node classification (Bosch et al., 2018; Tang et al., 2016a), node embeddings (Chiang et al., 2011; Derr et al., 2018; Kim et al., 2018; Wang et al., 2017; Yuan et al., 2017), node ranking (Chung et al., 2013; Shahriari & Jalili, 2014), and clustering (Chiang et al., 2012; Kunegis et al., 2010; Mercado et al., 2016; Sedoc et al., 2017; Doreian & Mrvar, 2009; Knyazev, 2018; Kirkley et al., 2018; Cucuringu et al., 2019; Cucuringu et al., 2018). See (Tang et al., 2016b; Gallier, 2016) for recent surveys on the topic.

In this paper we present a novel extension of spectral clustering for signed graphs. Spectral clustering (Luxburg, 2007) is a well established technique for non-signed graphs, which partitions the set of the nodes based on a k -dimensional node embedding obtained using the first eigenvectors of the graph Laplacian. **Our contributions are as follows:** We introduce the family of *Signed Power Mean (SPM) Laplacians*: a one-parameter family of graph matrices for signed graphs that blends the information from positive and negative interactions through the matrix power mean, a general class of matrix means that contains the arithmetic, geometric, and harmonic mean as special cases. This is inspired by recent extensions of spectral clustering which merge the information encoded by positive and negative interactions through different types of arithmetic (Chiang et al., 2012; Kunegis et al., 2010) and geometric (Mercado et al., 2016) means of the standard and signless graph Laplacians. We analyze the performance of the signed power mean Laplacian in a

¹Saarland University ²University of Tübingen ³University of Strathclyde. Correspondence to: Pedro Mercado <pedro@cs.uni-saarland.de>.

general Signed Stochastic Block Model. We first provide an analysis in expectation showing that the smaller is the parameter of the signed power mean Laplacian, the less restrictive are the conditions that ensure to recover the ground truth clusters. In particular, we show that the limit cases $+\infty$ and $-\infty$ are related to the boolean operators AND and OR, respectively, in the sense that for the limit case $+\infty$ clusters are recovered only if both positive *and* negative interactions are informative, whereas for $-\infty$ clusters are recovered if positive *or* negative interactions are informative. This is consistent with related work in the context of unsigned multilayer graphs (Mercado et al., 2018). Second, we show that the eigenvalues and eigenvectors of the signed power mean Laplacian concentrate around their mean, so that our results hold also for the case where one samples from the stochastic block model. Our result extends with minor changes to the unsigned multilayer graph setting considered in (Mercado et al., 2018), where just the expected case has been studied. To our knowledge these are the first concentration results for matrix power means under any stochastic block model for signed graphs. Finally, we show that the signed power mean Laplacian compares favorably with state-of-the-art approaches through extensive numerical experiments on diverse real world datasets. All the proofs have been moved to the supplementary material.

Notation. A signed graph is a pair $G^\pm = (G^+, G^-)$, where $G^+ = (V, W^+)$ and $G^- = (V, W^-)$ encode positive and negative edges, respectively, with positive symmetric adjacency matrices W^+ and W^- , and a common vertex set $V = \{v_1, \dots, v_n\}$. Note that this definition allows the simultaneous presence of both positive and negative interactions between the same two nodes. This is a major difference with respect to the alternative point of view where G^\pm is associated to a single symmetric matrix W with positive and negative entries. In this case $W = W^+ - W^-$, with $W_{ij}^+ = \max\{0, W_{ij}\}$ and $W_{ij}^- = -\min\{0, W_{ij}\}$, implying that every interaction is either positive or negative, but not both at the same time. We denote by $D_{ii}^+ = \sum_{j=1}^n w_{ij}^+$ and $D_{ii}^- = \sum_{j=1}^n w_{ij}^-$ the diagonal matrix of the degrees of G^+ and G^- , respectively, and $\bar{D} = D^+ + D^-$.

2. Related work

The study of clustering of signed graphs can be traced back to the theory of social balance (Cartwright & Harary, 1956; Harary, 1953; Davis, 1967), where a signed graph is called k -balanced if the set of vertices can be partitioned into k sets such that within the subsets there are only positive edges, and between them only negative.

Inspired by the notion of k -balance, different approaches for signed graph clustering have been introduced. In particular, many of them aim to extend spectral clustering to signed graphs by proposing novel signed graph Laplacians.

A related approach is correlation clustering (Bansal et al., 2004). Unlike spectral clustering, where the number of clusters is fixed a-priori, correlation clustering approximates the optimal number of clusters by identifying a partition that is as close as possible to be k -balanced. In this setting, the case where the number of clusters is constrained has been considered in (Giotis & Guruswami, 2006).

We briefly introduce the standard and signless Laplacian and review different definitions of Laplacians on signed graphs. The final clustering algorithm to find k clusters is the same for all of them: compute the smallest k eigenvectors of the corresponding Laplacian, use the eigenvectors to embed the nodes into \mathbb{R}^k , obtain the final clustering by doing k -means in the embedding space. However, we will see below that in some cases we have to slightly deviate from this generic principle by using the $k - 1$ smallest eigenvectors instead.

Laplacians of Unsigned Graphs: In the following all weight matrices are non-negative and symmetric. Given an assortative graph $G = (V, W)$, standard spectral clustering is based on the Laplacian and its normalized version defined as:

$$L = D - W \quad L_{\text{sym}} = D^{-1/2} L D^{-1/2}$$

where $D_{ii} = \sum_{j=1}^n w_{ij}$ is the diagonal matrix of the degrees of G . Both Laplacians are symmetric positive semidefinite and the multiplicity of the eigenvalue 0 is equal to the number of connected components in G .

For disassortative graphs, i.e. when edges carry only dissimilarity information, the goal is to identify clusters such that the amount of edges between clusters is larger than the one inside clusters. Spectral clustering is extended to this setting by considering the signless Laplacian matrix and its normalized version (see e.g. (Liu, 2015; Mercado et al., 2016)), defined as:

$$Q = D + W \quad Q_{\text{sym}} = D^{-1/2} Q D^{-1/2}$$

Both Laplacians are positive semi-definite, and the smallest eigenvalue is zero if and only if the graph has a bipartite component (Desai & Rao, 1994).

Laplacians of Signed Graphs: Signed graphs encode both positive and negative interactions. In the ideal k -balanced case positive interactions present an assortative behaviour, whereas negative interactions present a disassortative behaviour. With this in mind, several novel definitions of *signed Laplacians* have been proposed. We briefly review them for later reference.

In (Chiang et al., 2012) the balance ratio Laplacian and its normalized version are defined as:

$$L_{BR} = D^+ - W^+ + W^-, \quad L_{BN} = \bar{D}^{-1/2} L_{BR} \bar{D}^{-1/2}$$

whereas in (Kunegis et al., 2010) the signed ratio Laplacian

and its normalized version have been defined as:

$$L_{SR} = \bar{D} - W^+ + W^-, \quad L_{SN} = \bar{D}^{-1/2} L_{SR} \bar{D}^{-1/2}$$

The signed Laplacians L_{BR} and L_{BN} need not be positive semidefinite, while the signed Laplacians L_{SR} and L_{SN} are positive semidefinite with eigenvalue zero if and only if the graph is 2-balanced.

In the context of correlation clustering, in (Saade et al., 2015) the Bethe Hessian matrix is defined as:

$$H = (\alpha - 1)I - \sqrt{\alpha}(W^+ - W^-) + \bar{D}$$

where α is the average node degree $\alpha = \frac{1}{n} \sum_{i=1}^n \bar{D}_{ii}$. The Bethe Hessian H need not be positive definite. In fact, eigenvectors with negative eigenvalues bring information of clustering structure (Saade et al., 2014).

Let $L^+ = D^+ - W^+$ and $Q^- = D^- + W^-$ be the Laplacian and signless Laplacian of G^+ and G^- , respectively. As noted in (Mercado et al., 2016), $L_{SR} = L^+ + Q^-$ i.e. it coincides with twice the arithmetic mean of L^+ and Q^- . Note that the same holds for H when the average degree α is equal to one, i.e. $H = L_{SR}$ when $\alpha = 1$. In (Mercado et al., 2016), the arithmetic mean and geometric mean of the normalized Laplacian and its signless version are used to define new Laplacians for signed graphs:

$$L_{AM} = L_{\text{sym}}^+ + Q_{\text{sym}}^-, \quad L_{GM} = L_{\text{sym}}^+ \# Q_{\text{sym}}^-$$

where $A \# B = A^{-1/2} (A^{1/2} B A^{1/2})^{1/2} A^{-1/2}$ is the geometric mean of A and B , $L_{\text{sym}}^+ = (D^+)^{-1/2} L^+ (D^+)^{-1/2}$ and $Q_{\text{sym}}^- = (D^-)^{-1/2} Q^- (D^-)^{-1/2}$. While the computation of L_{GM} is more challenging, in (Mercado et al., 2016) it is shown that the clustering assignment obtained with the geometric mean Laplacian L_{GM} outperforms all other signed Laplacians.

Both the arithmetic and the geometric means are special cases of a much richer one-parameter family of means known as power means. Based on this observation, we introduce the *Signed Power Mean Laplacian* in Section 2.2, defined via a matrix version of the family of power means which we briefly review below.

2.1. Matrix Power Means

The scalar power mean of two non-negative scalars a, b is a one-parameter family of means defined for $p \in \mathbb{R}$ as $m_p(a, b) = \left(\frac{a^p + b^p}{2}\right)^{1/p}$. Particular cases are the arithmetic, geometric and harmonic means, as shown in Table 1. Moreover, the scalar power mean is monotone in the parameter p , i.e. $m_p(a, b) \leq m_q(a, b)$ when $p \leq q$ (see (Bullen, 2013), Ch. 3, Thm. 1), which yields the well known arithmetic-geometric-harmonic mean inequality $m_{-1}(a, b) \leq m_0(a, b) \leq m_1(a, b)$. As matrices do not commute, several matrix extensions of the scalar power

Table 1 Particular cases of scalar power means

p	$m_p(a, b)$	name
$p \rightarrow \infty$	$\max\{a, b\}$	maximum
$p = 1$	$(a + b)/2$	arithmetic mean
$p \rightarrow 0$	\sqrt{ab}	geometric mean
$p = -1$	$2\left(\frac{1}{a} + \frac{1}{b}\right)^{-1}$	harmonic mean
$p \rightarrow -\infty$	$\min\{a, b\}$	minimum

mean have been introduced, which typically agree if the matrices commute, see e.g. Chapter 4 in (Bhatia, 2009). We consider the following matrix extension of the scalar power mean:

Definition 1 ((Bhagwat & Subramanian, 1978)). *Let A, B be symmetric positive definite matrices, and $p \in \mathbb{R}$. The matrix power mean of A, B with exponent p is*

$$M_p(A, B) = \left(\frac{A^p + B^p}{2}\right)^{1/p}$$

where $Y^{1/p}$ is the unique positive definite solution of the matrix equation $X^p = Y$.

Please note that this definition can be extended to positive semidefinite matrices (Bhagwat & Subramanian, 1978) for $p > 0$, as $M_p(A, B)$ exists, whereas for $p < 0$ a diagonal shift is necessary to ensure that the matrices A, B are positive definite.

2.2. The Signed Power Mean Laplacian

Given a signed graph $G^\pm = (G^+, G^-)$ we define the Signed Power Mean (SPM) Laplacian L_p of G^\pm as

$$L_p = M_p(L_{\text{sym}}^+, Q_{\text{sym}}^-). \quad (1)$$

For the case $p < 0$ the matrix power mean requires positive definite matrices, hence we use in this case the matrix power mean of diagonally shifted Laplacians, i.e. $L_{\text{sym}}^+ + \varepsilon I$ and $Q_{\text{sym}}^- + \varepsilon I$. Our following theoretical analysis holds for all possible shifts $\varepsilon > 0$, whereas we discuss in the supplementary material the numerical robustness with respect to ε . The clustering algorithm for identifying k clusters in signed graphs is given in Algorithm 1. Please note that for $p \geq 1$ we deviate from the usual scheme and use the first $k - 1$ eigenvectors rather than the first k . The reason is a result of the analysis in the stochastic block model in Section 3. In general, the main influence of the parameter p of the power mean is on the ordering of the eigenvalues. In Section 3 we will see that this significantly influences the performance of different instances of SPM Laplacians, in particular, the arithmetic and geometric mean discussed in (Mercado et al., 2016) are suboptimal for the recovery of the ground truth clusters. For the computation of the matrix power mean we adapt the scalable Krylov subspace-based algorithm proposed in (Mercado et al., 2018).

Algorithm 1: Spectral clustering of signed graphs with L_p

Input: Symmetric matrices W^+ , W^- , number k of clusters to construct.

Output: Clusters C_1, \dots, C_k .

- 1 Let $k' = k - 1$ if $p \geq 1$ and $k' = k$ if $p < 1$.
 - 2 Compute eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_{k'}$ corresponding to the k' smallest eigenvalues of L_p .
 - 3 Set $U = (\mathbf{u}_1, \dots, \mathbf{u}_{k'})$ and cluster the rows of U with k -means into clusters C_1, \dots, C_k .
-

3. Stochastic Block Model Analysis of the Signed Power Mean Laplacian

In this section we analyze the signed power mean Laplacian L_p under a general Signed Stochastic Block Model. Our results here are twofold. First, we derive new conditions in expectation that guarantee that the eigenvectors corresponding to the smallest eigenvalues of L_p recover the ground truth clusters. These conditions reveal that, in this setting, the state-of-the-art signed graph matrices are suboptimal as compared to L_p for negative values of p . Second, we show that our result in expectation transfer to sampled graphs as we prove conditions that ensure that both eigenvalues and eigenvectors of L_p concentrate around their expected value with high probability. We verify our results by several experiments where the clustering performance of state-of-the-art matrices and L_p are compared on random graphs following the Signed Stochastic Block Model.

All proofs hold for an arbitrary diagonal shift $\varepsilon > 0$, whereas the shift is set to $\varepsilon = \log_{10}(1 + |p|) + 10^{-6}$ in the numerical experiments. Numerical robustness with respect to ε is discussed in the supplementary material.

The Stochastic Block Model (SBM) is a well-established generative model for graphs and a canonical tool for studying clustering methods (Holland et al., 1983; Rohe et al., 2011; Abbe, 2018). Graphs drawn from the SBM show a prescribed clustering structure, as the probability of an edge between two nodes depends only on the clustering membership of each node. We introduce our SBM for signed Graphs (SSBM): we consider k ground truth clusters C_1, \dots, C_k , all of them of size $|\mathcal{C}| = \frac{n}{k}$, and parameters $p_{\text{in}}^+, p_{\text{out}}^+, p_{\text{in}}^-, p_{\text{out}}^- \in [0, 1]$ where p_{in}^+ (resp. p_{in}^-) is the probability of observing an edge inside clusters in G^+ (resp. G^-) and p_{out}^+ (resp. p_{out}^-) is the probability of observing an edge between clusters in G^+ (resp. G^-). Calligraphic letters are used for the expected adjacency matrices: \mathcal{W}^+ and \mathcal{W}^- are the expected adjacency matrix of G^+ and G^- , respectively, where $\mathcal{W}_{i,j}^+ = p_{\text{in}}^+$ and $\mathcal{W}_{i,j}^- = p_{\text{in}}^-$ if v_i, v_j belong to the same cluster, whereas $\mathcal{W}_{i,j}^+ = p_{\text{out}}^+$ and $\mathcal{W}_{i,j}^- = p_{\text{out}}^-$ if v_i, v_j belong to different clusters.

Other extensions of the SBM to the signed setting have been considered. Particularly relevant examples are the Labelled Stochastic Block Model (LSBM) (Heimlicher et al.,

2012) and the Censored Block Model (CBM) (Abbe et al., 2014). In the context of signed graphs, both LSBM and CBM assume that an observed edge can be either positive or negative, but not both. Our SSBM, instead, allows the simultaneous presence of both positive and negative edges between the same pair of nodes, as the parameters $p_{\text{in}}^+, p_{\text{out}}^+, p_{\text{in}}^-, p_{\text{out}}^-$ in SSBM are independent. Moreover, the edge probabilities defining both the LSBM and the CBM can be recovered as special cases of the SSBM. In particular, the LSBM corresponds to the SSBM for the choices

$$\begin{aligned} p_{\text{in}}^+ &= \bar{p}_{\text{in}}\mu^+, & p_{\text{in}}^- &= \bar{p}_{\text{in}}\mu^- & (\text{within clusters}) \\ p_{\text{out}}^+ &= \bar{p}_{\text{out}}\nu^+, & p_{\text{out}}^- &= \bar{p}_{\text{out}}\nu^- & (\text{between clusters}) \end{aligned}$$

where \bar{p}_{in} and \bar{p}_{out} are edge probabilities within and between clusters, respectively, whereas μ^+ and $\mu^- = 1 - \mu^+$ (resp. ν^+ and $\nu^- = 1 - \nu^+$) are the probabilities of assigning a positive and negative label to an edge within (resp. between) clusters. Similarly, the CBM corresponds to the SSBM for the particular choices $\bar{p}_{\text{in}} = \bar{p}_{\text{out}}$, $\mu^+ = \nu^- = (1 - \eta)$ and $\mu^- = \nu^+ = \eta$ where η is a noise parameter.

Our goal is to identify conditions in terms of $k, p_{\text{in}}^+, p_{\text{out}}^+, p_{\text{in}}^-$, and p_{out}^- , such that C_1, \dots, C_k are recovered by the smallest eigenvectors of the signed power mean Laplacian. Consider the following k vectors:

$$\chi_1 = \mathbf{1}, \quad \chi_i = (k-1)\mathbf{1}_{C_i} - \mathbf{1}_{\bar{C}_i}.$$

$i = 2, \dots, k$. The node embedding given by $\{\chi_i\}_{i=1}^k$ is informative in the sense that applying k -means on $\{\chi_i\}_{i=1}^k$ trivially recovers the ground truth clusters C_1, \dots, C_k as all nodes of a cluster are mapped to the same point. Note that the constant vector χ_1 could be omitted as it does not add clustering information. We derive conditions for the SSBM such that $\{\chi_i\}_{i=1}^k$ are the smallest eigenvectors of the signed power mean Laplacian in expectation.

Theorem 1. Let $\mathcal{L}_p = M_p(\mathcal{L}_{\text{sym}}^+, \mathcal{Q}_{\text{sym}}^-)$ and let $\varepsilon > 0$ be the diagonal shift.

- If $p \geq 1$, then $\{\chi_i\}_{i=2}^k$ correspond to the $(k-1)$ -smallest eigenvalues of \mathcal{L}_p if and only if $m_p(\rho_\varepsilon^+, \rho_\varepsilon^-) < 1 + \varepsilon$;
 - If $p < 1$, then $\{\chi_i\}_{i=1}^k$ correspond to the k -smallest eigenvalues of \mathcal{L}_p if and only if $m_p(\rho_\varepsilon^+, \rho_\varepsilon^-) < 1 + \varepsilon$;
- with $\rho_\varepsilon^+ = 1 - (p_{\text{in}}^+ - p_{\text{out}}^+) / (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) + \varepsilon$ and $\rho_\varepsilon^- = 1 + (p_{\text{in}}^- - p_{\text{out}}^-) / (p_{\text{in}}^- + (k-1)p_{\text{out}}^-) + \varepsilon$.

Note that Theorem 1 is the reason why Alg. 1 uses only the first $k - 1$ eigenvectors for $p \geq 1$. The problem is that the constant eigenvector need not be among the first k eigenvectors in the SSBM for $p \geq 1$. However, as it is constant and thus uninformative in the embedding, this does not lead to any loss of information. The following Corollary shows that the limit cases of L_p are related to the boolean operators AND and OR.

Corollary 1. Let $\mathcal{L}_p = M_p(\mathcal{L}_{\text{sym}}^+, \mathcal{Q}_{\text{sym}}^-)$.

- $\{\chi_i\}_{i=2}^k$ correspond to the $(k-1)$ -smallest eigenvalues of

- \mathcal{L}_∞ if and only if $p_{\text{in}}^+ > p_{\text{out}}^+$ **and** $p_{\text{in}}^- < p_{\text{out}}^-$,
- $\{\chi_i\}_{i=1}^k$ correspond to the k -smallest eigenvalues of $\mathcal{L}_{-\infty}$ if and only if $p_{\text{in}}^+ > p_{\text{out}}^+$ **or** $p_{\text{in}}^- < p_{\text{out}}^-$.

The conditions for \mathcal{L}_∞ are the most conservative ones, as they require that $G^+ \underline{\text{and}} G^-$ are informative, i.e. G^+ has to be assortative $\underline{\text{and}} G^-$ disassortative. Under these conditions every clustering method for signed graphs should be able to identify the ground truth clusters in expectation. On the other hand, the less restrictive conditions for the recovery of the ground truth clusters correspond to the limit case $\mathcal{L}_{-\infty}$. If $G^+ \underline{\text{or}} G^-$ are informative, then the ground truth clusters are recovered, that is, $\mathcal{L}_{-\infty}$ only requires that G^+ is assortative $\underline{\text{or}} G^-$ is disassortative. In particular, the following corollary shows that smaller values of p require less restrictive conditions to ensure the identification of the informative eigenvectors.

Corollary 2. Let $q \leq p$. If $\{\chi_i\}_{i=\theta(p)}^k$ correspond to the k -smallest eigenvalues of \mathcal{L}_p , then $\{\chi_i\}_{i=\theta(q)}^k$ correspond to the k -smallest eigenvalues of \mathcal{L}_q , where $\theta(x) = 1$ if $x \leq 0$ and $\theta(x) = 2$ if $x > 0$.

To better understand the different conditions we have derived, we visualize them in Fig. 1, where the x -axis corresponds to how assortative G^+ is, while the y -axis corresponds to how disassortative G^- is. The conditions of the limit case \mathcal{L}_∞ , i.e. the case where $G^+ \underline{\text{and}} G^-$ have to be informative, correspond to the upper-right region, dark blue region in Fig. 1c, and correspond to the 25% of all possible configurations of the SBM. The conditions for the limit case $\mathcal{L}_{-\infty}$, i.e. the case where $G^+ \underline{\text{or}} G^-$ has to be informative, instead correspond to all possible configurations of the SBM except for the bottom-left region. This is depicted in Fig. 1b and corresponds to the 75% of all possible configurations under the SBM.

In Fig. 2 we present the corresponding conditions for recovery in expectation for the cases $p \in \{-10, -1, 0, 1, 10\}$. We can visually verify that the larger the value of p the smaller is the region where the conditions of Theorem 1 hold. In particular, one can compare the change of conditions as one moves from the signed harmonic (\mathcal{L}_{-1}), geometric (\mathcal{L}_0), to the arithmetic (\mathcal{L}_1) mean Laplacians verifying the ordering described in Corollary 2. Moreover, we clearly observe that \mathcal{L}_{-10} and \mathcal{L}_{10} are already quite close to the conditions necessary for the limit cases $\mathcal{L}_{-\infty}$ and \mathcal{L}_∞ , respectively.

In the middle row of Fig. 2 we show the average clustering error for each power mean Laplacian when sampling 50 times from the SSBM following the diagram presented in Fig. 1a and fixing the sparsity of G^+ and G^- by setting $p_{\text{in}}^+ + p_{\text{out}}^+ = 0.1$ and $p_{\text{in}}^- + p_{\text{out}}^- = 0.1$ with two clusters each of size 100. We observe that the areas with low clustering error qualitatively match the regions where in expectation we have recovery of the clusters. However, due to the sampling

which can make one of the graphs G^+ and G^- quite sparse and as we just consider graphs with 200 nodes, the area of low clustering error is smaller in comparison to the region of guaranteed recovery in expectation due to the sampling variance in the stochastic block model.

In the bottom row of Fig. 2 we show the clustering error for the state of the art methods L_{GM}, L_{SN}, L_{BM} and H . We can see that L_{GM} presents a similar performance as the signed power mean Laplacian L_0 . The next Theorem shows that the geometric mean Laplacian \mathcal{L}_{GM} and the limit $p \rightarrow 0$ of the signed power mean Laplacian agree in expectation for the SSBM. This implies via Corollary 2 that this operator is inferior to the signed power mean Laplacian for $p < 0$. This is why we use in the experiments on real world graphs later on always $p < 0$.

Theorem 2. Let $\mathcal{L}_{GM} = \mathcal{L}_{\text{sym}}^+ \# \mathcal{Q}_{\text{sym}}^-$ and \mathcal{L}_0 be the signed power mean Laplacian with $p \rightarrow 0$ of the expected signed graph. Then, $\mathcal{L}_0 = \mathcal{L}_{GM}$.

In the bottom row of Fig. 2 we can observe that L_{SN}, L_{BN} and H present a similar behaviour to the arithmetic mean Laplacian L_1 . A quick computation shows that for the case where both G^+, G^- have the same node degree in expectation, the conditions of Theorem 1 for \mathcal{L}_1 reduce to $p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$. It turns out that this condition is also required by $\mathcal{L}_{SN}, \mathcal{L}_{BN}$ and \mathcal{H} , as the following shows.

Theorem 3 ((Mercado et al., 2016)). Let \mathcal{L}_{BN} and \mathcal{L}_{SN} be the balanced normalized Laplacian and signed normalized Laplacian of the expected signed graph. The following statements are equivalent:

- $\{\chi_i\}_{i=1}^k$ are the eigenvectors corresponding to the k -smallest eigenvalues of \mathcal{L}_{BN} .
- $\{\chi_i\}_{i=1}^k$ are the eigenvectors corresponding to the k -smallest eigenvalues of \mathcal{L}_{SN} .
- inequalities $p_{\text{in}}^- + (k-1)p_{\text{out}}^- < p_{\text{in}}^+ + (k-1)p_{\text{out}}^+$ and $p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$ hold.

Finally, we present conditions in expectation for the Bethe Hessian to identify the ground truth clustering.

Theorem 4. Let \mathcal{H} be the Bethe Hessian of the expected signed graph. Then $\{\chi_i\}_{i=2}^k$ are the eigenvectors corresponding to the $(k-1)$ -smallest negative eigenvalues of \mathcal{H} if and only if the following conditions hold:

1. $\max\{0, \frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-} |C|}\} < (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)$
2. $p_{\text{out}}^+ < p_{\text{out}}^-$

Moreover, for the limit case $|V| \rightarrow \infty$ the first condition reduces to $p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$.

Please see the supplementary material for a further analysis in expectation. We can observe that the first condition in Theorem 4 is related to conditions of \mathcal{L}_1 and $\mathcal{L}_{SN}, \mathcal{L}_{BN}$ through the inequality $p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$. This explains





Figure 1: Stochastic Block Model (SBM) for signed graphs. From left to right: Fig. 1a SBM Diagram. Fig. 1b SBM for $L_{-\infty}$ (OR), Fig. 1c SBM for L_{∞} (AND). according to Corollary 1.

Recovery of Clusters in Expectation



Clustering Error

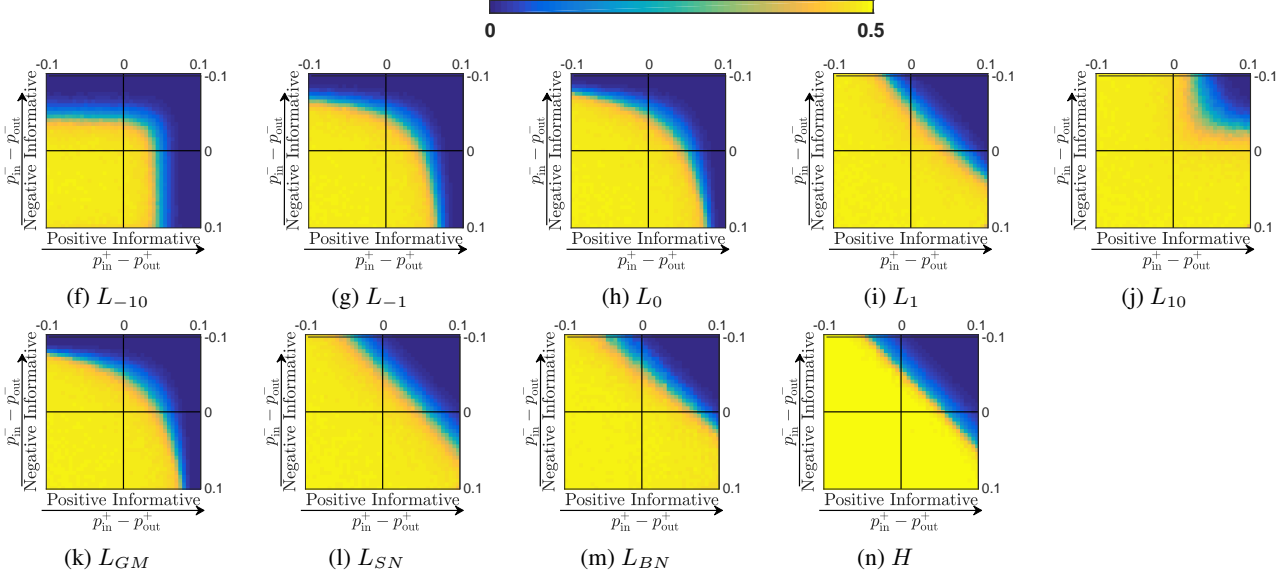


Figure 2: Performance visualization for two clusters for different parameters of the SBM. **Top row:** In dark blue the settings where the signed power mean Laplacians \mathcal{L}_p identify the ground truth clusters in expectation for the SBM, see Theorem 1, whereas yellow indicates failure. **Middle/Bottom row:** average clustering error (dark blue: small error, yellow: large error) of the signed power mean Laplacian L_p and L_{GM} , L_{SN} , L_{BN} , H for 50 samples from the SBM.

why the performance of the Bethe Hessian H resembles the one of arithmetic Laplacians L_{SN} , L_{BN} , L_1 . A more detailed comparison between the conditions of Theorems 1, 3 and 4 is detailed in the supplementary material.

Note that our analysis in expectation considers the dense regime where the average degree increases with the number of nodes and hence our results in expectation are verified under the SSBM setting here considered, showing that

L_{SN} , L_{BN} , L_1 , H have a similar performance. However, in the case of sparse graphs, it is known that the Bethe Hessian is asymptotically optimal in the information-theoretic transition limit (Saade et al., 2014; 2015). Please see the supplementary material for an evaluation under the CBM.

We now zoom in on a particular setting of Fig. 2. Namely, the case where G^+ (resp. G^-) is fixed to be informative, whereas the remaining graph transitions from informative to

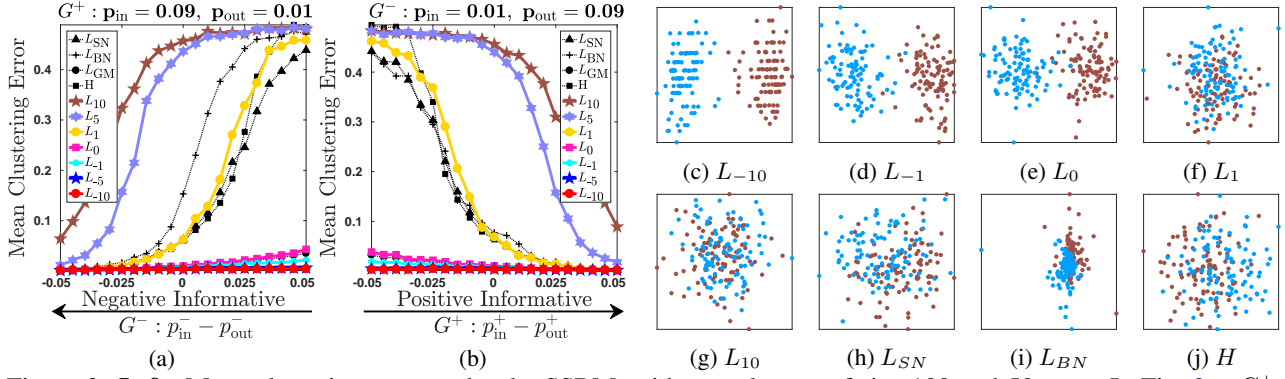


Figure 3: **Left:** Mean clustering error under the SSBM, with two clusters of size 100 and 50 runs. In Fig. 3a: G^+ is informative, i.e. assortative with $p_{in}^+ = 0.09$ and $p_{out}^+ = 0.01$. In Fig. 3b: G^- is informative, i.e. disassortative with $p_{in}^- = 0.01$ and $p_{out}^- = 0.09$. **Right:** Node embeddings induced by eigenvectors of different signed Laplacians for a random graph drawn from SSBM for 2 clusters of size 100, $p_{in}^+ = 0.025, p_{out}^+ = 0.075, p_{in}^- = 0.01, p_{out}^- = 0.09$.

uninformative. The corresponding results are in Fig. 3. In Fig. 3a we consider the case where G^+ is informative with parameters $p_{in}^+ = 0.09$ and $p_{out}^+ = 0.01$ (this corresponds to $p_{in}^+ - p_{out}^+ = 0.08$ in Fig. 2), and G^- goes from being informative ($p_{in}^- < p_{out}^-$) to non-informative ($p_{in}^- \geq p_{out}^-$). We confirm that the power mean Laplacian L_p presents smaller clustering errors for smaller values of p . Moreover, it is clear that in the case $p < 0$, L_p is able to recover clusters even in the case where G^- is not informative, whereas for $p > 0$, L_p requires both G^+ and G^- to be informative. We observe that the smallest (resp. largest) clustering errors correspond to L_{-10} (resp. L_{10}), corroborating Corollary 2. Further, we can observe that L_{GM} and L_0 have a similar performance, as well as L_{SN}, L_{BN}, L_1, H , as observed before, confirming Theorem 2 and Theorem 4, respectively. In Fig. 3b similar observations hold for the case where G^- is informative with parameters $p_{in}^- = 0.01$ and $p_{out}^- = 0.09$ (this corresponds to $p_{in}^- - p_{out}^- = -0.08$ in Fig. 2), and G^+ goes from being non-informative ($p_{in}^+ \leq p_{out}^+$) to informative ($p_{in}^+ > p_{out}^+$). Within this setting we present the eigenvector-based node embeddings of each method for the case $p_{in}^+ = 0.025, p_{out}^+ = 0.075, p_{in}^- = 0.01, p_{out}^- = 0.09$, in right hand side of Fig. 3. For L_{-10}, L_{-1}, L_0 the embeddings split the clusters properly, whereas remaining embeddings are not informative, verifying the effectivity of L_p with $p < 0$.

3.1. Consistency of the Signed Power Mean Laplacian for the Stochastic Block Model

In this section we prove two novel concentration bounds for signed power mean Laplacians of signed graphs drawn from the SSBM. The bounds show that, for large graphs, our previous results in expectation transfer to sampled graphs with high probability. We first show in Theorem 5 that L_p is close to \mathcal{L}_p . Then, in Theorem 6, we show that eigenvalues and eigenvectors of L_p are close to those of \mathcal{L}_p . We derive this result by tracing back the consistency of the matrix

power mean to the consistency of the standard and signless Laplacian established in (Chung & Radcliffe, 2011).

The consistency of spectral clustering on unsigned graphs for the SBM has been studied in (Lei & Rinaldo, 2015; Sarkar & Bickel, 2015; Rohe et al., 2011) and more recently consistency of several variants of spectral clustering has been shown (Qin & Rohe, 2013; Joseph & Yu, 2016; Chaudhuri et al., 2012; Le et al.; Fasino & Tudisco, 2018; Davis & Sethuraman, 2018). Moreover, while the case of multilayer graphs under the SBM has been previously analyzed (Han et al., 2015; Heimlicher et al., 2012; Jog & Loh, 2015; Paul & Chen, 2017; Xu et al., 2014; 2017; Yun & Proutiere, 2016), there are no consistency results for matrix power means for multilayer graphs as studied in (Mercado et al., 2018). While our main emphasis is on the analysis of the SPM Laplacian, our proofs are general enough to cover also the consistency of the matrix power means for unsigned multilayer graphs (Mercado et al., 2018). In Thm. 5 we show that the SPM Laplacian L_p for the SSBM is concentrated around \mathcal{L}_p , with high probability for large n . The following results hold for general shifts ϵ .

Theorem 5. *Let p be a non-zero integer; let*

$$C_p = \begin{cases} (2p)^{1/p}(2 + \epsilon)^{1-1/p} & p \geq 1 \\ |2p|^{1/|p|} \epsilon^{-(3+1/|p|)} & p \leq -1 \end{cases}$$

and choose $\epsilon > 0$. If $\frac{n}{k}(p_{in}^+ + (k-1)p_{out}^+) > 3 \ln(8n/\epsilon)$, and $\frac{n}{k}(p_{in}^- + (k-1)p_{out}^-) > 3 \ln(8n/\epsilon)$, then with probability at least $1 - \epsilon$, we have

$$\|L_p - \mathcal{L}_p\| \leq C_p m_{|p|}^{1/|p|} \left(\sqrt{\frac{3 \ln(8n/\epsilon)}{\frac{n}{k}(p_{in}^+ + (k-1)p_{out}^+)}} , \sqrt{\frac{3 \ln(8n/\epsilon)}{\frac{n}{k}(p_{in}^- + (k-1)p_{out}^-)}} \right)$$

In Thm 5 we take the spectral norm. A more general version of Theorem 5 for the inhomogeneous Erdős-Rényi model, where edges are formed independently with probabilities p_{ij}^+, p_{ij}^- is given in the supplementary material. Theorem 5 builds on top of concentration results of (Chung & Radcliffe, 2011) proven for the unsigned case $\|L_{sym}^+ - \mathcal{L}_{sym}^+\|$. We

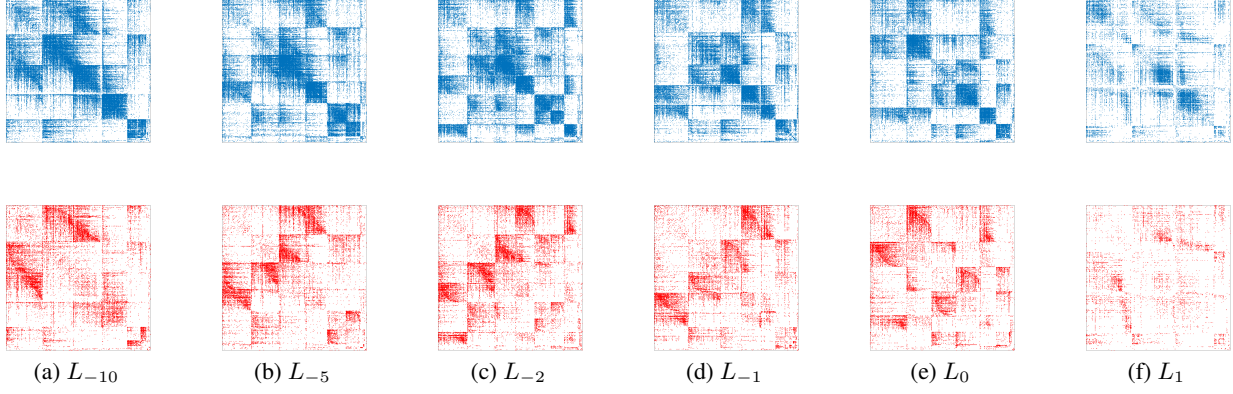


Figure 4: Adjacency matrices W^+ and W^- sorted through clustering of Wikipedia Elections dataset by the proposed Signed Power mean Laplacians $L_{-10}, L_{-5}, L_{-2}, L_{-1}, L_0, L_1$. **Top row:** zoom-in visualization of positive edges W^+ . **Bottom row:** zoom-in visualization of negative edges W^- . See supplementary material for more details.

can see that the deviation of L_p from \mathcal{L}_p depends on the power mean of the individual deviations of L_{sym}^+ and Q_{sym}^- from $\mathcal{L}_{\text{sym}}^+$ and $\mathcal{Q}_{\text{sym}}^-$, respectively. Note that the larger the size n of the graph is, the stronger is the concentration of L_p around \mathcal{L}_p .

The next Theorem shows that the eigenvectors corresponding to the smallest eigenvalues of L_p are close to the corresponding eigenvectors of \mathcal{L}_p . This is a key result showing consistency of our spectral clustering technique with L_p for signed graphs drawn from the SSBM.

Theorem 6. *Let $p \neq 0$ be an integer. Let $V_k, \mathcal{V}_k \in \mathbb{R}^{n \times k}$ be orthonormal matrices whose columns are the eigenvectors of the k smallest eigenvalues of L_p and \mathcal{L}_p , respectively. Let $\rho_\varepsilon^+, \rho_\varepsilon^-$ and C_p be defined as in Theorems 1 and 5, respectively. Define $\tilde{k} = k - 1$, if $p \geq 1$, and $\tilde{k} = k$, if $p \leq -1$ and choose $\varepsilon > 0$.*

If $m_p(\rho_\varepsilon^+, \rho_\varepsilon^-) < 1 + \varepsilon$, $\delta^+ := \frac{n}{\tilde{k}}(p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) > 3 \ln(8n/\varepsilon)$, and $\delta^- := \frac{n}{\tilde{k}}(p_{\text{in}}^- + (k-1)p_{\text{out}}^-) > 3 \ln(8n/\varepsilon)$, then there exists an orthogonal matrix $O_{\tilde{k}} \in \mathbb{R}^{\tilde{k} \times \tilde{k}}$ such that, with probability at least $1 - \varepsilon$, we have

$$\|V_{\tilde{k}} - \mathcal{V}_{\tilde{k}} O_{\tilde{k}}\| \leq \frac{\sqrt{8\tilde{k}} C_p m_{|p|}^{1/|p|} \left(\sqrt{\frac{3 \ln(8n/\varepsilon)}{\delta^+}}, \sqrt{\frac{3 \ln(8n/\varepsilon)}{\delta^-}} \right)}{(1 + \varepsilon) - m_p(\rho_\varepsilon^+, \rho_\varepsilon^-)}$$

Note that the main difference compared to Thm. 5 is the spectral gap $\gamma_p = (1 + \varepsilon) - m_p(\rho_\varepsilon^+, \rho_\varepsilon^-)$ of \mathcal{L}_p , which is the difference of the eigenvalues corresponding to the informative versus non-informative eigenvectors of \mathcal{L}_p . Thus the stronger the clustering structure the tighter is the concentration of the eigenvectors. Moreover, from the monotonicity of m_p we have $\gamma_p \geq \gamma_q$ for $p < q$, and thus for $p \leq -1$ the spectral gap increases with $|p|$, ensuring a stronger concentration of eigenvectors for smaller values of p .

4. Experiments on Wikipedia-Elections

We now evaluate the Signed Power Mean Laplacian L_p with $p \in \{-10, -5, -2, -1, 0, 1\}$ on Wikipedia-Elections dataset (Leskovec & Krevl, 2014). In this dataset each node represents an editor requesting to become administrator and positive (resp. negative) edges represent supporting (resp. against) votes to the corresponding admin candidate.

While (Chiang et al., 2012) conjectured that this dataset has no clustering structure, recent works (Mercado et al., 2016; Cucuringu et al., 2019) have shown that indeed there is clustering structure. As noted in (Mercado et al., 2016), using the geometric mean Laplacian L_{GM} and looking for k clusters unveils the presence of a large non-informative cluster and $k - 1$ remaining smaller clusters which show relevant clustering structure.

Our results verify these recent findings. We set the number of clusters to identify to $k = 30$ and in Fig. 4 we portray the portion of adjacency matrices of positive and negative edges W^+ and W^- corresponding to $k - 1$ clusters sorted according to the corresponding identified clusters. We can see that when $p \leq 0$ the Signed Power Mean Laplacian L_p identifies clustering structure, whereas this structure is overlooked by the arithmetic mean case $p = 1$. Moreover, we can see that different powers identify slightly different clusters: this happens as this dataset does not necessarily follow the Signed Stochastic Block Model, and hence we do not fully retrieve the same behaviour studied in Section 3.

Further experiments on UCI datasets are available in the supplementary material, suggesting that the L_{GM} together with L_p is a reasonable option under different settings.

Acknowledgments The work of F.T. has been funded by the Marie Curie Individual Fellowship MAGNET n. 744014.

References

- Abbe, E. Community detection and stochastic block models: Recent developments. *Journal of Machine Learning Research*, 18(177):1–86, 2018.
- Abbe, E., Bandeira, A. S., Bracher, A., and Singer, A. Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery. *IEEE Transactions on Network Science and Engineering*, 1(1):10–22, Jan 2014.
- Bansal, N., Blum, A., and Chawla, S. Correlation clustering. *Machine Learning*, 56(1):89–113, Jul 2004.
- Bhagwat, K. V. and Subramanian, R. Inequalities between means of positive operators. *Mathematical Proceedings of the Cambridge Philosophical Society*, 83(3):393401, 1978.
- Bhatia, R. *Matrix Analysis*. Springer New York, 1997.
- Bhatia, R. *Positive definite matrices*. Princeton University Press, 2009.
- Bosch, J., Mercado, P., and Stoll, M. Node classification for signed networks using diffuse interface methods. *arXiv:1809.06432*, 2018.
- Bullen, P. S. *Handbook of means and their inequalities*, volume 560. Springer Science & Business Media, 2013.
- Cartwright, D. and Harary, F. Structural balance: a generalization of Heider’s theory. *Psychological Review*, 63(5):277–293, 1956.
- Chaudhuri, K., Chung, F., and Tsias, A. Spectral clustering of graphs with general degrees in the extended planted partition model. In *COLT*, 2012.
- Chiang, K., Whang, J., and Dhillon, I. Scalable clustering of signed networks using balance normalized cut. *CIKM*, 2012.
- Chiang, K.-Y., Natarajan, N., Tewari, A., and Dhillon, I. S. Exploiting longer cycles for link prediction in signed networks. *CIKM*, 2011.
- Chung, F. and Radcliffe, M. On the spectra of general random graphs. *the electronic journal of combinatorics*, 18(1):P215, 2011.
- Chung, F., Tsias, A., and Xu, W. Dirichlet pagerank and ranking algorithms based on trust and distrust. *Internet Mathematics*, 9(1):113–134, 2013.
- Cucuringu, M., Pizzoferrato, A., and van Gennip, Y. An MBO scheme for clustering and semi-supervised clustering of signed networks. *arXiv:1901.03091*, 2018.
- Cucuringu, M., Davies, P., Glielmo, A., and Tyagi, H. SPONGE: A generalized eigenproblem for clustering signed networks. In *AISTATS*, 2019.
- Davis, E. and Sethuraman, S. Consistency of modularity clustering on random geometric graphs. *Ann. Appl. Probab.*, 28(4):2003–2062, 08 2018.
- Davis, J. A. Clustering and structural balance in graphs. *Human Relations*, 20:181–187, 1967.
- Derr, T., Ma, Y., and Tang, J. Signed Graph Convolutional Network. *arXiv:1808.06354*, 2018.
- Desai, M. and Rao, V. A characterization of the smallest eigenvalue of a graph. *Journal of Graph Theory*, 18(2):181–194, 1994.
- Doreian, P. and Mrvar, A. Partitioning signed social networks. *Social Networks*, 31(1):1–11, 2009.
- Falher, G. L., Cesa-Bianchi, N., Gentile, C., and Vitale, F. On the Troll-Trust Model for Edge Sign Prediction in Social Networks. In *AISTATS*, 2017.
- Fasino, D. and Tudisco, F. A modularity based spectral method for simultaneous community and anti-community detection. *Linear Algebra and its Applications*, 542:605–623, 2018.
- Fujita, A., Severino, P., Kojima, K., Sato, J. R., Patriota, A. G., and Miyano, S. Functional clustering of time series gene expression data by granger causality. *BMC systems biology*, 6(1):137, 2012.
- Gallier, J. Spectral theory of unsigned and signed graphs. applications to graph clustering: a survey. *arXiv:1601.04692*, 2016.
- Giotis, I. and Guruswami, V. Correlation clustering with a fixed number of clusters. In *SODA*, 2006.
- Han, Q., Xu, K. S., and Airolidi, E. M. Consistent estimation of dynamic and multi-layer block models. In *ICML*, 2015.
- Harary, F. On the notion of balance of a signed graph. *Michigan Mathematical Journal*, 2:143–146, 1953.
- Heimlicher, S., Lelarge, M., and Massoulié, L. Community detection in the labelled stochastic block model. *arXiv:1209.2910*, 2012.
- Holland, P. W., Laskey, K. B., and Leinhardt, S. Stochastic blockmodels: First steps. *Social Networks*, 5(2):109 – 137, 1983.
- Jog, V. and Loh, P.-L. Information-theoretic bounds for exact recovery in weighted stochastic block models using the Renyi divergence. *arXiv:1509.06418*, 2015.

- Joseph, A. and Yu, B. Impact of regularization on spectral clustering. *Ann. Statist.*, 44(4):1765–1791, 08 2016.
- Kim, J., Park, H., Lee, J.-E., and Kang, U. Side: Representation learning in signed directed networks. In *WWW*, 2018.
- Kirkley, A., Cantwell, G. T., and Newman, M. Balance in signed networks. *arXiv:1809.05140*, 2018.
- Knyazev, A. On spectral partitioning of signed graphs. In *SIAM Workshop on Combinatorial Scientific Computing*, 2018.
- Kumar, S., Spezzano, F., Subrahmanian, V., and Faloutsos, C. Edge weight prediction in weighted signed networks. In *ICDM*, 2016.
- Kunegis, J., Schmidt, S., Lommatzsch, A., Lerner, J., Luca, E., and Albayrak, S. Spectral analysis of signed graphs for clustering, prediction and visualization. In *ICDM*, 2010.
- Le, C. M., Levina, E., and Vershynin, R. Concentration and regularization of random graphs. *Random Structures & Algorithms*, 51(3):538–561.
- Lei, J. and Rinaldo, A. Consistency of spectral clustering in stochastic block models. *Ann. Statist.*, 43(1):215–237, 02 2015.
- Leskovec, J. and Krevl, A. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>, June 2014.
- Leskovec, J., Huttenlocher, D., and Kleinberg, J. Predicting positive and negative links in online social networks. In *WWW*, 2010a.
- Leskovec, J., Huttenlocher, D., and Kleinberg, J. Signed Networks in Social Media. In *CHI*, 2010b.
- Liu, S. Multi-way dual cheeger constants and spectral bounds of graphs. *Advances in Mathematics*, 268:306–338, 2015.
- Luxburg, U. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, December 2007.
- Mercado, P., Tudisco, F., and Hein, M. Clustering signed networks with the geometric mean of Laplacians. In *NIPS*. 2016.
- Mercado, P., Gautier, A., Tudisco, F., and Hein, M. The power mean laplacian for multilayer graph clustering. In *AISTATS*, 2018.
- Paul, S. and Chen, Y. Consistency of community detection in multi-layer networks using spectral and matrix factorization methods. *arXiv:1704.07353*, 2017.
- Pavlidis, N. G., Plagianakos, V. P., Tasoulis, D. K., and Vrahatis, M. N. Financial forecasting through unsupervised clustering and neural networks. *Operational Research*, 6(2):103–127, 2006.
- Qin, T. and Rohe, K. Regularized spectral clustering under the degree-corrected stochastic blockmodel. In *NIPS*. 2013.
- Rohe, K., Chatterjee, S., Yu, B., et al. Spectral clustering and the high-dimensional stochastic blockmodel. *The Annals of Statistics*, 39(4):1878–1915, 2011.
- Saade, A., Krzakala, F., and Zdeborová, L. Spectral clustering of graphs with the bethe hessian. In *NIPS*. 2014.
- Saade, A., Lelarge, M., Krzakala, F., and Zdeborov, L. Spectral detection in the censored block model. In *2015 IEEE International Symposium on Information Theory (ISIT)*, pp. 1184–1188, June 2015.
- Sarkar, P. and Bickel, P. J. Role of normalization in spectral clustering for stochastic blockmodels. *Ann. Statist.*, 43(3):962–990, 06 2015.
- Sedoc, J., Gallier, J., Foster, D., and Ungar, L. Semantic word clusters using signed spectral clustering. In *ACL*, 2017.
- Shahriari, M. and Jalili, M. Ranking nodes in signed social networks. *Social Network Analysis and Mining*, 4(1):172, Jan 2014. ISSN 1869-5469.
- Tang, J., Aggarwal, C., and Liu, H. Node classification in signed social networks. In *SDM*, 2016a.
- Tang, J., Chang, Y., Aggarwal, C., and Liu, H. A survey of signed network mining in social media. *ACM Comput. Surv.*, 49(3):42:1–42:37, August 2016b. ISSN 0360-0300.
- Tropp, J. A. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015. ISSN 1935-8237.
- Wang, S., Tang, J., Aggarwal, C., Chang, Y., and Liu, H. Signed network embedding in social media. In *SDM*, 2017.
- Xu, J., Massouli, L., and Lelarge, M. Edge label inference in generalized stochastic block models: from spectral theory to impossibility results. In *COLT*, 2014.
- Xu, M., Jog, V., and Loh, P.-L. Optimal rates for community estimation in the weighted stochastic block model. *arXiv:1706.01175*, 2017.
- Yu, Y., Wang, T., and Samworth, R. J. A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2015.

Yuan, S., Wu, X., and Xiang, Y. Sne: Signed network embedding. In *PAKDD*, 2017.

Yun, S.-Y. and Proutiere, A. Optimal cluster recovery in the labeled stochastic block model. In *NIPS*. 2016.

Ziegler, H., Jenny, M., Gruse, T., and Keim, D. A. Visual market sector analysis for financial time series data. In *Visual Analytics Science and Technology (VAST), 2010 IEEE Symposium on*, pp. 83–90. IEEE, 2010.

Appendices

This section contains all the proofs of results mention in the main paper. It is organized as follows:

- Section **A** contains the proof of Theorem 1,
- Section **B** contains the proof of Corollary 1,
- Section **C** contains the proof of Corollary 2,
- Section **D** contains the proof of Theorem 2,
- Section **I** contains the proof of Theorem 4
- Section **E** contains the proof of Theorem 5,
- Section **G** contains the proof of 6.
- Section **J** contains experiments on the Censored Block Model where the superiority of the Bethe Hessian on the sparse regime is verified,
- Section **K** contains experiments on UCI datasets,
- Section **L** contains an analysis on diagonal shift of the signed power mean Laplacian,
- Section **M** contains a description on the numerical scheme for the computation of eigenvectors without ever computing the power mean Laplacian matrix,
- Section **N** contains a further study on conditions of expectation of the power mean Laplacian and state of the art approaches.
- Section **O** contains a more detailed inspection of results on Wikipedia-Elections dataset presented in Section 4.

Before starting with the general results, we first mention a couple of basic results.

The following theorem states the monotonicity of the scalar power mean.

Theorem 7 ((Bullen, 2013), Ch. 3, Thm. 1). *Let $p < q$ then $m_p(a, b) \leq m_q(a, b)$ with equality if and only if $a = b$.*

The following lemma shows the effect of the matrix power mean when matrices have a common eigenvector.

Lemma 1 ((Mercado et al., 2018)). *Let \mathbf{u} be an eigenvector of both A and B , with corresponding eigenvalues α and β . Then \mathbf{u} is an eigenvector of $M_p(A, B)$ with eigenvalue $m_p(\alpha, \beta)$.*

A. Proof of Theorem 1

Proof. (Proof of Theorem 1) We first show that χ_1, \dots, χ_k are eigenvectors of \mathcal{W}^+ and \mathcal{W}^- . For χ_1 we have,

$$\mathcal{W}^+ \chi_1 = \mathcal{W}^+ \mathbf{1} = |\mathcal{C}| (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) \mathbf{1} = d^+ \mathbf{1} = \lambda_1^+ \mathbf{1}$$

For the remaining vectors χ_2, \dots, χ_k we have

$$\begin{aligned} \mathcal{W}^+ \chi_i &= \mathcal{W}^+ ((k-1)\mathbf{1}_{\mathcal{C}_i} - \mathbf{1}_{\overline{\mathcal{C}_i}}) \\ &= \mathcal{W}^+ (k\mathbf{1}_{\mathcal{C}_i} - (\mathbf{1}_{\mathcal{C}_i} + \mathbf{1}_{\overline{\mathcal{C}_i}})) \\ &= \mathcal{W}^+ (k\mathbf{1}_{\mathcal{C}_i} - \mathbf{1}) \end{aligned}$$

$$\begin{aligned}
 &= k |\mathcal{C}| (p_{\text{in}}^+ \mathbf{1}_{\mathcal{C}_i} + p_{\text{out}}^+ \mathbf{1}_{\overline{\mathcal{C}_i}}) - d^+ \mathbf{1} \\
 &= k |\mathcal{C}| (p_{\text{in}}^+ \mathbf{1}_{\mathcal{C}_i} + p_{\text{out}}^+ \mathbf{1}_{\overline{\mathcal{C}_i}}) - d^+ (\mathbf{1}_{\mathcal{C}_i} + \mathbf{1}_{\overline{\mathcal{C}_i}}) \\
 &= |\mathcal{C}| (k p_{\text{in}}^+ - d^+) \mathbf{1}_{\mathcal{C}_i} + |\mathcal{C}| (k p_{\text{out}}^+ - d^+) \mathbf{1}_{\overline{\mathcal{C}_i}} \\
 &= |\mathcal{C}| (k-1) (p_{\text{in}}^+ - p_{\text{out}}^+) \mathbf{1}_{\mathcal{C}_i} - |\mathcal{C}| (p_{\text{in}}^+ - p_{\text{out}}^+) \mathbf{1}_{\overline{\mathcal{C}_i}} \\
 &= |\mathcal{C}| (p_{\text{in}}^+ - p_{\text{out}}^+) ((k-1) \mathbf{1}_{\mathcal{C}_i} - \mathbf{1}_{\overline{\mathcal{C}_i}}) \\
 &= |\mathcal{C}| (p_{\text{in}}^+ - p_{\text{out}}^+) \chi_i \\
 &= \lambda_i \chi_i
 \end{aligned}$$

The same procedure holds for \mathcal{W}^- . Thus, we have shown that χ_1, \dots, χ_k are eigenvectors of both \mathcal{W}^+ and \mathcal{W}^- . In particular, we have seen that

$$\begin{aligned}
 \lambda_1^+ &= |\mathcal{C}| (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+), & \lambda_i^+ &= |\mathcal{C}| (p_{\text{in}}^+ - p_{\text{out}}^+) \\
 \lambda_1^- &= |\mathcal{C}| (p_{\text{in}}^- + (k-1)p_{\text{out}}^-), & \lambda_i^- &= |\mathcal{C}| (p_{\text{in}}^- - p_{\text{out}}^-)
 \end{aligned}$$

for $i = 2, \dots, k$. Further, as both matrices \mathcal{W}^+ and \mathcal{W}^- share all their eigenvectors, they are simultaneously diagonalizable, that is there exists a non-singular matrix Σ such that $\Sigma^{-1} \mathcal{W}^\pm \Sigma = \Lambda^\pm$, where Λ^+ and Λ^- are diagonal matrices $\Lambda^\pm = \text{diag}(\lambda_1^\pm, \dots, \lambda_k^\pm, 0, \dots, 0)$.

As we assume that all clusters are of the same size $|\mathcal{C}|$, the expected signed graph is a regular graph with degrees d^+ and d^- . Hence, the normalized Laplacian and normalized signless Laplacian of the expected signed graph can be expressed as

$$\begin{aligned}
 \mathcal{L}_{\text{sym}}^+ &= \Sigma (I - \frac{1}{d^+} \Lambda^+) \Sigma^{-1} \\
 \mathcal{Q}_{\text{sym}}^- &= \Sigma (I + \frac{1}{d^-} \Lambda^-) \Sigma^{-1}
 \end{aligned}$$

Thus, we can observe that

$$\begin{aligned}
 \lambda_1^+(\mathcal{L}_{\text{sym}}^+) &= 0, & \lambda_1^-(\mathcal{Q}_{\text{sym}}^-) &= 2 \\
 \lambda_i^+(\mathcal{L}_{\text{sym}}^+) &= 1 - \rho^+, & \lambda_i^-(\mathcal{Q}_{\text{sym}}^-) &= 1 + \rho^- \\
 \lambda_j^+(\mathcal{L}_{\text{sym}}^+) &= 1, & \lambda_j^-(\mathcal{Q}_{\text{sym}}^-) &= 1
 \end{aligned}$$

for $i = 2, \dots, k$, and $j = k+1, \dots, |V|$, where

$$\begin{aligned}
 \rho^+ &= (p_{\text{in}}^+ - p_{\text{out}}^+) / (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) \\
 \rho^- &= (p_{\text{in}}^- - p_{\text{out}}^-) / (p_{\text{in}}^- + (k-1)p_{\text{out}}^-)
 \end{aligned}$$

By obtaining the signed power mean Laplacian on diagonally shifted matrices,

$$\mathcal{L}_p = M_p(\mathcal{L}_{\text{sym}}^+ + \varepsilon I, \mathcal{Q}_{\text{sym}}^- + \varepsilon I)$$

we have by Lemma 1

$$\begin{aligned}
 \lambda_1(\mathcal{L}_p) &= m_p(\lambda_1^+ + \varepsilon, \lambda_1^- + \varepsilon) = m_p(\varepsilon, 2 + \varepsilon) \\
 \lambda_i(\mathcal{L}_p) &= m_p(1 - \rho^+ + \varepsilon, 1 + \rho^- + \varepsilon) \\
 \lambda_j(\mathcal{L}_p) &= m_p(\lambda_j^+ + \varepsilon, \lambda_j^- + \varepsilon) = 1 + \varepsilon
 \end{aligned} \tag{2}$$

Observe that $\lambda_j(\mathcal{L}_p)$, with $j = k+1, \dots, |V|$, corresponds to eigenvectors that do not yield an informative embedding. Hence, we do not want this eigenvalue to belong to the bottom of the spectrum of \mathcal{L}_p . Thus, for the case of χ_2, \dots, χ_k , we can see that they will be located at the bottom of the spectrum if the following condition holds:

$$\lambda_i(\mathcal{L}_p) = m_p(1 - \rho^+ + \varepsilon, 1 + \rho^- + \varepsilon) < 1 + \varepsilon = \lambda_j(\mathcal{L}_p)$$

It remains to analyze the case of the constant eigenvector χ_1 . Note that its associated eigenvalue $\lambda_1(\mathcal{L}_1)$ has the following relationship to the non-informative eigenvectors:

$$\lambda_1(\mathcal{L}_1) = m_1(\varepsilon, 2 + \varepsilon) = 1 + \varepsilon = \lambda_j(\mathcal{L}_p)$$

By Theorem 7 we know that the scalar power mean is monotone in its parameter p , and thus, for the case $p < 1$ we observe

$$\lambda_1(\mathcal{L}_p) = m_p(\varepsilon, 2 + \varepsilon) < m_1(\varepsilon, 2 + \varepsilon) = \lambda_1(\mathcal{L}_1) = \lambda_j(\mathcal{L}_p)$$

and for the case $p \geq 1$ we observe

$$\lambda_1(\mathcal{L}_p) = m_p(\varepsilon, 2 + \varepsilon) \geq m_1(\varepsilon, 2 + \varepsilon) = \lambda_1(\mathcal{L}_1) = \lambda_j(\mathcal{L}_p)$$

This means that for positive powers $p \geq 1$, the constant eigenvector χ_1 does not belong to the bottom of the spectrum, whereas for $p < 1$ it always does.

With this in mind, we reach the desired result, namely

- Let $p \geq 1$. $\{\chi_i\}_{i=2}^k$ correspond to the $(k-1)$ -smallest eigenvalues of \mathcal{L}_p if and only if $m_p(\boldsymbol{\mu} + \varepsilon \mathbf{1}) < 1 + \varepsilon$,
- Let $p < 1$. $\{\chi_i\}_{i=1}^k$ correspond to the k -smallest eigenvalues of \mathcal{L}_p if and only if $m_p(\boldsymbol{\mu} + \varepsilon \mathbf{1}) < 1 + \varepsilon$

□

B. Proof of Corollary 1

Proof of Corollary 1. Following the proof from Theorem 1, recall that $\lim_{p \rightarrow \infty} m_p(x) = \max\{x_1, \dots, x_T\}$ and $\lim_{p \rightarrow -\infty} m_p(x) = \min\{x_1, \dots, x_T\}$.

Thus, $m_\infty(\boldsymbol{\mu} + \varepsilon I) = \max(1 - \rho^+, 1 + \rho^-)$, and hence $m_\infty(\boldsymbol{\mu} + \varepsilon I) < 1 + \varepsilon$ if and only if $\rho^+ > 0$ and $\rho^- < 0$, yielding the desired conditions.

The case for $p \rightarrow -\infty$ is analogous: $m_{-\infty}(\boldsymbol{\mu} + \varepsilon I) = \min(1 - \rho^+, 1 + \rho^-)$ and thus $m_{-\infty}(\boldsymbol{\mu} + \varepsilon I) < 1 + \varepsilon$ if and only if $\rho^+ > 0$ or $\rho^- < 0$, yielding the desired conditions. □

C. Proof of Corollary 2

Proof of Corollary 2. If $\lambda_1, \dots, \lambda_k$ resp. $(\lambda_2, \dots, \lambda_k)$ are among the k (resp. $k-1$)-smallest eigenvalues of \mathcal{L}_p , then by Theorem 1, we have $m_p(\boldsymbol{\mu} + \varepsilon \mathbf{1}) < 1 + \varepsilon$. By Theorem 7 we have $m_q(\boldsymbol{\mu} + \varepsilon \mathbf{1}) \leq m_p(\boldsymbol{\mu} + \varepsilon \mathbf{1})$, Theorem 1 concludes the proof. □

D. Proof of Theorem 2

Proof of Theorem 2. Following the proof from Theorem 1 we can see that $\mathcal{L}_{\text{sym}}^+$ and $\mathcal{Q}_{\text{sym}}^-$ share all of their eigenvectors. Let \mathbf{u} be an eigenvector of $\mathcal{L}_{\text{sym}}^+$ and $\mathcal{Q}_{\text{sym}}^-$ with eigenvalues α and β , respectively.

By Lemma 1 we have

$$\mathcal{L}_0 \mathbf{u} = m_0(\alpha, \beta) \mathbf{u}$$

Moreover, from Theorem 1 in (Mercado et al., 2016) we know that

$$(\mathcal{L}_{\text{sym}}^+ \# \mathcal{Q}_{\text{sym}}^-) \mathbf{u} = \sqrt{\alpha\beta} \mathbf{u}$$

Further, $m_0(\alpha, \beta) = \sqrt{\alpha\beta}$. Hence, as \mathcal{L}_0 and $\mathcal{L}_{\text{sym}}^+ \# \mathcal{Q}_{\text{sym}}^-$ have in common all eigenvectors and eigenvalues, we conclude that $\mathcal{L}_0 = \mathcal{L}_{\text{sym}}^+ \# \mathcal{Q}_{\text{sym}}^-$. \square

E. Proof of Theorem 5

For the proof of Theorem 5, we first present Theorem 8 which is a general version that allows to choose different diagonal shifts of the Laplacians together with different edge probabilities.

Theorem 8. *Let G^+ and G^- be random graphs with independent edges $\mathbb{P}(W_{i,j}^+ = 1) = p_{ij}^+$ and $\mathbb{P}(W_{i,j}^- = 1) = p_{ij}^-$. Let δ^+ , δ^- be the minimum expected degrees of G^+ and G^- , respectively. Let $C_p^+ = p^{1/p} \beta^{1-1/p}$, and $C_p^- = |p|^{1/|p|} \alpha^{-(3+1/|p|)}$. Choose $\epsilon > 0$. Then there exist constants $k^+ = k^+(\epsilon/2)$ and $k^- = k^-(\epsilon/2)$ such that if $\delta^+ > k^+ \ln n$, and $\delta^- > k^- \ln n$ then with probability at least $1 - \epsilon$,*

$$\|L_p - \mathcal{L}_p\| \leq C_p^+ m_p \left(2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^+}}, 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^-}} \right)^{1/p}$$

for $p \geq 1$, with p integer and

$$\|L_p - \mathcal{L}_p\| \leq C_p^- m_{|p|} \left(2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^+}}, 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^-}} \right)^{1/|p|}$$

for $p \leq -1$, with p integer, where $L_p = M_p(L_{\text{sym}}^+ + \alpha I, Q_{\text{sym}}^- + \alpha I)$, and $\mathcal{L}_p = M_p(\mathcal{L}_{\text{sym}}^+ + \alpha I, \mathcal{Q}_{\text{sym}}^- + \alpha I)$.

Before starting the proof of Theorem 8, we present an upper bound on the matrix power mean.

Theorem 9. *Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric matrices where $\alpha \leq \lambda(A_i) \leq \beta$, $\alpha \leq \lambda(B_i) \leq \beta$ for $i = 1, \dots, T$ and $\alpha, \beta > 0$.*

Let $C_p^+ = p^{1/p} \beta^{1-1/p}$ and $C_p^- = |p|^{1/|p|} \alpha^{-(3+1/|p|)}$. Then, for $p \geq 1$, with p integer

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq C_p^+ m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/p} \end{aligned}$$

and, for $p \leq -1$, with p integer

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq C_p^- m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/|p|} \end{aligned}$$

Proof. The proof is contained in Section F. \square

Observe that the upper bound in Theorem 9 is general in the sense that it is suitable for symmetric definite matrices with bounded spectrum, and for an arbitrary number of matrices.

We are now ready to prove Theorem 8.

Proof of Theorem 8. Let

$$\begin{aligned} A_1 &= L_{\text{sym}}^+, & B_1 &= \mathcal{L}_{\text{sym}}^+ \\ A_2 &= Q_{\text{sym}}^-, & B_2 &= \mathcal{Q}_{\text{sym}}^- \end{aligned}$$

with the corresponding signed power mean Laplacian

$$L_p = M_p(L_{\text{sym}}^+, Q_{\text{sym}}^-), \quad \mathcal{L}_p = M_p(\mathcal{L}_{\text{sym}}^+, \mathcal{Q}_{\text{sym}}^-)$$

We start with the case $p \geq 1$, with p integer. Let $C_p^+ = p^{1/p} \beta^{1-1/p}$. By Theorem 9 we have

$$\|L_p - \mathcal{L}_p\| \leq C_p^+ m_p(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|, \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|)^{1/p}$$

Let $\gamma = (\gamma_1, \gamma_2)$ where

$$\gamma_1 = 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^+}} \quad \gamma_2 = 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^-}}$$

Define $a = c m_p(\gamma)$ and $c = C_p^+$. Then,

$$\begin{aligned} & \mathbb{P}(\|L_p - \mathcal{L}_p\| > a) \\ & \leq \mathbb{P}\left(c m_p(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|, \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|) > a\right) \\ & = \mathbb{P}\left(m_p(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|, \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|) > \frac{a}{c}\right) \\ & = \mathbb{P}\left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|^p + \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|^p > 2\left(\frac{a}{c}\right)^p\right) \\ & = \mathbb{P}\left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|^p + \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|^p > \sum_{i=1}^2 \gamma_i^p\right) \\ & \leq \mathbb{P}\left(\left\{\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|^p > \gamma_1^p\right\} \cup \left\{\|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|^p > \gamma_2^p\right\}\right) \\ & \leq \mathbb{P}\left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\|^p > \gamma_1^p\right) \\ & \quad + \mathbb{P}\left(\|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\|^p > \gamma_2^p\right) \end{aligned} \tag{3}$$

$$\begin{aligned}
 &= \mathbb{P} \left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\| > \gamma_1 \right) \\
 &\quad + \mathbb{P} \left(\|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\| > \gamma_2 \right) \\
 &= \mathbb{P} \left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\| > 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^+}} \right) \\
 &\quad + \mathbb{P} \left(\|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\| > 2\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^-}} \right) \\
 &= \mathbb{P} \left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\| > 2\sqrt{\frac{3 \ln(4n/\hat{\epsilon})}{\delta^+}} \right) \\
 &\quad + \mathbb{P} \left(\|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\| > 2\sqrt{\frac{3 \ln(4n/\hat{\epsilon})}{\delta^-}} \right) \\
 &= \mathbb{P} \left(\|L_{\text{sym}}^+ - \mathcal{L}_{\text{sym}}^+\| > 2\sqrt{\frac{3 \ln(4n/\hat{\epsilon})}{\delta^+}} \right) \\
 &\quad + \mathbb{P} \left(\|L_{\text{sym}}^- - \mathcal{L}_{\text{sym}}^-\| > 2\sqrt{\frac{3 \ln(4n/\hat{\epsilon})}{\delta^-}} \right) \\
 &\leq \hat{\epsilon} + \hat{\epsilon} \\
 &= \epsilon
 \end{aligned} \tag{4}$$

where $\hat{\epsilon} = \epsilon/2$. Inequality (3) follows from Boole's inequality. Inequality (4) comes from applying Theorem 12 from (Chung & Radcliffe, 2011) to G^+ and G^- , with corresponding minimum expected degree δ^+ , and δ^- , respectively, and $\hat{\epsilon}$, and

$$\begin{aligned}
 \|Q_{\text{sym}}^- - \mathcal{Q}_{\text{sym}}^-\| &= \|(I + T) - (I + \mathcal{T})\| \\
 &= \|(I - T) - (I - \mathcal{T})\| \\
 &= \|L_{\text{sym}}^- - \mathcal{L}_{\text{sym}}^-\|
 \end{aligned}$$

where

$$\begin{aligned}
 T &= (D^-)^{-1/2} W^- (D^-)^{-1/2} \\
 \mathcal{T} &= (\mathcal{D}^-)^{-1/2} \mathcal{W}^- (\mathcal{D}^-)^{-1/2}
 \end{aligned}$$

Thus,

$$\mathbb{P} \left(\|L_p - \mathcal{L}_p\| \geq a \right) < \epsilon$$

and hence

$$\mathbb{P} \left(\|L_p - \mathcal{L}_p\| \leq a \right) < 1 - \epsilon$$

completing the proof for the case $p \geq 1$.

For the proof of the case $p \leq -1$ with p integer, let $c = |p|^{1/|p|} \alpha^{-(3+1/|p|)}$, and proceed as for the previous case with $|p|$. \square

We now finally give the proof for Theorem 5.

Proof of Theorem 5. We will adapt to our particular case the general version presented in Theorem 8. We do this by showing that our Stochastic Block Model approach together with the shift of our model are particular cases of Theorem 8.

First, note that the spectrum of the normalized Laplacians L_{sym}^+ and Q_{sym}^- is upper bounded by two, i.e. $\lambda(L_{\text{sym}}^+), \lambda(Q_{\text{sym}}^-) \in [0, 2]$. Hence, by adding a diagonal shift we get $\lambda(L_{\text{sym}}^+ + \alpha I), \lambda(Q_{\text{sym}}^- + \alpha I) \in [\alpha, 2 + \alpha]$. Letting $\alpha = \epsilon$ and $\beta = 2 + \alpha$ we get the shift corresponding to the particular case from Theorem 5.

Further, observe that our SBM model is obtained by setting $p_{ij}^+ = p_{\text{in}}^+$ and $p_{ij}^- = p_{\text{in}}^-$ if v_i, v_j belong to the same cluster and $p_{ij}^+ = p_{\text{out}}^+$ and $p_{ij}^- = p_{\text{out}}^-$ if v_i, v_j belong to different clusters.

Moreover, under the Stochastic Block Model here considered, the induced expected graphs are regular, and thus all nodes have the same degree. Hence, the minimum expected degrees of G^+ and G^- are

$$\begin{aligned}
 \delta^+ &= \frac{n}{k} (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) \\
 \delta^- &= \frac{n}{k} (p_{\text{in}}^- + (k-1)p_{\text{out}}^-)
 \end{aligned}$$

Thus, taking these settings into Theorem 8 we get the desired result, except that the condition on the minimum expected degrees is that there exists constants $k^+ = k^+(\epsilon/2)$, and $k^- = k^-(\epsilon/2)$ such that the desired concentration holds.

To overcome this, observe in the proof of Theorem 12 (p.9) that the condition $\delta > k \ln(n)$ comes from the requirement

$$\sqrt{\frac{3 \ln(4n/\epsilon)}{\delta}} < 1$$

Thus, by setting $\delta > 3 \ln(4n/\epsilon)$ the condition is fulfilled. In our case, this yields to $\delta^+ > 3 \ln(4n/\hat{\epsilon}) = 3 \ln(8n/\epsilon)$ and $\delta^- > 3 \ln(4n/\hat{\epsilon}) = 3 \ln(8n/\epsilon)$, leading to the desired result. \square

F. Proof of Theorem 9

Before going into the proof, a set of preliminary results are necessary.

In what follows, for Hermitian matrices A and B we mean by $A \preceq B$ that $B - A$ is positive semidefinite (see (Bhatia, 1997), Ch. 5, and (Tropp, 2015), Ch. 2.1.8 for more details). We now proceed with the definition of a operator monotone function:

Definition 2 ((Tropp, 2015) Ch. 8.4.2, (Bhatia, 1997) Ch. 5.). *Let $f: I \rightarrow \mathbb{R}$ be a function on an interval I of the*

real line. The function f is operator monotone on I when $A \preceq B$ implies $f(A) \preceq f(B)$ for all Hermitian matrices A and B whose eigenvalues are contained in I .

The following result states that the negative inverse is operator monotone.

Proposition 1 ((Bhatia, 1997), Prop. V.1.6, (Tropp, 2015), Prop. 8.4.3). *The function $f(t) = -\frac{1}{t}$ is operator monotone on $(0, \infty)$.*

The following result states that the effect of operator monotone functions can be upper bounded in a helpful way.

Theorem 10 ((Bhatia, 1997), Theorem. X.3.8). *Let f be an operator monotone function on $(0, \infty)$ and let A, B be two positive definite matrices that are bounded below by a ; i.e. $A \geq aI$ and $B \geq aI$ for the positive number a . Then for every unitarily invariant norm*

$$\| \|f(A) - f(B)\| \| \leq f'(a) \| \|A - B\| \|$$

Applying this to the case of the negative inverse leads to the following Corollary.

Corollary 3. *Let A, B be two positive definite matrices that are bounded below by a ; i.e. $A \geq aI$ and $B \geq aI$ for the positive number a . Then for every unitarily invariant norm*

$$\| \|A^{-1} - B^{-1}\| \| \leq \frac{1}{a^2} \| \|A - B\| \|$$

Proof. Let $f(t) = -\frac{1}{t}$. Then, by Proposition 1 we know that f is operator monotone. Since $f'(t) = 1/t^2$, it follows from Theorem 10

$$\begin{aligned} \| \|A^{-1} - B^{-1}\| \| &= \| \|f(A) - f(B)\| \| \leq \\ & f'(a) \| \|A - B\| \| = \frac{1}{a^2} \| \|A - B\| \| \end{aligned}$$

□

The next results states a useful result on positive powers between zero and one.

Corollary 4 ((Bhatia, 1997), Eq. X.2). *Let A, B be two positive semidefinite matrices. Then, for $0 \leq r \leq 1$*

$$\| \|A^r - B^r\| \| \leq \| \|A - B\| \|^r$$

Its equivalent to positive integer powers is stated in the following result.

Proposition 2 (See (Bhatia, 1997), Eq. IX.4). *For any two matrices X, Y , and for $m = 1, 2, \dots$,*

$$\| \|X^m - Y^m\| \| \leq m M^{m-1} \| \|X - Y\| \|$$

where $M = \max(\| \|X\| \|, \| \|Y\| \|)$.

Next we show that the spectrum of the matrix power mean is well bounded for positive powers larger than one.

Proposition 3. *Let A_1, \dots, A_T be symmetric positive definite matrices that are bounded below and above by α and β ; i.e. $\alpha I \leq A_i \leq \beta I$ for positive numbers α and β . Then, for $p \geq 1$, with p integer*

$$\alpha \leq \lambda(M_p(A_1, \dots, A_T)) \leq \beta$$

Proof. Let $S_p(A_1, \dots, A_T) = \frac{1}{T} \sum_{i=1}^T A_i^p$. Then

$$\begin{aligned} \langle x, S_p(A_1, \dots, A_T)x \rangle &= \langle x, \left(\frac{1}{T} \sum_{i=1}^T A_i^p \right) x \rangle \\ &= \frac{1}{T} \sum_{i=1}^T \langle x, A_i^p x \rangle \end{aligned}$$

Thus, we obtain the following upper bound

$$\begin{aligned} \max_{\|x\|=1} \langle x, S_p(A_1, \dots, A_T)x \rangle &= \max_{\|x\|=1} \frac{1}{T} \sum_{i=1}^T \langle x, A_i^p x \rangle \\ &\leq \frac{1}{T} \sum_{i=1}^T \max_{\|x\|=1} \langle x, A_i^p x \rangle \\ &\leq \beta^p \end{aligned}$$

Hence, $\lambda_{max}(S_p(A_1, \dots, A_T)) \leq \beta^p$, and thus we obtain the corresponding upper bound $\lambda_{max}(M_p(A_1, \dots, A_T)) \leq \beta$.

In a similar way we obtain the following lower bound,

$$\begin{aligned} \min_{\|x\|=1} \langle x, S_p(A_1, \dots, A_T)x \rangle &= \min_{\|x\|=1} \frac{1}{T} \sum_{i=1}^T \langle x, A_i^p x \rangle \\ &\geq \frac{1}{T} \sum_{i=1}^T \min_{\|x\|=1} \langle x, A_i^p x \rangle \\ &\geq \alpha^p \end{aligned}$$

Hence, $\lambda_{min}(S_p(A_1, \dots, A_T)) \geq \alpha^p$, and thus we obtain the corresponding lower bound $\lambda_{min}(M_p(A_1, \dots, A_T)) \geq \alpha$.

Therefore, $\alpha \leq \lambda(M_p(A_1, \dots, A_T)) \leq \beta$. □

We now present results for $p \geq 1$ of Theorem 9.

E.1. Results for the case $p \geq 1$

The following two propositions are the main ingredients for the upper bound presented in Theorem 9 for the case $p \geq 1$.

Proposition 4. Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive semidefinite matrices. Then, for $p \geq 1$ with p integer

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq \|M_p^p(A_1, \dots, A_T) - M_p^p(B_1, \dots, B_T)\|^{\frac{1}{p}} \end{aligned}$$

Proof. Let $S_p(A_1, \dots, A_T) = \frac{1}{T} \sum_{i=1}^T A_i^p$ and $r = 1/p$. Then,

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & = \left\| S_p^{1/p}(A_1, \dots, A_T) - S_p^{1/p}(B_1, \dots, B_T) \right\| \\ & = \left\| S_p^r(A_1, \dots, A_T) - S_p^r(B_1, \dots, B_T) \right\| \\ & \leq \|S_p(A_1, \dots, A_T) - S_p(B_1, \dots, B_T)\|^r \\ & = \|S_p(A_1, \dots, A_T) - S_p(B_1, \dots, B_T)\|^{1/p} \\ & = \|M_p^p(A_1, \dots, A_T) - M_p^p(B_1, \dots, B_T)\|^{1/p} \end{aligned}$$

where the inequality comes from Corollary 4, giving the desired result. \square

Proposition 5. Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive semidefinite matrices such that $\lambda(A_i) \leq \beta$ and $\lambda(B_i) \leq \beta$ for $i = 1, \dots, T$. Then, for $p \geq 1$,

$$\begin{aligned} & \|M_p^p(A_1, \dots, A_T) - M_p^p(B_1, \dots, B_T)\| \\ & \leq p\beta^{p-1} m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \end{aligned}$$

Proof. Let $\beta_i = \max(\|A_i\|, \|B_i\|)$. Then,

$$\begin{aligned} & \|M_p^p(A_1, \dots, A_T) - M_p^p(B_1, \dots, B_T)\| \\ & = \left\| \left(\frac{1}{T} \sum_{i=1}^T A_i^p \right) - \left(\frac{1}{T} \sum_{i=1}^T B_i^p \right) \right\| \\ & = \left\| \frac{1}{T} \sum_{i=1}^T (A_i^p - B_i^p) \right\| \\ & \leq \frac{1}{T} \sum_{i=1}^T \|A_i^p - B_i^p\| \\ & \leq \frac{1}{T} \sum_{i=1}^T p(\beta_i)^{p-1} \|A_i - B_i\| \\ & \leq \frac{1}{T} \sum_{i=1}^T p\beta^{p-1} \|A_i - B_i\| \\ & = p\beta^{p-1} \left(\frac{1}{T} \sum_{i=1}^T \|A_i - B_i\| \right) \\ & = p\beta^{p-1} m_1(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \\ & \leq p\beta^{p-1} m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \end{aligned}$$

where: the first inequality follows from the triangular inequality, the second inequality follows from Proposition 2, the third inequality follows as $\beta_i \leq \beta$, and the last inequality comes from the monotonicity of the scalar power means. \square

The next Lemma contains the proof corresponding to the case of positive powers of Theorem 9.

Lemma 2 (Theorem 9 for the case $p \geq 1$). Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive semidefinite matrices where $\lambda(A_i) \leq \beta$ and $\lambda(B_i) \leq \beta$ for $i = 1, \dots, T$. Let $C_p^+ = p^{1/p} \beta^{1-1/p}$. Let $p \geq 1$, with p integer Then,

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq C_p^+ m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/p} \end{aligned}$$

Proof. We can see that

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq \|M_p^p(A_1, \dots, A_T) - M_p^p(B_1, \dots, B_T)\|^{\frac{1}{p}} \\ & \leq \left(p\beta^{p-1} m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \right)^{\frac{1}{p}} \\ & = C_p^+ m_p(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/p} \end{aligned}$$

where the first inequality comes from Proposition 4, and the second inequality comes from Proposition 5. \square

F.2. Results for the case $p \leq -1$

The following two propositions are the main ingredients for the upper bound presented in Theorem 9 for the case $p \leq -1$.

Proposition 6. Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive definite matrices where $\alpha \leq \lambda(A_i)$ and $\alpha \leq \lambda(B_i)$ for $i = 1, \dots, T$, and $\alpha > 0$. Then, for $p \leq -1$, with p integer

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & \leq \frac{1}{\alpha^2} \left\| M_{|p|}^{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}^{|p|}(B_1^{-1}, \dots, B_T^{-1}) \right\|^{1/|p|} \end{aligned}$$

Proof. Let $S_p(A_1, \dots, A_T) = \frac{1}{T} \sum_{i=1}^T A_i^p$. Then,

$$\begin{aligned} & \|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\ & = \left\| S_p^{1/p}(A_1, \dots, A_T) - S_p^{1/p}(B_1, \dots, B_T) \right\| \\ & = \left\| S_p^{1/|p|}(A_1, \dots, A_T)^{-1} - S_p^{1/|p|}(B_1, \dots, B_T)^{-1} \right\| \\ & = \left\| S_{|p|}^{1/|p|}(A_1^{-1}, \dots, A_T^{-1})^{-1} - S_{|p|}^{1/|p|}(B_1^{-1}, \dots, B_T^{-1})^{-1} \right\| \end{aligned}$$

$$\begin{aligned}
 &= \|M_{|p|}(A_1^{-1}, \dots, A_T^{-1})^{-1} - M_{|p|}(B_1^{-1}, \dots, B_T^{-1})^{-1}\| \\
 &\leq \frac{1}{\alpha^2} \|M_{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}(B_1^{-1}, \dots, B_T^{-1})\| \\
 &\leq \frac{1}{\alpha^2} \|M_{|p|}^{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}^{|p|}(B_1^{-1}, \dots, B_T^{-1})\|^{1/|p|}
 \end{aligned}$$

where the first inequality follows from Corollary 3 and Proposition 3, whereas the second inequality follows from Proposition 4. \square

Proposition 7. Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive definite matrices such that $\alpha \leq \lambda(A_i)$ and $\alpha \leq \lambda(B_i)$ for $i = 1, \dots, T$. Then, for $p \leq -1$, with p integer

$$\begin{aligned}
 &\|M_{|p|}^{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}^{|p|}(B_1^{-1}, \dots, B_T^{-1})\| \\
 &\leq |p| \alpha^{-(1+|p|)} m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)
 \end{aligned}$$

Proof. Let $\alpha_i = \min(\|A_i\|, \|B_i\|)$, then it clearly follows that $\frac{1}{\alpha_i} = \max(\|A_i^{-1}\|, \|B_i^{-1}\|)$. Thus,

$$\begin{aligned}
 &\|M_{|p|}^{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}^{|p|}(B_1^{-1}, \dots, B_T^{-1})\| \\
 &= \left\| \left(\frac{1}{T} \sum_{i=1}^T (A_i^{-1})^{|p|} \right) - \left(\frac{1}{T} \sum_{i=1}^T (B_i^{-1})^{|p|} \right) \right\| \\
 &= \left\| \frac{1}{T} \sum_{i=1}^T (A_i^{-1})^{|p|} - (B_i^{-1})^{|p|} \right\| \\
 &\leq \frac{1}{T} \sum_{i=1}^T \left\| (A_i^{-1})^{|p|} - (B_i^{-1})^{|p|} \right\| \\
 &\leq \frac{1}{T} \sum_{i=1}^T |p| \left(\frac{1}{\alpha_i} \right)^{|p|-1} \|A_i^{-1} - B_i^{-1}\| \\
 &\leq |p| \left(\frac{1}{\alpha} \right)^{|p|-1} \left(\frac{1}{T} \sum_{i=1}^T \|A_i^{-1} - B_i^{-1}\| \right) \\
 &\leq |p| \left(\frac{1}{\alpha} \right)^{|p|-1} \left(\frac{1}{T\alpha^2} \sum_{i=1}^T \|A_i - B_i\| \right) \\
 &= |p| \left(\frac{1}{\alpha} \right)^{|p|+1} \left(\frac{1}{T} \sum_{i=1}^T \|A_i - B_i\| \right) \\
 &= |p| \left(\frac{1}{\alpha} \right)^{|p|+1} m_1(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \\
 &\leq |p| \left(\frac{1}{\alpha} \right)^{|p|+1} m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \\
 &= |p| \alpha^{-(1+|p|)} m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)
 \end{aligned}$$

where: the first inequality follows from the triangular inequality, the second inequality follows from Proposition 2, the third inequality follows as $\alpha \leq \alpha_i$, and the fourth inequality follows as Corollary 3, and the last inequality comes from the monotonicity of the scalar power means. \square

The next Lemma contains the proof corresponding to the case of negative powers of Theorem 9.

Lemma 3 (Theorem 9 for the case $p \leq -1$). Let $A_1, \dots, A_T, B_1, \dots, B_T$ be symmetric positive definite matrices where $\alpha \leq \lambda(A_i)$ and $\alpha \leq \lambda(B_i)$ for $i = 1, \dots, T$. Let $C_p^- = |p|^{1/|p|} \alpha^{-(3+1/|p|)}$. Let $p \leq -1$ with p integer. Then,

$$\begin{aligned}
 &\|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\
 &\leq C_p^- m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/|p|}
 \end{aligned}$$

Proof.

$$\begin{aligned}
 &\|M_p(A_1, \dots, A_T) - M_p(B_1, \dots, B_T)\| \\
 &\leq \frac{1}{\alpha^2} \|M_{|p|}^{|p|}(A_1^{-1}, \dots, A_T^{-1}) - M_{|p|}^{|p|}(B_1^{-1}, \dots, B_T^{-1})\|^{1/|p|} \\
 &\leq \frac{1}{\alpha^2} \left(|p| \alpha^{-(1+|p|)} m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|) \right)^{1/|p|} \\
 &= C_p^- m_{|p|}(\|A_1 - B_1\|, \dots, \|A_T - B_T\|)^{1/|p|}
 \end{aligned}$$

where the first inequality comes from Proposition 6, and the second inequality comes from Proposition 7. \square

We are now ready to prove the result of Theorem 9.

Proof of Theorem 9. For the case $p \geq 1$ see Lemma 2. For the case $p \leq -1$ see Lemma 3. \square

G. Proof of Theorem 6

Before giving the proof of Theorem 6 we need to present two auxiliary results.

The following is an auxiliary technical result that extends an implicit result stated in (Rohe et al., 2011)(p.1908-1909) for the Frobenius norm to the case of the operator norm.

Lemma 4. Let $X, \mathcal{X} \in \mathbb{R}^{n \times k}$ be matrices with orthonormal columns. Let U, V be orthonormal matrices and Σ a diagonal matrix such that

$$\mathcal{X}^T X = U \Sigma V^T$$

where the diagonal entries of Σ are the cosines of the principal angles between the column space of X and the column space of \mathcal{X} . Let $O = UV^T$. Then,

$$\frac{1}{\sqrt{2}} \|X - \mathcal{X}O\| \leq \|\sin \Theta(\mathcal{X}, X)\|$$

Proof. For the proof we will make use of the identity $X^T \mathcal{X}O = (V \Sigma U^T) UV^T = V \Sigma V^T$, and the fact that $\|A\| = \sqrt{\lambda_{\max}(A^T A)}$. That is,

$$(X - \mathcal{X}O)^T (X - \mathcal{X}O)$$

$$\begin{aligned}
 &= (X^T - O^T \mathcal{X}^T)(X - \mathcal{X}O) && \leq \sum_i^p \sin^2 \Theta_i \\
 &= X^T X - X^T \mathcal{X}O - O^T \mathcal{X}^T X + O^T \mathcal{X}^T \mathcal{X}O \\
 &= I - X^T \mathcal{X}O - O^T \mathcal{X}^T X + O^T O \\
 &= I - X^T \mathcal{X}O - O^T \mathcal{X}^T X + I \\
 &= 2I - X^T \mathcal{X}O - O^T \mathcal{X}^T X \\
 &= 2I - V \Sigma V^T - V \Sigma V^T \\
 &= 2(I - V \Sigma V^T)
 \end{aligned}$$

Thus,

$$\begin{aligned}
 \|X - \mathcal{X}O\|^2 &= \lambda_{\max} \left((X - \mathcal{X}O)^T (X - \mathcal{X}O) \right) \\
 &= 2\lambda_{\max}(I - V \Sigma V^T) \\
 &= 2 \max_i (1 - \cos \Theta_i) \\
 &\leq 2 \max_i (1 - \cos^2 \Theta_i) \\
 &= 2 \max_i (\sin^2 \Theta_i) \\
 &= 2 \|\sin \Theta\|^2
 \end{aligned}$$

Hence, $\frac{1}{\sqrt{2}} \|X - \mathcal{X}O\| \leq \|\sin \Theta(\mathcal{X}, X)\|$ \square

The next result is a useful representation of the Davis-Kahan theorem. It is a technical adaption from the Frobenius norm to the operator norm based on Lemma 4 and Theorem 14.

Theorem 11. *Let $\Sigma, \hat{\Sigma} \in \mathbb{R}^{p \times p}$ be symmetric, with eigenvalues $\mu_1 \geq \dots \geq \mu_p$ and $\hat{\mu}_1 \geq \dots \geq \hat{\mu}_p$ respectively. Fix $1 \leq r \leq s \leq p$ and assume that $\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1}) > 0$, where $\mu_0 := \infty$ and $\mu_{p+1} := -\infty$. Let $d := s - r + 1$, and let $V = (v_r, v_{r+1}, \dots, v_s) \in \mathbb{R}^{p \times d}$ and $\hat{V} = (\hat{v}_r, \hat{v}_{r+1}, \dots, \hat{v}_s) \in \mathbb{R}^{p \times d}$ have orthonormal columns satisfying $\Sigma v_j = \mu_j v_j$ and $\hat{\Sigma} \hat{v}_j = \hat{\mu}_j \hat{v}_j$ for $j = r, r+1, \dots, s$. Then there exists an orthogonal matrix $O \in \mathbb{R}^{d \times d}$ such that*

$$\frac{1}{\sqrt{2}} \|V - \hat{V}O\| \leq \frac{2d^{1/2} \|\hat{\Sigma} - \Sigma\|}{\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1})}$$

Proof. By theorem 14 we have

$$\|\sin \Theta(\hat{V}, V)\|_F \leq \frac{2 \min(d^{1/2} \|\hat{\Sigma} - \Sigma\|, \|\hat{\Sigma} - \Sigma\|_F)}{\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1})}.$$

From lemma 4 we can see that

$$\frac{1}{\sqrt{2}} \|V - \hat{V}O\| \leq \|\sin \Theta(\hat{V}, V)\|$$

Moreover, as $\sin \Theta(\hat{V}, V)$ is a diagonal matrix, it holds that

$$\|\sin \Theta(\hat{V}, V)\|^2 = \max_i (\sin^2 \Theta_i)$$

Thus,

$$\frac{1}{\sqrt{2}} \|V - \hat{V}O\| \leq \|\sin \Theta(\hat{V}, V)\| \leq \|\sin \Theta(\hat{V}, V)\|_F$$

Further, it is straightforward to see that

$$\min(d^{1/2} \|\hat{\Sigma} - \Sigma\|, \|\hat{\Sigma} - \Sigma\|_F) \leq d^{1/2} \|\hat{\Sigma} - \Sigma\|$$

Thus, all in all, we have

$$\begin{aligned}
 \frac{1}{\sqrt{2}} \|V - \hat{V}O\| &\leq \|\sin \Theta(\hat{V}, V)\| \\
 &\leq \|\sin \Theta(\hat{V}, V)\|_F \\
 &\leq \frac{2d^{1/2} \|\hat{\Sigma} - \Sigma\|}{\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1})}
 \end{aligned}$$

which completes the proof. \square

We are now ready to give the proof of Theorem 6.

Proof of Theorem 6. The proof is an application of the Davis-Kahan theorem as presented in Theorem 11. Observe that in Theorem 11 the eigenvalues are sorted in a decreasing way i.e. $\mu_1 \geq \dots \geq \mu_n$, whereas in our case they are sorted in an increasing manner i.e. $\lambda_1 \leq \dots \leq \lambda_n$.

Notationally, let the variables p, s, r from Theorem 11 be defined as $p = s = n, r = p - k + 1$.

We first focus in the case for $p \leq -1$. For this case we are interested in the k -smallest eigenvalues, i.e. $\lambda_1, \dots, \lambda_k$, which correspond to μ_p, \dots, μ_r , where $\mu_p = \lambda_1$ and $\mu_r = \lambda_k$.

By definition, in Theorem 11 we have that $\mu_{p+1} = -\infty$. Thus, $\mu_p - \mu_{p+1} = \infty$. Further, we can see $\mu_{r-1} - \mu_r = \lambda_{k+1} - \lambda_k = (1 + \varepsilon) - m_p(1 - \rho^+ + \varepsilon, 1 + \rho^- + \varepsilon)$ and hence by Eq.2

$$\begin{aligned}
 \min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1}) &= \\
 &(1 + \varepsilon) - m_p(1 - \rho^+ + \varepsilon, 1 + \rho^- + \varepsilon)
 \end{aligned}$$

which by Theorem 11 leads to the following inequality

$$\|V_k - \mathcal{V}_k O_k\| \leq \frac{2^{3/2} k^{1/2}}{\gamma} \|L_p - \mathcal{L}_p\| = \frac{\sqrt{8k}}{\gamma} \|L_p - \mathcal{L}_p\|$$

By applying Theorem 5, we know that if

$$\delta^+ = \frac{n}{k} (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+) > 3 \ln(8n/\varepsilon), \text{ and}$$

$$\delta^- = \frac{n}{k} (p_{\text{in}}^- + (k-1)p_{\text{out}}^-) > 3 \ln(8n/\varepsilon)$$

then with probability at least $1 - \epsilon$

$$\|V_k - \mathcal{V}_k O_k\| \leq \frac{\sqrt{8k}}{\gamma} C_p^- m_{|p|}^{1/|p|} \left(\sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^+}}, \sqrt{\frac{3 \ln(8n/\epsilon)}{\delta^-}} \right)$$

yielding the desired result. The case for $p \geq 1$ is similar, where instead of k the value $k' = k - 1$ is used. \square

H. Main building block for our results

In this section present two results from (Chung & Radcliffe, 2011) that are the main building blocks for our results.

Theorem 12 ((Chung & Radcliffe, 2011)). *Let G be a random graph, where $\text{pr}(v_i \sim v_j) = p_{ij}$, and each edges is independent of each other edge. Let A be the adjacency matrix of G , so $A_{ij} = 1$ if $v_i \sim v_j$ and 0 otherwise, and $\bar{A} = E(A)$, so $\bar{A}_{ij} = p_{i,j}$. Let D be the diagonal matrix with $D_{ii} = \text{deg}(v_i)$, and $\bar{D} = E(D)$. Let δ be the minimum expected degree of G , and $L = I - D^{-1/2} A D^{-1/2}$ the (normalized) Laplacian matrix for G . Choose $\epsilon > 0$. Then there exists a constant $k = k(\epsilon)$ such that if $\delta > k \ln n$, then the probability at least $1 - \epsilon$, the eigenvalues of L and \bar{L} satisfy*

$$|\lambda_j(L) - \lambda_j(\bar{L})| \leq 2\sqrt{\frac{3 \ln(4n/\epsilon)}{\delta}}$$

for all $1 \leq j \leq n$, where $\bar{L} = I - \bar{D}^{-1/2} \bar{A} \bar{D}^{-1/2}$.

Although this theorem is presented as the main result, one can see in the proof of theorem 12 in (Chung & Radcliffe, 2011), that in deed what they proved was a concentration bound for $\|L - \bar{L}\|$.

Theorem 13 ((Chung & Radcliffe, 2011)). *Assume that conditions of Theorem 12 hold. Choose $\epsilon > 0$. Then there exists a constant $k = k(\epsilon)$ such that if $\delta > k \ln n$, then*

$$\mathbb{P}\left(\|L - \bar{L}\| \leq 2\sqrt{\frac{3 \ln(4n/\epsilon)}{\delta}}\right) > 1 - \epsilon \quad (5)$$

Theorem 14 (Yu et al., 2015)). *Let $\Sigma, \hat{\Sigma} \in \mathbb{R}^{p \times p}$ be symmetric, with eigenvalues $\mu_1 \geq \dots \geq \mu_p$ and $\hat{\mu}_1 \geq \dots \geq \hat{\mu}_p$ respectively. Fix $1 \leq r \leq s \leq p$ and assume that $\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1}) > 0$, where $\mu_0 := \infty$ and $\mu_{p+1} := -\infty$. Let $d := s - r + 1$, and let $V = (v_r, v_{r+1}, \dots, v_s) \in \mathbb{R}^{p \times d}$ and $\hat{V} = (\hat{v}_r, \hat{v}_{r+1}, \dots, \hat{v}_s) \in \mathbb{R}^{p \times d}$ have orthonormal columns satisfying $\Sigma v_j = \mu_j v_j$ and $\hat{\Sigma} \hat{v}_j = \hat{\mu}_j \hat{v}_j$ for $j = r, r+1, \dots, s$. Then*

$$\|\sin \Theta(\hat{V}, V)\|_F \leq \frac{2 \min(d^{1/2} \|\hat{\Sigma} - \Sigma\|, \|\hat{\Sigma} - \Sigma\|_F)}{\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1})}.$$

Moreover, there exists an orthogonal matrix $\hat{O} \in \mathbb{R}^{d \times d}$ such that

$$\|\hat{V} \hat{O} - V\|_F \leq \frac{2^{3/2} \min(d^{1/2} \|\hat{\Sigma} - \Sigma\|, \|\hat{\Sigma} - \Sigma\|_F)}{\min(\mu_{r-1} - \mu_r, \mu_s - \mu_{s+1})}.$$

I. Results on Bethe Hessian

The following Lemma 5 states that for the case where $\alpha = 1$ the Bethe Hessian is equal to the arithmetic mean of Laplacians, i.e. the signed ratio Laplacian L_{SR} .

Lemma 5. *Let $\alpha = 1$. Then the Bethe Hessian is two times the arithmetic mean of L^+ and Q^- .*

Proof of Lemma 5. Let J^+, J^- be the positive and negative part of J , i.e. $J_{ij}^+ = \max\{0, J_{ij}\}$ and $J_{ij}^- = -\min\{0, J_{ij}\}$. Let D^+ and D^- be degree diagonal matrices of J^+ and J^- respectively, i.e. $\bar{D} = D^+ + D^-$. Then,

$$\begin{aligned} H &= (\alpha - 1)I - \sqrt{\alpha}J + \bar{D} \\ &= -J + \bar{D} \\ &= -J^+ + J^- + D^+ + D^- \\ &= (D^+ - J^+) + (D^- + J^-) \\ &= L^+ + Q^- \\ &= L_{SR} \end{aligned}$$

\square

Lemma 6. *Let \mathcal{H} be the Bethe hessian of the expected signed graph. Then $\{\chi_i\}_{i=2}^k$ are the eigenvectors corresponding to the $(k-1)$ -smallest negative eigenvalues of \mathcal{H} if and only if the following conditions hold:*

1. $\max\{0, \frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-} |C|}\} < (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)$
2. $p_{\text{out}}^+ < p_{\text{out}}^-$

Proof. In our framework the we can see that $J = W^+ - W^-$. In Section A we can see that expected adjacency matrices \mathcal{W}^+ and \mathcal{W}^- have three distinct eigenvalues:

$$\begin{aligned} \lambda_1^+ &= |C| (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+), & \lambda_i^+ &= |C| (p_{\text{in}}^+ - p_{\text{out}}^+) \\ \lambda_1^- &= |C| (p_{\text{in}}^- + (k-1)p_{\text{out}}^-), & \lambda_i^- &= |C| (p_{\text{in}}^- - p_{\text{out}}^-) \end{aligned}$$

for $i = 2, \dots, k$, with corresponding eigenvectors χ_1, \dots, χ_k . Remaining eigenvalues are equal to zero. Further, as both matrices \mathcal{W}^+ and \mathcal{W}^- share all their eigenvectors, then the expected matrix $\mathcal{J} = \mathcal{W}^+ - \mathcal{W}^-$ has the same eigenvectors with eigenvalues being the difference between the positive and negative counterparts, i.e. $\mathcal{J} \chi_i = \mu_i \chi_i$ where

$$\mu_i = \lambda_i^+ - \lambda_i^-.$$

As we assume that all clusters are of the same size $|C|$, the expected signed graph is a regular graph with degrees d^+ and d^- . Thus, in expectation $\hat{\alpha} = d^+ + d^-$, where $d^+ = |C| (p_{\text{in}}^+ + (k-1)p_{\text{out}}^+)$ and $d^- = |C| (p_{\text{in}}^- + (k-1)p_{\text{out}}^-)$.

Hence, the Bethe hessian of the expected signed graph can be expressed as

$$\begin{aligned}\mathcal{H} &= (\widehat{\alpha} - 1)I - \sqrt{\widehat{\alpha}}\mathcal{J} + \overline{\mathcal{D}} \\ &= (\widehat{\alpha} - 1)I - \sqrt{\widehat{\alpha}}\mathcal{J} + \widehat{\alpha}I \\ &= (2\widehat{\alpha} - 1)I - \sqrt{\widehat{\alpha}}\mathcal{J}\end{aligned}$$

It is easy to see that the matrix \mathcal{H} is some sort of a diagonal shift of \mathcal{J} , and thus they have the same eigenvectors. In particular we can observe that:

$$\begin{aligned}\mathcal{H}\chi_i &= ((2\widehat{\alpha} - 1)I - \sqrt{\widehat{\alpha}}\mathcal{J})\chi_i \\ &= (2\widehat{\alpha} - 1)\chi_i - \sqrt{\widehat{\alpha}}\mathcal{J}\chi_i \\ &= (2\widehat{\alpha} - 1)\chi_i - \sqrt{\widehat{\alpha}}\mu_i\chi_i \\ &= ((2\widehat{\alpha} - 1) - \sqrt{\widehat{\alpha}}\mu_i)\chi_i\end{aligned}$$

Hence, the corresponding eigenvalues of \mathcal{H} are:

$$\lambda_i = (2\widehat{\alpha} - 1) - \sqrt{\widehat{\alpha}}\mu_i. \quad (6)$$

All in all, the corresponding eigenvalues of the expected Bethe hessian matrix \mathcal{H} are:

$$\begin{aligned}\lambda_1 &= (2\widehat{\alpha} - 1) - \sqrt{\widehat{\alpha}}(d^+ - d^-), \\ \lambda_i &= (2\widehat{\alpha} - 1) - \sqrt{\widehat{\alpha}}|\mathcal{C}|((p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)), \\ \lambda_j &= (2\widehat{\alpha} - 1).\end{aligned}$$

for $i = 2, \dots, k$ and $j = k + 1, \dots, n$.

We now focus on the conditions that are necessary so that eigenvectors χ_2, \dots, χ_k have the smallest negative eigenvalues. This is based on the fact that informative eigenvectors of the Bethe Hessian H have the smallest negative eigenvalue. From Eq.6 we can see that the general condition for eigenvalues of the Bethe Hessian in expectation \mathcal{H} to be negative is

$$\lambda_i < 0 \iff \frac{2\widehat{\alpha} - 1}{\sqrt{\widehat{\alpha}}} < \mu_i. \quad (7)$$

Hence the conditions to be analyzed are:

$$\begin{aligned}\lambda_i &< \lambda_1, \quad \text{for } i = 2, \dots, k \\ \lambda_i &< 0, \quad \text{for } i = 2, \dots, k \\ \lambda_i &< \lambda_j, \quad \text{for } i = 2, \dots, k \text{ and } j = k + 1, \dots, n\end{aligned}$$

Therefore we can easily see that the corresponding condition $\lambda_i < \lambda_1$ boils down to

$$p_{\text{out}}^+ < p_{\text{out}}^-$$

whereas condition $\lambda_i < 0$ is equivalent to

$$\frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-}|\mathcal{C}|} < ((p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-))$$

and for the remaining condition $\lambda_i < \lambda_j$ the equivalent condition is

$$0 < (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)$$

by putting together conditions for $\lambda_i < 0$ and $\lambda_i < \lambda_j$ we get the desired result. \square

Lemma 7. *Let \mathcal{H} be the Bethe hessian of the expected non-empty signed graph. Let $|V| \rightarrow \infty$. Then $\{\chi_i\}_{i=2}^k$ are the eigenvectors corresponding to the $(k-1)$ smallest negative eigenvalues of \mathcal{H} if and only if the following conditions hold:*

1. $p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$
2. $p_{\text{out}}^+ < p_{\text{out}}^-$

Proof. From Lemma 6 we have the following conditions for the recovery of informative eigenvectors on finite graphs:

1. $\max\{0, \frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-}|\mathcal{C}|}\} < (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)$
2. $p_{\text{out}}^+ < p_{\text{out}}^-$

Let

$$\begin{aligned}c_2 &= p_{\text{in}}^+ + (k-1)p_{\text{out}}^+ + p_{\text{in}}^- + (k-1)p_{\text{out}}^- \\ c_3 &= (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)\end{aligned}$$

For the first condition of Lemma 6 can be expressed as follows:

$$\begin{aligned}\frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-}|\mathcal{C}|} < ((p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)) \iff \\ \frac{2c_2^{1/2}}{|\mathcal{C}|^{1/2}} - \frac{1}{|\mathcal{C}|^{3/2}c_2^{1/2}} < c_3.\end{aligned}$$

Hence, in the limit where $|\mathcal{C}| \rightarrow \infty$ the above condition turns into

$$0 < c_3 \iff p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-. \quad (8)$$

yielding the desired conditions. \square

The following Lemma states the interesting fact that the Bethe Hessian works better for large graphs

Lemma 8. *Let \mathcal{H}_n be the Bethe Hessian of the expected signed graph under the SBM with n nodes. Let $\chi^n = \{\chi_i\}_{i=2}^k$ where $\chi_2, \dots, \chi_k \in \mathbb{R}^n$. Let $\frac{3}{2} < d^+ + d^-$. Let $n < m$. If χ^n are eigenvectors corresponding to the $(k-1)$ -smallest negative eigenvalues of \mathcal{H}_n , then χ^m are eigenvectors corresponding to the $(k-1)$ -smallest negative eigenvalues of \mathcal{H}_m .*

Proof. In this proof we show that if for a given signed graph with n nodes the conditions of Lemma 6, then conditions of Lemma 6 hold for expected signed graphs with a larger number of nodes.

By Lemma 6, we know that for a given graph in expectation with n nodes, eigenvectors $\chi^n = \{\chi_i\}_{i=2}^k$ correspond to the $(k-1)$ -smallest negative eigenvalues of \mathcal{H}_n if and only the following conditions hold:

1. $\max\{0, \frac{2(d^+ + d^-) - 1}{\sqrt{d^+ + d^-}|\mathcal{C}|}\} < (p_{\text{in}}^+ - p_{\text{out}}^+) - (p_{\text{in}}^- - p_{\text{out}}^-)$
2. $p_{\text{out}}^+ < p_{\text{out}}^-$

Observe that the right hand side of the above conditions does not depend on the number of nodes in the graph. We proceed by analyzing the left hand side of the first condition:

$$\frac{2(d^+ + d^-) - 1}{|\mathcal{C}| \sqrt{d^+ + d^-}}. \quad (9)$$

Note that under the Stochastic Block Model in consideration, all k clusters are of size $|\mathcal{C}| = \frac{n}{k}$. We now identify conditions such that the Equation 9 decreases with larger values of $|\mathcal{C}|$.

Let $x, \alpha \in \mathbb{R}$. Define the scalar function $g : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ as

$$g(x) = \frac{2\alpha x - 1}{\sqrt{\alpha x^3}}$$

Observe that we recover Equation 9 by letting $x = |\mathcal{C}|$ and $\alpha = p_{\text{in}}^+ + (k-1)p_{\text{out}}^+ + p_{\text{in}}^- + (k-1)p_{\text{out}}^-$ where $\alpha x = d^+ + d^-$.

The corresponding derivative is

$$g'(x) = \frac{3 - 2\alpha x}{2x\sqrt{\alpha x^3}}$$

Then

$$g'(x) < 0 \iff \frac{3}{2} < \alpha x. \quad (10)$$

Hence, if $\frac{3}{2} < \alpha x$ then $g(y) < g(x)$ if and only if $x < y$. We now apply this result to our setting.

Let $|\mathcal{C}_n| := |\mathcal{C}| = \frac{n}{k}$ and $|\mathcal{C}_m| = \frac{m}{k}$ denote the cluster size of the expected signed graphs with n and m nodes, respectively. Let $\alpha = p_{\text{in}}^+ + (k-1)p_{\text{out}}^+ + p_{\text{in}}^- + (k-1)p_{\text{out}}^-$. Let $\frac{3}{2} < d^+ + d^-$. Then

$$g(|\mathcal{C}_m|) < g(|\mathcal{C}_n|) = \frac{2(d^+ + d^-) - 1}{|\mathcal{C}| \sqrt{d^+ + d^-}} \quad (11)$$

if and only if $n < m$. Hence, if conditions 1 and 2 hold for the expected graph G with n nodes and its expected absolute degree is larger than $\frac{3}{2}$, i.e. $\frac{3}{2} < d^+ + d^-$, then conditions 1 and 2 hold for expected graphs with a larger number of nodes, leading to the desired result. \square




Figure 5: Mean clustering error under the Censored Block Model (Saade et al., 2015), with two clusters of size 500 and 20 runs. Fig. 5a: probability of observing an edge is fixed to $p = 0.03$, and $\epsilon \in [0, 0.5]$. Fig. 5b: probability of flipping sign of an edge is fixed to $\epsilon = 0.25$, and $p \in [0.001, 0.03]$

J. Experiments with the Censored Block Model

In this section we present a numerical evaluation of different methods under the Stochastic Block Model following the parameters corresponding to the Censored Block Model (CBM), following (Saade et al., 2015). Observe that the CBM is a particular case of the Stochastic Block Model for signed graphs as introduced in Section 3. Following (Saade et al., 2015), the CBM has two parameters: probability of observing an edge (\bar{p}), and the probability of flipping the sign of an edge (ϵ). The CBM can be recovered from the SSBM introduced in Section 3 by setting $p_{\text{in}}^+ = p_{\text{out}}^- = \bar{p}(1 - \epsilon)$ and $p_{\text{in}}^- = p_{\text{out}}^+ = \bar{p}\epsilon$. Observe that the parameter ϵ works as a noise parameter: the noiseless setting corresponds to $\epsilon = 0$, where positive and negative edges are only inside and between clusters, respectively. The case where $\epsilon = 0.5$ corresponds to the case where no clustering structure is conveyed by the sign of the edges.

We present a numerical evaluation under the SSBM with parameters from CBM in Fig. 5. We consider two clusters and fix a priori its size to be of 500 nodes each. We present the clustering error out of 20 realizations from the SSM with parameters following the CBM. We consider two settings: **First setting:** we fix the probability of observing an edge to $\bar{p} = 0.03$, and evaluate over different values of $\epsilon \in [0, 0.5]$. In Fig. 5a we can observe that there is no relevant difference in clustering error between methods. Further, as expected we can see that for small values of ϵ all methods perform well, and for larger values of ϵ the clustering error increases; **Second setting:** we fix the probability of flipping the sign of an edge to $\epsilon = 0.25$, and evaluate over different values of $\bar{p} \in [0.001, 0.03]$. In Fig. 5b we can observe that the performance of the Bethe Hessian is best for small values of \bar{p} , i.e. for sparser graphs. Following the Bethe Hessian are

Table 2 Experiments on UCI datasets. Positive edges generated by k -nearest neighbours, and negative edges generated by k -farthest neighbours. Reported is the percentage of cases where each method achieves the smallest and strictly smallest clustering error, and the average clustering error.

	iris	wine	ecoli	australian	cancer	vehicle	german	image	optdig	isolet	USPS	pendig	20new	MNIST	
# vertices	150	178	336	690	699	846	1000	2310	5620	7797	9298	10992	18846	70000	
# classes	3	3	8	2	2	4	2	7	10	26	10	10	20	10	
H	Best (%)	14.1	14.1	10.9	7.8	0.0	20.3	34.4	0.0	3.1	0.0	0.0	45.3	0.0	
	Str. best (%)	4.7	3.1	7.8	3.1	0.0	14.1	17.2	0.0	3.1	0.0	0.0	45.3	0.0	
	Avg. error	16.8	32.2	23.9	41.7	13.2	58.4	29.5	64.0	32.5	54.0	41.3	39.2	88.5	48.2
L_{SN}	Best (%)	10.9	10.9	14.1	4.7	15.6	12.5	12.5	17.2	26.6	7.8	14.1	7.8	26.6	14.1
	Str. best (%)	1.6	3.1	9.4	1.6	15.6	7.8	0.0	17.2	26.6	7.8	14.1	7.8	26.6	14.1
	Avg. error	17.5	32.2	24.5	42.8	8.8	57.2	29.9	53.9	24.9	51.2	38.6	37.8	89.0	45.8
L_{BN}	Best (%)	4.7	12.5	0.0	6.3	1.6	6.3	40.6	0.0	0.0	0.0	0.0	1.6	0.0	
	Str. best (%)	1.6	1.6	0.0	0.0	1.6	4.7	17.2	0.0	0.0	0.0	0.0	1.6	0.0	
	Avg. error	26.6	33.6	30.5	42.5	10.2	61.6	29.6	57.2	41.1	67.4	50.1	50.5	92.5	58.6
L_{AM}	Best (%)	6.3	20.3	7.8	6.3	0.0	20.3	15.6	9.4	0.0	0.0	0.0	4.7	0.0	
	Str. best (%)	1.6	9.4	6.3	1.6	0.0	7.8	1.6	6.3	0.0	0.0	0.0	4.7	0.0	
	Avg. error	19.0	32.7	24.4	42.7	11.6	58.1	29.7	47.7	33.5	49.6	44.7	48.3	89.7	56.1
L_{GM}	Best (%)	32.8	35.9	34.4	32.8	7.8	17.2	46.9	6.3	29.7	28.1	12.5	0.0	1.6	82.8
	Str. best (%)	1.6	7.8	21.9	23.4	6.3	14.1	25.0	6.3	28.1	28.1	9.4	0.0	1.6	82.8
	Avg. error	14.1	31.9	20.4	39.3	11.3	57.6	29.5	46.8	13.0	42.6	27.6	45.0	89.9	26.7
L_{-1}	Best (%)	25.0	45.3	39.1	42.2	0.0	12.5	15.6	39.1	4.7	37.5	4.7	9.4	12.5	1.6
	Str. best (%)	0.0	14.1	18.8	31.3	0.0	9.4	1.6	29.7	4.7	37.5	4.7	9.4	12.5	1.6
	Avg. error	13.8	29.8	20.3	38.2	8.3	56.2	29.8	39.7	16.3	42.1	25.2	32.9	88.3	32.3
L_{-10}	Best (%)	73.4	43.8	25.0	34.4	76.6	31.3	20.3	39.1	37.5	26.6	71.9	82.8	7.8	1.6
	Str. best (%)	42.2	7.8	10.9	18.8	75.0	25.0	4.7	31.3	35.9	26.6	68.8	82.8	7.8	1.6
	Avg. error	12.7	30.2	20.8	38.6	5.7	55.9	29.7	39.4	12.1	42.3	21.9	26.9	89.8	28.6

the arithmetic mean Laplacian L_1 together with the signed normalized Laplacian L_{SN} .

Hence we have observed that for sufficiently dense graphs following the Censored Block Model, the performance of different methods is rather similar, whereas for sparser graphs the Bethe Hessian performs best, confirming the analysis presented in (Saade et al., 2015).

K. Experiments on UCI datasets

We evaluate the signed power mean Laplacian with L_{-10}, L_{-1} against $L_{SN}, L_{BN}, L_{AM}, L_{GM}$ and H using datasets from the UCI repository. We build W^+ from the k^+ nearest neighbor graph, whereas W^- is obtained from the k^- farthest neighbor graph. For each dataset we evaluate all clustering methods over all possible choices of $k^+, k^- \in \{3, 5, 7, 10, 15, 20, 40, 60\}$, yielding in total 64 cases. We present the following statistics: Best(%): proportion of cases where a method yields the smallest clustering error. Strictly Best(%): proportion of cases where a method is the *only one* yielding the smallest clustering error. Results are shown in Table 2.

Observe that in 4 datasets H and L_{GM} present a competitive performance. For the remaining cases we can see that the best performance are obtained by the signed power mean Laplacians L_{-1}, L_{-10} . This verifies the superiority of negative powers ($p < 0$) to positive ($p > 0$) powers of L_p and related approaches like L_{SN}, L_{BN} . Moreover, although the Bethe Hessian is known to be optimal under the sparse transition theoretic limit under the Censored Block

Model (Saade et al., 2015), in the context where graphs unlikely follow a SBM distribution we can see that it is outperformed by the signed power mean Laplacian L_p .

We consider a second setting where we generate k^- noiseless negative edges via cannot link constraints between nodes of different classes. The corresponding results are shown in Table 3. We observe in this setting that the arithmetic mean Laplacian L_{AM} presents the best performance, followed by the geometric mean Laplacian L_{GM} and the Balance Normalized Laplacian L_{BN} . This suggests that from the family of non-arithmetic based Laplacians the case of L_{GM} is a reasonable option showing certain robustness to different signed graph regimes.

We emphasize that the eigenvectors of L_p are calculated without ever computing the matrix itself, by adapting the method proposed in (Mercado et al., 2018), described in Sec.M. Also, please see Section L for a performance comparison with respect to changes in the diagonal shift on UCI datasets.

L. On Diagonal Shift

In this section we briefly discuss the effect of the diagonal shift on the power mean Laplacian L_p for $p \leq 0$. In the definition of power mean Laplacian in Eq. 1 it is mentioned that for negative powers $p \leq 0$ a diagonal shift is necessary. To evaluate the influence of the magnitude of the diagonal shift we perform numerical evaluations on two different kinds of signed graphs: on one side we consider signed graphs generated through the Signed Stochastic Block Model introduced

Table 3 Experiments on UCI datasets. Positive edges generated by k -nearest neighbours. Negative edges generated are cannot links between nodes of different classes. Reported is the percentage of cases where each method achieves the smallest and stricly smallest clustering error, and the average clustering error.

		iris	wine	ecoli	australian	cancer	vehicle	german	image	optdig	isolet	USPS	pendigits	20new	MNIST
	# vertices	150	178	336	690	699	846	1000	2310	5620	7797	9298	10992	18846	70000
	# classes	3	3	8	2	2	4	2	7	10	26	10	10	20	10
H	Best (%)	54.7	51.6	20.3	43.8	40.6	26.6	45.3	12.5	4.7	3.1	4.7	1.6	15.6	4.7
	Str. best (%)	0.0	9.4	9.4	1.6	0.0	7.8	0.0	0.0	3.1	3.1	4.7	1.6	15.6	3.1
	Avg. error	3.6	14.3	15.9	8.6	0.8	32.0	8.5	16.9	6.5	47.8	15.7	10.4	87.2	10.7
L_{SN}	Best (%)	59.4	42.2	17.2	42.2	45.3	37.5	48.4	17.2	15.6	20.3	12.5	12.5	28.1	9.4
	Str. best (%)	0.0	4.7	9.4	0.0	0.0	17.2	0.0	3.1	10.9	18.8	12.5	12.5	28.1	7.8
	Avg. error	4.0	15.5	15.9	11.9	5.3	30.5	11.1	13.8	4.9	44.1	11.9	7.2	86.1	7.8
L_{BN}	Best (%)	68.8	51.6	45.3	42.2	53.1	50.0	50.0	50.0	37.5	28.1	35.9	35.9	35.9	42.2
	Str. best (%)	3.1	12.5	40.6	0.0	1.6	28.1	0.0	35.9	34.4	26.6	35.9	35.9	35.9	42.2
	Avg. error	2.5	14.6	14.9	11.0	0.8	30.1	10.1	13.6	6.6	47.0	12.5	8.8	84.6	9.2
L_{AM}	Best (%)	42.2	25.0	26.6	0.0	0.0	23.4	0.0	45.3	45.3	35.9	46.9	48.4	18.8	45.3
	Str. best (%)	21.9	21.9	23.4	0.0	0.0	21.9	0.0	43.8	45.3	35.9	46.9	48.4	18.8	45.3
	Avg. error	2.1	22.9	17.0	11.3	0.5	44.2	12.5	13.0	3.2	40.2	8.2	3.9	88.1	4.2
L_{GM}	Best (%)	3.1	4.7	0.0	87.5	87.5	1.6	100.0	0.0	0.0	7.8	0.0	0.0	1.6	0.0
	Str. best (%)	0.0	0.0	0.0	56.3	46.9	1.6	48.4	0.0	0.0	7.8	0.0	0.0	1.6	0.0
	Avg. error	5.6	28.8	19.4	2.2	4.3	55.9	0.0	35.9	9.4	41.9	21.1	25.0	89.5	25.4
L_{-1}	Best (%)	7.8	9.4	1.6	0.0	0.0	1.6	0.0	1.6	1.6	6.3	0.0	1.6	0.0	0.0
	Str. best (%)	0.0	4.7	0.0	0.0	0.0	1.6	0.0	1.6	1.6	6.3	0.0	1.6	0.0	0.0
	Avg. error	3.9	28.3	19.1	10.8	4.2	54.9	22.5	27.3	9.1	41.5	19.2	15.9	89.5	19.5
L_{-10}	Best (%)	6.3	0.0	4.7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	Str. best (%)	4.7	0.0	3.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	Avg. error	8.1	29.3	20.0	5.6	3.6	56.4	7.8	38.5	11.2	42.1	21.4	26.3	89.9	25.9

in Section 3, and on the otherside we consider signed graphs built from standard machine learning benchmark datasets following Section K.

Experiments with SSBM. We begin with experiments based on signed graphs following the SSBM. The corresponding results are presented in Fig. 6. We study the performance of the power mean Laplacians $L_{-1}, L_{-2}, L_{-5}, L_{-10}$ with diagonal shifts $\{10^{-10}, 10^{-9}, \dots, 10^3\}$. Moreover, the case where either G^+ or G^- are informative i.e. assortative and disassortative, respectively. In particular, in top (resp. bottom) row of Fig. 6 the results correspond to the case where G^+ (resp. G^-) is fixed to be assortative (resp. disassortative).

We can observe that the larger the value of p , the more robust the performance of the corresponding power mean Laplacian L_p to the values of the diagonal shift. For instance, we can see for L_{-1} (see Figs. 6a and 6e) that the smaller the diagonal shift, the better the smaller the clustering error, whereas for diagonal shifts $10^0, 10^1, 10^2, 10^3$ its performance clearly deteriorates.

On the other side we can see that the power mean Laplacian L_{-10} presents a high sensibility towards the value of the diagonal shift (see Figs. 6d and 6h) where the diagonal shift should be neither too large nor too small, being the values $\{10^{-2}, 10^{-1}, 10^0\}$ the more suitable for this particular case. This observations are confirmation for the setting with sparse graphs, as it is observed in Fig. 7.

Experiments with benchmark datasets. We now perform a numerical evaluation on different real world networks, following the procedure of Section K. Moreover, we perform this analysis for $p \in \{-1, -10\}$ and diagonal shifts $\{10^{-10}, 10^{-9}, \dots, 10^3\}$. The corresponding results are presented in Fig. 8, where we present the average clustering error taken across all values of k^+ and k^- (for more details on the construction of the corresponding signed graphs please see Section K).

We can observe a general behaviour for L_{-10} across datasets, where for a small diagonal shift, the clustering error is high, and decreases for larger shifts, generally reaching its minimum clustering error around diagonal shifts equal to one, to later present a slight increase in clustering error. This confirms the proposed approach to set the diagonal shift to $\log_{10}(1 + |p|) + 10^{-6}$ which for the case of $p = -10$ is ≈ 1.04 . For the case of the harmonic mean Laplacian L_{-1} we can observe that it presents a more stable behaviour that slightly resembles the one of L_{-10} . In particular, we can observe that there is a region from 10^{-6} to 10^{-1} where the smallest average clustering error is achieved. Hence, L_{-1} is relatively more robust to different diagonal shifts. This confirms the observations made based on signed graphs following the SBM.

On condition number. We now consider a condition number approach to study the effect of the diagonal shift. Recall that the eigenvalue computation scheme considered in this paper is described in Section M with the corresponding Al-

Spectral Clustering of Signed Graphs via Matrix Power Means




Figure 6: Mean clustering error under SBM for different diagonal shifts with sparsity 0.1. Details in Sec. L.




Figure 7: Mean clustering error under SBM for different diagonal shifts with sparsity 0.05. Details in Sec. L.




Figure 8: Mean clustering error of the power mean Laplacians L_{-1} and L_{-10} with diagonal shifts $\{10^{-10}, 10^{-9}, \dots, 10^3\}$.

Algorithm 2: PM applied to $M_p(L_{\text{sym}}^+, Q_{\text{sym}}^-)$.

Input: $\mathbf{x}_0, p < 0$

Output: Eigenpair (λ, \mathbf{x}) of $M_p(L_{\text{sym}}^+, Q_{\text{sym}}^-)$

```

1 repeat
2    $\mathbf{u}_k^{(1)} \leftarrow (L_{\text{sym}}^+)^p \mathbf{x}_k$  (Compute with Alg. 3)
3    $\mathbf{u}_k^{(2)} \leftarrow (Q_{\text{sym}}^-)^p \mathbf{x}_k$  (Compute with Alg. 3)
4    $\mathbf{y}_{k+1} \leftarrow \frac{1}{2}(\mathbf{u}_k^{(1)} + \mathbf{u}_k^{(2)})$ 
5    $\mathbf{x}_{k+1} \leftarrow \mathbf{y}_{k+1} / \|\mathbf{y}_{k+1}\|_2$ 
6 until tolerance reached
7  $\lambda \leftarrow (\mathbf{x}_{k+1}^T \mathbf{x}_k)^{1/p}$ ,  $\mathbf{x} \leftarrow \mathbf{x}_{k+1}$ 
    
```

gorithm 2. We can observe that the main computation steps are related to the matrix vector operations $(L_{\text{sym}}^+)^p \mathbf{x}_k$ and $(Q_{\text{sym}}^-)^p \mathbf{x}_k$ with $p < 0$. We highlight that this framework considers only the case where $p < 0$.

Observe that in the operation $(L_{\text{sym}}^+)^p \mathbf{x}_k$, with $p < 0$, the condition number plays a influential place due to the inverse operation implied by the negativity of p . Note that the eigenvalues of the normalized Laplacians L_{sym}^+ are contained in the interval $[0, 2]$, hence, it is a singular matrix. As mentioned in definition of the power mean Laplacian in Eq. 1, a suitable diagonal shift is necessary for the case where $p < 0$. Hence, the eigenvalues of the shifted Laplacian $L_{\text{sym}}^+ + \mu I$ are contained in the interval $[\mu, 2 + \mu]$, therefore, condition number is equal to $\frac{\lambda_{\max}(L_{\text{sym}}^+)}{\lambda_{\min}(L_{\text{sym}}^+)}$ which in this case reduces to $\frac{2+\mu}{\mu}$. Thus, it follows that the condition number of $(L_{\text{sym}}^+ + \mu I)^p$ is $g(\mu, p) := \left(\frac{2+\mu}{\mu}\right)^{|p|}$. It is easy to see that $\frac{2+\mu}{\mu} > 1$ and hence $g(\mu, p)$ grows with larger values of $|p|$, hence the condition number is larger for smaller values of the power mean Laplacian. Moreover, the growth rate of $g(\mu, p)$ is larger for smaller values of μ , suggesting that the shift μ should be set as large as possible. Yet, very large values of μ overcome the information contained in the Laplacian matrix. Hence, the diagonal shift should not be too small (due to numerical stability) and should not be too large (due to information ofuscation). This confirms the

Algorithm 3: PKSM for the computation of $A^p \mathbf{y}$

Input: $\mathbf{u}_0 = \mathbf{y}, V_0 = [\cdot], p < 0$

Output: $\mathbf{x} = A^p \mathbf{y}$

```

1  $\mathbf{v}_0 \leftarrow \mathbf{y} / \|\mathbf{y}\|_2$ 
2 for  $s = 0, 1, 2, \dots, n$  do
3    $\tilde{V}_{s+1} \leftarrow [V_s, \mathbf{v}_s]$ 
4    $V_{s+1} \leftarrow$  Orthogonalize columns of  $\tilde{V}_{s+1}$ 
5    $H_{s+1} \leftarrow V_{s+1}^T A V_{s+1}$ 
6    $\mathbf{x}_{s+1} \leftarrow V_{s+1} (H_{s+1})^p \mathbf{e}_1 \|\mathbf{y}\|_2$ 
7   if tolerance reached then break
8    $\mathbf{v}_{s+1} \leftarrow A \mathbf{v}_s$ 
9 end
10  $\mathbf{x} \leftarrow \mathbf{x}_{s+1}$ 
    
```

behaviour presented in Figs. 6, 7 and 8.

M. Computation Of the Smallest Eigenvalues and Eigenvectors of L_p

For the computation of the eigenvectors corresponding to the smallest eigenvalues of the signed power mean Laplacian L_p with $p < 0$, we take the Polynomial Krylov Subspace Method for multilayer graphs presented in (Mercado et al., 2018) and apply it to our case. The corresponding adaption is presented in Algorithms 2 and 3.

We briefly explain Algorithm 2. Let $\lambda_1 \leq \dots \leq \lambda_n$ be the eigenvalues of $L_p = M_p(L_{\text{sym}}^+, Q_{\text{sym}}^-)$. Let $p < 0$. Then the eigenvalues of L_p^p are $\lambda_1^p \geq \dots \geq \lambda_n^p$, that is, the eigenvectors corresponding to the smallest eigenvalues of L_p correspond to the largest eigenvalues of L_p^p . Thus, in order to obtain the eigenvectors corresponding to the smallest eigenvalues of L_p we have to apply the power method to L_p^p . This is depicted in Algorithm 2. However, the main computational task now is the matrix-vector multiplications $(L_{\text{sym}}^+)^p \mathbf{x}$ and $(Q_{\text{sym}}^-)^p \mathbf{x}$. This is approximated through the Polynomial Krylov Subspace Method (PKSM). This approximation method allows to obtain $(L_{\text{sym}}^+)^p \mathbf{x}$ and $(Q_{\text{sym}}^-)^p \mathbf{x}$ without ever computing the matrices $(L_{\text{sym}}^+)^p$ and $(Q_{\text{sym}}^-)^p$, respectively. This is depicted in Algorithm 3.

The main idea of PKSM s -step is to project a given ma-




Figure 9: Time execution analysis

trix A onto the space $\mathbb{K}^s(A, \mathbf{y}) = \{\mathbf{y}, A\mathbf{y}, \dots, A^{s-1}\mathbf{y}\}$ and solve the corresponding problem there. The projection on to $\mathbb{K}^s(A, \mathbf{y})$ is done by means of the Lanczos process, producing a sequence of matrices V_s with orthogonal columns where the first column of V_s is $\mathbf{y}/\|\mathbf{y}\|$ and $\text{range}(V_s) = \mathbb{K}^s(A, \mathbf{y})$. Moreover, at each step we have $AV_s = V_s H_s + \mathbf{v}_{s+1} \mathbf{e}_s^T$ where H_s is $s \times s$ symmetric tridiagonal, and \mathbf{e}_i is the i -th canonical vector. The matrix product vector $\mathbf{x} = A^p \mathbf{y}$ is approximated by $\mathbf{x}_s = V_s (H_s)^p \mathbf{e}_1 \|\mathbf{y}\| \approx A^p \mathbf{y}$.

Time Execution Analysis. We present a time execution analysis in Fig. 9. We depict the mean time execution out of 10 runs of the power mean Laplacian L_p with $p \in \{-1, -2, -5, -10\}$. In particular $L_{-1}(\text{ours})$, $L_{-2}(\text{ours})$, $L_{-5}(\text{ours})$ and $L_{-10}(\text{ours})$ depict the time execution using our proposed method based on Algorithm 2 together with the polynomial Krylov subspace method described in Algorithm 3. For comparison we consider $L_{-1}(\text{eigs})$ which is computed with the function `eigs` from MATLAB instead of using Algorithm 3. All experiments are performed using one thread. For evaluation random signed graphs following the SSBM are generated, with parameters $p_{\text{in}}^+ = p_{\text{out}}^- = 0.05$ and $p_{\text{in}}^- = p_{\text{out}}^+ = 0.025$ with two equal sized clusters, and graph size $|V| \in \{10000, 20000, 30000, 40000\}$. We can observe that our computational matrix-free approach based on the polynomial Krylov subspace method systematically outperforms the natural approach based on the explicit computation of power matrices per layer.

N. Proportion of cases where conditions hold

In order to understand how often the conditions from Theorems 1, 3 and 4, we perform a series of experiments.

For the Bethe Hessian we take the limit result when $|V| \rightarrow \infty$ as the corresponding conditions do not have as a parameter the size of graph. The corresponding results are depicted in Fig. 10. We discretize each of the parameters

$p_{\text{in}}^+, p_{\text{out}}^+, p_{\text{in}}^-, p_{\text{out}}^-$ in $[0, 1]$ in one hundred steps and count how many times the conditions of Theorems 1, 3 and 4 hold under different settings. In Fig. 10a we analyze the case when both G^+ and G^- are informative ($p_{\text{in}}^+ > p_{\text{out}}^+$ and $p_{\text{in}}^- < p_{\text{out}}^-$). We can see that the conditions for the signed power mean Laplacian \mathcal{L}_p are always fulfilled, whereas those of L_{SN} and L_{BN} hold in a significantly smaller fraction of cases, whereas the case of the Bethe Hessian H are closer to the power mean Laplacians than to L_{SN} and L_{BN} . In Fig. 10b we analyze the case when G^+ or G^- is informative ($p_{\text{in}}^+ > p_{\text{out}}^+$ or $p_{\text{in}}^- < p_{\text{out}}^-$). Now we see an ordering between different \mathcal{L}_p where the smaller the value of p the larger the proportion of cases leading to recovery of the clusters in expectation. In particular, $\mathcal{L}_{-\infty}$ always fulfills the conditions, whereas \mathcal{L}_{∞} realizes the smallest proportion of cases where its conditions hold comparable to the one of L_{SN} and L_{BN} , while the Bethe Hessian holds for 50% of the cases. In Fig. 10c we treat the case where on average G^+ and G^- are informative ($p_{\text{in}}^- + p_{\text{out}}^+ < p_{\text{in}}^+ + p_{\text{out}}^-$). We observe the same ordering as in the previous case and again all signed power mean Laplacians outperform L_{SN} and L_{BN} , while the Bethe Hessian holds for around 75% of the cases. In Figs 10b and 10c we observe that the difference between the signed power mean Laplacians with finite p gets smaller as the number of clusters k increases. The reason is that the eigenvalues of \mathcal{L}_{sym} and \mathcal{Q}_{sym} are of the form $1 \pm \rho(k)$, where $\lim_{k \rightarrow \infty} \rho(k) = 0$. Thus, as k increases the eigenvalues become equal and thus the gap vanishes.

O. On Wikipedia Experiments

We provide a more detailed inspection of the results from Sec. 4. In Fig. 11 we present the sorted adjacency matrices according to the identified clusters. In the first two columns, (left to right), we can see that there is a large cluster (upper-left corner of each adjacency matrix) that does not resemble any structure, whereas the remaining part of the graph does present certain clustering structure. The following third and fourth columns zoom in into this region, which corresponds to results presented in Fig. 4.

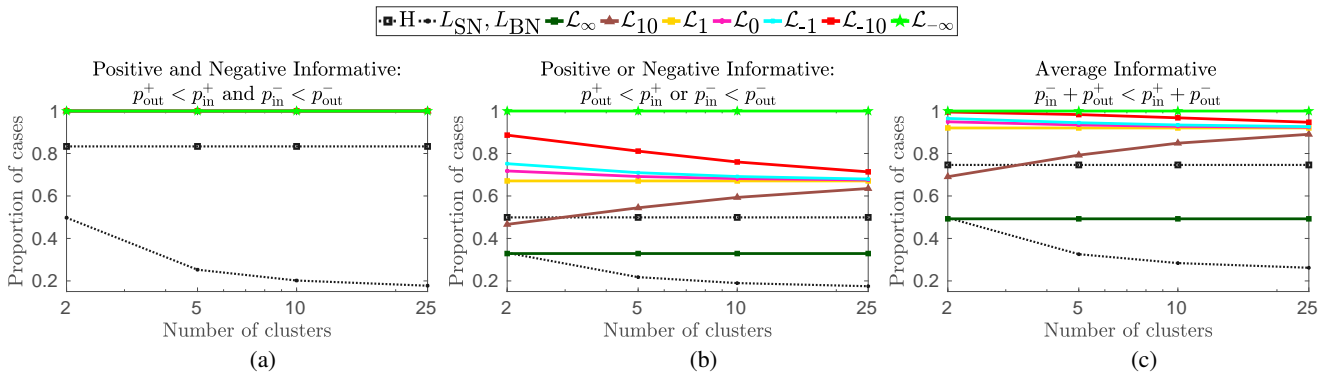


Figure 10: Proportion of cases where conditions of Theorems 1, 3 and 4 hold under different settings.

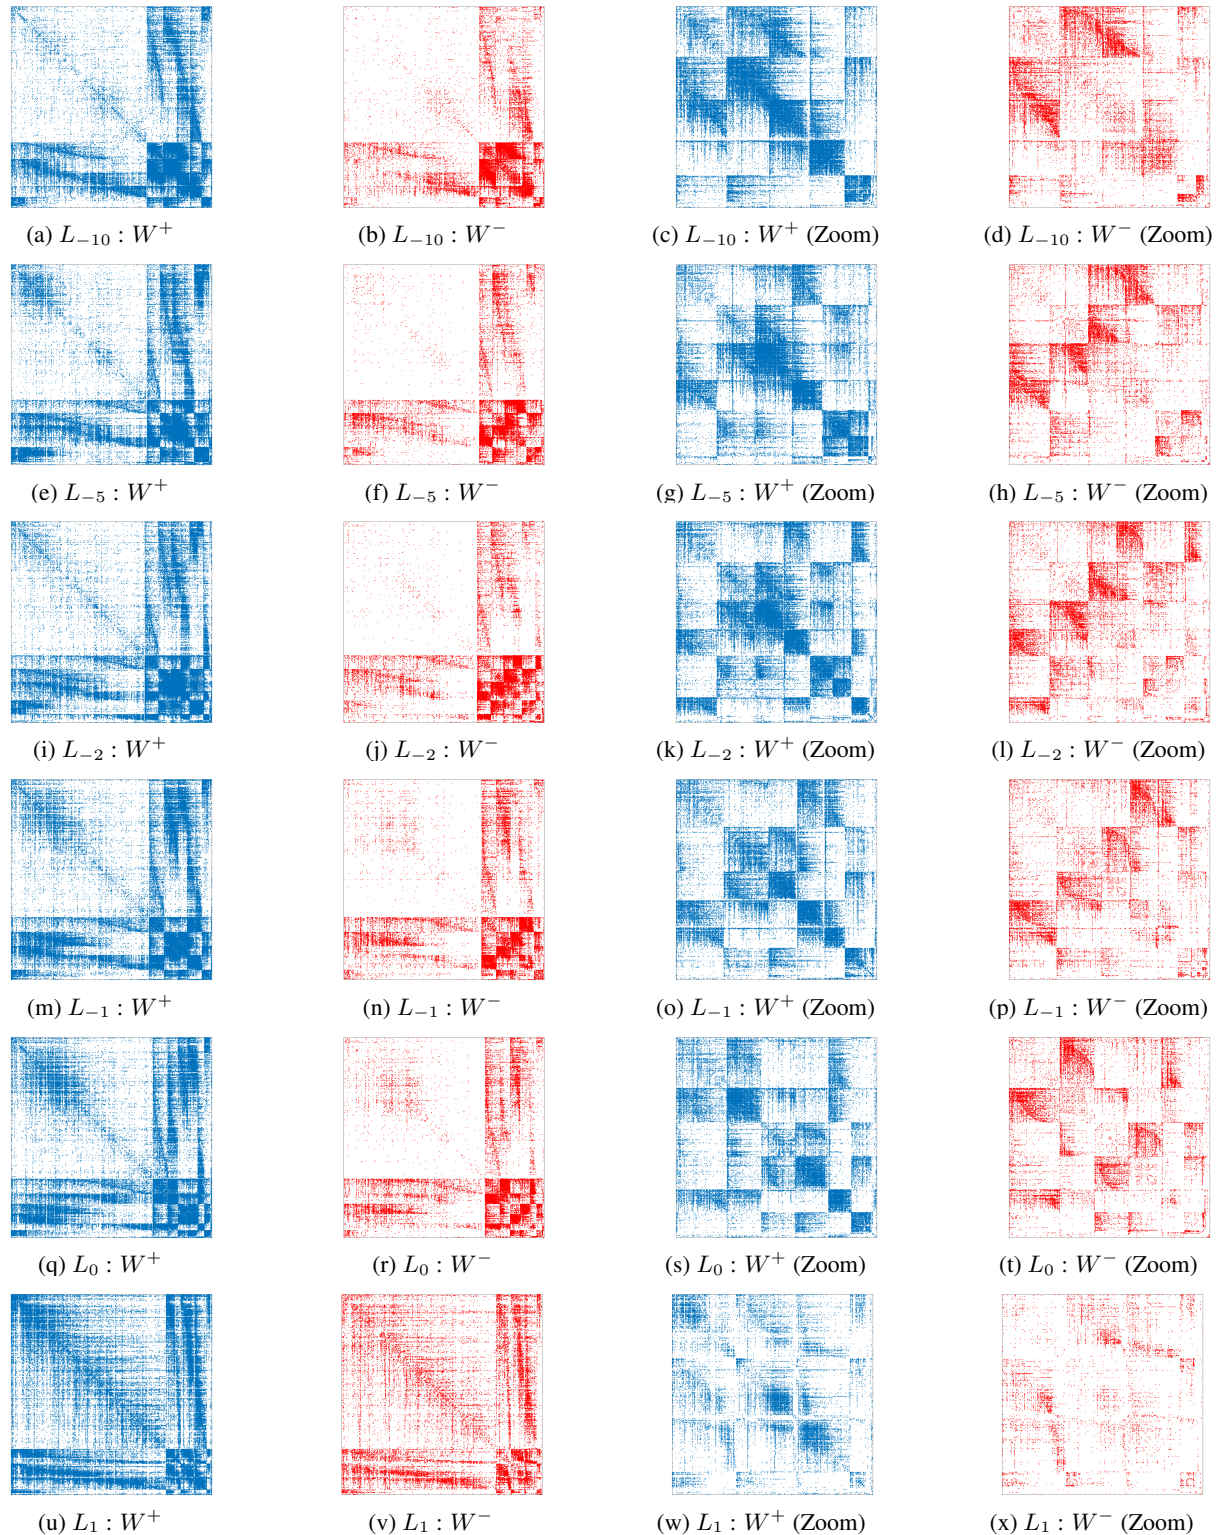


Figure 11: Sorted adjacency matrices according to clusters identified by the Power Mean Laplacian L_p with $p \in \{-10, -5, -2, -1, 0, 1\}$. Columns from left to right: First two columns depict adjacency matrices W^+ and W^- sorted through the corresponding clustering. Third and fourth columns depict the portion of adjacency matrices W^+ and W^- corresponding to the $k - 1$ identified clusters. Rows from top to bottom: Clustering corresponding to $L_{-10}, L_{-5}, L_{-2}, L_{-1}, L_0, L_1$.