A Additional Proofs for Subsection 3.1

A.1 Proof of Lemma 6

Denote $0 = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n$ and $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ to be the eigenvalues and eigenvectors of L, and let $0 = \tilde{\lambda}_1 \leq \tilde{\lambda}_2 \leq \cdots \leq \tilde{\lambda}_n$ and $\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2, \ldots, \tilde{\mathbf{v}}_n$ to be the eigenvalues and eigenvectors of $\tilde{L}$ (we note that a Laplacian is always positive semidefinite and has a 0 eigenvalue, see [25]).

By applying a classical eigenvector perturbation theorem due to Davis and Kahan (see section V in [24]), we have that if $[\mathbf{v}_1, \mathbf{v}_2]$ is the subspace spanned by the first two eigenvectors of L, and $\tilde{V}$ is the subspace spanned by the eigenvectors of $\tilde{L}$ whose eigenvalue is at most λ_2, then

$$\left\|\sin\left(\Theta\left([\mathbf{v}_1, \mathbf{v}_2], \tilde{V}\right)\right)\right\| \leq \frac{\|\tilde{L} - L\|}{\lambda_3 - \lambda_2},$$

where $\sin\left(\Theta\left([\mathbf{v}_1, \mathbf{v}_2], \tilde{V}\right)\right)$ is the diagonal matrix with the sines of the canonical angles between the subspaces $[\mathbf{v}_1, \mathbf{v}_2], \tilde{V}$ along the main diagonal. Moreover, by Weyl's inequality (see Corollary 4.9 in [24]), we have that $|\lambda_3 - \tilde{\lambda}_3| \leq \|\tilde{L} - L\|$. Therefore, if $\lambda_3 - \lambda_2 > \|\tilde{L} - L\|$, it guarantees us that $\tilde{\lambda}_3 > \lambda_2$, so the subspace $\tilde{V}$ is simply the one spanned by $\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2$. If $\lambda_3 - \lambda_2 \leq \|\tilde{L} - L\|$, we still trivially have $\|\sin\left(\Theta\left([\mathbf{v}_1, \mathbf{v}_2], [\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2]\right)\right)\| \leq 1$. So in any case, we get

$$\|\sin\left(\Theta\left([\mathbf{v}_1, \mathbf{v}_2], [\tilde{\mathbf{v}}_1, \tilde{\mathbf{v}}_2]\right)\right)\| \leq \min\left\{\frac{\|\tilde{L} - L\|}{\lambda_3 - \lambda_2}, 1\right\}.$$

For any Laplacian, the vector $\frac{1}{\sqrt{n}}\mathbf{1}$ is an eigenvector corresponding to the 0 eigenvalue (see [25]). So we may assume that $\tilde{\mathbf{v}}_1 = \mathbf{v}_1 = \frac{1}{\sqrt{n}}\mathbf{1}$, and by definition of canonical angles, it follows from the inequality above that

$$\sin(\arccos(\langle \mathbf{v}_2, \tilde{\mathbf{v}}_2\rangle)) \leq \min\left\{\frac{\|\tilde{L} - L\|}{\lambda_3 - \lambda_2}, 1\right\}.$$

Denoting $\|\tilde{L} - L\|/(\lambda_3 - \lambda_2)$ as ϵ, it is straightforward to show from this that

$$\langle \mathbf{v}_2, \tilde{\mathbf{v}}_2\rangle \geq \sqrt{1 - \min\{\epsilon^2, 1\}},$$

where we choose the sign of $\tilde{\mathbf{v}}_2$ so as to make the l.h.s. nonnegative. Therefore, $\|\tilde{\mathbf{v}}_2 - \mathbf{v}_2\|^2$ equals

$$2(1 - \langle \tilde{\mathbf{v}}_2, \mathbf{v}_2\rangle) \leq 2(1 - \sqrt{1 - \min\{\epsilon^2, 1\}}) \leq 2\min\{\epsilon^2, 1\},$$

from which the result follows.

A.2 Proof of Thm. 1

Let $c = \frac{n(n-1)}{2b}$, and let L be the Laplacian of the matrix A. By the triangle inequality,

$$\|c\tilde{L} - L\| \leq \|c\tilde{D} - D\| + \|c\tilde{A} - A\|. \tag{5}$$

We will treat each term separately. By Lemma 3, we know that $\tilde{A}$ consists of negatively dependent entries, and thus $c\tilde{A} - A$ is also a matrix with negatively dependent entries (this follows from Lemma 1). We now wish to apply Lemma 4 on this matrix, so we need to check that all of its conditions are fulfilled. First, note that each entry $\tilde{A}_{i,j}$ equals $A_{i,j}$ with probability $1/c$, and 0 otherwise. Thus, it is easy to verify that the $\mathbb{E}[c\tilde{A}_{i,j} - A_{i,j}] = 0$. Moreover, since $A_{i,j}$ is assumed to be bounded in $[0,1]$, it follows that $|c\tilde{A}_{i,j} - A_{i,j}| \leq c$. In addition, the variance of each entry $\mathbb{E}[(c\tilde{A}_{i,j} - A_{i,j})^2]$ is at most $(1/c)(c-1)^2 + (1-1/c)1^2 \leq c$. Finally, $c\tilde{A}_{i,j} - A_{i,j}$ takes only two values, in the manner assumed in Lemma 4. So applying Lemma 4 on the matrix $c\tilde{A} - A$, we have that with probability at least $1 - \delta$,

$$\|c\tilde{A} - A\| \leq 2\sqrt{cn} + 3\sqrt[3]{c^2 n}\log(2n/\delta). \tag{6}$$

Turning to analyze $\|c\tilde{D} - D\|$, we note that $c\tilde{D} - D$ is a diagonal matrix, hence the norm is equal to the absolute value of the largest entry on the diagonal:

$$\|c\tilde{D} - D\| = \max_i \left|\sum_{j=1}^{n}(c\tilde{a}_{i,j} - a_{i,j})\right|,$$

For any fixed i, the term in the absolute values is the sum of n zero mean, negatively dependent random variables, with absolute values and variances at most c. Applying Lemma 5, we have for any $\epsilon \in (0,1)$ that

$$\begin{aligned}
&\Pr\left(\left|\sum_{j=1}^{n}(c\tilde{a}_{i,j} - a_{i,j})\right| > n\epsilon\right) \\
&= \Pr\left(\left|\frac{1}{n}\sum_{j=1}^{n}(\tilde{a}_{i,j} - \frac{1}{c}a_{i,j})\right| > \frac{\epsilon}{c}\right) \\
&\leq 2\exp\left(-\frac{n\epsilon^2/c^2}{2(1/c + \epsilon/3c)}\right) \leq 2\exp\left(-\frac{n\epsilon^2}{3c}\right).
\end{aligned}$$

This implies that with probability at least $1 - \delta$, for any $\delta \geq 2\exp(-n/3c)$,

$$\left|\sum_{j=1}^{n}(c\tilde{a}_{i,j} - a_{i,j})\right| \leq \sqrt{3cn\log(2/\delta)}.$$

By a union bound, this implies that with probability at least $1 - \delta$, for any $\delta \geq 2\exp(-n/3c)$,

$$\|c\tilde{D} - D\| = \max_i \left|\sum_{j=1}^{n}(c\tilde{a}_{i,j} - a_{i,j})\right| \leq \sqrt{3cn\log(2n/\delta)}. \tag{7}$$

Plugging Eq. (6) and Eq. (7) into Eq. (5), and using again a union bound, we have with probability at least $1 - \delta$, for any $\delta \geq 2\exp(-n/3c)$ that

$$\|c\tilde{L} - L\| \leq 2\sqrt{cn} + 3\sqrt[3]{c^2 n}\log(4n/\delta) + \sqrt{3cn\log(4n/\delta)}. \tag{8}$$

Finally, applying Lemma 6, this event implies that $\|\tilde{\mathbf{v}}_2 - \mathbf{v}_2\|$ is at most

$$\sqrt{2}\,\min\left\{\frac{2\sqrt{cn} + 3\sqrt[3]{c^2 n}\log(4n/\delta) + \sqrt{3cn\log(4n/\delta)}}{\lambda_3 - \lambda_2}, 1\right\}. \tag{9}$$

In fact, this occurs with probability at least $1-\delta$ even when $\delta < 2\exp(-n/3c)$, because the bound can be shown to be vacuously true in this case (the expression inside the min will always be at least 1). Simplifying the bound a bit for readability, and using the fact that $c \leq n^2/2b$, we get the result stated in the theorem.

Finally, note that if one desires a fully empirical bound, which depends only on $\tilde{A}$ and not on the unknown matrix L, then it is possible to estimate $\lambda_3 - \lambda_2$. Indeed, if we let $\tilde{\lambda}_2, \tilde{\lambda}_3$ denote the 2nd and 3rd smallest eigenvalues of $c\tilde{L}$, then by Weyl's inequality (see corollary 4.9 in [24]) we have $|\tilde{\lambda}_3 - \lambda_3|, |\tilde{\lambda}_2 - \lambda_2| \leq \|c\tilde{L} - L\|$. Thus, by using the bound on $\|c\tilde{L} - L\|$ obtained above, one can bound $\lambda_3 - \lambda_2$ using the empirically obtainable quantity $\tilde{\lambda}_3 - \tilde{\lambda}_2$.

B Proof of Thm. 2

Before we begin, we will need the following lemma:

Lemma 7. *Let $\hat{A}$ be a matrix during some point in the run of algorithm 2, and let $\hat{A}'$ be the matrix obtained by setting some entry pairs $\hat{a}_{i,j}, \hat{a}_{j,i}$ to 0 in an arbitrary manner. Let $\hat{L}, \hat{L}'$ be the Laplacians of $\hat{A}, \hat{A}'$ respectively. Then it holds that*

$$\|\hat{L} - L\| \leq \|\hat{L}' - L\|.$$

Proof. Clearly, it is enough to prove the assertion for $\hat{A}'$ which is obtained from $\hat{A}$ by setting a particular entry pair $\hat{a}_{i,j}, \hat{a}_{j,i}$ to 0, and then repeating the argument. In this case, it is easy to verify that for any vector $\mathbf{v}$,

$$\mathbf{v}^\top(\hat{L} - L)\mathbf{v} - \mathbf{v}^\top(\hat{L}' - L)\mathbf{v} = \hat{a}_{i,j}(v_i - v_j)^2.$$

In particular, if we pick $\mathbf{v}$ to be the maximal eigenvector of $(\hat{L} - L)$ (e.g., such that $\mathbf{v}^\top(\hat{L} - L)\mathbf{v} = \|\hat{L} - L\|$), and $\mathbf{v}'$ to be the maximal eigenvector of $(\hat{L}' - L)$, then

$$\|\hat{L}' - L\| - \|\hat{L} - L\| = \mathbf{v}'^\top(\hat{L} - L)\mathbf{v}' - \mathbf{v}^\top(\hat{L}' - L)\mathbf{v}$$
$$\geq \mathbf{v}^\top(\hat{L} - L)\mathbf{v} - \mathbf{v}^\top(\hat{L}' - L)\mathbf{v} = \hat{a}_{i,j}(\hat{v}_i - \hat{v}_j)^2 \geq 0.$$

□

We now turn to prove Thm. 1. The basic idea is that after running Algorithm 2 with a budget size $2b$, the matrix $\hat{A}$ always includes b revealed entries which "look" as if they were sampled uniformly without replacement. Notice that these entries might not be the entries queried in the even iterations of Algorithm 2, since these were performed on a matrix with some entries already revealed in a non-random manner. The theorem then follows from an application of Lemma 7.

More precisely, suppose that we implement the random samples in the even iterations of Algorithm 2 as follows. Before the algorithm begins, we sample $2b$ matrix indices from $\{(i,j) \in \{1,\ldots,n\}^2 : i < j\}$ uniformly at random without replacement, and put them in an ordered list. At each even iteration, we pick the first index (i,j) from the list that wasn't sampled yet (in either the even or odd iterations). Clearly, this is a valid implementation of the random samples in Algorithm 2.

Since we pick an element from the list every even iteration, and all previous elements were necessarily already sampled, we have that by the end of the algorithm's run, all the first b elements in the list are already picked. Thus, the matrix $\hat{A}$ always include these b entries. Let $\hat{A}'$ be the matrix obtained from $\hat{A}$ by zeroing all sampled entries except those b entries. Notice that $\hat{A}'$ has the same distribution as if we picked b entries uniformly without replacement from $\{(i,j) \in \{1,\ldots,n\}^2 : i < j\}$. In other words, $\hat{A}'$ has the same distribution as the matrix obtained during the run of Algorithm 1. Therefore, the analysis performed in the proof of Thm. 1 applies. In particular, Eq. (8) applies, namely that with probability at least $1 - \delta$,

$$\|c\hat{L}' - L\| \leq 2\sqrt{cn} + 3\sqrt[3]{c^2 n}\log(4n/\delta) + \sqrt{3cn\log(4n/\delta)},$$

where $\hat{L}'$ is the Laplacian of $\hat{A}'$ and $c = (n-1)n/2b$ is a constant scaling factor. But from Lemma 7, we also know that

$$\|c\hat{L} - L\| \leq \|c\hat{L}' - L\|.$$

Combining the two inequalities, we get that

$$\|c\hat{L} - L\| \leq 2\sqrt{cn} + 3\sqrt[3]{c^2 n}\log(4n/\delta) + \sqrt{3cn\log(4n/\delta)}.$$

Repeating the end of the proof of Thm. 1 (following Eq. (8)), we get the same guarantee for Algorithm 2 as for Algorithm 1, only with a budget size of $2b$ rather than b.

C Algorithm 2 - Proof of Correctness

In this appendix, we show that Algorithm 2 computes the (squared) norm of the derivative of the 2nd eigenvector of the Laplacian, with respect to a symmetric perturbation of entries $(i,j),(j,i)$ in the similarity matrix.

The relevant general theorem, stated below, is based on techniques used in the proof of theorem 7 in section 8.8 of [19], and lemma 1 in [17]. In Corollary 1, we use it to compute the squared norm of the derivative.

Theorem 3. *Let $(\lambda_1\ldots,\lambda_n)$ and $(\mathbf{v}_1,\ldots,\mathbf{v}_n)$ be the n eigenvalues and eigenvectors of the Laplacian L of a given $n \times n$ matrix A, and assume that λ_k for some $k > 1$ is a simple eigenvalue (i.e. has multiplicity 1).*

Fix some indices $i,j \in \{1,\ldots,n\}$, $i \neq j$, and let E^{ij} be the $n \times n$ matrix with $E^{ij}_{i,j} = E^{ij}_{j,i} = 1$ and zeros otherwise. Finally, define $A(t) = A + tE^{ij}$, where $t \in \mathbb{R}$.

Then it is possible to define functions $\mu_k(t)$ and $\mathbf{u}_k(t)$ in an open set T around 0, such that $\mu_k(t), \mathbf{u}_k(t)$ are the k-th eigenvalue and eigenvector of $A(t)$, and it holds that

$$\left.\frac{d\mathbf{u}_k(t)}{dt}\right|_{t=0} = \left(\sum_{l\neq k}\frac{\mathbf{v}_l\mathbf{v}_l^\top}{\lambda_k - \lambda_l}\right)P^{ij}\mathbf{v}_k,$$

where $P^{ij}_{i,j} = P^{ij}_{j,i} = -1, P^{ij}_{i,i} = P^{ij}_{j,j} = 1$, and zeros otherwise.

Intuitively, $A(t)$ tracks a perturbation of A by adding a small positive element at the symmetric indices $(i,j),(j,i)$. We note that the assumption of λ_k being simple is reasonable for a general matrix (this can always be ensured with probability 1 by adding an arbitrarily small random perturbation to the matrix, but we did not bother to do this in our algorithm).

Proof. Let $M(t)$ be the Laplacian of $A(t)$. Note that $M(0) = L, \mu_k(0) = \lambda_k$, and $\mathbf{u}_k(0) = \mathbf{v}_k$. For simplicity, we will drop the indexing by t from now on, as it should be clear from context.

We begin by proving that $\mathbf{u}_k, \mu_k$ are well-defined, differentiable functions of t in some region around $t = 0$.

Consider the vector function $f : \mathbb{R}^{n+2} \mapsto \mathbb{R}^{n+1}$ defined as

$$f(\mathbf{u}, \mu, t) = \begin{pmatrix} (\mu I - M(t))\mathbf{u} \\ \mathbf{u}^\top \mathbf{u} - 1 \end{pmatrix}.$$

It is easy to see that the function is continuously differentiable everywhere, and that $f(\mathbf{v}_k, \lambda_k, 0) = \mathbf{0}$. Moreover, the Jacobian of f w.r.t. $\mathbf{u}, \mu$ at that point equals

$$\left.\frac{\partial f}{\partial(\mathbf{u}, \mu)}\right|_{\mathbf{u}=\mathbf{v}_k, \mu=\lambda_k, t=0} = \begin{bmatrix} \lambda_k I - L & \mathbf{v}_k \\ 2\mathbf{v}_k^\top & 0 \end{bmatrix}. \qquad (10)$$

We claim that the Jacobian is non-singular. To see why, notice that the matrix in Eq. (10) has $n - 1$ eigenvalues of the form $(\lambda_k - \lambda_l)$ for all $l \neq k$ (which are all non-zero by the assumption that λ_k is a simple eigenvalue of L), with corresponding eigenvectors $(\mathbf{v}_l^\top, 0)$, as well as the two eigenvalues $-\sqrt{2}, \sqrt{2}$ with corresponding eigenvectors $(\mathbf{v}_k^\top, -\sqrt{2}), (\mathbf{v}_k^\top, \sqrt{2})$. Thus, the determinant of the Jacobian in Eq. (10), which equals the product of its eigenvalues, is non-zero.

By the implicit function theorem, this implies that there is some open set $T \subseteq \mathbb{R}$, $0 \in T$, with well defined differentiable functions $\mu_k : T \mapsto \mathbb{R}, \mathbf{u}_k : T \mapsto \mathbb{R}^n$, such that

$$\forall\ t \in T, \quad (\mu_k I - M)\mathbf{u}_k = 0. \qquad (11)$$

Differentiating Eq. (11) at $t = 0$, and using the definition of a Laplacian, we get

$$\left(\frac{d\mu_k}{dt} I - P^{ij}\right)\mathbf{v}_k + (\lambda_k I - L)\frac{d\mathbf{v}_k}{dt} = 0.$$

Multiplying from the left by $\mathbf{v}_l$ for any $j \neq k$, and using the fact that $\mathbf{v}_l^\top \mathbf{v}_i = 0$ and $L\mathbf{v}_l = \lambda_l \mathbf{v}_l$, we get that

$$-\mathbf{v}_l^\top P^{ij}\mathbf{v}_k + (\lambda_k - \lambda_l)\mathbf{v}_l^\top \frac{d\mathbf{u}_k}{dt} = 0,$$

or

$$(\lambda_k - \lambda_l)\mathbf{v}_l^\top \frac{d\mathbf{u}_k}{dt} = \mathbf{v}_l^\top P^{ij}\mathbf{v}_k. \qquad (12)$$

Also, since $\mathbf{v}_1, \ldots, \mathbf{v}_n$ is an orthonormal basis for $\mathbb{R}^n$, and $d\mathbf{u}_k/dt$ at $t = 0$ is orthogonal to $\mathbf{v}_k$ (since the eigenvectors are forced to lie on the unit sphere), we have

$$\frac{d\mathbf{u}_k}{dt} = \sum_{l=1}^{n}\langle \mathbf{v}_l, \frac{d\mathbf{u}_k}{dt}\rangle \mathbf{v}_l = \sum_{l \neq k}\langle \mathbf{v}_l, \frac{d\mathbf{u}_k}{dt}\rangle \mathbf{v}_l.$$

Since we assume that λ_k is a simple eigenvector, the above is equal to

$$\sum_{l \neq k}(\lambda_k - \lambda_l)^{-1}(\lambda_k - \lambda_l)\left(\mathbf{v}_l^\top \frac{d\mathbf{u}_k}{dt}\right)\mathbf{v}_l,$$

which by Eq. (12) equals

$$\sum_{l \neq k}(\lambda_k - \lambda_l)^{-1}\left(\mathbf{v}_l^\top P^{ij}\mathbf{v}_k\right)\mathbf{v}_l.$$

Slightly rearranging, the theorem follows. ∎

The following corollary to Thm. 3 provides an expression for the squared norm of the derivative of the k-th eigenvector of the Laplacian. Its derivation is a simple algebraic exercise, utilizing the orthogonality of the eigenvectors.

Corollary 1. *Under the conditions and notation of Thm. 3, define*

$$\tilde{\mathbf{v}}_l = \begin{cases} \mathbf{v}_l/(\lambda_k - \lambda_l) & \lambda_k \neq \lambda_l \\ 0 & \lambda_k = \lambda_l \end{cases}$$

Then $\left\| \left.\frac{d\mathbf{u}_k(t)}{dt}\right|_{t=0} \right\|^2$*, where the derivative is with respect to the perturbation* E^{ij}*, equals*

$$(v_{k,i} - v_{k,j})^2 \sum_{l=1}^{n}(\tilde{v}_{l,i} - \tilde{v}_{l,j})^2$$

Proof. Define the matrix G as

$$G = \sum_{l:\lambda_l \neq \lambda_k}\frac{\mathbf{v}_l \mathbf{v}_l^\top}{\lambda_k - \lambda_l}.$$

By Thm. 3, the vector $d\mathbf{u}_k(t)/dt$ at $t = 0$ w.r.t. perturbation E^{ij} can be written as $(v_{k,i} - v_{k,j})(G_i - G_j)$, where G_i, G_j are the i-th and j-th column of G. Therefore, $\|d\mathbf{u}_k(t)/dt\|^2$ equals

$$(v_{k,i} - v_{k,j})^2 \left\| \sum_{l:\lambda_l \neq \lambda_k}\frac{(v_{l,i} - v_{l,j})\mathbf{v}_l}{\lambda_k - \lambda_l} \right\|^2$$
$$= (v_{k,i} - v_{k,j})^2 \sum_{l:\lambda_l \neq \lambda_k}\left(\frac{v_{l,i} - v_{l,j}}{\lambda_k - \lambda_l}\right)^2,$$

since the eigenvectors are orthogonal to each other. This equals the expression in the corollary statement. ∎