10 Synchronization Problems and Alignment

10.1 Synchronization-type problems

This section will focuses on synchronization-type problems.[36] These are problems where the goal is to estimate a set of parameters from data concerning relations or interactions between pairs of them. A good example to have in mind is an important problem in computer vision, known as structure from motion: the goal is to build a three-dimensional model of an object from several two-dimensional photos of it taken from unknown positions. Although one cannot directly estimate the positions, one can compare pairs of pictures and gauge information on their relative positioning. The task of estimating the camera locations from this pairwise information is a synchronization-type problem. Another example, from signal processing, is multireference alignment, which is the problem of estimating a signal from measuring multiple arbitrarily shifted copies of it that are corrupted with noise.

We will formulate each of these problems as an estimation problem on a graph $G = (V, E)$. More precisely, we will associate each data unit (say, a photo, or a shifted signal) to a graph node $i \in V$. The problem can then be formulated as estimating, for each node $i \in V$, a group element $g_i \in \mathcal{G}$, where the group $\mathcal{G}$ is a group of transformations, such as translations, rotations, or permutations. The pairwise data, which we identify with edges of the graph $(i, j) \in E$, reveals information about the ratios $g_i(g_j)^{-1}$. In its simplest form, for each edge $(i, j) \in E$ of the graph, we have a noisy estimate of $g_i(g_j)^{-1}$ and the synchronization problem consists of estimating the individual group elements $g : V \to \mathcal{G}$ that are the most consistent with the edge estimates, often corresponding to the Maximum Likelihood (ML) estimator. Naturally, the measure of "consistency" is application specific. While there is a general way of describing these problems and algorithmic approaches to them [BCS15, Ban15a], for the sake of simplicity we will illustrate the ideas through some important examples.

10.2 Angular Synchronization

The angular synchronization problem [Sin11, BSS13] consist in estimating n unknown angles $\theta_1, \dots, \theta_n$ from m noisy measurements of their offsets $\theta_i - \theta_j \mod 2\pi$. This problem easily falls under the scope of synchronization-type problem by taking a graph with a node for each θ_i, an edge associated with each measurement, and taking the group to be $\mathcal{G} \cong SO(2)$, the group of in-plane rotations. Some of its applications include time-synchronization of distributed networks [GK06], signal reconstruction from phaseless measurements [ABFM12], surface reconstruction problems in computer vision [ARC06] and optics [RW01].

Let us consider a particular instance of this problem (with a particular noise model).

Let $z_1, \dots, z_n \in \mathbb{C}$ satisfying $|z_a| = 1$ be the signal (angles) we want to estimate ($z_a = \exp(i\theta_a)$). Suppose for every pair (i, j) we make a noisy measurement of the angle offset

$$Y_{ij} = z_i \overline{z_j} + \sigma W_{ij},$$

where $W_{ij} \sim \mathcal{N}(0, 1)$. The maximum likelihood estimator for z is given by solving (see [Sin11, BBS14])

$$\max_{|x_i|^2 = 1} x^* Y x. \tag{103}$$

[36] And it will follow somewhat the structure in Chapter 1 of [Ban15a]

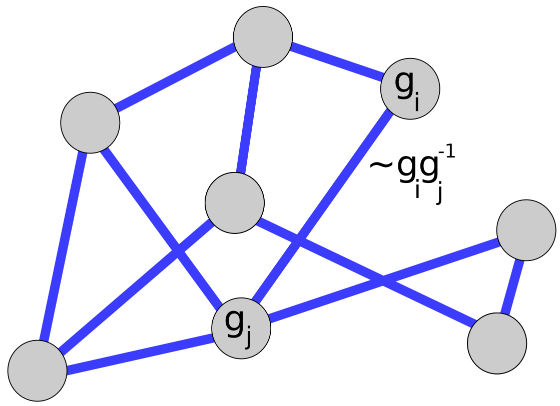

Figure 22: Given a graph $G = (V, E)$ and a group $\mathcal{G}$, the goal in synchronization-type problems is to estimate node labels $g : V \to \mathcal{G}$ from noisy edge measurements of offsets $g_i g_j^{-1}$.

There are several approaches to try to solve (103). Using techniques very similar to the study of the spike model in PCA on the first lecture one can (see [Sin11]), for example, understand the performance of the spectral relaxation of (103) into

$$\max_{\|x\|^2 = n} x^* Y x. \tag{104}$$

Notice that, since the solution to (104) will not necessarily be a vector with unit-modulus entries, a rounding step will, in general, be needed. Also, to compute the leading eigenvector of A one would likely use the power method. An interesting adaptation to this approach is to round after each iteration of the power method, rather than waiting for the end of the process, more precisely:

Algorithm 10.1 *Given Y. Take a original (maybe random) vector $x^{(0)}$. For each iteration k (until convergence or a certain number of iterations) take $x^{(k+1)}$ to be the vector with entries:*

$$\left(x^{(k+1)}\right)_i = \frac{\left(Yx^{(k)}\right)_i}{\left|\left(Yx^{(k)}\right)_i\right|}.$$

Although this method appears to perform very well in numeric experiments, its analysis is still an open problem.

Open Problem 10.1 *In the model where $Y = zz^* + \sigma W$ as described above, for which values of σ will the Projected Power Method (Algorithm 10.1) converge to the optimal solution of* (103) *(or at least to a solution that correlates well with z), with high probability?*[37]

[37] We thank Nicolas Boumal for suggesting this problem.

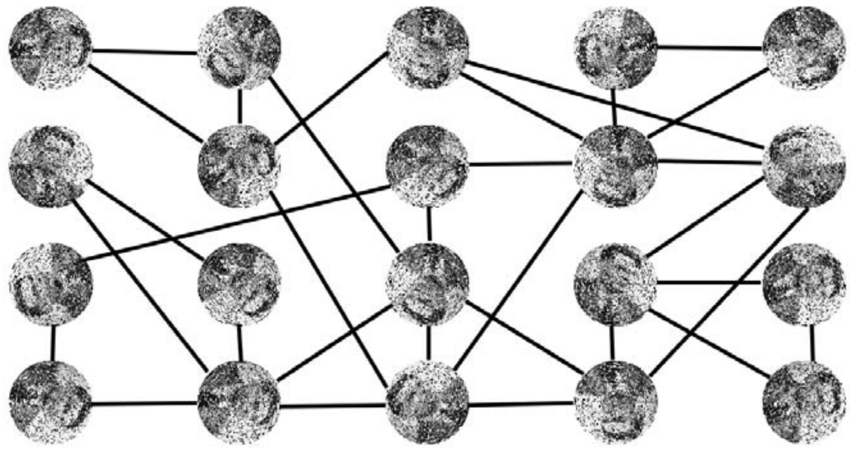

Figure 23: An example of an instance of a synchronization-type problem. Given noisy rotated copies of an image (corresponding to vertices of a graph), the goal is to recover the rotations. By comparing pairs of images (corresponding to edges of the graph), it is possible to estimate the relative rotations between them. The problem of recovering the rotation of each image from these relative rotation estimates is an instance of Angular synchronization.

We note that Algorithm 10.1 is very similar to the Approximate Message Passing method presented, and analyzed, in [MR14] for the positive eigenvector problem.

Another approach is to consider an SDP relaxation similar to the one for Max-Cut and minimum bisection.

$$\begin{array}{rc} \max & \mathrm{Tr}(YX) \\ \text{s.t.} & X_{ii} = 1, \forall_i \\ & X \succeq 0. \end{array} \tag{105}$$

In [BBS14] it is shown that, in the model of $Y = zz^* zz^* + \sigma W$, as long as $\sigma = \tilde{\mathcal{O}}(n^{1/4})$ then (105) is tight, meaning that the optimal solution is rank 1 and thus it corresponds to the optimal solution of (103).[38]. It is conjecture [BBS14] however that $\sigma = \tilde{\mathcal{O}}(n^{1/2})$ should suffice. It is known (see [BBS14]) that this is implied by the following conjecture:

If $x^\natural$ is the optimal solution to (103), then with high probability $\|Wx^\natural\|_\infty = \tilde{\mathcal{O}}(n^{1/2})$. This is the content of the next open problem.

Open Problem 10.2 *Prove or disprove: With high probability the SDP relaxation* (105) *is tight as long as* $\sigma = \tilde{\mathcal{O}}(n^{1/2})$*. This would follow from showing that, with high probability* $\|Wx^\natural\|_\infty = \tilde{\mathcal{O}}(n^{1/2})$*, where* $x^\natural$ *is the optimal solution to* (103)*.*

[38]Note that this makes (in this regime) the SDP relaxation a Probably Certifiably Correct algorithm [Ban15b]

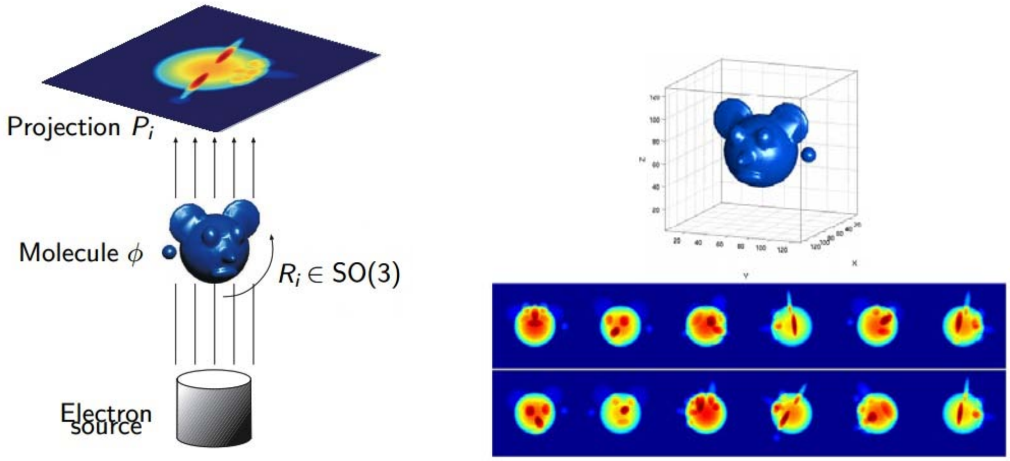

Image courtesy of Prof. Amit Singer, Princeton University. Used with permission.

Figure 24: Illustration of the Cryo-EM imaging process: A molecule is imaged after being frozen at a random (unknown) rotation and a tomographic 2-dimensional projection is captured. Given a number of tomographic projections taken at unknown rotations, we are interested in determining such rotations with the objective of reconstructing the molecule density. Images courtesy of Amit Singer and Yoel Shkolnisky [SS11].

We note that the main difficulty seems to come from the fact that W and $x^\natural$ are not independent random variables.

10.2.1 Orientation estimation in Cryo-EM

A particularly challenging application of this framework is the orientation estimation problem in Cryo-Electron Microscopy [SS11].

Cryo-EM is a technique used to determine the three-dimensional structure of biological macromolecules. The molecules are rapidly frozen in a thin layer of ice and imaged with an electron microscope, which gives 2-dimensional projections. One of the main difficulties with this imaging process is that these molecules are imaged at different unknown orientations in the sheet of ice and each molecule can only be imaged once (due to the destructive nature of the imaging process). More precisely, each measurement consists of a tomographic projection of a rotated (by an unknown rotation) copy of the molecule. The task is then to reconstruct the molecule density from many such measurements. As the problem of recovering the molecule density knowing the rotations fits in the framework of classical tomography—for which effective methods exist— the problem of determining the unknown rotations, the orientation estimation problem, is of paramount importance. While we will not go into details here, there is a mechanism that, from two such projections, obtains information between their orientation. The problem of finding the orientation of each projection from such pairwise information naturally fits in the framework of synchronization and some of the techniques described here can be adapted to this setting [BCS15].

10.2.2 Synchronization over $\mathbb{Z}_2$

This particularly simple version already includes many applications of interest. Similarly to before, given a graph $G = (V, E)$, the goal is recover unknown node labels $g : V \to \mathbb{Z}_2$ (corresponding to memberships to two clusters) from pairwise information. Each pairwise measurement either suggests the two involved nodes are in the same cluster or in different ones (recall the problem of recovery in the stochastic block model). The task of clustering the graph in order to agree, as much as possible, with these measurements is tightly connected to *correlation clustering* [BBC04] and has applications to determining the orientation of a manifold [SW11].

In the case where all the measurements suggest that the involved nodes belong in different communities, then this problem essentially reduces to the `Max-Cut` problem.

10.3 Signal Alignment

In signal processing, the multireference alignment problem [BCSZ14] consists of recovering an unknown signal $u \in \mathbb{R}^L$ from n observations of the form

$$y_i = R_{l_i} u + \sigma \xi_i, \tag{106}$$

where R_{l_i} is a circulant permutation matrix that shifts u by $l_i \in \mathbb{Z}_L$ coordinates, ξ_i is a noise vector (which we will assume standard gaussian i.i.d. entries) and l_i are unknown shifts.

If the shifts were known, the estimation of the signal u would reduce to a simple denoising problem. For that reason, we will focus on estimating the shifts $\{l_i\}_{i=1}^n$. By comparing two observations y_i and y_j we can obtain information about the relative shift $l_i - l_j \mod L$ and write this problem as a Synchronization problem

10.3.1 The model bias pitfall

In some of the problems described above, such as the multireference alignment of signals (or the orientation estimation problem in Cryo-EM), the alignment step is only a subprocedure of the estimation of the underlying signal (or the 3d density of the molecule). In fact, if the underlying signal was known, finding the shifts would be nearly trivial: for the case of the signals, one could simply use match-filtering to find the most likely shift l_i for measurement y_i (by comparing all possible shifts of it to the known underlying signal).

When the true signal is not known, a common approach is to choose a reference signal z that is not the true template but believed to share some properties with it. Unfortunately, this creates a high risk of model bias: the reconstructed signal $\hat{u}$ tends to capture characteristics of the reference z that are not present on the actual original signal u (see Figure 10.3.1 for an illustration of this phenomenon). This issue is well known among the biological imaging community [SHBG09, Hen13] (see, for example, [Coh13] for a particularly recent discussion of it). As the experiment shown on Figure 10.3.1 suggests, the methods treated in this paper, based solely on pairwise information between observations, do not suffer from model bias as they do not use any information besides the data itself.

In order to recover the shifts l_i from the shifted noisy signals (106) we will consider the following estimator

$$\operatorname{argmin}_{l_1,\ldots,l_n \in \mathbb{Z}_L} \sum_{i,j \in [n]} \left\| R_{-l_i} y_i - R_{-l_j} y_j \right\|^2, \tag{107}$$

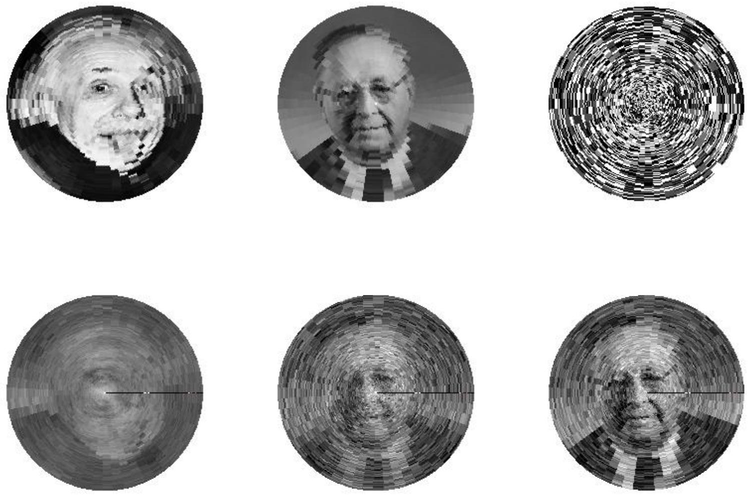

Figure 25: A simple experiment to illustrate the model bias phenomenon: Given a picture of the mathematician Hermann Weyl (second picture of the top row) we generate many images consisting of random rotations (we considered a discretization of the rotations of the plane) of the image with added gaussian noise. An example of one such measurements is the third image in the first row. We then proceeded to align these images to a reference consisting of a famous image of Albert Einstein (often used in the model bias discussions). After alignment, an estimator of the original image was constructed by averaging the aligned measurements. The result, first image on second row, clearly has more resemblance to the image of Einstein than to that of Weyl, illustration the model bias issue. One the other hand, the method based on the synchronization approach produces the second image of the second row, which shows no signs of suffering from model bias. As a benchmark, we also include the reconstruction obtained by an oracle that is given the true rotations (third image in the second row).

which is related to the maximum likelihood estimator of the shifts. While we refer to [Ban15a] for a derivation we note that it is intuitive that if l_i is the right shift for y_i and l_j for y_j then $R_{-l_i}y_i - R_{-l_j}y_j$ should be random gaussian noise, which motivates the estimator considered.

Since a shift does not change the norm of a vector, (107) is equivalent to

$$\underset{l_1,\dots,l_n\in\mathbb{Z}_L}{\operatorname{argmax}} \sum_{i,j\in[n]} \langle R_{-l_i}y_i, R_{-l_j}y_j\rangle, \tag{108}$$

we will refer to this estimator as the quasi-MLE.

It is not surprising that solving this problem is NP-hard in general (the search space for this optimization problem has exponential size and is nonconvex). In fact, one can show [BCSZ14] that, conditioned on the Unique Games Conjecture, it is hard to approximate up to any constant.

10.3.2 The semidefinite relaxation

We will now present a semidefinite relaxation for (108) (see [BCSZ14]).

Let us identify R_l with the $L \times L$ permutation matrix that cyclicly permutes the entries fo a vector by l_i coordinates:

$$R_l \begin{bmatrix} u_1 \\ \vdots \\ u_L \end{bmatrix} = \begin{bmatrix} u_{1-l} \\ \vdots \\ u_{L-l} \end{bmatrix}.$$

This corresponds to an L-dimensional representation of the cyclic group. Then, (108) can be rewritten:

$$\begin{aligned} \sum_{i,j\in[n]} \langle R_{-l_i} y_i, R_{-l_j} y_j \rangle &= \sum_{i,j\in[n]} \left(R_{-l_i} y_i\right)^T R_{-l_j} y_j \\ &= \sum_{i,j\in[n]} \operatorname{Tr}\left[\left(R_{-l_i} y_i\right)^T R_{-l_j} y_j\right] \\ &= \sum_{i,j\in[n]} \operatorname{Tr}\left[y_i^T R_{-l_i}^T R_{-l_j} y_j\right] \\ &= \sum_{i,j\in[n]} \operatorname{Tr}\left[\left(y_i y_j^T\right)^T R_{l_i} R_{l_j}^T\right]. \end{aligned}$$

We take

$$X = \begin{bmatrix} R_{l_1} \\ R_{l_2} \\ \vdots \\ R_{l_n} \end{bmatrix} \begin{bmatrix} R_{l_1}^T & R_{l_2}^T & \cdots & R_{l_n}^T \end{bmatrix} \in \mathbb{R}^{nL\times nL}, \tag{109}$$

and can rewrite (108) as

$$\begin{array}{ll} \max & \operatorname{Tr}(CX) \\ \text{s. t.} & X_{ii} = I_{L\times L} \\ & X_{ij} \text{ is a circulant permutation matrix} \\ & X \succeq 0 \\ & \operatorname{rank}(X) \leq L, \end{array} \tag{110}$$

where C is the rank 1 matrix given by

$$C = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} \begin{bmatrix} y_1^T & y_2^T & \cdots & y_n^T \end{bmatrix} \in \mathbb{R}^{nL\times nL}, \tag{111}$$

with blocks $C_{ij} = y_i y_j^T$.

The constraints $X_{ii} = I_{L\times L}$ and $\text{rank}(X) \le L$ imply that $\text{rank}(X) = L$ and $X_{ij} \in O(L)$. Since the only doubly stochastic matrices in $O(L)$ are permutations, (110) can be rewritten as

$$\begin{array}{ll} \max & \text{Tr}(CX) \\ \text{s. t.} & X_{ii} = I_{L\times L} \\ & X_{ij}\mathbf{1} = \mathbf{1} \\ & X_{ij} \text{ is circulant} \\ & X \ge 0 \\ & X \succeq 0 \\ & \text{rank}(X) \le L. \end{array} \tag{112}$$

Removing the nonconvex rank constraint yields a semidefinite program, corresponding to (**??**),

$$\begin{array}{ll} \max & \text{Tr}(CX) \\ \text{s. t.} & X_{ii} = I_{L\times L} \\ & X_{ij}\mathbf{1} = \mathbf{1} \\ & X_{ij} \text{ is circulant} \\ & X \ge 0 \\ & X \succeq 0. \end{array} \tag{113}$$

Numerical simulations (see [BCSZ14, BKS14]) suggest that, below a certain noise level, the semidefinite program (113) is tight with high probability. However, an explanation of this phenomenon remains an open problem [BKS14].

Open Problem 10.3 *For which values of noise do we expect that, with high probability, the semidefinite program* (113) *is tight? In particular, is it true that for any σ by taking arbitrarily large n the SDP is tight with high probability?*

10.3.3 Sample complexity for multireference alignment

Another important question related to this problem is to understand its sample complexity. Since the objective is to recover the underlying signal u, a larger number of observations n should yield a better recovery (considering the model in (**??**)). Another open question is the consistency of the quasi-MLE estimator, it is known that there is some bias on the power spectrum of the recovered signal (that can be easily fixed) but the estimates for phases of the Fourier transform are conjecture to be consistent [BCSZ14].

Open Problem 10.4
1. *Is the quasi-MLE (or the MLE) consistent for the Multireference alignment problem? (after fixing the power spectrum appropriately).*
2. *For a given value of L and σ, how large does n need to be in order to allow for a reasonably accurate recovery in the multireference alignment problem?*

Remark 10.2 *One could design a simpler method based on angular synchronization: for each pair of signals take the best pairwise shift and then use angular synchronization to find the signal shifts from these pairwise measurements. While this would yield a smaller SDP, the fact that it is not* *using all of the information renders it less effective [BCS15]. This illustrates an interesting trade-off between size of the SDP and its effectiveness. There is an interpretation of this through dimensions of representations of the group in question (essentially each of these approaches corresponds to a different representation), we refer the interested reader to [BCS15] for more one that.*

References

[AABS15] E. Abbe, N. Alon, A. S. Bandeira, and C. Sandon. Linear boolean classification, coding and "the critical problem". *Available online at arXiv:1401.6528v3 [cs.IT]*, 2015.

[ABC$^+$15] P. Awasthi, A. S. Bandeira, M. Charikar, R. Krishnaswamy, S. Villar, and R. Ward. Relax, no need to round: integrality of clustering formulations. *6th Innovations in Theoretical Computer Science (ITCS 2015)*, 2015.

[ABFM12] B. Alexeev, A. S. Bandeira, M. Fickus, and D. G. Mixon. Phase retrieval with polarization. *available online*, 2012.

[ABG12] L. Addario-Berry and S. Griffiths. The spectrum of random lifts. *available at arXiv:1012.4097 [math.CO]*, 2012.

[ABH14] E. Abbe, A. S. Bandeira, and G. Hall. Exact recovery in the stochastic block model. *Available online at arXiv:1405.3267 [cs.SI]*, 2014.

[ABKK15] N. Agarwal, A. S. Bandeira, K. Koiliaris, and A. Kolla. Multisection in the stochastic block model using semidefinite programming. *Available online at arXiv:1507.02323 [cs.DS]*, 2015.

[ABS10] S. Arora, B. Barak, and D. Steurer. Subexponential algorithms for unique games related problems. 2010.

[AC09] Nir Ailon and Bernard Chazelle. The fast Johnson-Lindenstrauss transform and approximate nearest neighbors. *SIAM J. Comput*, pages 302–322, 2009.

[AGZ10] G. W. Anderson, A. Guionnet, and O. Zeitouni. *An introduction to random matrices.* Cambridge studies in advanced mathematics. Cambridge University Press, Cambridge, New York, Melbourne, 2010.

[AJP13] M. Agarwal, R. Jaiswal, and A. Pal. k-means++ under approximation stability. *The 10th annual conference on Theory and Applications of Models of Computation*, 2013.

[AL06] N. Alon and E. Lubetzky. The shannon capacity of a graph and the independence numbers of its powers. *IEEE Transactions on Information Theory*, 52:21722176, 2006.

[ALMT14] D. Amelunxen, M. Lotz, M. B. McCoy, and J. A. Tropp. Living on the edge: phase transitions in convex programs with random data. 2014.

[Alo86] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6:83–96, 1986.

[Alo03] N. Alon. Problems and results in extremal combinatorics i. *Discrete Mathematics*, 273(1–3):31–53, 2003.

[AM85] N. Alon and V. Milman. Isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory*, 38:73–88, 1985.

[AMMN05] N. Alon, K. Makarychev, Y. Makarychev, and A. Naor. Quadratic forms on graphs. *Invent. Math*, 163:486–493, 2005.

[AN04] N. Alon and A. Naor. Approximating the cut-norm via Grothendieck's inequality. In *Proc. of the 36 th ACM STOC*, pages 72–80. ACM Press, 2004.

[ARC06] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In A. Leonardis, H. Bischof, and A. Pinz, editors, *Computer Vision – ECCV 2006*, volume 3951 of *Lecture Notes in Computer Science*, pages 578–591. Springer Berlin Heidelberg, 2006.

[AS15] E. Abbe and C. Sandon. Community detection in general stochastic block models: fundamental limits and efficient recovery algorithms. *to appear in FOCS 2015, also available online at arXiv:1503.00609 [math.PR]*, 2015.

[B+11] J. Bourgain et al. Explicit constructions of RIP matrices and related problems. *Duke Mathematical Journal*, 159(1), 2011.

[Bai99] Z. D. Bai. Methodologies in spectral analysis of large dimensional random matrices, a review. *Statistics Sinica*, 9:611–677, 1999.

[Ban15a] A. S. Bandeira. *Convex relaxations for certain inverse problems on graphs*. PhD thesis, Program in Applied and Computational Mathematics, Princeton University, 2015.

[Ban15b] A. S. Bandeira. A note on probably certifiably correct algorithms. *Available at arXiv:1509.00824 [math.OC]*, 2015.

[Ban15c] A. S. Bandeira. Random Laplacian matrices and convex relaxations. *Available online at arXiv:1504.03987 [math.PR]*, 2015.

[Ban15d] A. S. Bandeira. Relax and Conquer BLOG: Ten Lectures and Forty-two Open Problems in Mathematics of Data Science. 2015.

[Bar14] B. Barak. Sum of squares upper bounds, lower bounds, and open questions. *Available online at http://www.boazbarak.org/sos/files/all-notes.pdf*, 2014.

[BBAP05] J. Baik, G. Ben-Arous, and S. Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697, 2005.

[BBC04] N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004.

[BBRV01] S. Bandyopadhyay, P. O. Boykin, V. Roychowdhury, and F. Vatan. A new proof for the existence of mutually unbiased bases. *Available online at arXiv:quant-ph/0103162*, 2001.

[BBS14] A. S. Bandeira, N. Boumal, and A. Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *Available online at arXiv:1411.3272 [math.OC]*, 2014.

[BCS15] A. S. Bandeira, Y. Chen, and A. Singer. Non-unique games over compact groups and orientation estimation in cryo-em. *Available online at arXiv:1505.03840 [cs.CV]*, 2015.

[BCSZ14] A. S. Bandeira, M. Charikar, A. Singer, and A. Zhu. Multireference alignment using semidefinite programming. *5th Innovations in Theoretical Computer Science (ITCS 2014)*, 2014.

[BDMS13] A. S. Bandeira, E. Dobriban, D.G. Mixon, and W.F. Sawin. Certifying the restricted isometry property is hard. *IEEE Trans. Inform. Theory*, 59(6):3448–3450, 2013.

[BFMM14] A. S. Bandeira, M. Fickus, D. G. Mixon, and J. Moreira. Derandomizing restricted isometries via the Legendre symbol. *Available online at arXiv:1406.4089 [math.CO]*, 2014.

[BFMW13] A. S. Bandeira, M. Fickus, D. G. Mixon, and P. Wong. The road to deterministic matrices with the restricted isometry property. *Journal of Fourier Analysis and Applications*, 19(6):1123–1149, 2013.

[BGN11] F. Benaych-Georges and R. R. Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 2011.

[BGN12] F. Benaych-Georges and R. R. Nadakuditi. The singular values and vectors of low rank perturbations of large rectangular random matrices. *Journal of Multivariate Analysis*, 2012.

[BKS13a] A. S. Bandeira, C. Kennedy, and A. Singer. Approximating the little grothendieck problem over the orthogonal group. *Available online at arXiv:1308.5207 [cs.DS]*, 2013.

[BKS13b] B. Barak, J. Kelner, and D. Steurer. Rounding sum-of-squares relaxations. *Available online at arXiv:1312.6652 [cs.DS]*, 2013.

[BKS14] A. S. Bandeira, Y. Khoo, and A. Singer. Open problem: Tightness of maximum likelihood semidefinite relaxations. In *Proceedings of the 27th Conference on Learning Theory*, volume 35 of *JMLR W&CP*, pages 1265–1267, 2014.

[BLM15] A. S. Bandeira, M. E. Lewis, and D. G. Mixon. Discrete uncertainty principles and sparse signal processing. *Available online at arXiv:1504.01014 [cs.IT]*, 2015.

[BMM14] A. S. Bandeira, D. G. Mixon, and J. Moreira. A conditional construction of restricted isometries. *Available online at arXiv:1410.6457 [math.FA]*, 2014.

[Bou14] J. Bourgain. An improved estimate in the restricted isometry problem. *Lect. Notes Math.*, 2116:65–70, 2014.

[BR13] Q. Berthet and P. Rigollet. Complexity theoretic lower bounds for sparse principal component detection. *Conference on Learning Theory (COLT)*, 2013.

[BRM13] C. Bachoc, I. Z. Ruzsa, and M. Matolcsi. Squares and difference sets in finite fields. *Available online at arXiv:1305.0577 [math.CO]*, 2013.

[BS05] J. Baik and J. W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. 2005.

[BS14] B. Barak and D. Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. *Survey, ICM 2014*, 2014.

[BSS13] A. S. Bandeira, A. Singer, and D. A. Spielman. A Cheeger inequality for the graph connection Laplacian. *SIAM J. Matrix Anal. Appl.*, 34(4):1611–1630, 2013.

[BvH15] A. S. Bandeira and R. v. Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Annals of Probability, to appear*, 2015.

[Che70] J. Cheeger. A lower bound for the smallest eigenvalue of the Laplacian. *Problems in analysis (Papers dedicated to Salomon Bochner, 1969), pp. 195–199. Princeton Univ. Press*, 1970.

[Chi15] T.-Y. Chien. *Equiangular lines, projective symmetries and nice error frames.* PhD thesis, 2015.

[Chu97] F. R. K. Chung. *Spectral Graph Theory.* AMS, 1997.

[Chu10] F. Chung. Four proofs for the cheeger inequality and graph partition algorithms. *Fourth International Congress of Chinese Mathematicians, pp. 331–349*, 2010.

[Chu13] M. Chudnovsky. The erdos-hajnal conjecture – a survey. 2013.

[CK12] P. G. Casazza and G. Kutyniok. *Finite Frames: Theory and Applications.* 2012.

[Coh13] J. Cohen. Is high-tech view of HIV too good to be true? *Science*, 341(6145):443–444, 2013.

[Coh15] G. Cohen. Two-source dispersers for polylogarithmic entropy and improved ramsey graphs. *Electronic Colloquium on Computational Complexity*, 2015.

[Con09] David Conlon. A new upper bound for diagonal ramsey numbers. *Annals of Mathematics*, 2009.

[CR09] E.J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[CRPW12] V. Chandrasekaran, B. Recht, P.A. Parrilo, and A.S. Willsky. The convex geometry of linear inverse problems. *Foundations of Computational Mathematics*, 12(6):805–849, 2012.

[CRT06a] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52:489–509, 2006.

[CRT06b] E. J. Candès, J. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Comm. Pure Appl. Math.*, 59:1207–1223, 2006.

[CT] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley-Interscience.

[CT05] E. J. Candès and T. Tao. Decoding by linear programming. *IEEE Trans. Inform. Theory*, 51:4203–4215, 2005.

[CT06] E. J. Candès and T. Tao. Near optimal signal recovery from random projections: universal encoding strategies? *IEEE Trans. Inform. Theory*, 52:5406–5425, 2006.

[CT10] E. J. Candes and T. Tao. The power of convex relaxation: Near-optimal matrix completion. *Information Theory, IEEE Transactions on*, 56(5):2053–2080, May 2010.

[CW04] M. Charikar and A. Wirth. Maximizing quadratic programs: Extending grothendieck's inequality. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '04, pages 54–60, Washington, DC, USA, 2004. IEEE Computer Society.

[CZ15] E. Chattopadhyay and D. Zuckerman. Explicit two-source extractors and resilient functions. *Electronic Colloquium on Computational Complexity*, 2015.

[DG02] S. Dasgupta and A. Gupta. An elementary proof of the johnson-lindenstrauss lemma. Technical report, 2002.

[DKMZ11] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys. Rev. E*, 84, December 2011.

[DM13] Y. Deshpande and A. Montanari. Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time. *Available online at arXiv:1304.7047 [math.PR]*, 2013.

[DMS15] A. Dembo, A. Montanari, and S. Sen. Extremal cuts of sparse random graphs. *Available online at arXiv:1503.03923 [math.PR]*, 2015.

[Don06] D. L. Donoho. Compressed sensing. *IEEE Trans. Inform. Theory*, 52:1289–1306, 2006.

[Dor43] R. Dorfman. The detection of defective members of large populations. 1943.

[Duc12] J. C. Duchi. Commentary on "towards a noncommutative arithmetic-geometric mean inequality" by b. recht and c. re. 2012.

[Dur06] R. Durrett. *Random Graph Dynamics (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press, New York, NY, USA, 2006.

[DVPS14] A. G. D'yachkov, I. V. Vorob'ev, N. A. Polyansky, and V. Y. Shchukin. Bounds on the rate of disjunctive codes. *Problems of Information Transmission*, 2014.

[EH89] P. Erdos and A. Hajnal. Ramsey-type theorems. *Discrete Applied Mathematics*, 25, 1989.

[F$^+$14] Y. Filmus et al. Real analysis in computer science: A collection of open problems. *Available online at http://simons.berkeley.edu/sites/default/files/openprobsmerged.pdf*, 2014.

[Fei05] U. Feige. On sums of independent random variables with unbounded variance, and estimating the average degree in a graph. 2005.

[FP06] D. Féral and S. Péché. The largest eigenvalue of rank one deformation of large wigner matrices. *Communications in Mathematical Physics*, 272(1):185–228, 2006.

[FR13] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Birkhauser, 2013.

[Fuc04] J. J. Fuchs. On sparse representations in arbitrary redundant bases. *Information Theory, IEEE Transactions on*, 50(6):1341–1344, 2004.

[Fur96] Z. Furedia. On r-cover-free families. *Journal of Combinatorial Theory, Series A*, 1996.

[Gil52] E. N. Gilbert. A comparison of signalling alphabets. *Bell System Technical Journal*, 31:504–522, 1952.

[GK06] A. Giridhar and P.R. Kumar. Distributed clock synchronization over wireless networks: Algorithms and analysis. In *Decision and Control, 2006 45th IEEE Conference on*, pages 4915–4920. IEEE, 2006.

[GLV07] N. Gvozdenovic, M. Laurent, and F. Vallentin. Block-diagonal semidefinite programming hierarchies for 0/1 programming. *Available online at arXiv:0712.3079 [math.OC]*, 2007.

[Gol96] G. H. Golub. *Matrix Computations*. Johns Hopkins University Press, third edition, 1996.

[Gor85] Y. Gordon. Some inequalities for gaussian processes and applications. *Israel J. Math*, 50:109–110, 1985.

[Gor88] Y. Gordon. On milnan's inequality and random subspaces which escape through a mesh in $\mathbb{R}^n$. 1988.

[GRS15] V. Guruswami, A. Rudra, and M. Sudan. *Essential Coding Theory*. Available at: http://www.cse.buffalo.edu/faculty/atri/courses/coding-theory/book/, 2015.

[GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefine programming. *Journal of the Association for Computing Machinery*, 42:1115–1145, 1995.

[GZC+15] Amir Ghasemian, Pan Zhang, Aaron Clauset, Cristopher Moore, and Leto Peel. Detectability thresholds and optimal algorithms for community structure in dynamic networks. *Available online at arXiv:1506.06179 [stat.ML]*, 2015.

[Haa87] U. Haagerup. A new upper bound for the complex Grothendieck constant. *Israel Journal of Mathematics*, 60(2):199–224, 1987.

[Has02] J. Hastad. Some optimal inapproximability results. 2002.

[Hen13] R. Henderson. Avoiding the pitfalls of single particle cryo-electron microscopy: Einstein from noise. *Proceedings of the National Academy of Sciences*, 110(45):18037–18041, 2013.

[HJ85] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

[HMPW] T. Holenstein, T. Mitzenmacher, R. Panigrahy, and U. Wieder. Trace reconstruction with constant deletion probability and related results. *In Proceedings of the Nineteenth Annual ACM-SIAM*.

[HMT09] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. *Available online at arXiv:0909.4061v2 [math.NA]*, 2009.

[HR] I. Haviv and O. Regev. The restricted isometry property of subsampled fourier matrices. *SODA 2016*.

[HWX14] B. Hajek, Y. Wu, and J. Xu. Achieving exact cluster recovery threshold via semidefinite programming. *Available online at arXiv:1412.6156*, 2014.

[HWX15] B. Hajek, Y. Wu, and J. Xu. Achieving exact cluster recovery threshold via semidefinite programming: Extensions. *Available online at arXiv:1502.07738*, 2015.

[IKW14] A. Israel, F. Krahmer, and R. Ward. An arithmetic-geometric mean inequality for products of three matrices. *Available online at arXiv:1411.0333 [math.SP]*, 2014.

[IMPV15a] T. Iguchi, D. G. Mixon, J. Peterson, and S. Villar. On the tightness of an sdp relaxation of k-means. *Available online at arXiv:1505.04778 [cs.IT]*, 2015.

[IMPV15b] T. Iguchi, D. G. Mixon, J. Peterson, and S. Villar. Probably certifiably correct k-means clustering. *Available at arXiv*, 2015.

[JL84] W. Johnson and J. Lindenstrauss. Extensions of Lipschitz mappings into a Hilbert space. In *Conference in modern analysis and probability (New Haven, Conn., 1982)*, volume 26 of *Contemporary Mathematics*, pages 189–206. American Mathematical Society, 1984.

[Joh01] I. M. Johnston. On the distribution of the largest eigenvalue in principal components analysis. *The Annals of Statistics*, 29(2):295–327, 2001.

[Kar05] N. E. Karoui. Recent results about the largest eigenvalue of random covariance matrices and statistical application. *Acta Physica Polonica B*, 36(9), 2005.

[Kho02] S. Khot. On the power of unique 2-prover 1-round games. *Thiry-fourth annual ACM symposium on Theory of computing*, 2002.

[Kho10] S. Khot. On the unique games conjecture (invited survey). In *Proceedings of the 2010 IEEE 25th Annual Conference on Computational Complexity*, CCC '10, pages 99–121, Washington, DC, USA, 2010. IEEE Computer Society.

[KKMO05] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? 2005.

[KV13] S. A. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into l1. *Available online at arXiv:1305.4581 [cs.CC]*, 2013.

[KW92] J. Kuczynski and H. Wozniakowski. Estimating the largest eigenvalue by the power and lanczos algorithms with a random start. *SIAM Journal on Matrix Analysis and Applications*, 13(4):1094–1122, 1992.

[Las01] J. B. Lassere. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001.

[Lat05] R. Latała. Some estimates of norms of random matrices. *Proc. Amer. Math. Soc.*, 133(5):1273–1282 (electronic), 2005.

[LGT12] J.R. Lee, S.O. Gharan, and L. Trevisan. Multi-way spectral partitioning and higher–order cheeger inequalities. *STOC '12 Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, 2012.

[Llo82] S. Lloyd. Least squares quantization in pcm. *IEEE Trans. Inf. Theor.*, 28(2):129–137, 1982.

[LM00] B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.*, 2000.

[Lov79] L. Lovasz. On the shannon capacity of a graph. *IEEE Trans. Inf. Theor.*, 25(1):1–7, 1979.

[LRTV12] A. Louis, P. Raghavendra, P. Tetali, and S. Vempala. Many sparse cuts via higher eigenvalues. *STOC*, 2012.

[Lyo14] R. Lyons. Factors of IID on trees. *Combin. Probab. Comput.*, 2014.

[Mas00] P. Massart. About the constants in Talagrand's concentration inequalities for empirical processes. *The Annals of Probability*, 28(2), 2000.

[Mas14] L. Massoulié. Community detection thresholds and the weak ramanujan property. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 694–703, New York, NY, USA, 2014. ACM.

[Mek14] R. Meka. Windows on Theory BLOG: Discrepancy and Beating the Union Bound. *http://windowsontheory.org/2014/02/07/discrepancy-and-beating-the-union-bound/*, 2014.

[Mit09] M. Mitzenmacher. A survey of results for deletion channels and related synchronization channels. *Probability Surveys*, 2009.

[Mix14a] D. G. Mixon. Explicit matrices with the restricted isometry property: Breaking the square-root bottleneck. *available online at arXiv:1403.3427 [math.FA]*, 2014.

[Mix14b] D. G. Mixon. Short, Fat matrices BLOG: Gordon's escape through a mesh theorem. 2014.

[Mix14c] D. G. Mixon. Short, Fat matrices BLOG: Gordon's escape through a mesh theorem. 2014.

[Mix15] D. G. Mixon. Applied harmonic analysis and sparse approximation. *Short, Fat Matrices Web blog*, 2015.

[MM15] C. Musco and C. Musco. Stronger and faster approximate singular value decomposition via the block lanczos method. *Available at arXiv:1504.05477 [cs.DS]*, 2015.

[MNS14a] E. Mossel, J. Neeman, and A. Sly. A proof of the block model threshold conjecture. *Available online at arXiv:1311.4115 [math.PR]*, January 2014.

[MNS14b] E. Mossel, J. Neeman, and A. Sly. Stochastic block models and reconstruction. *Probability Theory and Related Fields (to appear)*, 2014.

[Mon14] A. Montanari. Principal component analysis with nonnegativity constraints. *http://sublinear.info/index.php?title=Open_Problems:62*, 2014.

[Mos11] M. S. Moslehian. Ky Fan inequalities. *Available online at arXiv:1108.1467 [math.FA]*, 2011.

[MP67] V. A. Marchenko and L. A. Pastur. Distribution of eigenvalues in certain sets of random matrices. *Mat. Sb. (N.S.)*, 72(114):507–536, 1967.

[MR14] A. Montanari and E. Richard. Non-negative principal component analysis: Message passing algorithms and sharp asymptotics. *Available online at arXiv:1406.4775v1 [cs.IT]*, 2014.

[MS15] A. Montanari and S. Sen. Semidefinite programs on sparse random graphs. *Available online at arXiv:1504.05910 [cs.DM]*, 2015.

[MSS15a] A. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families i: Bipartite ramanujan graphs of all degrees. *Annals of Mathematics*, 2015.

[MSS15b] A. Marcus, D. A. Spielman, and N. Srivastava. Interlacing families ii: Mixed characteristic polynomials and the kadison-singer problem. *Annals of Mathematics*, 2015.

[MZ11] S. Mallat and O. Zeitouni. A conjecture concerning optimality of the karhunen-loeve basis in nonlinear reconstruction. *Available online at arXiv:1109.0489 [math.PR]*, 2011.

[Nel] J. Nelson. Johnson-lindenstrauss notes. *http://web.mit.edu/minilek/www/jl_notes.pdf*.

[Nes00] Y. Nesterov. Squared functional systems and optimization problems. *High performance optimization*, 13(405-440), 2000.

[Nik13] A. Nikolov. The komlos conjecture holds for vector colorings. *Available online at arXiv:1301.4039 [math.CO]*, 2013.

[NN] J. Nelson and L. Nguyen. Osnap: Faster numerical linear algebra algorithms via sparser subspace embeddings. *Available at arXiv:1211.1002 [cs.DS]*.

[NPW14] J. Nelson, E. Price, and M. Wootters. New constructions of RIP matrices with fast multiplication and fewer rows. *SODA*, pages 1515–1528, 2014.

[NSZ09] B. Nadler, N. Srebro, and X. Zhou. Semi-supervised learning with the graph laplacian: The limit of infinite unlabelled data. 2009.

[NW13] A. Nellore and R. Ward. Recovery guarantees for exemplar-based clustering. *Available online at arXiv:1309.3256v2 [stat.ML]*, 2013.

[Oli10] R. I. Oliveira. The spectrum of random k-lifts of large graphs (with possibly large k). *Journal of Combinatorics*, 2010.

[Par00] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, 2000.

[Pau] D. Paul. Asymptotics of the leading sample eigenvalues for a spiked covariance model. *Available online at http://anson.ucdavis.edu/~debashis/techrep/eigenlimit.pdf*.

[Pau07] D. Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistics Sinica*, 17:1617–1642, 2007.

[Pea01] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine, Series 6*, 2(11):559–572, 1901.

[Pis03] G. Pisier. *Introduction to operator space theory*, volume 294 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 2003.

[Pis11] G. Pisier. Grothendieck's theorem, past and present. *Bull. Amer. Math. Soc.*, 49:237–323, 2011.

[PW15] W. Perry and A. S. Wein. A semidefinite program for unbalanced multisection in the stochastic block model. *Available online at arXiv:1507.05605 [cs.DS]*, 2015.

[QSW14] Q. Qu, J. Sun, and J. Wright. Finding a sparse vector in a subspace: Linear sparsity using alternating directions. *Available online at arXiv:1412.4659v1 [cs.IT]*, 2014.

[Rag08] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 245–254. ACM, 2008.

[Ram28] F. P. Ramsey. On a problem of formal logic. 1928.

[Rec11] B. Recht. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12:3413–3430, 2011.

[RR12] B. Recht and C. Re. Beneath the valley of the noncommutative arithmetic-geometric mean inequality: conjectures, case-studies, and consequences. *Conference on Learning Theory (COLT)*, 2012.

[RS60] I. S. Reed and G. Solomon. Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics (SIAM)*, 8(2):300–304, 1960.

[RS10] P. Raghavendra and D. Steurer. Graph expansion and the unique games conjecture. *STOC*, 2010.

[RS13] S. Riemer and C. Schütt. On the expectation of the norm of random matrices with non-identically distributed entries. *Electron. J. Probab.*, 18, 2013.

[RST09] V. Rokhlin, A. Szlam, and M. Tygert. A randomized algorithm for principal component analysis. *Available at arXiv:0809.2274 [stat.CO]*, 2009.

[RST12] P. Raghavendra, D. Steurer, and M. Tulsiani. Reductions between expansion problems. *IEEE CCC*, 2012.

[RV08] M. Rudelson and R. Vershynin. On sparse reconstruction from Fourier and Gaussian measurements. *Comm. Pure Appl. Math.*, 61:1025–1045, 2008.

[RW01] J. Rubinstein and G. Wolansky. Reconstruction of optical surfaces from ray data. *Optical Review*, 8(4):281–283, 2001.

[Sam66] S. M. Samuels. On a chebyshev-type inequality for sums of independent random variables. *Ann. Math. Statist.*, 1966.

[Sam68] S. M. Samuels. More on a chebyshev-type inequality. 1968.

[Sam69] S. M. Samuels. The markov inequality for sums of independent random variables. *Ann. Math. Statist.*, 1969.

[Sch12] K. Schmudgen. Around hilbert's 17th problem. *Documenta Mathematica - Extra Volume ISMP*, pages 433–438, 2012.

[Seg00] Y. Seginer. The expected norm of random matrices. *Combin. Probab. Comput.*, 9(2):149–166, 2000.

[SG10] A. J. Scott and M. Grassl. Sic-povms: A new computer study. *J. Math. Phys.*, 2010.

[Sha56] C. E. Shannon. The zero-error capacity of a noisy channel. *IRE Transactions on Information Theory*, 2, 1956.

[SHBG09] M. Shatsky, R. J. Hall, S. E. Brenner, and R. M. Glaeser. A method for the alignment of heterogeneous macromolecules from electron microscopy. *Journal of Structural Biology*, 166(1), 2009.

[Sho87] N. Shor. An approach to obtaining global extremums in polynomial mathematical programming problems. *Cybernetics and Systems Analysis*, 23(5):695–700, 1987.

[Sin11] A. Singer. Angular synchronization by eigenvectors and semidefinite programming. *Appl. Comput. Harmon. Anal.*, 30(1):20 – 36, 2011.

[Spe75] J. Spencer. Ramsey's theorem – a new lower bound. *J. Combin. Theory Ser. A*, 1975.

[Spe85] J. Spencer. Six standard deviations suffice. *Trans. Amer. Math. Soc.*, (289), 1985.

[Spe94] J. Spencer. *Ten Lectures on the Probabilistic Method: Second Edition.* SIAM, 1994.

[SS11] A. Singer and Y. Shkolnisky. Three-dimensional structure determination from common lines in Cryo-EM by eigenvectors and semidefinite programming. *SIAM J. Imaging Sciences*, 4(2):543–572, 2011.

[Ste74] G. Stengle. A nullstellensatz and a positivstellensatz in semialgebraic geometry. *Math. Ann. 207*, 207:87–97, 1974.

[SW11] A. Singer and H.-T. Wu. Orientability and diffusion maps. *Appl. Comput. Harmon. Anal.*, 31(1):44–58, 2011.

[SWW12] D. A Spielman, H. Wang, and J. Wright. Exact recovery of sparsely-used dictionaries. *COLT*, 2012.

[Tal95] M. Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Inst. Hautes Etudes Sci. Publ. Math.*, (81):73–205, 1995.

[Tao07] T. Tao. What's new blog: Open question: deterministic UUP matrices. 2007.

[Tao12] T. Tao. *Topics in Random Matrix Theory.* Graduate studies in mathematics. American Mathematical Soc., 2012.

[TdSL00] J. B. Tenenbaum, V. de Silva, and J. C. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[TP13] A. M. Tillmann and M. E. Pfefsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. 2013.

[Tre11] L. Trevisan. in theory BLOG: CS369G Llecture 4: Spectral Partitionaing. 2011.

[Tro05] J. A. Tropp. Recovery of short, complex linear combinations via ℓ_1 minimization. *IEEE Transactions on Information Theory*, 4:1568–1570, 2005.

[Tro12] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012.

[Tro15a] J. A. Tropp. The expected norm of a sum of independent random matrices: An elementary approach. *Available at arXiv:1506.04711 [math.PR]*, 2015.

[Tro15b] J. A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 2015.

[Tro15c] J. A. Tropp. Second-order matrix concentration inequalities. *In preparation*, 2015.

[Var57] R. R. Varshamov. Estimate of the number of signals in error correcting codes. *Dokl. Acad. Nauk SSSR*, 117:739–741, 1957.

[VB96] L. Vanderberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38:49–95, 1996.

[VB04] L. Vanderberghe and S. Boyd. *Convex Optimization*. Cambridge University Press, 2004.

[vH14] R. van Handel. Probability in high dimensions. *ORF 570 Lecture Notes, Princeton University*, 2014.

[vH15] R. van Handel. On the spectral norm of inhomogeneous random matrices. *Available online at arXiv:1502.05003 [math.PR]*, 2015.

[Yam54] K. Yamamoto. Logarithmic order of free distributive lattice. *Journal of the Mathematical Society of Japan*, 6:343–353, 1954.

[ZB09] L. Zdeborova and S. Boettcher. Conjecture on the maximum cut and bisection width in random regular graphs. *Available online at arXiv:0912.4861 [cond-mat.dis-nn]*, 2009.

[Zha14] T. Zhang. A note on the non-commutative arithmetic-geometric mean inequality. *Available online at arXiv:1411.5058 [math.SP]*, 2014.

[ZMZ14] Pan Zhang, Cristopher Moore, and Lenka Zdeborova. Phase transitions in semisupervised clustering of sparse networks. *Phys. Rev. E*, 90, 2014.