Notes on Spectral Partitioning

(DRAFT: Do not circulate, quote, or cite. Bugs likely.)

Andrew Y. Ng, Michael Jordan and Yair Weiss

1 The Algorithm

For simplicity, this text discusses Spectral Partioning only in the case of three clusters. The generalization to k clusters is straightforward.

Let S_1, S_2 and S_3 be three sets (clusters) of points of sizes n_1, n_2 and n_3 respectively, and $S = S_1 \cup S_2 \cup S_3$ be the total set of $n = n_1 + n_2 + n_3$ points. Let x_j be the j-th point in S, and with some abuse of notation, we will write either "$x_j \in S_i$" or "$j \in S_i$" to mean x_j is in cluster i. We assume that we are given the set S, and wish to characterize conditions under which spectral partitioning will reliably recover the three subsets.

Let $\tilde{A}$ be the nonnegative symmetric square matrix whose (i, j) element is the similarity of points i and j in set S if $i \neq j$ (for example, $\tilde{A}_{ij} = \exp(-||x_i - x_j||_2^2/2\sigma^2)$), and 0 if $i = j$. Assume for convenience that the similarity scores $\tilde{A}_{ij}$ lie in $[0, 1]$. Also, we assume that the points in S are ordered in a way that respects their cluster memberships. Note that $\tilde{A}$ can therefore be partitioned as follows:

$$\tilde{A} = \begin{bmatrix} \tilde{A}^{(11)} & \tilde{A}^{(12)} & \tilde{A}^{(13)} \\ \tilde{A}^{(21)} & \tilde{A}^{(22)} & \tilde{A}^{(23)} \\ \tilde{A}^{(31)} & \tilde{A}^{(22)} & \tilde{A}^{(33)} \end{bmatrix} \tag{1}$$

where $\tilde{A}^{(ii)} \in \mathbb{R}^{n_i \times n_i}$ is the matrix of similarities between all the points in S_i, etc. In the sequel, we will continue to use the notational convention of using parenthesized superscripts to index into the subblocks of partitioned matrices and vectors.

Let $\tilde{D}$ be the diagonal matrix whose i-th diagonal element equals the sum of $\tilde{A}$'s i-th row. Note that it can be partitioned as follows:

$$\tilde{D} = \begin{bmatrix} \tilde{D}^{(11)} & 0 & 0 \\ 0 & \tilde{D}^{(22)} & 0 \\ 0 & 0 & \tilde{D}^{(33)} \end{bmatrix} \tag{2}$$

We then construct the Laplacian $\tilde{L}$ according to[1]

$$\begin{aligned} \tilde{L} &= \tilde{D}^{-1/2}\tilde{A}\tilde{D}^{-1/2} \\ &= \begin{bmatrix} (\tilde{D}^{(11)})^{-1/2}\tilde{A}^{(11)}(\tilde{D}^{(11)})^{-1/2} & (\tilde{D}^{(11)})^{-1/2}\tilde{A}^{(12)}(\tilde{D}^{(22)})^{-1/2} & (\tilde{D}^{(11)})^{-1/2}\tilde{A}^{(13)}(\tilde{D}^{(33)})^{-1/2} \\ (\tilde{D}^{(22)})^{-1/2}\tilde{A}^{(21)}(\tilde{D}^{(11)})^{-1/2} & (\tilde{D}^{(22)})^{-1/2}\tilde{A}^{(22)}(\tilde{D}^{(22)})^{-1/2} & (\tilde{D}^{(22)})^{-1/2}\tilde{A}^{(23)}(\tilde{D}^{(33)})^{-1/2} \\ (\tilde{D}^{(33)})^{-1/2}\tilde{A}^{(31)}(\tilde{D}^{(11)})^{-1/2} & (\tilde{D}^{(33)})^{-1/2}\tilde{A}^{(32)}(\tilde{D}^{(22)})^{-1/2} & (\tilde{D}^{(33)})^{-1/2}\tilde{A}^{(33)}(\tilde{D}^{(33)})^{-1/2} \end{bmatrix} \end{aligned}$$

[1] Note Chung (1991) defines the Laplacian differently, as the identity minus this. This just changes each eigenvalue λ to $1 - \lambda$, without changing the eigenvectors.

We take the first 3 eigenvectors x_1, x_2, x_3 of $\tilde{L}$ and line them up in columns to form the n-by-3 matrix $\tilde{X} = [x_1 x_2 x_3]$. We then renormalize each row of $\tilde{X}$ to have unit (ℓ_2) length, resulting in a matrix $\tilde{Y}$. (Note that this renormalizing step just corresponds to premultiplying $\tilde{X}$ by some appropriate diagonal matrix to get $\tilde{Y}$.) We now think of the i-th point in S as being represented by the 3-dimensional vector that is the i-th row of $\tilde{Y}$.

To make the notation reflect this more clearly, let $\tilde{Y}$ be partitioned into 3 blocks of sizes n_i-by-3 ($i = 1, 2, 3$) conformally with the clusters, and let $(\tilde{y}_j^{(i)})^T$ be the j-th row of the i-th block of $\tilde{Y}$. i.e.

$$\tilde{Y} = \begin{bmatrix} (\tilde{y}_1^{(1)})^T \\ \vdots \\ (\tilde{y}_{n_1}^{(1)})^T \\ (\tilde{y}_1^{(2)})^T \\ \vdots \\ (\tilde{y}_{n_2}^{(2)})^T \\ (\tilde{y}_1^{(3)})^T \\ \vdots \\ (\tilde{y}_{n_3}^{(3)})^T \end{bmatrix} \tag{3}$$

We wish to understand the properties of these points $\tilde{y}_j^{(i)}$ lying on the surface of the 3-dimensional sphere, and in particular when they will "cluster" together on the surface of the sphere in a way that allows us to reliably separate out the three clusters S_1, S_2 and S_3.[2]

2 Overview of analysis

Rather than directly analyzing the algorithm described above, we first consider an alternative algorithm that does not directly normalize $\tilde{X}$'s rows to unit length to get $\tilde{Y}$. Instead, we consider an alternative algorithm that "cheats" and uses information about the true clustering of the data to construct a certain diagonal matrix Λ, and then uses, instead of $\tilde{Y}$, the matrix

$$\tilde{V} = \Lambda \tilde{X}. \tag{4}$$

In the manner of $\tilde{y}_j^{(i)}$ being the (transposed) rows of $\tilde{Y}$ (as in Equation (3)), also let $\tilde{v}_j^{(i)}$ be the (transposed) rows of $\tilde{V}$, where specifically $(\tilde{v}_j^{(i)})^T$ is the j-th row of the i-th block of $\tilde{V}$ ($i = 1, 2, 3$, $j = 1, \ldots, n_i$).

Our analysis then proceeds as follows. In Section 3, the details of the construction of Λ are given, and we analyze its behavior in the setting of the clusters

[2]Though we will not explicitly draw attention to this again in the sequel, these methods also apply straightforwardly to characterizing the behavior of a version of the algorithm that does not normalize $\tilde{X}$'s rows to have unit length, but instead uses $\tilde{X}$ directly.

being "infinitely" far apart, i.e. if we had perfectly clean separation between the clusters. Section 4 then gives the second part of the analysis, in which we generalize the result to the setting where the clusters are finite distances from each other. Finally, Section 5 makes use of our analysis of the alternate algorithm's $\tilde{V}$ to characterize the behavior of the original algorithm's $\tilde{Y}$.

We conclude this section with a few more definitions. Let $\tilde{d}$ be the n-dimensional vector containing $\tilde{D}$'s diagonal entries. For $i = 1, 2, 3$, let $\tilde{d}^{(i)}$ be the n_i-dimensional vector containing $\tilde{D}^{(ii)}$'s diagonal entries. i.e. $\tilde{D}^{(ii)} = \mathrm{diag}(\tilde{d}_1^{(i)}, \ldots, \tilde{d}_{n_i}^{(i)})$, and

$$\tilde{d} = \begin{bmatrix} \tilde{d}^{(1)} \\ \tilde{d}^{(2)} \\ \tilde{d}^{(3)} \end{bmatrix} \tag{5}$$

For each $i = 1, 2, 3$, also define $\tilde{Z}_i = \sqrt{\sum_{j=1}^{n_i} \tilde{d}_j^{(i)}}$.

3 Analysis part I: The case of "infinite" separation

For the first part of this analysis, assume that the clusters are "infinitely" far from each other. Thus, let A be a matrix that is the same as $\tilde{A}$, except that its off-diagonal blocks are zero:

$$A = \begin{bmatrix} A^{(11)} & 0 & 0 \\ 0 & A^{(22)} & 0 \\ 0 & 0 & A^{(33)} \end{bmatrix} \tag{6}$$

where $A^{(ii)} = \tilde{A}^{(ii)}$. Also let all the tilde-less quantities D, d, L, V, v, Y, y, and Z be defined similarly as in the preceding discussion, but with $\tilde{A}$ replaced by A. Informally, we think of this as the case where the three clusters' "internal structure" is kept the same as before (so the within-cluster similarity matrices are the same as before), but where the clusters are moved so far away from each other that the inter-cluster similarities become 0.

We then define the n-by-n diagonal matrix Λ as follows:

$$\Lambda = \begin{bmatrix} Z_1 (D^{(11)})^{-1/2} & 0 & 0 \\ 0 & Z_2 (D^{(22)})^{-1/2} & 0 \\ 0 & 0 & Z_3 (D^{(33)})^{-1/2} \end{bmatrix} \tag{7}$$

i.e. If we let $\Lambda^{(ii)}$ denote the i-th diagonal subblock of Λ ($i = 1, 2, 3$), then $\Lambda_{jj}^{(ii)}$ is $Z_i / \sqrt{d_j^{(i)}}$.

Let us now consider what happens when we calculate X, the matrix containing the first three eigenvectors of L. Note that L can be written

$$L = \begin{bmatrix} (D^{(11)})^{-1/2} A^{(11)} (D^{(11)})^{-1/2} & 0 & 0 \\ 0 & (D^{(22)})^{-1/2} A^{(22)} (D^{(22)})^{-1/2} & 0 \\ 0 & 0 & (D^{(33)})^{-1/2} A^{(33)} (D^{(33)})^{-1/2} \end{bmatrix},$$

so that it is also block diagonal. The spectrum and eigenvectors of L are therefore the union of the spectra and eigenvectors (padded appropriately with zeros) of its blocks.

For each $i = 1, 2, 3$, let $\lambda_1^{(i)} \geq \lambda_2^{(i)} \geq \cdots \geq \lambda_{n_i}^{(i)}$ be the n_i eigenvalues of $L^{(ii)} = (D^{(ii)})^{-1/2} A^{(ii)} (D^{(ii)})^{-1/2}$. Letting $\mathbf{1}$ denote the vector of 1's, we also know that $(D^{(ii)})^{1/2}\mathbf{1} = \left[\sqrt{d_1^{(i)}}, \ldots, \sqrt{d_{n_i}^{(i)}}\right]^T$ is the principal eigenvector of $L^{(ii)}$, with eigenvalue $\lambda_1^{(i)} = 1$. (This is easily seen, as in Meila and Shi (2001), using the fact that $(D^{(ii)})^{-1}A^{(ii)}$ is a stochastic matrix and hence has as its principal right eigenvector the vector $\mathbf{1}$, with eigenvalue 1; premultiplying $(D^{(ii)})^{-1}A^{(ii)}\mathbf{1} = \mathbf{1}$ on both sides by $(D^{(ii)})^{1/2}$ then gives this.) Now assume that each cluster is connected.[3] Then the next eigenvalue $\lambda_2^{(i)}$ of the matrix $L^{(ii)}$ will be strictly less than 1. [[ayn: Double-check that Chung says this....]]

Thus, L's first few eigenvalues are $\lambda_1 = \lambda_2 = \lambda_3 = 1$, and $\lambda_4 = \max_{i=1,2,3} \lambda_2^{(i)} < 1$. Moreover, the first three eigenvectors,[4] summarized in a matrix $X = [x_1 x_2 x_3]$, are therefore given by:

$$X = \begin{bmatrix} \sqrt{d_1^{(1)}}/Z_1 & 0 & 0 \\ \vdots & \vdots & \vdots \\ \sqrt{d_{n_1}^{(1)}}/Z_1 & 0 & 0 \\ 0 & \sqrt{d_{n_2}^{(2)}}/Z_2 & 0 \\ \vdots & \vdots & \vdots \\ 0 & \sqrt{d_{n_2}^{(2)}}/Z_2 & 0 \\ 0 & 0 & \sqrt{d_{n_1}^{(3)}}/Z_3 \\ \vdots & \vdots & \vdots \\ 0 & 0 & \sqrt{d_{n_3}^{(3)}}/Z_3 \end{bmatrix} \tag{8}$$

where we remind the reader that $Z_i = \sqrt{\sum_{j=1}^{n_i} d_j^{(i)}}$, so that each of these eigenvectors does indeed have unit length.

Note that X (and x_1, x_2, x_3) is, of course, uniquely defined only up to a rotation: X could just as well have been replaced by XU, for any orthogonal (rotation) matrix $U \in \mathbb{R}^{3\times 3}$ ($UU^T = U^TU = I$). Ignoring for now this complication, and proceeding as in Equation (4), the "alternative algorithm" then

[3]This is certainly satisfied if each of the blocks $A^{(ii)}$ has strictly positive elements, which is true for most similarity functions.

[4]Here and in the sequel, we assume in the case of repeated eigenvalues that the corresponding eigenvectors are orthogonal to each other.

constructs

$$V = \Lambda X = \begin{bmatrix} 1 & 0 & 0 \\ \vdots & \vdots & \vdots \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ \vdots & \vdots & \vdots \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ \vdots & \vdots & \vdots \\ 0 & 0 & 1 \end{bmatrix} \tag{9}$$

This suggests that all the points in cluster i will have their correponding rows in V lying at exactly e_i (the i-th basis vector, here in $\mathbb{R}^3$), though we again remind the reader that, just as X could easily have been replaced by any XU, the matrix V could also have been replaced by VU for an arbitrary orthogonal matrix U. Thus, these points can actually lie anywhere on the unit sphere, at 3 points that that $90°$ from each other.

Recall that a 3-by-3 matrix $U = [u_1 u_2 u_3]$ is orthogonal if and only if its columns u_i all lie on the unit sphere are are at $90°$ from each other. Thus, our first our result can be stated as follows.

Proposition 1 *Suppose A's off-diagonal blocks $A^{(ij)}$, $i \neq j$, are zero. Also assume that each cluster S_i is connected. Then there exists an orthogonal matrix $U = [u_1 u_2 u_3]$ such that V's rows satisfy*

$$v_j^{(i)} = u_i \tag{10}$$

for all $i = 1, 2, 3$, $j = 1, \ldots, n_i$.

In anticipation of eventually relating these results back to the behavior of the original algorithm that uses Y instead of V, note that the following ways of defining Y in terms of X are equivalent: (i) Normalize the rows of X to have unit length to get Y. (ii) First premultiply X by a diagonal matrix Λ (with a strictly positive diagonal) to get $V = \Lambda X$, and then normalize V's rows to have unit length to get Y. The equivalence of (i) and (ii) stems from the observation that ΛX is just X with its rows rescaled in some way, and that this rescaling is cancelled out by the second, renormalization, step. But we see above that V (or VU)'s rows already have unit length, so renormalizing it to unit length will not change anything. Hence, $Y = V$, and we therefore have the following Corollary to Proposition 1.

Corollary 2 *Suppose A's off-diagonal blocks A_{ij}, $i \neq j$, are zero. Also assume that each cluster S_i is connected. Then there exists an orthogonal matrix $U = [u_1 u_2 u_3]$ such that Y's rows satisfy*

$$y_j^{(i)} = u_i \tag{11}$$

for all $i = 1, 2, 3$, $j = 1, \ldots, n_i$.

4 Analysis part II: The case of finite cluster separation

We now continue to analyze the behavior of the "alternative" algorithm that uses $\tilde{V}$ instead of $\tilde{Y}$, and now generalize our result to the setting where the clusters are no longer infinitely far apart.

Our analysis proceeds as follows. In the previous section, we showed how, with zero off-diagonal blocks, the eigenvectors of L led to clean separation between the three clusters. Here, we view $\tilde{A}$ as a perturbation of A, and $\tilde{L}$ as a perturbation of L, and seek to characterize when the resulting "perturbed" eigenvectors will still reliably separate the clusters. Section 4.1 first gives the assumptions under which our analysis will be carried out. Section 4.2 then gives and proves our main theorem on the behavior of $\tilde{V}$.

4.1 Assumptions

For it to be reasonable to expect an algorithm to reliably separate S_1, S_2 and S_3, clearly every cluster should be "well separated" from every other cluster. Specifically, we hope that, for all $j \in S_{i_1}$ and $k \in S_{i_2}$ $(i_1 \neq i_2)$ in different clusters, we have that $\tilde{A}_{jk}$, the similarity score between points j and k, is small.

Assumption A1. There is some fixed $\epsilon_1 > 0$ so that for every $i_1, i_2 \in \{1,2,3\}$, $i_1 \neq i_2$, we have that

$$\sum_{j \in S_{i_1}} \sum_{k \in S_{i_2}} \frac{\tilde{A}_{jk}^2}{d_j d_k} \leq \epsilon_1. \tag{12}$$

To gain intuition about this assumption, consider the case of two "densely connected" clusters i_1 and i_2 of size $\Omega(n)$ each. Recall that that for $j \in S_i$, we had $d_j = \sum_{k:k\in S_i} A_{jk}$. Thus, d_j measures how "connected" point j is to other points in the same cluster. Specifically, if it is close to every other point in its own cluster, then $d_j = \Omega(n)$. The sum above is a sum over $O(n^2)$ terms, and is in our example divided by $d_j d_k = \Omega(n^2)$. Thus, as long as the individual $\tilde{A}_{jk}$'s are small, the overall sum will also be small, and the assumption will hold true with small ϵ_1.

Whereas d_j measures the connectedness of point j to the rest of the same cluster, the quantity $\tilde{d}_j - d_j = \sum_{k:k\notin S_i} \tilde{A}_{jk}$ measures how connected j is to points in *other* clusters. The next of our assumptions states that all points must be more connected to points in the same cluster compared to how connected it is to points in other clusters; specifically, it says that the ratio between these two quantities should be small.

Assumption A2. There is some fixed $\epsilon_2 > 0$ so that for each $i = 1,2,3$, $j \in S_i$, we have that

$$\frac{\sum_{k:k\notin S_i} \tilde{A}_{jk}}{d_j} \leq \epsilon_2 \left(\sum_{k,l \in S_i} \frac{A_{kl}^2}{d_k d_l} \right)^{-1/2}$$

For intuition about this assumption, it is again worth considering the case of densely connected clusters (as above), in which case the quantity in parentheses on the right hand side is $\Omega(1)$, so that all this is demanding is that

$$\frac{\sum_{k:k\notin S_i} \tilde{A}_{jk}}{d_j} = \frac{\sum_{k:k\notin S_i} \tilde{A}_{jk}}{\sum_{k:k\in S_i} \tilde{A}_{jk}} = O(\epsilon_2)$$

The next assumption is a subtle one whose rationale is similar to that of Assumption A2. Informally, it extends that idea that if we want an algorithm to find the clusters S_1, S_2 and S_3, then we require that each of these sets S_i really look like a "tight" cluster. Consider an example in which $S_1 = S_{1.1} \cup S_{1.2}$, where $S_{1.1}$ and $S_{1.2}$ are themselves two well-separated clusters, then $S = S_1 \cup S_2 \cup S_3$ really looks like (at least) four well-separated clusters, and it would be unreasonable to expect an algorithm to correctly guess what partition of the four clusters into three subsets we had in mind.

The "tightness" of a cluster is characterized by the second eigenvalue of the Laplacian for that cluster. Reminding the reader that we had previously defined $\lambda_j^{(i)}$ to be the j-th eigenvalue of the matrix $L^{(ii)}$, we therefore have the following assumption.

Assumption A3. There is some fixed $\delta > 0$ so that for all $i = 1, 2, 3$, we have $\lambda_2^{(i)} \leq 1 - \delta$.

This also takes care of the "connected clusters" assumption we previously had in Proposition 1, and gives exactly a precondition that we will later need to invoke some results from Matrix Perturbation theory. More precisely, to guarantee stability of (a subspace spanned by several) eigenvectors, we require that the *eigengap* be large. We had previously shown L's first few eigenvalues to be $\lambda_1 = \lambda_2 = \lambda_3 = 1$, and $\lambda_4 = \max_i \lambda_2^{(i)}$. So, Assumption A3 is exactly that λ_4 be at most $1-\delta$, which guarantees that the eigengap between λ_3 and the rest of the eigenvalues is at least δ. Assumption A3 can also be easily guaranteed in several ways that are easier to interpret in terms of the clusters S_i being "tight" clusters. Here are three such ways.

Assumption A3.1 Something about Cheeger constants.

Assumption A3.2 Something about mixing times of the Markov chain on each cluster.

Assumption A3.3 Something about paths through the clusters.

The last of our assumptions is as follows.

Assumption A4. There is some constant C so that for every $i = 1, 2, 3$, we have that

$$\frac{\max_{j=1,\ldots,n_i} d_j^{(i)}}{\min_{j=1,\ldots,n_i} d_j^{(i)}} \leq C. \tag{13}$$

[[ayn: This can be considerably weakened if we switch to a "no more than $o(1)$ points more than $o(1)$ away" type of result; decide later if should do this. Also, $d_j^{(i)} \geq C(\sum_{k=1}^{n_i} d_k^{(i)})/n_i$ works too, and is weaker and probably more palatable?]]

This assumption is that each cluster has a fairly "uniform" degree of interconnectedness, and in particular that there aren't some points in the cluster that have far more "nearby neighbors" that other parts of the same cluster. For example, the assumption will generally be satisfied with a small C if the cluster is either a tight "ball" (in which case almost every point has about the same, $\Omega(n_i)$, number of "neighbors" as every other point), or if the cluster is a long "chain" (in which case almost every point will have, e.g., about 2 neighbors). Note that, since similarity scores are in $[0,1]$, this assumption is always satisfied by picking $C = (\max_i\{n_i\})/B \leq n/B$, if B is a lower-bound on the d_j's.

4.2 Behavior of $\tilde{V}$

We first consider how L changes to become $\tilde{L}$. Let $E = \tilde{L} - L$ be the perturbation to L.

Lemma 3 *Let assumptions A1 and A2 hold, and $E = \tilde{L} - L$. Then*

$$||E||_F \leq \sqrt{6\epsilon_1 + 3\epsilon_2^2} \tag{14}$$

Proof. Let E be partitioned conformally with the clusters as follows.

$$E = \begin{bmatrix} E^{(11)} & E^{(12)} & E^{(13)} \\ E^{(21)} & E^{(22)} & E^{(23)} \\ E^{(31)} & E^{(22)} & E^{(33)} \end{bmatrix} \tag{15}$$

We then have

$$||E||_F^2 = \sum_{i_1=1}^{3} \sum_{i_2=1}^{3} ||E^{(i_1 i_2)}||_F^2. \tag{16}$$

We first bound the norms of the off-diagonal blocks, and then the norms of the on-diagonal blocks.

Let any $i_1, i_2 \in \{1,2,3\}$, $i_1 \neq i_2$, be fixed, and consider any $j \in S_{i_1}, k \in S_{i_2}$. Then

$$\begin{aligned} E_{jk}^2 &= (\tilde{L}_{jk} - L_{jk})^2 & (17) \\ &= \tilde{L}_{jk}^2 & (18) \\ &= \left(\frac{\tilde{A}_{jk}}{\sqrt{\tilde{d}_j \tilde{d}_k}} \right)^2 & (19) \\ &= \frac{\tilde{A}_{jk}^2}{\tilde{d}_j \tilde{d}_k} & (20) \\ &\leq \frac{\tilde{A}_{jk}^2}{d_j d_k}. & (21) \end{aligned}$$

Thus, using Assumption A1, we have

$$\begin{aligned}
||E^{(i_1 i_2)}||_F^2 &= \sum_{j \in S_{i_1}} \sum_{k \in S_{i_2}} E_{jk}^2 && (22)\\
&\leq \sum_{j \in S_{i_1}} \sum_{k \in S_{i_2}} \frac{\tilde{A}_{jk}^2}{d_j d_k}. && (23)\\
&\leq \epsilon_1 && (24)
\end{aligned}$$

We next bound the norms of the on-diagonal blocks of E. Let any $i \in \{1, 2, 3\}$, be fixed, and consider any $j, k \in S_i$. Note $\tilde{A}_{jk} = A_{jk}$ for such j, k. For each j, also let $a_j = \sum_{k:k \notin S_i} A_{jk}$, so that $\tilde{d}_j = d_j + a_j$. We have

$$\begin{aligned}
E_{jk}^2 &= (\tilde{L}_{jk} - L_{jk})^2 && (25)\\
&= \left(\frac{\tilde{A}_{jk}}{\sqrt{\tilde{d}_j \tilde{d}_k}} - \frac{\tilde{A}_{jk}}{\sqrt{d_j d_k}} \right)^2 && (26)\\
&= \left(\frac{\tilde{A}_{jk}}{\sqrt{d_j d_k}} \left(\sqrt{\frac{d_j d_k}{\tilde{d}_j \tilde{d}_k}} - 1 \right) \right)^2 && (27)\\
&= \frac{\tilde{A}_{jk}^2}{d_j d_k} \left(1 - \sqrt{\frac{d_j}{d_j + a_j}} \sqrt{\frac{d_k}{d_k + a_k}} \right)^2 && (28)\\
&= \frac{A_{jk}^2}{d_j d_k} \left(1 - \sqrt{1 - \frac{a_j}{d_j + a_j}} \sqrt{1 - \frac{a_k}{d_k + a_k}} \right)^2 && (29)\\
&\leq \frac{A_{jk}^2}{d_j d_k} \left(1 - \sqrt{1 - \frac{a_j}{d_j}} \sqrt{1 - \frac{a_k}{d_k}} \right)^2 && (30)\\
&\leq \frac{A_{jk}^2}{d_j d_k} \left(1 - \left(\sqrt{1 - \epsilon_2 \left(\sum_{l,m \in S_i} \frac{A_{lm}^2}{d_l d_m} \right)^{-1/2}} \right)^2 \right)^2 && (31)\\
&= \frac{A_{jk}^2}{d_j d_k} \left(\epsilon_2^2 \left(\sum_{l,m \in S_i} \frac{A_{lm}^2}{d_l d_m} \right)^{-1} \right) && (32)
\end{aligned}$$

where to get Equation (31), we used Assumption A2. [[ayn: Was slightly sloppy above in that not all the operands of $\sqrt{\cdot}$ in every step were all necessarily non-negative, but the conclusion is correct either way, since if $x, y \leq \min\{1, B\}$, then $(1 - \sqrt{1-x}\sqrt{1-y})^2 \leq B^2$. Clean up the argument later.]] Hence, we can bound the norm of $E^{(ii)}$ as follows.

$$
\begin{aligned}
||E^{(ii)}||_F^2 &= \sum_{j,k \in S_i} E_{jk}^2 && (33)\\
&\leq \sum_{j,k \in S_i} \frac{A_{jk}^2}{d_j d_k} \left(\epsilon_2^2 \left(\sum_{l,m \in S_i} \frac{A_{lm}^2}{d_l d_m} \right)^{-1} \right) && (34)\\
&= \epsilon_2^2. && (35)
\end{aligned}
$$

Finally, substituting the bounds from Equations (24) and (35) into Equation (16), we find that

$$||E||_F^2 \leq 6\epsilon_1 + 3\epsilon_2^2, \tag{36}$$

which proves the Lemma. □

Lemma 4 *Let $Y \in \mathbb{R}^{n \times m}$ have orthonormal columns ($m \leq n$), and $P \in \mathbb{R}^{m \times p}$. Then*

$$||YP||_F = ||P||_F. \tag{37}$$

Proof. By using the invariance of $||\cdot||_F$ to unitary transforms, and considering $[XY] \left[\begin{array}{c} \mathbf{0} \\ P \end{array} \right]$ where $[XY] \in \mathbb{R}^{n \times n}$ is unitary. □

Lemma 5 *Let $M \in \mathbb{R}^{n \times n}$ be symmetric positive semidefinate, with norm bounded by $||M||_F \leq t < 1$. Then*

$$||I - (I + M)^{-1/2}||_F \leq t \tag{38}$$

Proof. Let $\sigma_1 \geq \ldots \geq \sigma_n$ be M's n singular values. Since the Frobenius norm of a matrix is just the square-root of the sum of the squares of its singular values, from $||M||_F \leq t$, we also know

$$\sum_{i=1}^n \sigma_i^2 \leq t^2. \tag{39}$$

Also, we have that $1 + \sigma_1, \ldots, 1 + \sigma_n$ are the singular values of $I + M$. Lastly, $1/\sqrt{1 + \sigma_n}, \ldots, 1/\sqrt{1 + \sigma_1}$ are the singular values of $(I + M)^{-1/2}$, and $1 - 1/\sqrt{1 + \sigma_1}, \ldots, 1 - 1/\sqrt{1 + \sigma_n}$ the singular values of $I - (I + M)^{-1/2}$. Thus, to show the result, it suffices to show that

$$
\begin{aligned}
||I - (I + M)^{-1/2}||_F^2 &= \sum_{i=1}^n \left(1 - \frac{1}{\sqrt{1 + \sigma_i}} \right)^2 && (40)\\
&\leq t^2 && (41)
\end{aligned}
$$

Now, it can be verified by standard means that for any $x \geq 0$, it holds true that $1 - 1/\sqrt{1 + x} \leq x/2$ [[ayn: Double-check this. Also below we only used "$\leq x$," giving a looser bound.]] Hence,.

$$\sum_{i=1}^{n}\left(1-\frac{1}{\sqrt{1+\sigma_i}}\right)^2 \leq \sum_{i=1}^{n}\sigma_i^2 \quad (42)$$

$$\leq t^2 \quad (43)$$

This proves the Lemma. □

Lemma 6 *Let assumption A4 hold. Then for each $i = 1, 2, 3$, $j = 1, \ldots, n_i$, we have that $\Lambda_{jj}^{(ii)} \leq \sqrt{Cn}$.*

Proof. $\Lambda_{jj}^{(ii)} = Z_i/\sqrt{d_j^{(i)}} = \sqrt{(\sum_{k=1}^{n_i} d_k^{(i)})/d_j^{(i)}} \leq \sqrt{Cn_i} \leq \sqrt{Cn}$. □

The next Lemma is derived by putting together several results from Matrix Perturbation Theory (Stewart and Sun, 1991), and will be used to prove the main Theorem of this section.

Lemma 7 *Let a symmetric matrix $L \in \mathbb{R}^{n\times n}$ be given, and let $x_1, x_2, \ldots, x_n$ be its n eigenvectors, with corresponding eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_n$. Let the matrices $X \in \mathbb{R}^{n\times 3}$ and $Y \in \mathbb{R}^{n\times(n-3)}$ be defined according to*

$$X = [x_1 x_2 x_3], \quad Y = [x_4 x_5 \cdots x_n]. \quad (44)$$

so that the matrix $[XY] \in \mathbb{R}^{n\times n}$ is orthogonal.

Let $\tilde{L} = L + E$ be a perturbation of A, where E is symmetric. Let $\epsilon = ||E||_F$, and also define

$$\delta = \lambda_3 - \lambda_4 \quad (45)$$

If

$$\delta > (2+\sqrt{2})\epsilon \quad (46)$$

then there is a matrix P satisfying

$$||P||_F \leq 2\frac{\epsilon}{\delta - \sqrt{2}\epsilon} \quad (47)$$

so that the columns of

$$\check{X} = (X + YP)(I + P^T P)^{-1/2} \quad (48)$$

form an orthonormal basis for the subspace spanned by $\tilde{x}_1, \tilde{x}_2$ and $\tilde{x}_3$, the first three eigenvectors of $\tilde{L}$ (which will have eigenvalues $\tilde{\lambda}_1 \geq \cdots \geq \tilde{\lambda}_n$ that satisfy $\tilde{\lambda}_3 > \tilde{\lambda}_4$).

Proof. Let the matrix E be written

$$[XY]^T E[XY] = \begin{bmatrix} E^{[11]} & E^{[12]} \\ E^{[21]} & E^{[22]} \end{bmatrix}, \quad (49)$$

(where we have used square-brackets rather than parentheses in the superscripts to acknowledge that this is not a partioned form of E, but instead of $[XY]^T E[XY]$). Also, we have the following diagonalization of L:

$$[XY]^T L[XY] = \begin{bmatrix} L_1 & 0 \\ 0 & L_2 \end{bmatrix}, \tag{50}$$

where $L_1 = \mathrm{diag}(\lambda_1, \lambda_2, \lambda_3)$, and $L_2 = \mathrm{diag}(\lambda_4, \ldots, \lambda_n)$.

Now, define $\tilde{\delta} = \delta - \sqrt{2}\epsilon$. Equation (46) implies

$$\tilde{\delta} > 0 \tag{51}$$

and

$$\frac{\epsilon}{\tilde{\delta}} < \frac{1}{2}. \tag{52}$$

Since $\epsilon = ||E||_F$, it holds true that

$$||E^{[21]}||_F = ||E^{[12]}||_F \leq \epsilon. \tag{53}$$

Moreover, since $||E^{[11]}||_F + ||E^{[22]}||_F \leq \sqrt{2}||E||_F$, we have that

$$\tilde{\delta} = \delta - \sqrt{2}\epsilon \leq \delta - ||E^{[11]}||_F - ||E^{[22]}||_F. \tag{54}$$

Thus, Theorem V.2.7 of Stewart and Sun (1991) guarantees that there exists a matrix P satisfying

$$||P||_F \leq 2\frac{\epsilon}{\tilde{\delta}} \tag{55}$$

so that the columns of

$$\breve{X} = (X + YP)(I + P^T P)^{-1/2} \tag{56}$$

form an orthonormal basis for the subspace spanned by three eigenvectors of $\tilde{L}$. The eigenvalues associated with these three eigenvectors are the eigenvalues of

$$\tilde{L}_1 = (I + P^T P)^{1/2}(L_1 + E^{[11]} + E^{[12]}P)(I + P^T P)^{-1/2} \tag{57}$$

and the eigenvalues associated with the rest of $\tilde{L}$'s eigenvectors are the eigenvalues of

$$\tilde{L}_2 = (I + PP^T)^{-1/2}(L_2 + E^{[22]} - PE^{[12]})(I + PP^T)^{1/2} \tag{58}$$

Thus, to complete the proof of the Lemma, it suffices to verify that the eigenvalues of $\tilde{L}_1$ are all (strictly) larger than the eigenvalues of $\tilde{L}_2$. But since eigenvalues are invariant similarity transformation, this is the same as showing that the eigenvalues of

$$\hat{L}_1 = L_1 + E^{[11]} + E^{[12]}P \tag{59}$$

are all (strictly) larger than the eigenvalues of

$$\hat{L}_2 = L_2 + E^{[22]} - PE^{[12]}. \tag{60}$$

Now, since $||P||_F \leq 2\epsilon/\tilde{\delta} < 1$, we have that

$$\begin{aligned}
||E^{[11]} + E^{[12]}P||_F &\leq ||E^{[11]}||_F + ||E^{[12]}P||_F && (61)\\
&\leq ||E^{[11]}||_F + ||E^{[12]}||_F||P||_F && (62)\\
&< ||E^{[11]}||_F + ||E^{[12]}||_F && (63)\\
&\leq \sqrt{2}||E||_F && (64)\\
&= \sqrt{2}\epsilon. && (65)
\end{aligned}$$

By the same argument, we also have that

$$||E^{[22]} - PE^{[12]}||_F \leq \sqrt{2}\epsilon. \quad (66)$$

But since the Frobenius norm upperbounds the Spectral norm, this also shows

$$\begin{aligned}
||E^{[11]} + E^{[12]}P||_2 &\leq \sqrt{2}\epsilon, && (67)\\
||E^{[22]} - PE^{[12]}||_2 &\leq \sqrt{2}\epsilon. && (68)
\end{aligned}$$

The eigenvalues of L_1 are $\lambda_1, \lambda_2, \lambda_3$, and those of L_2 are $\lambda_4, \ldots, \lambda_n$. By Corollary IV.3.4 of Stewart and Sun (1991), the above conditions ensure that that eigenvalues of $\hat{L}_1$ lie in the interval $[\lambda_3 - \sqrt{2}\epsilon, \lambda_1 + \sqrt{2}\epsilon]$, and that those of $\hat{L}_2$ lie in the interval $[\lambda_n - \sqrt{2}\epsilon, \lambda_4 + \sqrt{2}\epsilon]$. But since, from Equation (46), we have that

$$\begin{aligned}
\lambda_3 - \lambda_4 &= \delta\\
&> (2 + \sqrt{2})\epsilon\\
&> 2\sqrt{2}\epsilon
\end{aligned}$$

we also have that

$$\lambda_3 - \sqrt{2}\epsilon > \lambda_4 + \sqrt{2}\epsilon, \quad (69)$$

which implies that all of $\hat{L}_1$'s eigenvalues are strictly larger than all of $\hat{L}_2$'s eigenvalues. □

Theorem 8 *Let assumptions A1, A2, A3 and A4 hold. Set*

$$\epsilon = \sqrt{6\epsilon_1 + 3\epsilon_2^2}. \quad (70)$$

If

$$\delta > (2 + \sqrt{2})\epsilon \quad (71)$$

then there exists an orthogonal matrix $U = [u_1 u_2 u_3]$ *so that*

$$\frac{1}{n}\sum_{i=1}^{3}\sum_{j=1}^{n_i} ||\tilde{v}_j^{(i)} - u_i||_2^2 \leq C\left(4 + 2\sqrt{3}\right)^2 \frac{\epsilon^2}{(\delta - \sqrt{2}\epsilon)^2} \quad (72)$$

Proof. Let $Y \in \mathbb{R}^{n\times(n-3)}$ have its columns contain the eigenvectors of L not already in X, so that that $[XY] \in \mathbb{R}^{n\times n}$ is orthogonal. By Lemma 3, we know that $||E||_F \leq \epsilon$.

Define $\tilde{\delta} = \delta - \sqrt{2}\epsilon$. Note that Equation (71) certainly implies

$$\frac{\epsilon}{\tilde{\delta}} < \frac{1}{2}. \tag{73}$$

By Lemma 7, there exists a matrix P satisfying

$$||P||_F \leq 2\frac{\epsilon}{\tilde{\delta}} \tag{74}$$

so that the columns of

$$\breve{X} = (X + YP)(I + P^T P)^{-1/2} \tag{75}$$

form an orthonormal basis for the subspace spanned by the first 3 eigenvectors of $\tilde{L}$. Hence, $\tilde{X} = \breve{X}U$ for some orthogonal matrix $U \in \mathbb{R}^{3\times 3}$.

We now bound $||\breve{X} - X||_F$. Let $M = P^T P$. Then $||M||_F \leq ||P||_F^2 \leq 4\epsilon^2/\tilde{\delta}^2 < 1$.

$$\begin{aligned}
||\breve{X} - X||_F &= ||(X + YP)(I + P^T P)^{-1/2} - X||_F && (76)\\
&= ||(X + YP)(I - I + (I + P^T P)^{-1/2}) - X||_F && (77)\\
&= ||X + YP - (X + YP)(I - (I + P^T P)^{-1/2}) - X||_F && (78)\\
&\leq ||YP||_F + ||(X + YP)(I - (I + M)^{-1/2})||_F && (79)\\
&= ||P||_F + ||(X + YP)(I - (I + M)^{-1/2})||_F && (80)\\
&\leq ||P||_F + (||X||_F + ||YP||_F) \cdot ||(I - (I + M)^{-1/2})||_F && (81)\\
&\leq ||P||_F + (||X||_F + ||YP||_F)\frac{4\epsilon^2}{\tilde{\delta}^2} && (82)\\
&= ||P||_F + (||X||_F + ||P||_F)\frac{4\epsilon^2}{\tilde{\delta}^2} && (83)\\
&\leq \frac{2\epsilon}{\tilde{\delta}} + (||X||_F + ||P||_F)\frac{4\epsilon^2}{\tilde{\delta}^2} && (84)\\
&\leq \frac{2\epsilon}{\tilde{\delta}} + (\sqrt{3} + \frac{2\epsilon}{\tilde{\delta}})\frac{4\epsilon^2}{\tilde{\delta}^2} && (85)\\
&\leq \frac{2\epsilon}{\tilde{\delta}} + (\sqrt{3} + 1)\frac{2\epsilon}{\tilde{\delta}} && (86)\\
&= \left(4 + 2\sqrt{3}\right)\frac{\epsilon}{\tilde{\delta}} && (87)
\end{aligned}$$

where to derive (82), we used Lemma 5, and the derivation also used Lemma 4 and Equations (74) and (73).

Now, consider $||\Lambda\breve{X} - \Lambda X||_F^2 = ||\Lambda(\breve{X} - X)||_F^2$. This is the sum of the squares of the ℓ_2 norms of the differences between the rows of $\Lambda\breve{X}$ and ΛX. Equation (87) implies that the sum of the squares of the differences between the rows of $\breve{X}$ and X is $||\breve{X}-X||_F^2 \leq \left(4+2\sqrt{3}\right)^2 \epsilon^2/\tilde{\delta}^2$. The effect of premultiplying $\breve{X}$ and X by Λ is that the i-th row of each matrix—and hence the difference between the i-rows—is scaled by a Λ_{ii}, which is at most $\sqrt{Cn}$ (by Lemma 6). Hence, their squared difference is scaled by at most Cn, and we have

$$||\Lambda\breve{X} - \Lambda X||_F^2 \leq Cn\left(4+2\sqrt{3}\right)^2 \frac{\epsilon^2}{\tilde{\delta}^2} \tag{88}$$

Recall that $\tilde{X} = \breve{X}U$ for some orthogonal U. Thus, we also have

$$\begin{aligned} ||\Lambda\tilde{X}U - \Lambda X||_F^2 &= ||\tilde{V}U - V||_F^2 & (89) \\ &= ||\tilde{V} - VU^T||_F^2 & (90) \\ &\leq Cn\left(4+2\sqrt{3}\right)^2 \frac{\epsilon^2}{\tilde{\delta}^2} & (91) \end{aligned}$$

where $U = [u_1 u_2 u_3]$ is orthogonal. But referring to Equation (9), we see that the j-th row of V is just e_i^T if $j \in S_i$, where e_i is the i-th basis vector in $\mathbb{R}^3$. Hence, the j-th row of VU^T is simply u_i^T (if $j \in S_i$).

Letting $\tilde{v}_j^{(i)}$ denote the (transpose of) the j-th row of the i-subblock of $\tilde{V}$ (as specified in Section 2), Equation (91) therefore implies that

$$\sum_{i=1}^{3}\sum_{j=1}^{n_i} ||\tilde{v}_j^{(i)} - u_i||_2^2 \leq Cn\left(4+2\sqrt{3}\right)^2 \frac{\epsilon^2}{\tilde{\delta}^2}, \tag{92}$$

which completes the proof. □

5 Analysis part III: Original algorithm

Lemma 9 *Let any vector $x \in \mathbb{R}^k$ be given, and let $\hat{x} = x/||x||_2$ be the unit-length vector in the x-direction (or let $\hat{x}$ be any vector on the surface of the unit k-sphere if $x = 0$). Then for any other vector $y \in \mathbb{R}^k, ||y||_2 = 1$, we have $||x - \hat{x}||_2 \leq ||x - y||_2$.*

Proof. If $||x||_2 = 1$, the result is obvious. We consider the cases of $||x||_2 < 1$ and $||x||_2 > 1$ seperately. If $||x||_2 < 1$, then

$$||0 - x||_2 + ||x - y||_2 \geq ||0 - y||_2 = ||y||_2 = 1 \tag{93}$$

by the triangle inequality, so that

$$||x - y||_2 \geq 1 - ||x||_2, \tag{94}$$

which we see holds with equality if $y = \hat{x}$. If, on the other hand, $||x||_2 > 1$, then

$$1 + ||y - x||_2 = ||0 - y||_2 + ||y - x||_2 \geq ||0 - x||_2 = ||x||_2. \tag{95}$$

This implies

$$||y - x||_2 \geq ||x||_2 - 1, \tag{96}$$

which holds with equality if $y = \hat{x}$. □

Lemma 10 *Let any $x, u \in \mathbb{R}^k$ be given, with $||u||_2 = 1$. Define $\hat{x} = x/||x||_2$ (or let $\hat{x}$ be any vector on the surface of the k-sphere if $x = 0$). Then*

$$||\hat{x} - u||_2 \leq 2||x - u||_2. \tag{97}$$

Proof. By Lemma 9, we have that $||\hat{x} - x||_2 \leq ||x - u||_2$. Using the triangle inequality and this, we have

$$\begin{aligned} ||\hat{x} - u||_2 &\leq ||\hat{x} - x||_2 + ||x - u||_2 && (98) \\ &\leq 2||x - u||_2 && (99) \end{aligned}$$

This proves the Lemma. □

Theorem 11 *Let assumptions A1, A2, A3 and A4 hold. Set*

$$\epsilon = \sqrt{6\epsilon_1 + 3\epsilon_2^2}. \tag{100}$$

If

$$\delta > (2 + \sqrt{2})\epsilon \tag{101}$$

then there exists an orthogonal matrix $U = [u_1 u_2 u_3]$ so that

$$\frac{1}{n} \sum_{i=1}^{3} \sum_{j=1}^{n_i} ||\tilde{y}_j^{(i)} - u_i||_2^2 \leq 4C \left(4 + 2\sqrt{3}\right)^2 \frac{\epsilon^2}{(\delta - \sqrt{2}\epsilon)^2} \tag{102}$$

Proof. Each $\tilde{y}_j^{(i)}$ is just the corresponding $\tilde{v}_j^{(i)}$ renormalized to have unit length. Also, u_i lies on the surface of the unit sphere. Thus by Lemma 10, we have that

$$||\tilde{y}_j^{(i)} - u_i||_2 \leq 2||\tilde{v}_j^{(i)} - u_i||_2. \tag{103}$$

The result then immediately follows from Theorem 8. □