

ANGULAR SYNCHRONIZATION BY EIGENVECTORS AND SEMIDEFINITE PROGRAMMING: ANALYSIS AND APPLICATION TO CLASS AVERAGING IN CRYO-ELECTRON MICROSCOPY

A. SINGER* AND Y. SHKOLNISKY†

Abstract. The angular synchronization problem is to obtain an accurate estimation (up to a constant additive phase) for a set of unknown angles $\theta_1, \dots, \theta_n$ from m noisy measurements of their offsets $\theta_i - \theta_j \pmod{2\pi}$. Of particular interest is angle recovery in the presence of many outlier measurements that are uniformly distributed in $[0, 2\pi)$ and carry no information on the true offsets.

We introduce an efficient recovery algorithm for the unknown angles from the top eigenvector of a specially designed Hermitian matrix. The eigenvector method is extremely stable and succeeds even when the number of outliers is exceedingly large. For example, we successfully estimate $n = 400$ angles from a full set of $m = \binom{400}{2}$ offset measurements of which 90% are outliers in less than a second on a commercial laptop. We use random matrix theory to prove that the eigenvector method gives meaningful results whenever the proportion of good offset measurements is greater than $\sqrt{\frac{n}{2m}}$. We show that the eigenvector method is asymptotically nearly optimal in the sense that it achieves the information theoretic Shannon bound up to a multiplicative factor that depends on the discretization error of the measurements $2\pi/L$, but not on m and n .

The angular synchronization problem is related to the combinatorial optimization problem MAX-2-LIN MOD L for maximizing the number of satisfied linear equations mod L with exactly 2 variables in each equation. There already exist known polynomial-time semidefinite programming (SDP) approximation algorithms to MAX-2-LIN MOD L , but in practice we find such algorithms to be limited to relatively small size problems. We also present other SDP relaxations for angle recovery, drawing similarities with the Goemans-Williamson algorithm for finding the maximum cut in a weighted graph. Our experiments show that the angle recovery by the eigenvector method and by the different SDP relaxations are comparable in their quality, making the eigenvector method preferable due to its much faster running time.

We formulate and analyze the problem of finding class averages for the three-dimensional structure determination of macromolecules from cryo-electron microscopy as a particular angular synchronization problem. The angular synchronization problem in class averaging is special in the sense that the underlying graph of offset measurements is a “small-world” graph on the real projective plane \mathbb{RP}^2 , a case we analyze in detail.

In the angular synchronization problem, the angles can be viewed as elements of the rotation group $SO(2)$ and the offsets as relations among the group elements. We discuss extensions of the eigenvector method to other synchronization problems that involve different group structures and their applications, such as the time synchronization problem in distributed networks and the surface reconstruction problems in computer vision and optics.

1. Introduction. The angular synchronization problem is to estimate n unknown angles $\theta_1, \dots, \theta_n \in [0, 2\pi)$ from m noisy measurements δ_{ij} of their offsets $\theta_i - \theta_j \pmod{2\pi}$. In general, only a subset of all possible $\binom{n}{2}$ offsets are measured. The set E of pairs $\{i, j\}$ for which offset measurements exist can be realized as the edge set of a graph $G = (V, E)$ with vertices corresponding to angles and edges corresponding to measurements.

When all offset measurements are exact with zero measurement error, it is possible to solve the angular synchronization problem iff the graph G is connected. Indeed, if G is connected then it contains a spanning tree and all angles are sequentially determined by traversing the tree while summing the offsets modulo 2π . The angles are uniquely determined up to an additive phase, e.g., the angle of the root. On the

*Department of Mathematics and PACM, Princeton University, Fine Hall, Washington Road, Princeton NJ 08544-1000 USA, email: amits@math.princeton.edu

†Department of Mathematics, Program in Applied Mathematics, Yale University, 10 Hillhouse Ave. PO Box 208283, New Haven, CT 06520-8283 USA. E-mail: yoel.shkolnisky@yale.edu

other hand, if G is disconnected then it is impossible to determine the offset between angles that belong to disjoint components of the graph.

Sequential algorithms that integrate the measured offsets over a particular spanning tree of the graph are very sensitive to measurement errors, due to accumulation of the errors. It is therefore desirable to integrate all offset measurements in a globally consistent way. The need for such a globally consistent integration method comes up in a variety of applications. One such application is the time synchronization of distributed networks [1, 2], where clocks measure noisy time offsets $t_i - t_j$ from which the determination of $t_1, \dots, t_n \in \mathbb{R}$ is required. Other applications include the surface reconstruction problems in computer vision [3, 4] and optics [5], where the surface is to be reconstructed from noisy measurements of the gradient to the surface and the graph of measurements is typically the two-dimensional regular grid. The most common approach in the above mentioned applications for a self consistent global integration is the least squares approach. The least squares solution is most suitable when the offset measurements have a small Gaussian additive error. The least squares solution can be efficiently computed and also mathematically analyzed in terms of the Laplacian of the underlying measurement graph.

There are many possible models for the measurement errors, and we are mainly interested in models that allow many outliers. An outlier is an offset measurement that has a uniform distribution on $[0, 2\pi)$ regardless of the true value for the offset. In addition to outliers that carry no information on the true angle values, there also exist of course good measurements whose errors are relatively small. We have no a-priori knowledge, however, which measurements are good and which are bad (outliers).

The many-outlier error model arises naturally in the class averaging problem in cryo-electron microscopy (EM) [6]. Cryo-EM is an imaging technique for determining three-dimensional structures of macromolecules that defy crystallization. A short exposition of cryo-EM and the class averaging problem is the subject of Section 2. In Section 3 we show that the measurement graph in the class averaging problem is a “small-world” graph [7], while in Section 4 we formulate the class averaging problem as an angular synchronization problem.

The least squares method is not suitable for problems with a large number of outliers, such as the class averaging problem in cryo-EM, because the sum of squared errors will be dominated by the many outliers. The recently popular and more robust ℓ_1 convex minimization approach should perform better than the ℓ_2 minimization of least squares, but it is also expected to be limited to a relatively small number of outliers. Another approach that turns out to be problematic is the maximum likelihood approach, because it leads to a non-convex optimization problem that cannot be efficiently solved for large scale problems.

In this paper we take a different approach and introduce two different estimators for the angles, the first one is based on an eigenvector computation while the second one on a semidefinite program (SDP) [8]. Our eigenvector estimator $\hat{\theta}_1, \dots, \hat{\theta}_n$ is obtained by the following two-step recipe. In the first step, we construct an $n \times n$ complex-valued matrix H whose entries are

$$H_{ij} = \begin{cases} e^{\iota \delta_{ij}} & \{i, j\} \in E \\ 0 & \{i, j\} \notin E \end{cases}, \quad (1.1)$$

where $\iota = \sqrt{-1}$. The matrix H is Hermitian, i.e. $H_{ij} = \bar{H}_{ji}$, because the offsets are skew-symmetric $\delta_{ij} = -\delta_{ji} \pmod{2\pi}$. As H is Hermitian, its eigenvalues are real. The second step is to compute the top eigenvector v_1 of H with maximal eigenvalue, and

to define the estimator in terms of this top eigenvector as

$$e^{i\hat{\theta}_i} = \frac{v_1(i)}{|v_1(i)|}, \quad i = 1, \dots, n. \quad (1.2)$$

The philosophy leading to the eigenvector method is explained in Section 5.

The second estimator is based on the following SDP

$$\max_{\Theta \in \mathbb{C}^{n \times n}} \text{trace}(\bar{H}\Theta) \quad (1.3)$$

$$s.t. \Theta \succeq 0 \quad (1.4)$$

$$\Theta_{ii} = 1 \quad i = 1, 2, \dots, n, \quad (1.5)$$

where $\Theta \succeq 0$ is a shorthand notation for Θ being a Hermitian semidefinite positive matrix. The only difference between this SDP and the Goemans-Williamson algorithm for finding the maximum cut in a weighted graph [9] is that the maximization is taken over all semidefinite positive Hermitian matrices with complex-valued entries rather than just the real-valued symmetric matrices. The SDP-based estimator $\hat{\theta}_1, \dots, \hat{\theta}_n$ is derived from the normalized top eigenvector v_1 of Θ by the same rounding procedure (1.2). Our numerical experiments show that the accuracy of the eigenvector method and the SDP method are comparable. Since the eigenvector method is much faster, we prefer using it for large scale problems. The eigenvector method is also numerically appealing, because in the useful case the spectral gap is large, rendering the simple power method an efficient and numerically stable way of computing the top eigenvector.

In Section 7 we use random matrix theory to analyze the eigenvector method for two different measurement graphs: the small-world graph in the cryo-EM class averaging problem and the complete graph. Our analysis shows that the top eigenvector of H has a non-trivial correlation with the vector of true angles as soon as the proportion p of good offset measurements becomes greater than $\sqrt{\frac{n}{2m}}$. In particular, the correlation goes to 1 as $\frac{2mp^2}{n} \rightarrow \infty$, meaning a successful recovery of the angles. Our numerical simulations confirm these results and demonstrate the robustness of the estimator (1.2) to outliers.

In Section 8 we prove that the eigenvector method is asymptotically nearly optimal in the sense that it achieves the information theoretic Shannon bound up to a multiplicative factor that depends only on the discretization error of the measurements $2\pi/L$, but not on m and n . In other words, no method whatsoever can accurately estimate the angles if the proportion of good measurements is $o(\sqrt{\frac{n}{m}})$. The connection between the angular synchronization problem and MAX-2-LIN MOD L [10] is explored in Section 9. Finally, Section 10 is a summary and discussion of further applications of the eigenvector method to other synchronization problems over different groups.

2. The Class Averaging Problem in Cryo-Electron Microscopy. Our interest in the angular synchronization problem originates in structural biology, specifically the class averaging problem in cryo-electron microscopy (EM) [6]. The goal in cryo-EM is to determine 3D macromolecular structures from noisy projection images taken at unknown random orientations by an electron microscope, i.e., a random Computational Tomography (CT). Determining 3D macromolecular structures for large biological molecules remains vitally important, as witnessed, for example, by the 2003 Chemistry Nobel Prize, co-awarded to R. MacKinnon for resolving the 3D structure of the Shaker K^+ channel protein [11, 12]. The standard procedure for structure determination of large molecules is X-ray crystallography; the challenge in this method is

often more in the crystallization itself than in the interpretation of the X-ray results: many large proteins have so far withstood all attempts to crystallize them.

In cryo-EM, an alternative to X-ray crystallography, the sample of macromolecules is rapidly frozen in an ice layer so thin that their tomographic projections are typically disjoint; this seems the most promising alternative for molecules that defy crystallization. The cryo-EM imaging process produces a large collection of tomographic projections of the same molecule, corresponding to different and unknown projection orientations. The goal is to reconstruct the three-dimensional structure of the molecule from such unlabeled projection images, where data sets typically range from 10^4 to 10^5 projection images whose size is roughly 100×100 pixels. The intensity of the pixels in a given projection image is proportional to the line integrals of the electric potential induced by the molecule along the path of the imaging electrons (see Figure 2.1). The highly intense electron beam destroys the frozen molecule and it is therefore impractical to take projection images of the same molecule at known different directions as in the case of classical CT. In other words, a single molecule can be imaged only once, rendering an extremely low signal-to-noise ratio (SNR) for the images (see Figure 2.2 for a sample of real microscope images), mostly due to shot noise induced by the maximal allowed electron dose (other sources of noise include the varying width of the ice layer and partial knowledge of the contrast function of the microscope). In the basic homogeneity setting considered hereafter, all imaged molecules are assumed to have the exact same structure; they differ only by their spatial rotation. Every image is a projection of the same molecule but an unknown random three-dimensional rotation and the cryo-EM problem is to find the three-dimensional structure of the molecule from a collection of noisy projection images.

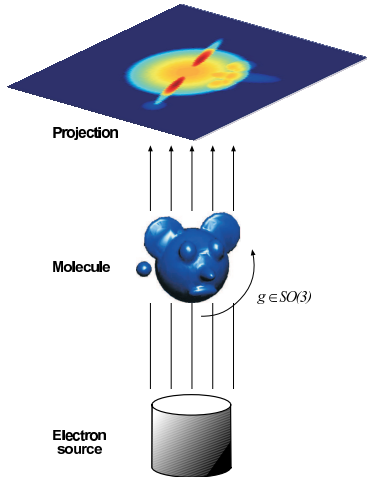


FIG. 2.1. Schematic drawing of the imaging process: every projection image corresponds to some unknown 3D rotation of the unknown molecule.

Any 3D rotation can be expressed as a rotation about some axis. The axis is a three-dimensional unit vector ν (unique except for sign) which remains unchanged by the rotation. The magnitude θ of the rotation angle is also unique, with its sign being determined by the sign of the rotation axis. We refer to the axis ν as the viewing angle, and since it is normalized, it has only two degrees of freedom and can be realized as a point on the unit sphere S^2 . The angle θ between 0 to 2π is




FIG. 2.2. A collection of four real electron microscope images of the *E. coli* ribosome.

a point on the unit circle S^1 that adds the third degree of freedom to this rotation representation. Axis-angle gives parameters in $S^2 \times S^1$; However, as ν and $-\nu$ give the same axis line, the set of viewing angles becomes \mathbb{RP}^2 , the real projective plane rather than S^2 .

As projection images in cryo-EM have extremely low SNR, a crucial initial step in all reconstruction methods is “class averaging” [6]. Class averaging is the grouping of a large data set of noisy raw projection images P_1, \dots, P_n into clusters, such that images within a single cluster have similar viewing angles. Averaging rotationally aligned noisy images within each cluster results in “class averages”; these are images that enjoy a higher SNR and are used in later cryo-EM procedures such as the angular reconstitution procedure that requires better quality images. Finding consistent class averages is challenging due to the high level of noise in the raw images as well as their large number. A sketch of the class averaging procedure is shown in Figure 2.3.




FIG. 2.3. The class averaging problem is to find, align and average images with similar viewing angles: (a) A clean simulated projection image of the *E. coli* ribosome generated from its known volume; (b) Noisy instance of (a), obtained by the addition of white Gaussian noise. The SNR is much better than that of experimental images in order for image features to be clearly visible; (c) Noisy projection of the ribosome at the same viewing angle but with a different in-plane rotation ($\theta = \pi/2$); 2.3(d) Averaging the noisy images (b) and (c) after in-plane rotational alignment by $\pi/2$. The class average of the two images has a higher SNR than that of the noisy images (b) and (c), and it has better similarity with the clean image (a).

Penczek, Zhu and Frank [13] introduced the rotationally invariant K-means clustering procedure to identify images that have similar viewing angles. Their invariant distance d_{ij} between image P_i and image P_j is defined as the Euclidean distance between the images when they are optimally aligned with respect to in-plane rotations (assuming the images are centered)

$$d_{ij} = \min_{\theta \in [0, 2\pi)} \|P_i - R(\theta)P_j\|, \quad (2.1)$$

where $R(\theta)$ is the rotation operator of an image by an angle θ . Prior to computing the invariant distances of (2.1), a common practice is to center all images by correlating them with their total average $\frac{1}{n} \sum_{i=1}^n P_i$, which is approximately radial due to the randomness in the rotations. The resulting centers usually miss the true centers by only a few pixels (as can be validated in simulations during the refinement procedure). Therefore, like [13], we also choose to focus first on the more challenging problem of rotational alignment. Later in the paper we will address some possible improvements to the centering procedure.

It is worth noting that the specific choice of metric to measure proximity between images can make a big difference in class averaging. The cross-correlation or Euclidean distance (2.1) are by no means optimal measures of proximity. In practice, it is common to denoise the images prior to computing their pairwise distances. A popular smoothing scheme is to convolve the images with a Gaussian kernel, and other linear and non-linear filters are also used in practice. Although the discussion which follows is independent of the particular choice of filter or distance metric, we stress again that filtering can have a dramatic effect on finding meaningful class averages.

The invariant distance (2.1) is invariant to in-plane rotations and as such it factors out the circle S^1 from the three-dimensional parameterization of rotations given by $\mathbb{RP}^2 \times S^1$. It follows that the invariant distance induces a metric on the two-dimensional viewing angle space \mathbb{RP}^2 . The invariant distance between images that share the same viewing angle (with perhaps a different in-plane rotation) is expected to be small. Ideally, all neighboring images of some reference image P_i in a small invariant distance ball centered at P_i should have similar viewing angles, and averaging such neighboring images (after proper rotational alignment) would amplify the signal and diminish the noise.

Unfortunately, due to the low SNR, it often happens that two images of completely different viewing angles have a small invariant distance. This can happen when the realizations of the noise match well for some random aligning angle, leading to spurious neighbor identification. Therefore, averaging the nearest neighbor images can sometimes yield a poor estimate of the true signal in the reference image.

Clustering algorithms, such as the K-means algorithm, perform much better than this naïve nearest neighbors averaging, because they take into account all pairwise distances, not just distances to the reference image. Such clustering procedures are based on the philosophy that images that share a similar viewing angle with the reference image are expected to have a small invariant distance not only to the reference image but also to all other images with similar viewing angles. This observation was utilized in the rotationally invariant K-means clustering algorithm [13]. Such clustering algorithms make it harder for spurious neighbors to sneak their way into the neighborhood. Still, noise is our enemy, and the rotationally invariant K-means clustering algorithm may suffer from misidentifications at the low SNR values present in the experimental data.

Is it possible to further improve the detection of neighboring images at even lower SNR values? In this paper we provide a positive answer to this question. First, we note that the rotationally invariant distance neglects an important piece of information, namely, the optimal angle that realizes the best rotational alignment in (2.1). Second, we observe that the optimal rotation angles must satisfy a global system of consistency relations that can be formulated as an angular synchronization problem, for which our eigenvector-based estimator (1.2) is mostly suitable due to the large amount of false detections of neighboring images.

3. A small world graph on \mathbb{RP}^2 . As mentioned earlier, the information in the optimal rotation angles has yet to be used in existing class averaging algorithms. We incorporate this additional information as follows. When computing the optimal alignment of images P_i and P_j and their invariant distance d_{ij} , we also record the rotation angle δ_{ij} that brings the distance between the two images to a minimum

$$\delta_{ij} = \operatorname{argmin}_{\theta \in [0, 2\pi)} \|P_i - R(\theta)P_j\|, \quad i, j = 1, \dots, n. \quad (3.1)$$

Note that

$$\delta_{ij} = -\delta_{ji} \pmod{2\pi}, \quad (3.2)$$

as the optimal rotation from P_j to P_i is in the opposite direction to that from P_i to P_j . By making a histogram of all $\binom{n}{2}$ distances d_{ij} , one can choose some threshold value ε , such that $d_{ij} \leq \varepsilon$ is indicative that perhaps P_i and P_j share similar viewing angles. The threshold ε defines an undirected graph $G = (V, E)$ with n vertices corresponding to the projection images, with an edge between nodes i and j iff their invariant distance is smaller than ε :

$$\{i, j\} \in E \iff d_{ij} \leq \varepsilon. \quad (3.3)$$

In an ideal noiseless world the topology of the graph is that of \mathbb{RP}^2 : the viewing angle of every image is a unit vector in three-space realized as a point on the sphere, where antipodal points correspond to the same viewing angle and are therefore being identified; if all invariant distances were trustworthy such that small distances imply similar viewing angles, then the edges of G would link neighboring points on \mathbb{RP}^2 . It is much easier to visualize the sphere S^2 rather than \mathbb{RP}^2 , as the sphere can be isometrically embedded in \mathbb{R}^3 while the projective plane requires the much harder to visualize \mathbb{R}^5 . The drawing of such a graph in five-dimensional space would show scattered points (vertices) on the two dimensional manifold of \mathbb{RP}^2 connected by short chords (edges). The analogy to the sphere is instructive because it is easier to imagine scattered points on the sphere in three-space connected by short edges. The experimental world, however, is far from ideal and is ruled by noise, giving rise to false edges that shortcut the manifold by long chords. Such graphs are known as “small-world” graphs [7], a popular model to describe social network phenomena such as the six degrees of separation: our social network consists of people living in our own town (neighboring edges), but also of some other family and friends that live across the world (shortcut edges). Planar drawings of a ring graph and its corresponding small world graph are given in Figure 3.1.

Can we tell the good edges (short chords) from the bad edges (long chords)? It is possible to denoise small world graphs based on the fact that they have many more “triangles” than random graphs: two images P_i and P_j that have similar viewing angles should have common neighboring images P_k whose viewing angles are close to theirs. All three edges $\{i, j\}$, $\{j, k\}$, and $\{k, i\}$ are in E forming a triangle (i, j, k) . On the other hand, shortcut edges are not expected to be sides of as many triangles. This cliquishness property of small-world graphs was used by Goldberg and Roth [14] to denoise protein-protein interaction maps by thresholding edges that appear in only a few triangles.

4. Class averaging as an angular synchronization problem. In the class averaging problem of cryo-EM, we can further test for the consistency of the triangles.




FIG. 3.1. (a) A ring graph with 20 vertices each of which is connected to its 6 nearest neighbors with short edges; (b) A small-world graph obtained by randomly rewiring the edges of the ring graph with probability 0.2 leading to about 20% of shortcut edges.

Indeed, if the three images P_i , P_j and P_k have similar viewing angles then the three corresponding rotation angles δ_{ij} , δ_{jk} and δ_{ki} must satisfy

$$\delta_{ij} + \delta_{jk} + \delta_{ki} = 0 \pmod{2\pi}, \quad (4.1)$$

because rotating first from P_i to P_j , followed by a rotation from P_j to P_k , and finally a rotation from P_k to P_i together complete a full circle. Equation (4.1) is a consistency relation that enables us to detect image triplets with similar viewing angles and to identify good triangles. Similarly, we may write consistency relations that involve four or more images.

The triplet consistency relation is a byproduct of an underlying angular synchronization problem. Theoretically, all projection images can be initially rotated such that they are optimally rotationally aligned. Let θ_i be the rotation angle of image P_i that brings it in sync with all other images. The mutual rotation angles δ_{ij} should satisfy the difference equations

$$\theta_i - \theta_j = \delta_{ij} \pmod{2\pi}, \quad (4.2)$$

from which the consistency relation (4.1) immediately follows:

$$\delta_{ij} + \delta_{jk} + \delta_{ki} = \theta_i - \theta_j + \theta_j - \theta_k + \theta_k - \theta_i = 0 \pmod{2\pi}.$$

There is no hope to properly align two images of completely different viewing angles, thus, the angle difference equations (4.2) make sense only for images of similar viewing angles. We therefore write a linear system of difference angle equations only for pairs $\{i, j\} \in E$

$$\theta_i - \theta_j = \delta_{ij} \pmod{2\pi}, \quad \text{for } \{i, j\} \in E, \quad (4.3)$$

as the condition $d_{ij} \leq \varepsilon$ is a proxy for the proximity of the viewing angles. Still, in practice, many of the inferred equations in (4.3) are faulty, due to noise and false detection of neighbors (shortcut edges). The edges of E can be split into a set of good edges and a set of bad edges. A good edge $\{i, j\}$ corresponds to a successful rotational alignment of P_i and P_j and a reliable measurement of δ_{ij} in the sense that (4.2) is satisfied, whereas a bad edge corresponds to a rotational alignment of P_i and P_j at

some spurious angle. We denote the set of good edges by E_{good} and the set of bad edges by E_{bad} , and their sizes by $m_{good} = |E_{good}|$ and $m_{bad} = |E_{bad}|$, respectively. Clearly, E is the disjoint union of E_{good} and E_{bad} , that is, $E = E_{good} \cup E_{bad}$ and $m = m_{good} + m_{bad}$, with $m = |E|$ being the total number of edges. Assuming that the spurious angle offsets for the bad edges are uniformly distributed on $[0, 2\pi)$, we write

$$\begin{aligned} \delta_{ij} &= \theta_i - \theta_j && \text{for } \{i, j\} \in E_{good} \\ \delta_{ij} &\sim Uniform([0, 2\pi)) && \text{for } \{i, j\} \in E_{bad}. \end{aligned} \quad (4.4)$$

Perhaps it would be more realistic to allow a small discretization error for the good offsets, for example, by letting them have the wrapped normal distribution on the circle with mean $\theta_i - \theta_j$ and variance σ^2 (where σ is a typical discretization error). This discretization error can be incorporated into the mathematical analysis of Section 7 with a little extra difficulty. However, the effect of the discretization error is negligible compared to that of the outliers, so we choose to ignore it in order to make the presentation as simple as possible.

It is trivial to find a solution to (4.3) if some oracle whispers to our ears which equations are good and which are bad (in fact, all we need in that case is that E_{good} contains a spanning tree of G). In reality, we have to be able to tell the good from the bad on our own.

The overdetermined system of equations modulo 2π in (4.3) can be solved by the method of least squares as follows. Introducing the complex-valued variables $z_i = e^{i\theta_i}$, the system (4.3) becomes

$$z_i - e^{i\delta_{ij}} z_j = 0, \quad \{i, j\} \in E, \quad (4.5)$$

which is an overdetermined system of homogeneous linear equations over \mathbb{C} . To prevent the solution from collapsing to the trivial solution $z_1 = z_2 = \dots = z_n = 0$, we set $z_1 = 1$ (recall that the angles are determined up to a global additive phase, so we may choose $\theta_1 = 0$), and look for the solution with minimal ℓ_2 -norm residual. However, it is expected that the sum of squares errors would be overwhelmingly dominated by the large number of outlier equations, making least squares least favorable to succeed (see numerical results involving least squares in Table 7.3). We therefore have to look for a solution method which is more robust to outliers.

Maximum likelihood is an obvious step in that direction. The maximum likelihood solution to (4.4) is simply the set of angles $\theta_1, \dots, \theta_n$ that satisfies as many equations of (4.3) as possible. We may therefore define the self consistency error (SCE) of $\theta_1, \dots, \theta_n$ as the number of equations not being satisfied

$$SCE(\theta_1, \dots, \theta_n) = \#\{\{i, j\} \in E : \theta_i - \theta_j \neq \delta_{ij} \pmod{2\pi}\}. \quad (4.6)$$

As even the good equations contain some error (due to angular discretization and noise), a more suitable self consistency error is SCE_f that incorporates some penalty function f

$$SCE_f(\theta_1, \dots, \theta_n) = \sum_{\{i, j\} \in E} f(\theta_i - \theta_j - \delta_{ij}), \quad (4.7)$$

where $f : [0, 2\pi) \rightarrow \mathbb{R}$ is a smooth periodic function with $f(0) = 0$ and $f(\theta) = 1$ for $|\theta| > \theta_0$, where θ_0 is the allowed discretization error. The minimization of (4.7) is equivalent to maximizing the log likelihood with a different probabilistic error model.

The maximum likelihood approach suffers from a major drawback though. It is virtually impossible to find the global minimizer $\theta_1, \dots, \theta_n$ when dealing with large scale problems ($n \gg 1$), because the minimization of either (4.6) or (4.7) is a non-convex optimization problem in a huge parameter space. It is like finding a needle in a haystack. We therefore take an alternative approach in which instead of maximizing the likelihood function we maximize a different function that measures the self-consistency of the solution. The new maximization has the advantage that it can be efficiently and effectively computed.

5. The Eigenvector Method. Our approach to finding the self consistent solution for $\theta_1, \dots, \theta_n$ starts with forming the following $n \times n$ matrix H

$$H_{ij} = \begin{cases} e^{i\delta_{ij}} & \{i, j\} \in E \\ 0 & \{i, j\} \notin E \end{cases}, \quad (5.1)$$

where $i = \sqrt{-1}$. Note that $H_{ij} = \bar{H}_{ji}$ due to (3.2), where for any complex number $z = a + ib$ we denote by $\bar{z} = a - ib$ its complex conjugate. The matrix H is therefore Hermitian, i.e., $H^* = H$.

Next, we consider the maximization problem

$$\max_{\theta_1, \dots, \theta_n \in [0, 2\pi)} \sum_{i, j=1}^n e^{-i\theta_i} H_{ij} e^{i\theta_j}, \quad (5.2)$$

and explain the philosophy behind it. For the correct set of angles $\theta_1, \dots, \theta_n$, each good edge contributes

$$e^{-i\theta_i} e^{i(\theta_i - \theta_j)} e^{i\theta_j} = 1$$

to the sum in (5.2). The total contribution of the good edges is just the sum of ones, piling up to be exactly the total number of good edges m_{good} . On the other hand, the contribution of each bad edge will be uniformly distributed on the unit circle in the complex plane. Adding up the terms due to bad edges can be thought of as a discrete planar random walk where each bad edge corresponds to a unit size step at a uniformly random direction. These random steps mostly cancel out each other, such that the total contribution of the m_{bad} edges is only $O(\sqrt{m_{bad}})$. It follows that the objective function in (5.2) has the desired property of diminishing the contribution of the faulty alignments by a square root relative to the linear contribution of the correct alignments.

Still, the maximization problem (5.2) is a non-convex maximization problem which is quite difficult to solve in practice. We therefore introduce the following relaxation of the problem

$$\max_{\substack{z_1, \dots, z_n \in \mathbb{C} \\ \sum_{i=1}^n |z_i|^2 = n}} \sum_{i, j=1}^n z_i^* H_{ij} z_j. \quad (5.3)$$

That is, we replace the previous n individual constraints for each of the variables $z_i = e^{i\theta_i}$ to have a unit magnitude, by a single and much weaker constraint, requiring the sum of squared magnitudes to be n . The maximization problem (5.3) is that of a quadratic form whose solution is simply given by the top eigenvector of the Hermitian matrix H . Indeed, the spectral theorem implies that the eigenvectors v_1, v_2, \dots, v_n of

H form an orthonormal basis for \mathbb{C}^n with corresponding real eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ satisfying $Hv_i = \lambda_i v_i$. Rewriting the constrained maximization problem (5.3) as

$$\max_{\|z\|^2=n} z^* H z, \quad (5.4)$$

it becomes clear that the maximizer z is given by $z = v_1$, where v_1 is the normalized top eigenvector satisfying $Hv_1 = \lambda_1 v_1$ and $\|v_1\|^2 = n$, with λ_1 being the largest eigenvalue. The components of the eigenvector v_1 are not necessarily of unit magnitude, so we normalize them and define the estimated rotation angles by

$$e^{i\hat{\theta}_i} = \frac{v_1(i)}{|v_1(i)|}, \quad \text{for } i = 1, \dots, n \quad (5.5)$$

(see also equation (1.2)).

The top eigenvector can be efficiently computed by the power iteration method that starts from a randomly chosen vector b_0 and iterates $b_{n+1} = \frac{Hb_n}{\|Hb_n\|}$. Each iteration requires just a matrix-vector multiplication that takes $O(n^2)$ operations for dense matrices, but only $O(m)$ operations for sparse matrices, where $m = |E|$ is the number of non-zero entries of H corresponding to edges in the graph. In our case, the number of edges is controlled by the threshold ε . The number of iterations required by the power method decreases with the spectral gap that indeed exists and is analyzed in detail in Section 7.

A closer look into the power iteration method reveals that multiplying the matrix H by itself integrates the information in the consistency relation of the triplets (4.1), while higher order iterations exploit consistency relations of longer cycles. Indeed,

$$H_{ij}^2 = \sum_{k=1}^n H_{ik} H_{kj} = \sum_{k:\{i,k\},\{j,k\} \in E} e^{i\delta_{ik}} e^{i\delta_{kj}} = \sum_{k:\{i,k\},\{j,k\} \in E} e^{-i(\delta_{jk} + \delta_{ki})} \quad (5.6)$$

$$\begin{aligned} &= \# \{k : \{i, k\} \text{ and } \{j, k\} \in E_{good}\} e^{i(\theta_i - \theta_j)} \\ &+ \sum_{k:\{i,k\} \text{ or } \{j,k\} \in E_{bad}} e^{-i(\delta_{jk} + \delta_{ki})}, \end{aligned} \quad (5.7)$$

where we employed (3.2) in (5.6), and (4.1) in (5.7).

The top eigenvector therefore gives a self-consistent rotational alignment of the raw images. The computation of the eigenvector is not only efficient but also takes all the global information into account.

6. The semidefinite program approach. A different natural relaxation of the optimization problem (5.2) is using SDP. Indeed, the objective function in (5.2) can be written as

$$\sum_{i,j=1}^n e^{-i\theta_i} H_{ij} e^{i\theta_j} = \text{trace}(\bar{H}\Theta), \quad (6.1)$$

where Θ is the $n \times n$ complex-valued rank-one Hermitian matrix

$$\Theta_{ij} = e^{i(\theta_i - \theta_j)}. \quad (6.2)$$

Note that Θ has ones on its diagonal

$$\Theta_{ii} = 1, \quad i = 1, 2, \dots, n. \quad (6.3)$$

Except for the non-convex rank-one constraint implied by (6.2), all other constraints are convex and lead to the natural SDP relaxation (1.3)-(1.5). This program is almost identical to the Goemans-Williamson SDP for finding the maximum cut in a weighted graph. The only difference is that here we maximize over all possible complex-valued Hermitian matrices, not just the symmetric real matrices. The SDP-based estimator corresponding to (1.3)-(1.5) is then obtained from the best rank-one approximation of the optimal matrix Θ using the Cholesky decomposition.

The SDP method may seem favorable to the eigenvector method as it explicitly imposes the unit magnitude constraint for $e^{i\theta_i}$. Our numerical experiments show that the two methods give similar results (see Table 7.3). Since the eigenvector method is much faster, it is also the method of choice for large scale problems.

7. Connections with random matrix theory and spectral graph theory.

In this section we analyze the eigenvector method using tools from random matrix theory and spectral graph theory.

7.1. Analysis of the complete graph angular synchronization problem.

We first consider the angular synchronization problem in which all $\binom{n}{2}$ angle differences of the form (4.2) are given, so that the corresponding graph is the complete graph of n vertices. We also assume that the probability for each edge to be good is p , independently of all other edges. This probabilistic model for the graph of good edges is known as the Erdős-Rényi random graph $G(n, p)$ [15]. We refer to this model as the complete graph angular synchronization model.

Although the angular synchronization problem in class averaging is different, as it involves a particular small-world graph rather than the complete graph, this related model allows us to gain some mathematical insight and understanding of the eigenvector method, before analyzing the more realistic small world graph model.

The elements of H in the complete graph angular synchronization model are random variables given by the following mixture model. With probability p the edge $\{i, j\}$ is good and $H_{ij} = e^{i(\theta_i - \theta_j)}$, whereas with probability $1 - p$ the edge is bad and $H_{ij} \sim \text{Uniform}(S^1)$. It is convenient to define the diagonal elements as $H_{ii} = p$.

The matrix H is Hermitian and the expected value of its elements is

$$\mathbb{E}H_{ij} = p e^{i(\theta_i - \theta_j)}. \quad (7.1)$$

In other words, the expected value of H is the rank-one matrix

$$\mathbb{E}H = npzz^*, \quad (7.2)$$

where z is the normalized vector ($\|z\| = 1$) given by

$$z_i = \frac{1}{\sqrt{n}} e^{i\theta_i}, \quad i = 1, \dots, n. \quad (7.3)$$

The matrix H can be decomposed as

$$H = npzz^* + R, \quad (7.4)$$

where $R = H - \mathbb{E}H$ is a random matrix whose elements have zero mean, with $R_{ii} = 0$, and for $i \neq j$

$$R_{ij} = \begin{cases} (1-p)e^{i(\theta_i - \theta_j)} & \text{with probability } p \\ e^{i\varphi} - pe^{i(\theta_i - \theta_j)} & \text{w.p. } 1-p \text{ and } \varphi \sim \text{Uniform}([0, 2\pi)) \end{cases}. \quad (7.5)$$

The variance of R_{ij} is

$$\mathbb{E}|R_{ij}|^2 = (1-p)^2p + (1+p^2)(1-p) = 1-p^2 \quad (7.6)$$

for $i \neq j$ and 0 for the diagonal elements. Note that for $p = 1$ the variance vanishes as all edges become good.

The distribution of the eigenvalues of the random matrix R follows Wigner's semi-circle law [16, 17] whose support is $[-2\sqrt{n(1-p^2)}, 2\sqrt{n(1-p^2)}]$. The largest eigenvalue of R , denoted $\lambda_1(R)$, is concentrated near the right edge of the support [18] and the universality of the edge of the spectrum [19] implies that it follows the Tracy-Widom distribution [20] even when the entries of R are non-Gaussian. For our purposes, the approximation

$$\lambda_1(R) \approx 2\sqrt{n(1-p^2)} \quad (7.7)$$

will suffice, with the probabilistic error bound given in [18].

The matrix $H = npzz^* + R$ can be considered as a rank-one perturbation to a random matrix. The distribution of the largest eigenvalue of such perturbed random matrices was investigated in [21, 22, 23] for the particular case where z is proportional to the all-ones vector $(1 \ 1 \cdots 1)^T$. Although our vector z given by (7.3) is different, without loss of generality, we can assume $\theta_1 = \theta_2 = \dots = \theta_n = 0$, since this assumption does not change the statistical properties of the random matrix R . Thus, adopting [22, Theorem 1.1] to H gives that for

$$np > \sqrt{n(1-p^2)} \quad (7.8)$$

the largest eigenvalue $\lambda_1(H)$ jumps outside the support of the semi-circle law and is normally distributed with mean μ and variance σ^2 given by

$$\lambda_1(H) \sim \mathcal{N}(\mu, \sigma^2), \quad \mu = \frac{np}{\sqrt{1-p^2}} + \frac{\sqrt{1-p^2}}{p}, \quad \sigma^2 = \frac{(n+1)p^2 - 1}{np^2}(1-p^2), \quad (7.9)$$

whereas for $np < \sqrt{n(1-p^2)}$, $\lambda_1(H)$ still tends to the right edge of the semicircle given at $2\sqrt{n(1-p^2)}$.

Note that the factor of 2 that appears in (7.7) has disappeared from (7.8), which is perhaps somewhat non-intuitive: it is expected that $\lambda_1(H) > \lambda_1(R)$ whenever $np > \lambda_1(R)$, but the theorem guarantees that $\lambda_1(H) > \lambda_1(R)$ also for $\frac{1}{2}\lambda_1(R) < np < \lambda_1(R)$.

The condition (7.8) also implies a lower bound on the correlation between the normalized top eigenvector v_1 of H and the vector z . To that end, consider the eigenvector equation satisfied by v_1 :

$$\lambda_1(H)v_1 = Hv_1 = (npzz^* + R)v_1. \quad (7.10)$$

Taking the dot product with v_1 yields

$$\lambda_1(H) = np|\langle z, v_1 \rangle|^2 + v_1^*Rv_1. \quad (7.11)$$

From $v_1^*Rv_1 \leq \lambda_1(R)$ we obtain the lower bound

$$|\langle z, v_1 \rangle|^2 \geq \frac{\lambda_1(H) - \lambda_1(R)}{np}, \quad (7.12)$$

with $\lambda_1(H)$ and $\lambda_1(R)$ given by (7.7) and (7.9). Thus, if the spectral gap $\lambda_1(H) - \lambda_1(R)$ is large enough then v_1 must be close to z , in which case the eigenvector method successfully recovers the unknown angles. Since the variance of the correlation of two random unit vectors in \mathbb{R}^n is $1/n$, the eigenvector method would give above random correlation values whenever

$$\frac{\lambda_1(H) - \lambda_1(R)}{np} > \frac{1}{n}. \quad (7.13)$$

Replacing in (7.13) $\lambda_1(H)$ by μ from (7.9) and $\lambda_1(R)$ by (7.7) and multiplying by $p\sqrt{n}$ yields the condition

$$\frac{\sqrt{np}}{\sqrt{1-p^2}} + \frac{\sqrt{1-p^2}}{\sqrt{np}} - 2\sqrt{1-p^2} > \frac{p}{\sqrt{n}}. \quad (7.14)$$

Since $\frac{\sqrt{np}}{\sqrt{1-p^2}} + \frac{\sqrt{1-p^2}}{\sqrt{np}} \geq 2$, it follows that (7.14) is satisfied for

$$p > \frac{1}{\sqrt{n}}. \quad (7.15)$$

Thus, already for $p > \frac{1}{\sqrt{n}}$ we should obtain above random correlations between the vector of angles z and the top eigenvector v_1 . We therefore define the threshold probability p_c as

$$p_c = \frac{1}{\sqrt{n}}. \quad (7.16)$$

When $np \gg \lambda_1(R)$, the correlation between v_1 and z can be predicted by using regular perturbation theory for solving the eigenvector equation (7.10) in an asymptotic expansion with the small parameter $\epsilon = \frac{\lambda_1(R)}{np}$. Such perturbations are derived in standard textbooks on quantum mechanics aiming to find approximations to the energy levels and eigenstates of perturbed time-independent Hamiltonians (see, e.g., [24, Chapter 6]). In our case, the resulting asymptotic expansions of the non-normalized eigenvector v_1 and of the eigenvalue $\lambda_1(H)$ are given by

$$v_1 \sim z + \frac{1}{np} [Rz - (z^* Rz)z] + \dots, \quad (7.17)$$

and

$$\lambda_1(H) \sim np + z^* Rz + \dots \quad (7.18)$$

Note that the first order term in (7.17) is perpendicular to the leading order term z , from which it follows that the angle α between the eigenvector v_1 and the vector of true angles z satisfies the asymptotic relation

$$\tan^2 \alpha \sim \frac{\|Rz\|^2 - (z^* Rz)^2}{(np)^2} + \dots, \quad (7.19)$$

because $\|Rz - (z^* Rz)z\|^2 = \|Rz\|^2 - (z^* Rz)^2$. The expected values of the numerator terms in (7.19) are given by

$$\mathbb{E}\|Rz\|^2 = \mathbb{E} \sum_{i=1}^n \left| \sum_{j=1}^n R_{ij} z_j \right|^2 = \sum_{i,j=1}^n \text{Var}(R_{ij} z_j) = \sum_{i=1}^n \sum_{j \neq i}^n |z_j|^2 (1-p^2) = (n-1)(1-p^2), \quad (7.20)$$

and

$$\begin{aligned}\mathbb{E}(z^* R z)^2 &= \mathbb{E} \left[\sum_{i,j=1}^n R_{ij} \bar{z}_i z_j \right]^2 = \sum_{i,j=1}^n \text{Var}(R_{ij} \bar{z}_i z_j) = (1-p^2) \sum_{i \neq j} |z_i|^2 |z_j|^2 \\ &= (1-p^2) \left[\left(\sum_{i=1}^n |z_i|^2 \right)^2 - \sum_{i=1}^n |z_i|^4 \right] = (1-p^2) \left(1 - \frac{1}{n} \right),\end{aligned}\quad (7.21)$$

where we used that R_{ij} are i.i.d zero mean random variables with variance given by (7.6) and that $|z_i|^2 = \frac{1}{n}$. Substituting (7.20)-(7.21) into (7.19) results in

$$\mathbb{E} \tan^2 \alpha \sim \frac{(n-1)^2(1-p^2)}{n^3 p^2} + \dots, \quad (7.22)$$

which for $p \ll 1$ and $n \gg 1$ reads

$$\mathbb{E} \tan^2 \alpha \sim \frac{1}{np^2} + \dots \quad (7.23)$$

This expression shows that as np^2 goes to infinity, the angle between v_1 and z goes to zero and the correlation between them goes to 1. For $np^2 \gg 1$, the leading order term in the expected squared correlation $\mathbb{E} \cos^2 \alpha$ is given by

$$\mathbb{E} \cos^2 \alpha = \mathbb{E} \frac{1}{1 + \tan^2 \alpha} \sim \frac{1}{1 + \frac{1}{np^2}} + \dots \quad (7.24)$$

We conclude that even for very small p values, the eigenvector method successfully recovers the angles if there are enough equations, that is, if np^2 is large enough.

Figure 7.1 shows the distribution of the eigenvalues of the matrix H for $n = 400$ and different values of p . The spectral gap decreases as p is getting smaller. From (7.8) we expect a spectral gap for $p \geq p_c$ where the critical value is $p_c = \frac{1}{\sqrt{400}} = 0.05$. The experimental values of $\lambda_1(H)$ also agree with (7.9). For example, for $n = 400$ and $p = 0.15$, the expected value of the largest eigenvalue is $\mu = 67.28$ and its standard deviation is $\sigma = 0.93$, while for $p = 0.1$ we get $\mu = 50.15$ and $\sigma = 0.86$; these value are in full agreement with the location of the largest eigenvalues in Figures 7.1(a)-7.1(b). Note that the right edge of the semi-circle is smaller than $2\sqrt{n} = 40$, so the spectral gap is significant even when $p = 0.1$.




FIG. 7.1. Histogram of the eigenvalues of the matrix H in the complete graph model for $n = 400$ and different values of p : (a) $p = 0.15$; (b) $p = 0.1$; (c) $p = 0.05$.

The skeptical reader may wonder whether the existence of a visible spectral gap necessarily implies that the normalized top eigenvector v_1 correctly recovers the original set of angles $\theta_1, \dots, \theta_n$ (up to a constant phase). To that end, we compute the

following two measures of correlation ρ_1 and ρ_2 for the correlation between the vector of true angles z and the computed normalized top eigenvector v_1 :

$$\rho_1 = \left| \frac{1}{n} \sum_{i=1}^n e^{-i\theta_i} \frac{v_1(i)}{|v_1(i)|} \right|, \quad \rho_2 = \left| \frac{1}{\sqrt{n}} \sum_{i=1}^n e^{-i\theta_i} v_1(i) \right| = |\langle z, v_1 \rangle|. \quad (7.25)$$

The correlation ρ_1 takes into account the rounding procedure (5.5), while ρ_2 is simply the dot product between v_1 and z without applying any rounding. Clearly, $\rho_1, \rho_2 \leq 1$ (Cauchy-Schwartz), and $\rho_1 = 1$ iff the two sets of angles are the same up to a rotation. Note that it is possible to have $\rho_1 = 1$ with $\rho_2 < 1$. This happens when the angles implied by $v_1(i)$ are all correct, but the magnitudes $|v_1(i)|$ are not all the same. Table 7.1 summarizes the experimentally obtained correlations ρ_1, ρ_2 for different values of p with $n = 100$ (Table 7.0(a)) and $n = 400$ (Table 7.0(b)). The experimental results show that for large values of np^2 the correlation is very close to 1, indicating a successful recovery of the angles. The third column, indicating the values of $\left(1 + \frac{1}{np^2}\right)^{-\frac{1}{2}}$ is motivated by the asymptotic expansion (7.24) and seems to provide a very good approximation for ρ_2 when $np^2 \gg 1$, with deviations attributed to higher order terms of the asymptotic expansion and to statistical fluctuations around the mean value. Below the threshold probability (ending rows of Tables 7.0(a) and 7.0(b) with $np^2 < 1$), the correlations take values near $\frac{1}{\sqrt{n}}$, as expected from the correlation of two unit random vectors in \mathbb{R}^n ($\frac{1}{\sqrt{100}} = 0.1$ and $\frac{1}{\sqrt{400}} = 0.05$).

From the practical point of view, most important is the fact that the eigenvector method successfully recovers the angles even when a large portion of the offset measurements consists of just outliers. For example, for $n = 400$, the correlation obtained when 85% of the offset measurements were outliers (only 15% are good measurements) was $\rho_1 = 0.97$.

(a) $n = 100$					(b) $n = 400$				
p	np^2	$\left(1 + \frac{1}{np^2}\right)^{-\frac{1}{2}}$	ρ_1	ρ_2	p	np^2	$\left(1 + \frac{1}{np^2}\right)^{-\frac{1}{2}}$	ρ_1	ρ_2
0.4	16	0.97	0.99	0.98	0.2	16	0.97	0.99	0.97
0.3	9	0.95	0.97	0.95	0.15	9	0.95	0.97	0.95
0.2	4	0.89	0.90	0.88	0.1	4	0.89	0.90	0.87
0.15	2.25	0.83	0.75	0.81	0.075	2.25	0.83	0.77	0.76
0.1	1	0.71	0.34	0.35	0.05	1	0.71	0.28	0.32
0.05	0.25	0.45	0.13	0.12	0.025	0.25	0.45	0.06	0.07

TABLE 7.1

Correlations between the top eigenvector v_1 of H and the vector z of true angles for different values of p in the complete graph model.

7.2. Analysis of the small world graph angular synchronization problem. The underlying graph in class averaging is a small-world graph obtained by randomly rewiring the edges of a neighborhood graph on the projective plane \mathbb{RP}^2 . This small-world graph is different from the complete graph considered earlier and requires a separate mathematical analysis. In the complete graph case, equation (7.4) gives the decomposition of H into a sum of a rank-one matrix and a Wigner random matrix. The small-world graph leads to a different decomposition that we derive be-

low. This decomposition is general in the sense that it can be applied to any graph, not just to small world graphs.

Let A be the adjacency matrix for the set of good edges E_{good} :

$$A_{ij} = \begin{cases} 1 & \{i, j\} \in E_{good} \\ 0 & \{i, j\} \notin E_{good} \end{cases} . \quad (7.26)$$

As the matrix A is symmetric, it has a complete set of real eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and corresponding real orthonormal eigenvectors ψ_1, \dots, ψ_n such that

$$A = \sum_{l=1}^n \lambda_l \psi_l \psi_l^T . \quad (7.27)$$

Let Z be an $n \times n$ diagonal matrix whose diagonal elements are $Z_{ii} = e^{i\theta_i}$. Clearly, Z is a unitary matrix ($ZZ^* = I$). Define the Hermitian matrix B by conjugating A with Z

$$B = ZAZ^* . \quad (7.28)$$

It follows that the eigenvalues of B are equal to the eigenvalues $\lambda_1, \dots, \lambda_n$ of A , and the corresponding eigenvectors $\{\phi_l\}_{l=1}^n$ of B are given by

$$\phi_l = Z\psi_l, \quad l = 1, \dots, n. \quad (7.29)$$

From (7.28) it follows that

$$B_{ij} = \begin{cases} e^{i(\theta_i - \theta_j)} & \{i, j\} \in E_{good} \\ 0 & \{i, j\} \notin E_{good} \end{cases} . \quad (7.30)$$

We are now ready to decompose the matrix H defined in (5.1) as

$$H = B + R, \quad (7.31)$$

where R is a random matrix whose elements are given by

$$R_{ij} = \begin{cases} e^{i\delta_{ij}} & \{i, j\} \in E_{bad} \\ 0 & \{i, j\} \notin E_{bad} \end{cases} , \quad (7.32)$$

where $\delta_{ij} \sim Uniform([0, 2\pi])$ for $\{i, j\} \in E_{bad}$. The decomposition (7.31) is extremely useful, because it reveals the eigen-structure of H in terms of the much simpler eigen-structures of B and R .

First, consider the matrix B defined in (7.28), which shares the same spectrum with A and whose eigenvectors ϕ_1, \dots, ϕ_n are phase modulations of the eigenvectors ψ_1, \dots, ψ_n of A . If the graph of good measurements is connected, as it must be in order to have a unique solution for the angular synchronization problem (see second paragraph of Section 1), then the Perron-Frobenius theorem (see, e.g., [25, Chapter 8]) for the non-negative matrix A implies that the entries of ψ_1 are all positive

$$\psi_1(i) > 0, \quad \text{for all } i = 1, 2, \dots, n, \quad (7.33)$$

and therefore the complex phases of the coordinates of the top eigenvector $\phi_1 = Z\psi_1$ of B are identical to the true angles, that is, $e^{i\theta_i} = \frac{\phi_1(i)}{|\phi_1(i)|}$. Hence, if the top eigenvector of

H is highly correlated with the top eigenvector of B then the angles will be recovered correctly. We will shortly derive the precise condition that guarantees such a high correlation between the eigenvectors of H and B .

The largest eigenvalue $\lambda_1(A) = \lambda_1(B)$ is squeezed between the maximal degree and the averaged degree of the good graph, that is

$$\bar{d} \leq \lambda_1(B) \leq d_{max}, \quad (7.34)$$

where

$$d_i = \sum_{j=1}^n A_{ij}, \quad \bar{d} = \frac{1}{n} \sum_{i=1}^n d_i, \quad d_{max} = \max_{i=1, \dots, n} d_i. \quad (7.35)$$

Indeed, the lower bound in (7.34) follows from the min-max Courant-Fischer theorem using the Rayleigh-Ritz quotient of the normalized all-ones vector $\frac{1}{\sqrt{n}}(1 \ 1 \ \dots \ 1)^T$

$$\lambda_1(B) = \max_{\|x\|_2=1} x^* A x \geq \frac{1}{n} \sum_{i,j=1}^n A_{ij} = \bar{d}, \quad (7.36)$$

while the upper bound is a consequence of the Perron-Frobenius theorem. From the lower bound in (7.34) we conclude that

$$\lambda_1(B) \geq \frac{2m_{good}}{n}. \quad (7.37)$$

An even better understanding of the eigen-structure of B can be achieved by looking at the normalized Laplacian matrix L of the good graph, defined as $L = I - D^{-1}A$ where I is the $n \times n$ identity matrix and D is a diagonal matrix with the vertex degrees on its diagonal, i.e., $D_{ii} = d_i$ [26]. In the particular case of the class averaging angular synchronization problem, the good graph is a neighborhood graph on \mathbb{RP}^2 and the n vertices of the graph can be viewed as n sample points on \mathbb{RP}^2 . If the vertices are sampled from the uniform distribution on \mathbb{RP}^2 , as it happens to be the case when the molecule has no preferred orientation, then in the limit of a large number of vertices/projection images $n \rightarrow \infty$, the discrete normalized Laplacian matrix L converges to the continuous Laplacian over the two dimensional manifold \mathbb{RP}^2 , denoted $\Delta_{\mathbb{RP}^2}$ [27, 28, 29, 30]. In particular, the eigenvectors and eigenvalues of the graph Laplacian matrix L converge to the eigenfunctions and eigenvalues of the Laplacian on \mathbb{RP}^2 which are the spherical harmonics of even degree. Indeed, the eigenfunctions of the (positively defined) Laplacian operator on the sphere S^2 are known to be the spherical harmonics Y_l^m [31, p.195] (also known as the eigenstates of the angular momentum operator in quantum mechanics)

$$\Delta_{S^2} Y_l^m = l(l+1)Y_l^m, \quad l = 0, 1, 2, \dots, \quad m = -l, -l+1, \dots, l. \quad (7.38)$$

The (non-normalized) spherical harmonics are given in terms of the associated Legendre polynomials of the zenith angle and trigonometric polynomials of the azimuthal angle. The eigenspaces are degenerated so that the eigenvalue $l(l+1)$ has multiplicity $2l+1$. Alternatively, the l 'th eigenspace corresponds to homogeneous polynomials of degree l restricted to S^2 (i.e., modulo $x^2 + y^2 + z^2 = 1$). Since the projective plane \mathbb{RP}^2 is obtained by identifying antipodal points on S^2 , only the even-degree homogeneous polynomials are well defined on \mathbb{RP}^2 (the odd-degree polynomials change their

sign under the transformation $(x, y, z) \mapsto (-x, -y, -z)$. It follows that the eigenfunctions of the Laplacian on \mathbb{RP}^2 are the even spherical harmonics ($l = 0, 2, 4, \dots$) with multiplicities $1, 5, 9, 13, \dots$

Now that we have a good understanding of the spectrum of B we turn to analyze the spectrum of the random matrix R given in (7.32). There are only m_{bad} nonzero elements in R , which makes R a sparse matrix with an average number of $2m_{bad}/n$ nonzero entries per row. The nonzero entries of R have zero mean and unit variance. The spectral norm of such sparse random matrices was studied in [32, 33] where it was shown that with probability 1, $\limsup_{n \rightarrow \infty} \frac{\sqrt{n}}{\sqrt{2m_{bad}}} \lambda_1(R) \leq 2$ as long as $\frac{m_{bad}}{n \log n} \rightarrow \infty$ as $n \rightarrow \infty$. The implication of this result is that we can approximate $\lambda_1(R)$ with

$$\lambda_1(R) \approx 2 \frac{\sqrt{2m_{bad}}}{\sqrt{n}}. \quad (7.39)$$

Similar to the spectral gap condition (7.8), requiring

$$\lambda_1(B) > \frac{1}{2} \lambda_1(R), \quad (7.40)$$

and employing equations (7.37) and (7.39), lead to the condition

$$\frac{2m_{good}}{n} > \frac{\sqrt{2m_{bad}}}{\sqrt{n}}, \quad (7.41)$$

or equivalently

$$m_{good} > \sqrt{\frac{nm_{bad}}{2}}. \quad (7.42)$$

Again, the condition (7.42) signifies the square-root diminishing effect of the outliers compared to the linear contribution of the good measurements.

As the original m edges of the small world graph are rewired with probability $1 - p$, the expected number of bad edges $\mathbb{E}m_{bad}$ and the expected number of good edges $\mathbb{E}m_{good}$ are given by

$$\mathbb{E}m_{good} = pm, \quad \mathbb{E}m_{bad} = (1 - p)m, \quad (7.43)$$

with relatively small fluctuations of $O(\sqrt{mp(1-p)})$. Plugging (7.43) in (7.42) results in the following condition for a non-trivial correlation between the top eigenvectors of H and B :

$$2mp^2 > n(1 - p). \quad (7.44)$$

The threshold probability p_c^{eig} satisfies the quadratic equation

$$2m(p_c^{eig})^2 = n(1 - p_c^{eig}), \quad (7.45)$$

and in the limit $m/n \gg 1$, p_c^{eig} is asymptotically given by

$$p_c^{eig} = \sqrt{\frac{n}{2m}} + O(n/m). \quad (7.46)$$

Note that (7.46) reduces to the threshold probability (7.16) of the complete graph case by setting $m = \binom{n}{2}$.

Figure 7.2 shows the histogram of the eigenvalues for the matrix H in the small-world \mathbb{RP}^2 class averaging problem. We generated random small-world graphs on \mathbb{RP}^2 with n vertices as follows. First, we sampled n points ν_1, \dots, ν_n on the unit sphere S^2 in \mathbb{R}^3 from the uniform distribution. We put an edge between i and j iff $|\nu_i \cdot \nu_j| > 1 - \varepsilon$, where ε is a small parameter that determines the connectivity (average degree) of the graph. The resulting graph is a neighborhood graph on \mathbb{RP}^2 . We then rewire edges with probability $1 - p$, so that the expected proportion of good edges is p . We used two different sets of parameters: $n = 400$ with $\varepsilon = 0.2$ for which there were about $m = 16250$ edges in the graph, and $n = 100$ with $\varepsilon = 0.3$ for which there were about $m = 1500$ edges.

The histograms of Figure 7.2 for the eigenvalues of H seem to be much more exotic than the ones obtained in the complete graph case shown in Figure 7.1. In particular, there seems to be a long tail of large eigenvalues, rather than a single eigenvalue that stands out from all the others. But now we understand that these eigenvalues are nothing but the top eigenvalues of the adjacency matrix of the good graph, related to the even spherical harmonics. This behavior is better visible in Figure 7.3.

The expected threshold probability (7.45) for $n = 400$ and $m = 16250$ is $p_c^{eg} \approx \sqrt{\frac{400}{2 \cdot 16250}} = 0.105$, while for $n = 100$ and $m = 1500$ it is $p_c^{eg} \approx 0.167$. This threshold values are in agreement with the experimental correlations given in Table 7.2, indicated by jumps in correlation that occurs between $p = 0.15$ and $p = 0.2$ for $n = 100$ and between $p = 0.1$ and $p = 0.12$ for $n = 400$. Also evident from Table 7.2 is that the correlation goes to 1 as $2mp^2/n \rightarrow \infty$. We remark that using regular perturbation theory and the relation of the eigenstructure of B to the even spherical harmonics, it is possible to obtain an asymptotic series for the correlation in terms of the large parameter $2mp^2/n$, similar to the asymptotic expansion (7.24).

(a) $n = 100, m = 1500$			(b) $n = 400, m = 16250$		
p	$\frac{2mp^2}{n}$	ρ_1	p	$\frac{2mp^2}{n}$	ρ_1
0.8	19.2	0.994	0.8	52	0.999
0.6	10.8	0.971	0.4	13	0.977
0.4	4.8	0.86	0.2	3.25	0.75
0.3	2.7	0.84	0.15	1.82	0.5
0.2	1.2	0.32	0.12	1.17	0.34
0.15	0.68	0.06	0.1	0.81	0.09

TABLE 7.2

Correlations between the top eigenvector of H and the vector of true angles for different values of p in the small-world \mathbb{RP}^2 model.




FIG. 7.2. Histogram of the eigenvalues of the matrix H in the small-world \mathbb{RP}^2 class averaging model for $n = 400$, $\varepsilon = 0.2$ and different values of p : (a) $p = 1$; (b) $p = 0.7$; (c) $p = 0.4$; (d) $p = 0.1$.




FIG. 7.3. Bar plot of the 30 largest eigenvalues of the matrix H in the small-world \mathbb{RP}^2 class averaging model for $n = 400$, $\varepsilon = 0.2$ and different values of p . The multiplicities $1, 5, 9, 13, \dots$ corresponding to the even spherical harmonics are evident as long as p is not too small. As p decreases the high-oscillatory spherical harmonics are getting “swallowed” by the semi-circle. (a) $p = 1$; (b) $p = 0.7$; (c) $p = 0.4$; (d) $p = 0.1$.

We also compared the correlations obtained by the eigenvector method and the SDP method. To solve the SDP (1.3)-(1.5) we used CVX, a package for specifying and solving convex programs [34, 35], which is mainly used for prototyping and not suitable for large scale problems that exploit the sparse block structure of the constraint (1.5). The comparison was therefore performed with $n = 200$ and $\varepsilon = 0.3$ (number of edges $m \approx 6000$), with an expected threshold $p_c^{eig} \approx 0.12$. The results are summarized in Table 7.3. Although the SDP is slightly more accurate, the eigenvector method runs much faster.

p	ρ_{lsqr}	ρ_{eig}	ρ_{sdp}	rank Θ
1	1	1	1	1
0.7	0.879	0.997	0.997	1
0.4	0.515	0.977	0.977	1
0.3	0.396	0.922	0.933	2
0.2	0.097	0.767	0.803	4
0.15	0.068	0.418	0.521	4
0.1	0.059	0.203	0.213	5

TABLE 7.3

Comparison between the correlations obtained by the eigenvector method ρ_{eig} , by the SDP method ρ_{sdp} and by the least squares method ρ_{lsqr} for different values of p (small world graph on \mathbb{RP}^2 , $n = 200$, $\varepsilon = 0.3$, $m \approx 6000$). The SDP tends to find low-rank matrices despite the fact that the rank-one constraint on Θ is not included in the SDP. The rightmost column gives the rank of the Θ matrices that were found by the SDP. The least squares solution was obtained using MATLAB’s `lsqr` function. As expected, the least squares method yields poor correlations compared to the eigenvector and the SDP methods.

8. Information Theoretic Analysis. The optimal solution to the angular synchronization problem can be considered as the set of angles that maximizes the log-likelihood. Unfortunately, the log-likelihood is a non-convex function and the maximum likelihood cannot be found in a polynomial time. Both the eigenvector method and the SDP method are polynomial-time relaxations of the maximum log-likelihood problem. In the previous section we showed that the eigenvector method fails to recover the true angles when p is below the threshold probability p_c^{eig} . It is clear that even the maximum likelihood solution would fail to recover the correct set of angles below some (perhaps lower) threshold. It is therefore natural to ask if the threshold value of the polynomial eigenvector method gets close to the optimal threshold value of the exponential-time maximum likelihood exhaustive search. In this section we provide a positive answer to this question using the information theoretic Shannon

bound [36]. Specifically, we show that the threshold probability for the eigenvector method is asymptotically larger by just a multiplicative factor compared to the threshold probability of the optimal recovery algorithm. The multiplicative factor is a function of the angular discretization resolution, but not a function of n and m . The eigenvector method becomes less optimal as the discretization resolution improves.

We start the analysis by recalling that from the information theoretic point of view, the uncertainty in the values of the angles is measured by their entropy. The noisy offset measurements carry some bits of information on the angle values, therefore decreasing their uncertainty, which is measured by the conditional entropy that we need to estimate.

The angles $\theta_1, \dots, \theta_n$ can take any real value in the interval $[0, 2\pi)$. However, an infinite number of bits is required to describe real numbers, and so we cannot hope to determine the angles with an arbitrary precision. Moreover, the offset measurements are also discretized, as each of the cryo-EM projection images is made of a finite number of pixels. We therefore seek to determine the angles only up to some discretization precision $\frac{2\pi}{L}$, where L is the number of subintervals of $[0, 2\pi)$ obtained by dividing the unit circle into L equally sized pieces.

Before observing any of the offset measurements, the angles are uniformly distributed on $\{0, 1, \dots, L-1\}$, that is, each of them falls with equal probability $1/L$ to any of the L subintervals. It follows that the entropy of the i 'th angle θ_i is given by

$$H(\theta_i) = - \sum_{l=0}^{L-1} \frac{1}{L} \log_2 \frac{1}{L} = \log_2 L, \quad \text{for } i = 1, 2, \dots, n. \quad (8.1)$$

We denote by $\boldsymbol{\theta}^n = (\theta_1, \dots, \theta_n)$ the vector of angles. Since $\theta_1, \dots, \theta_n$ are independent (the orientations of the molecules are random), their joint entropy $H(\boldsymbol{\theta}^n)$ is given by

$$H(\boldsymbol{\theta}^n) = \sum_{i=1}^n H(\theta_i) = n \log_2 L, \quad (8.2)$$

reflecting the fact that the configuration space is of size $L^n = 2^{n \log_2 L}$.

Let δ_{ij} be the random variable for the outcome of the noisy offset measurement of θ_i and θ_j . The random variable δ_{ij} is also discretized and takes values in $\{0, 1, \dots, L-1\}$. We denote by $\boldsymbol{\delta}^m = (\delta_{i_1 j_1}, \dots, \delta_{i_m j_m})$ the vector of all offset measurements. Conditioned on the values of θ_i and θ_j , the random variable δ_{ij} has the following conditional probability distribution

$$\Pr\{\delta_{ij} | \theta_i, \theta_j\} = \begin{cases} \frac{1-p}{L} & \delta_{ij} \neq \theta_i - \theta_j \pmod{L}, \\ p + \frac{1-p}{L} & \delta_{ij} = \theta_i - \theta_j \pmod{L}, \end{cases} \quad (8.3)$$

because with probability $1-p$ the measurement δ_{ij} is an outlier that takes each of the L possibilities with equal probability $\frac{1}{L}$, and with probability p it is a good measurement that equals $\theta_i - \theta_j$. It follows that the conditional entropy $H(\delta_{ij} | \theta_i, \theta_j)$ is

$$H(\delta_{ij} | \theta_i, \theta_j) = -(L-1) \frac{1-p}{L} \log_2 \frac{1-p}{L} - \left(p + \frac{1-p}{L} \right) \log_2 \left(p + \frac{1-p}{L} \right). \quad (8.4)$$

We denote this entropy by $H(L, p)$ and its deviation from $\log_2 L$ by $I(L, p)$, that is,

$$H(L, p) \equiv -(L-1) \frac{1-p}{L} \log_2 \frac{1-p}{L} - \left(p + \frac{1-p}{L} \right) \log_2 \left(p + \frac{1-p}{L} \right). \quad (8.5)$$

and

$$I(L, p) \equiv \log_2 L - H(L, p). \quad (8.6)$$

Without conditioning, the random variable δ_{ij} is uniformly distributed on $\{0, \dots, L-1\}$ and has entropy

$$H(\delta_{ij}) = \log_2 L. \quad (8.7)$$

It follows that the mutual information $I(\delta_{ij}; \theta_i, \theta_j)$ between the offset measurement δ_{ij} and the angle values θ_i and θ_j is

$$I(\delta_{ij}; \theta_i, \theta_j) = H(\delta_{ij}) - H(\delta_{ij} | \theta_i, \theta_j) = \log_2 L - H(L, p) = I(L, p). \quad (8.8)$$

This mutual information measures the reduction in the uncertainty of the random variable δ_{ij} from knowledge of θ_i and θ_j . Due to the symmetry of the mutual information,

$$I(\delta_{ij}; \theta_i, \theta_j) = H(\delta_{ij}) - H(\delta_{ij} | \theta_i, \theta_j) = H(\theta_i, \theta_j) - H(\theta_i, \theta_j | \delta_{ij}), \quad (8.9)$$

the mutual information is also the reduction in uncertainty of the angles θ_i and θ_j given the noisy measurement of their offset δ_{ij} . Thus,

$$H(\theta_i, \theta_j | \delta_{ij}) = H(\theta_i, \theta_j) - I(\delta_{ij}; \theta_i, \theta_j). \quad (8.10)$$

Similarly, given all m offset measurements $\boldsymbol{\delta}^m$, the uncertainty in $\boldsymbol{\theta}^n$ is given by

$$H(\boldsymbol{\theta}^n | \boldsymbol{\delta}^m) = H(\boldsymbol{\theta}^n) - I(\boldsymbol{\delta}^m; \boldsymbol{\theta}^n), \quad (8.11)$$

with

$$I(\boldsymbol{\delta}^m; \boldsymbol{\theta}^n) = H(\boldsymbol{\delta}^m) - H(\boldsymbol{\delta}^m | \boldsymbol{\theta}^n). \quad (8.12)$$

A simple upper bound for this mutual information is obtained by explicit evaluation of the conditional entropy $H(\boldsymbol{\delta}^m | \boldsymbol{\theta}^n)$ combined with a simple upper bound on the joint entropy term $H(\boldsymbol{\delta}^m)$. First, note that given the values of $\theta_1, \dots, \theta_n$, the offsets become independent random variables. That is, knowledge of $\delta_{i_1 j_1}$ (given $\theta_{i_1}, \theta_{j_1}$) does not give any new information on the value of $\delta_{i_2 j_2}$ (given $\theta_{i_2}, \theta_{j_2}$). The conditional probability distribution of the offsets is completely determined by (8.3), and the conditional entropy is therefore the sum of m identical entropies of the form (8.4)

$$H(\boldsymbol{\delta}^m | \boldsymbol{\theta}^n) = mH(L, p). \quad (8.13)$$

Next, bounding the joint entropy $H(\boldsymbol{\delta}^m)$ by the logarithm of its configuration space size L^m yields

$$H(\boldsymbol{\delta}^m) \leq m \log_2 L. \quad (8.14)$$

Note that this simple upper bound ignores the dependencies among the offsets which we know to exist, as implied, for example, by the triplet consistency relation (4.1). As such, (8.14) is certainly not a tight bound, but still good enough to prove our claim about the nearly optimal performance of the eigenvector method.

Plugging (8.13) and (8.14) in (8.12) yields the desired upper bound on the mutual information

$$I(\boldsymbol{\delta}^m; \boldsymbol{\theta}^n) \leq m \log_2 L - mH(L, p) = mI(L, p). \quad (8.15)$$

Now, substituting the bound (8.15) and the equality (8.2) in (8.11) gives a lower bound for the conditional entropy

$$H(\boldsymbol{\theta}^n | \boldsymbol{\delta}^m) \geq n \log_2 L - mI(L, p). \quad (8.16)$$

We may interpret this bound in the following way. Before seeing any offset measurement the entropy of the angles is $n \log_2 L$, and each of the m offset measurements can decrease the conditional entropy by at most $I(L, p)$, the information that it carries.

The bound (8.16) demonstrates, for example, that for fixed n , p and L , the conditional entropy is bounded from below by a linear decreasing function of m . It follows that unless m is large enough, the uncertainty in the angles would be too large. Information theory says that a successful recovery of all $\theta_1, \dots, \theta_n$ is possible only when their uncertainty, as expressed by the conditional entropy, is small enough. The last statement can be made precise by Fano's inequality and Wolfowitz' converse, also known as the weak and strong converse theorems to the coding theorem that provide a lower bound for the probability of the error probability in terms of the conditional entropy, see, e.g., [36, Chapter 8.9, pages 204-207] and [37, Chapter 5.8, pages 173-176].

In the language of coding, we may think of $\boldsymbol{\theta}^n$ as a codeword that we are trying to decode from the noisy vector of offsets $\boldsymbol{\delta}^m$ which is probabilistically related to $\boldsymbol{\theta}^n$. The codeword $\boldsymbol{\theta}^n$ is originally uniformly distributed on $\{1, 2, \dots, 2^{n \log_2 L}\}$ and from $\boldsymbol{\delta}^m$ we estimate $\boldsymbol{\theta}^n$ as one of the $2^{n \log_2 L}$ possibilities. Let the estimate be $\hat{\boldsymbol{\theta}}^n$ and define the probability of error as $P_e = \Pr\{\hat{\boldsymbol{\theta}}^n \neq \boldsymbol{\theta}^n\}$. Fano's inequality [36, Lemma 8.9.1, page 205] gives the following lower bound on the error probability

$$H(\boldsymbol{\theta}^n | \boldsymbol{\delta}^m) \leq 1 + P_e n \log_2 L. \quad (8.17)$$

Combining (8.17) with the lower bound for the conditional entropy (8.16) we obtain a weak lower bound on the error probability

$$P_e \geq 1 - \frac{m I(L, p)}{n \log_2 L} - \frac{1}{n \log_2 L}. \quad (8.18)$$

This lower bound for the probability of error is applicable to all decoding algorithms, not just for the eigenvector method. For large n , we see that for any $\beta < 1$,

$$\frac{m I(L, p)}{n \log_2 L} < \beta \implies P_e \geq 1 - \beta + o(1). \quad (8.19)$$

We are mainly interested in the limit $m, n \rightarrow \infty$ and $p \rightarrow 0$ with L being fixed. The Taylor expansion of $I(L, p)$ (given by (8.5)-(8.6)) near $p = 0$ reads

$$I(L, p) = \frac{1}{2}(L-1)p^2 + O(p^3). \quad (8.20)$$

Combining (8.19) and (8.20) we obtain that

$$p = \sqrt{\frac{n}{m} \frac{2 \log_2 L}{L-1}} \beta \implies P_e \geq 1 - \beta + o(1), \quad \text{as } n, m \rightarrow \infty, n/m \rightarrow 0. \quad (8.21)$$

Note that $n/m \rightarrow 0$, because $m \geq n \log n$ in order to ensure with high probability the connectivity of the measurement graph G . The bound (8.21) was derived using the weak converse theorem (Fano's inequality). It is also possible to show that the probability of error goes exponentially to 1 (using the Wolfowitz' converse and Chernoff bound, see [37, Theorem 5.8.5, pages 173-176]).

The above discussion shows that there does not exist a decoding algorithm with a small probability for the error for values of p below the threshold probability p_c^{inf} given by

$$p_c^{inf} = \sqrt{\frac{n}{m} \frac{2 \log_2 L}{L-1}}. \quad (8.22)$$

Note that for $L = 2$, the threshold probability of the eigenvector method (7.46) is 2 times smaller than p_c^{inf} . This is not a violation of information theory, because the fact that the top eigenvector has a non-trivial correlation with the vector of true angles does not mean that all angles are recovered correctly by the eigenvector. The fact that the eigenvector method gives non-trivial correlations below the information theoretic bound is just another evidence to its effectiveness.

We turn to shed some light on why it is possible to partially recover the angles below the information theoretic bound. The main issue here is that it is perhaps too harsh to measure the success of the decoding algorithm by $P_e = \Pr\{\hat{\theta}^n \neq \theta^n\}$. For example, when the decoding algorithm decodes 999 angles out of $n = 1000$ correctly while making just a single mistake, we still count it as a failure. It may be more natural to consider the probability of error in the estimation of the individual angles. We proceed to show that this measure of error leads to a threshold probability which is smaller than (8.22) by a constant factor.

Let $P_e^{(1)} = \Pr\{\hat{\theta}_1 \neq \theta_1\}$ be the probability of error in the estimation of θ_1 . Again, we want to use Fano's inequality to bound the probability of the error by bounding the conditional entropy. A simple lower bound to the conditional entropy $H(\theta_1|\delta^m)$ is obtained by conditioning on the remaining $n - 1$ angles

$$H(\theta_1|\delta^m) \geq H(\theta_1|\delta^m, \theta_2, \theta_3, \dots, \theta_n). \quad (8.23)$$

Suppose that there are d_1 noisy offset measurements of the form $\theta_1 - \theta_j$, that is, d_1 is the degree of node 1 in the measurement graph G . Let the neighbors of node 1 be j_1, j_2, \dots, j_{d_1} with corresponding offset measurements $\delta_{1j_1}, \dots, \delta_{1j_{d_1}}$. Given the values of all other angles $\theta_2, \dots, \theta_n$, and in particular the values of $\theta_{j_1}, \dots, \theta_{j_{d_1}}$, these d_1 equations become noisy equations for the single variable θ_1 . We denote these transformed equations for θ_1 alone by $\tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}$. All other $m - d_1$ equations do not involve θ_1 and therefore do not carry any information on its value. It follows that

$$H(\theta_1|\delta^m, \theta_2, \theta_3, \dots, \theta_n) = H(\theta_1|\tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}). \quad (8.24)$$

We have

$$H(\tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}|\theta_1) = d_1 H(L, p), \quad (8.25)$$

because given θ_1 these d_1 equations are i.i.d random variables with entropy $H(L, p)$. Also, a simple upper bound on the d_1 equations (without conditioning) is given by

$$H(\tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}) \leq d_1 \log_2 L, \quad (8.26)$$

ignoring possible dependencies among the outcomes. From (8.25)-(8.26) we get an upper bound for the mutual information between θ_1 and the transformed equations

$$I(\theta_1; \tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}) \leq d_1 [\log_2 L - H(L, p)] = d_1 I(L, p). \quad (8.27)$$

Combining (8.23),(8.24) (8.27) and (8.1) we get

$$\begin{aligned} H(\theta_1 | \boldsymbol{\delta}^m) &\geq H(\theta_1 | \boldsymbol{\delta}^m, \theta_2, \theta_3, \dots, \theta_n) \\ &= H(\theta_1 | \tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}) \\ &= H(\theta_1) - I(\theta_1; \tilde{\delta}_1, \dots, \tilde{\delta}_{d_1}) \\ &\geq \log_2 L - d_1 I(L, p). \end{aligned} \quad (8.28)$$

This lower bound on the conditional entropy translates, via Fano's inequality, to a lower bound on the probability of error $P_e^{(1)}$, and it follows that

$$d_1 I(L, p) > \log_2 L \quad (8.29)$$

is a necessary condition for having a small $P_e^{(1)}$. Similarly, the condition for a small probability of error in decoding θ_i is

$$d_i I(L, p) > \log_2 L, \quad (8.30)$$

where d_i is the degree of vertex i in the measurement graph. This condition suggests that we should have more success in decoding angles of high degree. The average degree \bar{d} in a graph with n vertices and m edges is $\bar{d} = \frac{2m}{n}$. The condition for successful decoding of angles with degree \bar{d} is

$$\frac{2m}{n} I(L, p) > \log_2 L. \quad (8.31)$$

In particular, this would be the condition for all vertices in a regular graph, or in a graph whose degree distribution is concentrated near \bar{d} . Recall that in class averaging of cryo-EM images, we can simply discard images with a small vertex degree, thus making the graph G more balanced. In light of (8.30), it may make sense to trim vertices whose degree is smaller than $\frac{\log_2 L}{I(L, p)}$.

Substituting the Taylor expansion (8.20) into (8.31) results in the condition

$$p > \sqrt{\frac{n \log_2 L}{m L - 1}}. \quad (8.32)$$

This means that successful decoding of the individual angles may be possible already for $p > p_c^{ind}$, where

$$p_c^{ind} = \sqrt{\frac{n \log_2 L}{m L - 1}}, \quad (8.33)$$

but the estimation of the individual angles must contain some error when $p < p_c^{ind}$. Note that $p_c^{ind} < p_c^{inf}$, so while for p values between p_c^{ind} and p_c^{inf} it is impossible to successfully decode all angles, it may still be possible to decode some angles.

Comparing the threshold probability of the eigenvector method p_c^{eig} given by (7.46) and the information theoretic threshold probability p_c^{ind} (8.33) below which

no algorithm can successfully recover individual angles, we find that their ratio is asymptotically independent of n and m :

$$\frac{p_c^{eig}}{p_c^{ind}} = \sqrt{\frac{L-1}{2\log_2 L}} + o(1). \quad (8.34)$$

Note that the threshold probability p_c^{eig} is smaller than p_c^{ind} for $L \leq 6$. Thus, we may regard the eigenvector method as an optimal recovery algorithm for offset equations with a small modulo L .

For $L \geq 7$, equation (8.34) implies a gap between the threshold probabilities p_c^{eig} and p_c^{ind} , suggesting that the exhaustive exponential search for the maximum likelihood would perform better than the polynomial time eigenvector method. Note, however, that the gap would be significant only for very large values of L that correspond to very fine angular resolutions. For example, the angular resolution of 100×100 pixel images is about $L = 100$, and the threshold probability of the eigenvector method would only be $\sqrt{\frac{99}{2\log_2 100}} \approx 2.73$ times larger than that of the maximum likelihood. The exponential complexity of $O(mL^n)$ of the exhaustive search for the maximum likelihood makes it impractical even for moderate-scale problems. On the other hand, the eigenvector method has a polynomial running time and it can handle large scale problems with relative ease.

9. Connection with Max-2-Lin mod L and Unique Games. The angular synchronization problem is related to the combinatorial optimization problem MAX-2-LIN MOD L for maximizing the number of satisfied linear equations mod L with exactly 2 variables in each equation, because the discretized offset equations $\theta_i - \theta_j = \delta_{ij}$ mod L are exactly of this form. MAX-2-LIN MOD L is a problem mainly studied in theoretical computer science, where we prefer using the notation “MOD L ” instead of the more common “MOD p ”, to avoid confusion between the size of the modulus and the proportion of good measurements.

Note that a random assignment of the angles would satisfy a $\frac{1}{L}$ fraction of the offset equations. Andersson, Engebretsen, and Håstad [10] considered SDP based algorithms for MAX-2-LIN MOD L , and showed that they could obtain an $\frac{1}{L}(1 + \kappa(L))$ -approximation algorithm, where $\kappa(L) > 0$ is a constant that depends on L . In particular, they gave a very weak proven performance guarantee of $\frac{1}{L}(1 + 10^{-8})$, though they concluded that it is most likely that their bounds can be improved significantly. Moreover, for $L = 3$ they numerically find the approximation ratio to be $\frac{1}{1.27} \approx 0.79$, and later Goemans and Williamson [38] proved a 0.793733-approximation. The SDP based algorithms in [10] are similar in their formulation to the SDP based algorithm of Frieze and Jerrum for MAX-k-CUT [39], but with a different rounding procedure. In these SDP models, L vectors are assigned to each of the n angle variables, so that the total number of vectors is nL . The resulting $nL \times nL$ matrix of inner products is required to be semidefinite positive, along with another set of $O(n^2L^2)$ linear and inequality constraints. Due to the large size of the inner product matrix and the large number of constraints, our numerical experiments with these SDP models were limited to relatively small size problems (such as $n = 20$ and $L = 7$) from which it was difficult to get a good understanding of their performance. In the small scale problems that we did manage to test, we did not find any supporting evidence that these SDP algorithms perform consistently better than the eigenvector method, despite their extensive running times and memory requirements. For our SDP experiments we used the software SDPT3 [40, 41] and SDPLR [42] in MATLAB. In [10] it is also shown that it

is NP-hard to approximate MAX-2-LIN MOD L within a constant ratio, independent of L . Thus, we should expect an L -dependent gap similar to (8.34) for any polynomial time algorithm, not just for the eigenvector method.

MAX-2-LIN is an instance of what is known as unique games [43], described below. One distinguishing feature of the offset equations is that every constraint corresponds to a bijection between the values of the associated variables. That is, for every possible value of θ_i , there is a unique value of θ_j that satisfies the constraint $\theta_i - \theta_j = \delta_{ij}$. Unique games are systems of constraints, a generalization of the offset equations, that have this uniqueness property, so that every constraint corresponds to some permutation.

As in the setting of offset equations, instances of unique games where all constraints are satisfiable are easy to handle. Given an instance where $1 - \varepsilon$ fraction of constraints are satisfiable, the Unique Games Conjecture (UGC) of Khot [44] says that it is hard to satisfy even a δ fraction of the constraints. The UGC has been shown to imply a number of inapproximability results for fundamental problems that seem difficult to obtain by more standard complexity assumptions. Note that in our angular synchronization problem the fraction of constraints that are satisfiable is $1 - \varepsilon = p + \frac{1-p}{L}$.

Charikar, Makarychev and Makarychev [45] presented improved approximation algorithms for unique games. For instances with domain size L where the optimal solution satisfies $1 - \varepsilon$ fraction of all constraints, their algorithms satisfy roughly $L^{-\varepsilon/(2-\varepsilon)}$ and $1 - O(\sqrt{\varepsilon \log L})$ fraction of all constraints. Their algorithms are based on SDP, also with an underlying inner products matrix of size $nL \times nL$, but their constraints and rounding procedure are different than those of [10]. Given the results of [46], the algorithms in [45] are near optimal if the UGC is true, that is, any improvement (beyond low order terms) would refute the conjecture. We have not tested their SDP based algorithm in practice, because, like the SDP of [10] it is also expected to be limited to relatively small scale problems.

10. Summary and Further Applications. In this paper we presented an eigenvector method and an SDP approach for solving the angular synchronization problem, and in particular for finding a self consistent rotational alignment of a large data set of noisy cryo-EM projection images. We used random matrix theory to prove that the eigenvector method finds an accurate estimate for the rotation angles even in the presence of a large number of outlier measurements.

The idea of synchronization by eigenvectors goes beyond the class averaging problem of cryo-EM images, and can be applied to other problems exhibiting a group structure and noisy measurements of relations of group elements. In this paper we specialized the class averaging problem of cryo-EM, in which images that share the same view angle were viewed as points on the unit circle S^1 which has the group structure of $SO(2)$. If two projections P_i and P_j share the same viewing angle then the optimal rotation angle δ_{ij} satisfies the group relation

$$e^{i\delta_{ij}} = e^{i\theta_i} e^{-i\theta_j}. \quad (10.1)$$

If, however, P_i and P_j do not share a similar viewing angle then the measurement of δ_{ij} is just a random angle.

In the general case we may consider a group G other than $SO(2)$ for which we have good and bad measurements of the group relations

$$g_{ij} = g_i g_j^{-1}, \quad g_i, g_j \in G. \quad (10.2)$$

For example, in the general case, the triplet consistency relation (4.1) simply reads

$$g_{ij}g_{jk}g_{ki} = g_i g_j^{-1} g_j g_k^{-1} g_k g_i^{-1} = e, \quad (10.3)$$

where e is the identity element of G .

Whenever the group G is compact and has a complex or real representation (for example, the rotation group $\text{SO}(3)$ has a real representation using 3×3 rotation matrices), we may construct an Hermitian matrix whose ij element is the matrix representation of the measurement g_{ij} and look for its top eigenvectors. In some cases the eigenvector method can be applied even when there is only partial information about the group relations as we demonstrated in constructing a self-consistent solution to the angular reconstitution problem in cryo-EM that admits the group structure of $\text{SO}(3)$ and where the common-lines give only partial information about the group relations [47]. The generalization of the eigenvector method to the distance geometry problem in NMR spectroscopy [48, 49], to localization of sensor networks [50, 51] and to the problem of low-rank matrix completion [52], as well as the formulation of such problems as d -to-1 games [45], will be discussed in a separate publication.

The eigenvector method can also be applied to non-compact groups that can be compactified. For example, consider the group of real numbers \mathbb{R} with addition. One may consider the synchronization problem of clocks that measure noisy time differences of the form

$$t_i - t_j = t_{ij}, \quad t_i, t_j \in \mathbb{R}. \quad (10.4)$$

We compactify the group \mathbb{R} by mapping it to the unit circle $t \mapsto e^{i\omega t}$, where $\omega \in \mathbb{R}$ is a parameter to be chosen not too small and not too big, as we now explain. There may be two kinds of measurement errors in (10.4). The first kind of error is a small discretization error (e.g., a small Gaussian noise) of typical size Δ . The second type of error is a large error that can be regarded as an outlier. For example, in some practical application an error of size 10Δ may be considered as an outlier. We therefore want ω to satisfy $\omega \gg (1/10)\Delta^{-1}$ (not too small) and $\omega \ll \Delta^{-1}$ (not too large), so that when constructing the matrix

$$H_{ij} = \begin{cases} e^{i\omega t_{ij}} & \{i, j\} \in E, \\ 0 & \{i, j\} \notin E, \end{cases} \quad (10.5)$$

each good equation will contribute approximately 1, while the contribution of the bad equations will be uniformly distributed on the unit circle. One may even try several different values for the “frequency” ω in analogy to the Fourier transform. An overdetermined linear system of the form (10.4) can also be solved using least squares. However, the contribution of outlier equations will dominate the sum of squares error. For example, each outlier equation with error 10Δ contributes to the sum of squares error the same as 100 good equations with error Δ . The compactification of the group combined with the eigenvector method has the appealing effect of reducing the impact of the outlier equations. This technique can also be used to center the raw projection images with respect to in-plane shifts in both directions. Each pair of images can be optimally aligned not only with respect to rotations but also with respect to translations in the x and y in-plane directions, yielding optimal shifts of the form x_{ij} and y_{ij} , each of which can be used to construct a matrix of the form (10.5). We do not expect the top eigenvectors of the translation “ H ” matrices to be as significant as the top eigenvector of the rotation H matrix, since images are

already roughly centered, and translations span only a few pixels in each direction (left-right, up-down). The eigenvector method based on (10.5) may also be useful for the surface reconstruction problems in computer vision [3, 4] and optics [5] in which current methods succeed only in the presence of a limited number of outliers.

Acknowledgments. We would like to thank Fred Sigworth and Ronald Coifman for introducing us to the cryo-electron microscopy problem and for many stimulating discussions. We are indebt to Boaz Barak for introducing us to vast literature on MAX-2-LIN MOD L and unique games. Amir Bennatan is acknowledged for referring us to the weak and strong converse theorems to the coding theorem. We also thank Robert Ghrist and Michael Robinson for valuable discussions at UPenn and for referring us to [1]. Last but not least, we thank Ben Sunday and Yosi Keller for reviewing an earlier version of the manuscript and for their helpful suggestions.

REFERENCES

- [1] Giridhar, A. and Kumar, P.R. (2006) Distributed Clock Synchronization over Wireless Networks: Algorithms and Analysis, *45th IEEE Conference on Decision and Control 2006*, pp. 4915–4920.
- [2] Karp, R., Elson, J., Estrin, D. and Shenker, S. (2003) Optimal and global time synchronization in sensor networks, *Technical Report*, Center for Embedded Networked Sensing, University of California, Los Angeles.
- [3] Frankot, R. T. and Chellappa, R. (1988) A method for enforcing integrability in shape from shading algorithms, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **10** (4): 439-451.
- [4] Agrawal, A. K., Raskar, R., and Chellappa, R. (2006) What is the range of surface reconstructions from a gradient field? in *Computer Vision – ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7-13, 2006, Proceedings, Part IV (Lecture Notes in Computer Science)*, pp. 578–591.
- [5] Rubinstein, J. and Wolansky, G. (2001) Reconstruction of optical surfaces from ray data, *Optical Review*, **8** (4), pp. 281–283.
- [6] Frank, J. (2006) *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*, Oxford.
- [7] Watts, D. J. and Strogatz, S. H. (1998) Collective dynamics of small-world networks. *Nature* **393**, pp. 440–442.
- [8] Vandenberghe, L., and Boyd, S. (1996) Semidefinite programming. *SIAM Review*, **38**(1):49–95.
- [9] Goemans, M. X. and Williamson, D. P. (1995) Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)* **42** (6), pp. 1115–1145.
- [10] Andersson, G., Engebretsen, L., and Håstad, J. (1999) A new way to use semidefinite programming with applications to linear equations mod p . *Proceedings 10th annual ACM-SIAM symposium on Discrete algorithms*, pp. 41–50.
- [11] Doyle, D. A., Cabral, J. M., Pfuetzner, R. A., Kuo, A., Gulbis, J. M., Cohen, S. L., Chait, B. T., MacKinnon, R. (1998) The Structure of the Potassium Channel: Molecular Basis of K⁺ Conduction and Selectivity, *Science* 3 April 1998 **280** (5360): pp. 69–77.
- [12] MacKinnon, R. (2004) Potassium Channels and the Atomic Basis of Selective Ion Conduction, 8 December 2003, Nobel Lecture, in *Bioscience Reports*, Springer Netherlands, **24** (2) pp. 75–100.
- [13] Penczek, P.A., Zhu, J., and Frank, J. (1996) A common-lines based method for determining orientations for $N > 3$ particle projections simultaneously. *Ultramicroscopy* **63**, pp. 205–218.
- [14] Goldberg, D. S. and Roth, F. P. (2003) Assessing experimentally derived interactions in a small world. *Proceedings of the National Academy of Sciences* **100** (8) pp. 4372–4376.
- [15] Erdős, P. and Rényi, A. (1959) On random graphs. *Publicationes Mathematicae* **6**, pp. 290–297.
- [16] Wigner, E. P. (1955) Characteristic vectors of bordered matrices with infinite dimensions, *Annals of Mathematics* **62** pp. 548–564.
- [17] Wigner, E. P. (1958) On the distribution of tile roots of certain symmetric matrices, *Annals of Mathematics* **67** pp. 325–328.

- [18] Alon, N., Krivelevich, M., and Vu, V. H. (2002) On the concentration of eigenvalues of random symmetric matrices, *Israel Journal of Mathematics* **131** (1) pp. 259–267.
- [19] Soshnikov, A. (1999) Universality at the edge of the spectrum in Wigner random matrices. *Comm. Math. Phys.* **207**, pp. 697–733.
- [20] Tracy, C. A., and H. Widom (1994) Level-spacing distributions and the Airy kernel. *Communications in Mathematical Physics* **159** (1), pp. 151–174.
- [21] Péché, S. (2006) The largest eigenvalues of small rank perturbations of Hermitian random matrices, *Prob. Theo. Rel. Fields* **134** (1): 127–174 .
- [22] Féral, D. and Péché, S. (2007) The Largest Eigenvalue of Rank One Deformation of Large Wigner Matrices, *Communications in Mathematical Physics* **272** (1): 185–228.
- [23] Füredi, Z., and Komlós, J. (1981) The eigenvalues of random symmetric matrices. *Combinatorica*, **1**, pp. 233–241.
- [24] Griffiths, D. J. (1994) *Introduction to Quantum Mechanics*, Prentice Hall, NJ, 416 pages.
- [25] Horn, R. A., and Johnson, C. R. (1990) *Matrix Analysis*. Cambridge University Press, 575 pages.
- [26] Chung, F. R. K. (1997) *Spectral Graph Theory*, AMS Bookstore, pp. 1–203.
- [27] Coifman, R. R. and Lafon, S. (2006) Diffusion maps. *Applied and Computational Harmonic Analysis*, **21** (1), pp. 5–30.
- [28] Belkin, M. and Niyogi, P. (2005) Towards a theoretical foundation for laplacian-based manifold methods. In *Proc. 18th Conf. Learning Theory (COLT)*, ser. Lecture Notes Computer Science, P. Auer and R. Meir, Eds. Berlin: Springer-Verlag, vol. 3559, pp. 486–500.
- [29] Hein, M., Audibert, J., and von Luxburg, U. (2005) From Graphs to Manifolds - weak and strong pointwise consistency of graph laplacians. *Proceedings of the 18th Conference on Learning Theory (COLT)*, pp. 470–485.
- [30] Singer, A. (2006) From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, **21**, pp. 128–134.
- [31] Natterer, F. (2001) *The Mathematics of Computerized Tomography*, SIAM: Society for Industrial and Applied Mathematics, Classics in Applied Mathematics.
- [32] Khorunzhy, A. (2001) Sparse Random Matrices: Spectral Edge and Statistics of Rooted Trees, *Advances in Applied Probability* **33** (1) pp. 124–140.
- [33] Khorunzhiy, O. (2003) Rooted trees and moments of large sparse random matrices, *Discrete Mathematics and Theoretical Computer Science AC*, pp. 145–154.
- [34] Grant, M. and Boyd, S. (2009) CVX: Matlab software for disciplined convex programming (web page and software). <http://stanford.edu/~boyd/cvx>.
- [35] Grant, M. and Boyd, S. (2008) Graph implementations for nonsmooth convex programs. Recent Advances in Learning and Control (a tribute to M. Vidyasagar), V. Blondel, S. Boyd, and H. Kimura, editors, pages 95–110, Lecture Notes in Control and Information Sciences, Springer. http://stanford.edu/~boyd/graph_dcp.html.
- [36] Cover, T. M. and Thomas, J. A. (1991) *Elements of Information Theory*, Wiley, New York.
- [37] Gallager, R. G. (1968) *Information Theory and Reliable Communication*, Wiley, New York.
- [38] Goemans, M. X. and Williamson, D. P. (2001) Approximation algorithms for MAX-3-CUT and other problems via complex semidefinite programming. *Proceedings 33rd annual ACM symposium on Theory of computing*, pp. 443–452.
- [39] Frieze A., and Jerrum, M. (1997) Improved Approximation Algorithms for MAX k-CUT and MAX BISECTION. *Algorithmica* **18** (1), pp. 67–81.
- [40] Toh, K.C., Todd, M.J., and Tutuncu, R.H. (1999) SDPT3 — a Matlab software package for semidefinite programming, *Optimization Methods and Software*, **11**, pp. 545–581.
- [41] Tutuncu, R.H., Toh, K.C., and Todd, M.J. (2003) Solving semidefinite-quadratic-linear programs using SDPT3, *Mathematical Programming Ser. B*, **95**, pp. 189–217.
- [42] Burer, S. and Monteiro, R.D.C. (2003) A Nonlinear Programming Algorithm for Solving Semidefinite Programs Via Low-Rank Factorization. *Mathematical Programming (series B)*, **95** (2):329–357.
- [43] Feige, U., and Lovász, L. (1992) Two-prover one round proof systems: Their power and their problems. In *Proceedings of the 24th ACM Symposium on Theory of Computing*, pp. 733–741.
- [44] Khot, S. (2002) On the power of unique 2-prover 1-round games. In *Proceedings of the ACM Symposium on the Theory of Computing*, **34**, pp. 767–775.
- [45] Charikar, M., Makarychev, K., and Makarychev, Y. (2006) Near-Optimal Algorithms for Unique Games. *Proceedings 38th annual ACM symposium on Theory of computing*, pp. 205–214.
- [46] Khot, S., Kindler, G., Mossel, E., and O’Donnell, R. (2007) Optimal inapproximability results for MAX-CUT and other two-variable CSPs? *SIAM Journal of Computing* **37** (1), pp. 319–357 (2007).

- [47] Coifman, R. R., Shkolnisky, Y., Sigworth, F. J., and Singer, A. (2009) *in preparation*.
- [48] Wuthrich, K. (2003). NMR studies of structure and function of biological macromolecules (Nobel Lecture). *J Biomol NMR* **27**, 13-39.
- [49] Havel, T. F., and Wuthrich, K. (1985). An evaluation of the combined use of nuclear magnetic resonance and distance geometry for the determination of protein conformation in solution. *J Mol Biol* **182**, 281-294.
- [50] Biswas, P., Liang, T. C., Toh, K.C., Wang, T. C., and Ye, Y. (2006) Semidefinite programming approaches for sensor network localization with noisy distance measurements. *IEEE Transactions on Automation Science and Engineering*, **3**(4):360-371.
- [51] Singer, A. (2008) A Remark on Global Positioning from Local Distances. *Proceedings of the National Academy of Sciences*, **105** (28):9507-9511.
- [52] Candès, E.J. and Recht, B. Exact Matrix Completion via Convex Optimization, *submitted for publication*.