

Statistical Estimation in the Presence of Group Actions

by

Alexander Spence Wein

Submitted to the Department of Mathematics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018

© Alexander Spence Wein, MMXVIII. All rights reserved.

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Author
Department of Mathematics
April 27, 2018

Certified by
Ankur Moitra
Associate Professor of Mathematics
Thesis Supervisor

Accepted by
Jonathan A. Kelner
Chairman, Applied Mathematics Committee

Statistical Estimation in the Presence of Group Actions

by

Alexander Spence Wein

Submitted to the Department of Mathematics
on April 27, 2018, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Imagine we want to recover an unknown vector given many noisy copies of it, except each copy is cyclically shifted by an unknown offset (this is “multi-reference alignment”). Or imagine we want to reconstruct an unknown 3D structure (e.g. a molecule) given many noisy pictures of it taken from different unknown angles (this is “cryo-EM”). These problems (and many others) involve the action of unknown group elements drawn randomly from a compact group such as \mathbb{Z}/p or $\text{SO}(3)$.

In this thesis we study two statistical models for estimation in the presence of group actions. The first is the *synchronization* model in which we attempt to learn an unknown collection of group elements based on noisy pairwise comparisons. The second is the *orbit recovery* model in which we observe noisy copies of a hidden signal, each of which is acted upon by a random group element. For both of these models, we explore the fundamental statistical limits as well as the fundamental computational limits (i.e. how well can a polynomial-time algorithm perform?). We use methods from a wide variety of areas, including statistical physics, approximate message passing, representation theory, contiguity and the associated second moment method, invariant theory, algebraic geometry, and the sum-of-squares hierarchy.

Thesis Supervisor: Ankur Moitra

Title: Associate Professor of Mathematics

In Memoriam

Amelia Perry

1991 – 2018



Amelia was a fellow student with whom I worked very closely throughout graduate school. Nearly everything I worked on during my PhD was in collaboration with her (including basically everything in this thesis), and I could not have done it without her. Amelia was truly exceptional, both as a mathematician and as a friend. I will remember her amazing ability to work out difficult problems on the fly, her passion and excitement for talking about math, and her sense of humor. This picture shows her in our office, where she would write all over our desks in chalk. She passed away during our final year of grad school, following a long struggle with anxiety and depression.

Acknowledgments

Thank you to the many great mentors I was fortunate to have throughout graduate school, including my advisor Ankur Moitra along with Afonso Bandeira, Michel Goemans, and Philippe Rigollet. Thank you to all of my collaborators for everything they taught me. Thank you to my fellow students (both in the math department and the CSAIL theory group) for all the fun times during grad school. Finally, thank you to my family and to Melissa (my fiancée) for their support.

Contents

1	Introduction	13
1.1	Motivation	13
1.2	Models	15
1.2.1	Synchronization	15
1.2.2	Orbit recovery	16
1.3	Summary of contributions	17
1.3.1	Synchronization: analysis via statistical physics	18
1.3.2	Synchronization: contiguity and rigorous bounds	19
1.3.3	Orbit recovery: statistical limits	19
1.3.4	Orbit recovery: computational limits	20
1.4	Preliminaries on groups and representations	20
2	Synchronization: analysis via statistical physics	23
2.1	Introduction	23
2.2	Intuition: iterative methods for $\mathbb{Z}/2$ and $U(1)$ synchronization	31
2.2.1	$\mathbb{Z}/2$ synchronization	31
2.2.2	Belief propagation and approximate message passing	34
2.2.3	AMP for Gaussian $U(1)$ synchronization with one frequency	36
2.2.4	AMP for Gaussian $U(1)$ synchronization with multiple frequencies	38
2.3	AMP over general compact groups	43
2.3.1	Representation theory preliminaries	44

2.3.2	Graphical model formulation	46
2.3.3	AMP algorithm	48
2.3.4	Gaussian observation model	50
2.3.5	Representation theory of some common examples	51
2.4	Experimental results	52
2.5	Derivation of AMP from belief propagation	56
2.5.1	Onsager correction	58
2.6	MMSE derivation and state evolution	61
2.6.1	MMSE estimator	61
2.6.2	AMP update step	62
2.6.3	State evolution	63
2.6.4	Simplified AMP update step	65
2.6.5	Reduction to single parameter (per frequency)	65
2.6.6	Threshold at $\lambda = 1$	67
2.7	Correctness of state evolution?	68
2.7.1	Rigorous work on state evolution	68
2.7.2	Experiments on state evolution	69
2.8	Statistical-to-computational gaps	69
2.8.1	Free energy	70
2.8.2	Examples	71
3	Synchronization: contiguity and rigorous bounds	77
3.1	Introduction	77
3.2	Contiguity and the second moment method	80
3.3	The truth-or-Haar model	82
3.3.1	Main results	82
3.3.2	Second moment computation	84
3.3.3	The conditioning method	85
3.3.4	Upper bound via exhaustive search	89

3.4	The Gaussian synchronization model	90
3.4.1	The model	90
3.4.2	Second moment computation	94
3.4.3	The sub-Gaussian method	95
3.4.4	Applications of the sub-Gaussian method	98
3.4.5	The conditioning method for finite groups	100
3.4.6	Upper bound via exhaustive search	104
4	Orbit recovery: statistical limits	109
4.1	Introduction	109
4.1.1	Prior work	110
4.1.2	Our contributions	113
4.1.3	Motivating examples	114
4.1.4	Problem statement	115
4.1.5	Method of moments	117
4.1.6	Extensions: projection and heterogeneity	120
4.1.7	Outline of remainder of chapter	121
4.2	General problem statement	121
4.3	Algebraic results	124
4.3.1	Invariant theory basics	124
4.3.2	Generic list recovery	126
4.3.3	Generic unique recovery	132
4.3.4	Worst-case unique recovery	132
4.3.5	Worst-case list recovery	133
4.4	Examples	134
4.4.1	Learning a bag of numbers	134
4.4.2	Learning a rigid body	135
4.4.3	Multi-reference alignment (MRA)	135
4.4.4	S^2 registration	139

4.4.5	Cryo-EM	142
4.5	Open questions	145
4.6	Proofs for Section 4.3: algebraic results	146
4.6.1	Algorithm for generators of $\mathbb{R}[\mathbf{x}]^G$	146
4.6.2	Bounding the list size for generic signals	147
4.6.3	Generic list recovery converse	148
4.6.4	Hilbert series and Hironaka decomposition	149
4.6.5	Transcendence degree for heterogeneity	150
4.6.6	Gröbner bases	151
4.7	Proofs for S^2 registration	155
4.7.1	Formula for Hilbert series of $\mathbb{R}[\mathbf{x}]^G$	155
4.7.2	Formula for dimension of $\mathbb{R}[\mathbf{x}]_d^G$	156
5	Orbit recovery: computational limits	159
5.1	Introduction	159
5.2	Prior work: MRA	160
5.3	Conjectures based on tensor completion	162
5.4	Heterogeneous MRA	163
5.4.1	Preliminaries	164
5.4.2	Basic facts	166
5.4.3	Main result	167
5.4.4	Spectral bounds	170
5.4.5	SoS proof	175
A	Additional proofs for Chapter 2	179
A.1	Log-likelihood expansion for the Gaussian observation model	179
A.2	Proof of Lemma 2.6.1	181
B	Additional proofs for Chapter 3	183
B.1	Proof of Proposition 3.3.5	183

C	Additional proofs for Chapter 4	189
C.1	Spherical harmonics and $\text{SO}(3)$ invariants	189
C.1.1	Spherical harmonics	189
C.1.2	Wigner D-matrices	190
C.1.3	Moment tensor	191
C.1.4	Projection	192
C.1.5	Explicit construction of invariants	193
C.1.6	Counting the number of invariants	195

Chapter 1

Introduction

1.1 Motivation

Many computational problems throughout the sciences exhibit rich symmetry and geometry, especially in fields such as signal and image processing, computer vision, and microscopy. One example is *multi-reference alignment* (MRA) [17, 13, 118], a problem from signal processing [155, 123] with further relevance to structural biology [60, 144]. In this problem, there is an unknown signal $\theta \in \mathbb{R}^p$ and we observe many noisy copies of it, each with its coordinates cyclically shifted by a random unknown offset. More formally, we observe samples of the form $y_i = R_{\ell_i} \theta + \xi_i$ where ξ_i is noise, $\ell_i \in \{0, 1, \dots, p-1\}$ is a random offset, and R_ℓ is the cyclic shift operator $(R_\ell \theta)_i = \theta_{i-\ell}$ where the subscript $i - \ell$ is taken modulo p . Thus we have a statistical estimation problem involving a group action: namely, the cyclic group \mathbb{Z}/p acts on \mathbb{R}^p via cyclic shift.

Another problem involving a group action is the reconstruction problem in *cryo-electron microscopy* (cryo-EM) [5, 142, 115], an imaging technique in structural biology that was recently awarded the 2017 Nobel Prize in Chemistry. This is a technique for determining the 3-dimensional structure of a large molecule, such as a protein. The idea is to freeze many copies of the molecule and take a 2-dimensional image (tomographic projection) of each one. In each image, the molecule is rotated 3-dimensionally to a random unknown orientation.

To make matters worse, each image is extremely noisy. The core computational challenge in cryo-EM is to take this data and build a denoised 3-dimensional model of the molecule. To mathematically abstract this problem, we take our unknown signal to be the density θ of the molecule, considered as a function $\mathbb{R}^3 \rightarrow \mathbb{R}$. We have access to observations of the following form: our microscopy sample contains many rotated copies $R_i\theta$ of the molecule, where $R_i \in \text{SO}(3)$ are random unknown 3D rotations, and we observe the noisy projections $\Pi(R_i\theta) + \xi_i$, where Π denotes tomographic projection (in a fixed direction) and ξ_i is noise. Thus we have a statistical estimation problem involving a group action by the 3-dimensional rotation group $\text{SO}(3)$.

Computational problems involving group actions arise in many other settings, including community detection [56], time synchronization in networks [71], sensor network localization [51], simultaneous localization and mapping in robotics [130], surface reconstruction in computer vision [6], phase alignment in signal processing [17], and many other areas (see e.g. [14]). These problems exhibit a range of group structure, including rotation groups, Euclidean groups, and cyclic groups.

While various methods have been proposed to solve these types of problems in practice, they often lack provable guarantees or strong theoretical justifications. The aim of this thesis is to build a theoretical foundation for statistical estimation in the presence of group actions. In particular, this includes (i) defining statistical models that capture the core difficulties of such problems, (ii) determining the fundamental statistical limits of these models, and (iii) finding efficient (polynomial-time) algorithms that achieve these limits. In some cases we will see that (iii) is likely impossible due to inherent *statistical-to-computational tradeoffs*; in such cases we aim to understand the fundamental limits of efficient algorithms. Our results are in high generality as we often work over an arbitrary compact group and allow a wide variety of observation models.

1.2 Models

In this thesis, we define and study two different statistical models: the *Gaussian synchronization model* and the *orbit recovery model*.

1.2.1 Synchronization

In the context of group actions, the *synchronization approach* is to estimate the unknown group elements, e.g. the rotation of the molecule in each image. In a *synchronization problem* (see e.g. [14]), there is an unknown vector of group elements (g_1, \dots, g_n) and for each pair $i < j$, we are given a noisy measurement of the relative group element $g_i g_j^{-1}$. The goal is to use this weak pairwise information to recover all the group elements g_i (up to a global right-multiplication by some group element, since we can only hope to recover the *relative* group elements).

For example, in cryo-EM we have an unknown rotation $g_i \in \text{SO}(3)$ for each image. One can compare two images to obtain weak information about their relative angle $g_i g_j^{-1}$ [147, 149, 142]. (A more precise explanation is as follows. By the Fourier projection-slice theorem, the Fourier transforms of the tomographic projections are 2D slices of the Fourier transform of the molecule density. Given a hypothesis as to the angles of two projections (slices), we can predict a 1D line of intersection along which those slices should agree. By measuring correlation along that common line, we obtain some weak information by which to confirm or refute our hypothesized angles. Indeed, this test only depends on the relative angle of the slices, thus providing weak information about the value of $g_i g_j^{-1}$.) We can then use a synchronization algorithm to recover the g_i using this information. Once the g_i are known, it is straightforward to reconstruct the molecule.

Synchronization problems have been studied previously and various methods for solving them have been proposed, including spectral methods [141, 142] and semidefinite programming [141, 142, 17, 18, 14, 29]. We define the first statistical observation model for a large class of synchronization problems, allowing us to investigate the fundamental statistical limits of these problems.

For intuition, consider the following extremely simple Gaussian model for synchronization over the group $\mathbb{Z}/2 = \{\pm 1\}$, which was introduced by [56] as a simplification of community detection. One observes the $n \times n$ matrix

$$Y = \frac{\lambda}{n} g g^\top + \frac{1}{\sqrt{n}} W, \quad (1.1)$$

where $g \in \{\pm 1\}^n$ is the signal to be recovered, W is a GOE matrix¹, and $\lambda > 0$ is a signal-to-noise parameter. Each entry Y_{ij} represents a noisy measurement of the pairwise relative alignment $g_i g_j$. (Note that in $\mathbb{Z}/2$, $g_j = g_j^{-1}$.)

In Chapters 2 and 3 we present and analyze our statistical model for synchronization, which generalizes the above model to all compact groups. For each pair of group elements, we observe $g_i g_j^{-1}$ plus Gaussian noise. Note that in general, it is not obvious how to add Gaussian noise to a group element; our solution, which we believe is the most natural one, uses *representation theory* to represent group elements as matrices. Our model captures a wide variety of observation models, allowing for different signal strengths on different *frequency channels* (which correspond to irreducible representations of the group).

1.2.2 Orbit recovery

For some applications, such as cryo-EM and MRA, the synchronization model has some shortcomings. For instance, in cryo-EM there is independent noise on each image (group element), whereas our Gaussian synchronization model has independent noise on each pairwise comparison of group elements. An even more problematic flaw is that when the noise level is large, no consistent estimation of the group elements g_i is possible [7]. This is because even if we knew the true molecule structure, each individual image is too noisy for us to be able to reliably determine the associated rotation. It is the high-noise regime that is practically relevant for many applications, including cryo-EM, where the presence of large noise is a primary obstruction to current techniques [140]. Thus, we should not aim to estimate the

¹Gaussian orthogonal ensemble: a random symmetric matrix with off-diagonal entries $\mathcal{N}(0, 1)$, diagonal entries $\mathcal{N}(0, 2)$, and all entries independent (up to symmetry).

group elements but instead to directly estimate the signal of interest (the molecule). This idea originates from [13] in the context of MRA.

The considerations above motivate studying the following *orbit recovery* model which more directly captures problems like cryo-EM. Fix a compact group G acting (by orthogonal transformations) on a finite-dimensional vector space $V = \mathbb{R}^p$. Let $\theta \in V$ be the unknown signal. We receive noisy measurements of its orbit as follows: for $i = 1, \dots, n$ we observe a sample of the form

$$y_i = g_i \cdot \theta + \xi_i$$

where g_i is drawn randomly from G (according to *Haar measure*, the “uniform distribution” on the group) and $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$ is noise. Since e.g. we cannot hope to distinguish between different rotations of the same molecule, the goal is to recover the *orbit* $\{g \cdot \theta : g \in G\}$ of θ , i.e. to output a vector in (or close to) the orbit.

This model is a straightforward generalization of a popular model for MRA (which, recall, is the special case where G is the cyclic group \mathbb{Z}/p acting via cyclic shifts) [17, 13, 118, 35].

The above model, already a rich object of study, neglects the tomographic projection in cryo-EM; we will also study a generalization of the problem which allows such a projection. We will also consider the additional extension of *heterogeneity* [83, 93, 94, 35], where mixtures of signals are allowed: we have K signals $\theta_1, \dots, \theta_K$, and each sample $y_i = g_i \cdot \theta_{k_i} + \xi_i$ comes from a random choice $1 \leq k_i \leq K$ of which signal is observed. This extension is of paramount importance for cryo-EM in practice, since the laboratory samples often contain one protein in multiple conformations, and understanding the range of conformations is key to understanding the function of the protein.

1.3 Summary of contributions

In this section we summarize the main results in each chapter of this thesis.

1.3.1 Synchronization: analysis via statistical physics

In Chapter 2 we define our Gaussian synchronization model and analyze its statistical and computational limits using methods from statistical physics, including the approximate message passing framework.

It is well known that there are deep connections between Bayesian inference and statistical physics (see e.g. [154]). The core connection is that in inference problems, the posterior distribution (of the unknown signal given the observed data) often follows a Gibbs (or Boltzmann) distribution and thus behaves similarly to a disordered physical system (such as a magnet or a *spin glass*). Various tools from statistical physics, such as the replica and cavity methods, can be applied to Bayesian inference problems. These often come in the form of non-rigorous heuristics that give extremely precise predictions about the behavior of the system. Similar to physical systems, Bayesian inference problems often exhibit *phase transitions*, i.e. sharp boundaries in parameter space that separate regions in which the problem is computationally easy, computationally hard, or statistically impossible.

Ideas from statistical physics have inspired a powerful framework for algorithm design, known as *approximate message passing* (AMP). The first AMP algorithm was proposed by [63] and later rigorously analyzed by [23, 81]. Since then, AMP algorithms have been derived in numerous settings and have often been shown to achieve optimal statistical performance. In particular, for the simple Gaussian model for $\mathbb{Z}/2$ synchronization (1.1), AMP is known to achieve statistically optimal mean squared error (in the limit $n \rightarrow \infty$) for every value of the signal-to-noise parameter λ [56].

Recall that our Gaussian synchronization model is a generalization of (1.1). It is therefore natural for us to attempt to extend the AMP algorithm of [56] to our more general setting. We do this in Chapter 2, leading to a sharp analysis of both the statistical and computational limits of the model. We expect that our AMP algorithm achieves optimal performance among all polynomial-time algorithms. However, unlike the $\mathbb{Z}/2$ case, we predict that under certain conditions there are *statistical-to-computational gaps*, i.e. an inefficient estimator outperforms AMP. Most of the results in this chapter are non-rigorous, but are based on well-established

ideas from statistical physics, many of which have been rigorously verified in related settings.

1.3.2 Synchronization: contiguity and rigorous bounds

In Chapter 3 we complement the above non-rigorous results with some rigorous (albeit less sharp) statistical lower and upper bounds. A central concept to this chapter is the notion of *contiguity* [92] which captures whether two (sequences of) distributions cannot be consistently distinguished. Associated with contiguity is a particular second moment method that we use to show that the Gaussian synchronization model is statistically impossible in certain regimes. We also give statistical upper bounds by analyzing an inefficient exhaustive search algorithm.

1.3.3 Orbit recovery: statistical limits

In Chapter 4 we determine the statistical sample complexity of the orbit recovery problem in the high-noise regime, i.e. we determine how the number of samples n needs to scale with the noise variance σ^2 in the limit $\sigma^2 \rightarrow \infty$. Here we generalize prior work on the special case of MRA (cyclic shifts) which shows that the *method of invariants* achieves optimal sample complexity [13]. The idea behind the method of invariants is to use the samples to estimate *invariant features* of the signal which are unaffected by the group action. This leaves us with an algebraic question of determining how many invariants are needed in order to uniquely determine the signal (up to orbit). Our main contribution is a method to answer this question using tools from invariant theory and algebraic geometry. One result of this is that similarly to MRA, cryo-EM requires invariants up to degree 3 and thus has sample complexity $n = \Theta(\sigma^6)$.

There are some caveats to this result. For instance, instead of unique recovery of the signal, we only show *list recovery* wherein we output a finite list of candidate solutions, one of which is close to the true orbit. Furthermore, our recovery procedure is inefficient, leaving open the question of finding a polynomial-time algorithm for cryo-EM.

1.3.4 Orbit recovery: computational limits

In Chapter 5 we discuss issues of computational efficiency for solving the orbit recovery problem. We first survey existing results on provable efficient recovery, which have been restricted to the case of MRA. We then give a general conjecture about what we expect are the fundamental computational limits of orbit recovery. Finally we prove a result on heterogeneous MRA, showing that if the signals are random then polynomial-time recovery is possible up to the conjectured threshold.

1.4 Preliminaries on groups and representations

In this section we review some basic concepts from group theory and representation theory that will be essential in the coming chapters. Further preliminaries will be covered as needed in the individual chapters.

For a standard reference on this material, we refer the reader to e.g. [40]. We assume the reader is familiar with the algebraic notion of a *group*. We will restrict our study to *compact groups*, which are algebraically well-behaved in many ways.

Definition 1.4.1. A *topological group* is a group G along with a topology on G for which the group's binary operator and inverse function are continuous. A *compact group* is a topological group that is compact with respect to its topology.

Examples of compact groups include all finite groups (such as the cyclic group \mathbb{Z}/p for any positive integer p), and compact Lie groups such as $\mathrm{SO}(2)$ (2×2 rotation matrices), $\mathrm{SO}(3)$ (3×3 rotation matrices), and $\mathrm{U}(1)$ (the unit circle in \mathbb{C} ; note that $\mathrm{U}(1)$ is isomorphic to $\mathrm{SO}(2)$).

One crucial property of compact groups is that they admit a *Haar measure*. This can be thought of as the ‘uniform distribution’ on the group. (For finite groups, Haar measure coincides with the uniform distribution.)

Definition 1.4.2. For a compact group G , *Haar measure* is the unique positive measure μ on G that is invariant under left and right translation by any group element, normalized

so that $\mu(G) = 1$. Formally, for every $g \in G$ and every Borel subset $S \subseteq G$, we have $\mu(gS) = \mu(S) = \mu(Sg)$.

We will use some basic notions from *representation theory*.

Definition 1.4.3. A *(linear) representation* of a group G over a field k (e.g. \mathbb{R} or \mathbb{C}) is a homomorphism $\rho : G \rightarrow GL(V)$ where V is a vector space over k . A representation is denoted by (V, ρ) , or simply by ρ .

We will be exclusively concerned with finite-dimensional representations in which V is a finite-dimensional vector space k^p for some positive integer p (called the *dimension* of the representation). In this case $GL(V)$ is the set of invertible $p \times p$ matrices with entries in k . Thus we should think of a representation as a way to assign a matrix to each group element (in a way that respects group multiplication and inverse).

Definition 1.4.4. A representation (V, ρ) gives rise to a *linear group action*. For $g \in G$ and $x \in V$, the action of g on x is given by $g \cdot x \triangleq \rho(g)x$.

Definition 1.4.5. A *subrepresentation* of a representation (V, ρ) is given by a subspace W of V for which $g \cdot x \in W$ for all $g \in G$ and all $x \in W$. This is a representation (W, φ) where $\varphi(g) \in GL(W)$ is the restriction of $\rho(g)$ to W .

Definition 1.4.6. A representation is *irreducible* if it has only two subrepresentations, namely $\{0\}$ and itself. Otherwise it is *reducible*.

Definition 1.4.7. The *trivial representation* is the 1-dimensional representation in which every group element acts as the constant 1.

We will often restrict ourselves to working with representations that are unitary (or orthogonal), which is justified by the following.

Fact 1.4.8. *Let G be a compact group. For any finite-dimensional representation (V, ρ) of G over \mathbb{C} , there is a basis for V such that the representation is unitary, i.e. $\rho(g)$ is a unitary matrix for every $g \in G$. Similarly, if the representation is over \mathbb{R} , there is a basis in which it is orthogonal ($\rho(g)$ is an orthogonal matrix).*

Chapter 2

Synchronization: analysis via statistical physics

This chapter is adapted (with minor modifications) from joint work with Amelia Perry, Afonso Bandeira, and Ankur Moitra [121].

2.1 Introduction

Among the most common data problems in the sciences and machine learning is that of recovering low-rank signal present in a noisy matrix. The standard tool for such problems is principal component analysis (PCA), which estimates the signal by the top eigenvectors. One example out of many is in macroeconomics, where large, noisy correlation matrices reveal useful volatility and yield predictions in their top eigenvectors [99, 65]. However, many particular applications involve extra structure such as sparsity in the signal, and this structure is ignored by conventional PCA, leading to sub-optimal estimates. Thus, a major topic of recent interest in machine learning is to devise efficient algorithms for sparse PCA [10, 28], non-negative PCA [110], general Bayesian PCA with a prior [30], and other variants. These problems pose a major computational challenge. While significant advances have appeared, it is also expected that there are fundamental gaps between what is statistically

possible and what can be done efficiently [27, 96, 97, 101], and thus carried out in practice on very large datasets that are now prevalent.

A number of low-rank recovery problems involve symmetry and group structure in an integral way, and have been studied together under the general heading of *synchronization problems*. Broadly, the goal of such problems is to recover a list of group elements g_u from noisy pairwise measurements of the relative alignments $g_u g_v^{-1}$. Such problems arise in community detection [56], cryo-electron microscopy [142], time synchronization in networks [71], sensor network localization [51], simultaneous localization and mapping in robotics [130], surface reconstruction in computer vision [6], phase alignment in signal processing [17], and many other settings [14]. These problems exhibit a range of group structure, including rotation groups, Euclidean groups, and cyclic groups. Among the most common approaches to such problems is to linearize the observations into a matrix and then take top eigenvectors (“spectral methods”); thus, synchronization is often studied as a low rank recovery problem, with a great richness of extra structure to be exploited. We now examine a few of these problems in detail, along with the algorithmic challenges that they pose.

Community detection as $\mathbb{Z}/2$ synchronization. The problem of partitioning a graph into two well-connected subcommunities can be viewed as synchronization over the group $\{\pm 1\} \cong \mathbb{Z}/2$: each vertex has a latent group element $g_u \in \{\pm 1\}$, its community identity, and each edge is a noisy measurement of the relative status $g_u g_v^{-1}$ [14]. Spectral methods have a long history of use in such community detection and minimum cut problems (e.g. [104]); here one hopes to recover the community structure as the second eigenvector of an adjacency or Laplacian matrix. However, this approach breaks down in sparse graphs: localized noise eigenvectors associated to high-degree vertices dominate the spectrum. This is essentially a failure of PCA to adequately exploit the problem structure, as these localized eigenvectors lie far from the constraint that the truth is entrywise $\{\pm 1\}$. A number of more structured approaches have been shown to improve over PCA, including modified spectral methods [103, 112, 132, 91, 90] and semidefinite programming [2, 75, 74, 1, 111]. A major algorithmic challenge in this problem is to obtain an efficient algorithm that optimally exploits this

structure, to obtain the minimum possible estimation error. It is widely believed that belief propagation achieves this limit, with significant progress in this direction [53, 113, 56].

Gaussian $\mathbb{Z}/2$ synchronization. The following Gaussian model of $\mathbb{Z}/2$ synchronization was introduced by [56] as a simplification of community detection. One observes

$$Y = \frac{\lambda}{n}xx^\top + \frac{1}{\sqrt{n}}W,$$

where $x \in \{\pm 1\}^n$ is the signal to be recovered, W is a GOE matrix¹, and $\lambda > 0$ is a signal-to-noise parameter. Each entry of Y represents a noisy measurement of the pairwise relative alignment $x_u x_v$.

This estimation problem may be approached through ordinary PCA (top eigenvector), which one might perform by initializing with a small random guess v and repeatedly assigning $v \leftarrow Yv$; this is the method of power iteration. Random matrix theory implies that in the limit $n \rightarrow \infty$, this method achieves a nontrivial result as soon as $\lambda > 1$ [69, 25]; it is known that this threshold is tight in the sense that nontrivial estimation is information-theoretically impossible when $\lambda \leq 0$ [56, 120]. However, PCA does not achieve the minimum possible estimation error when $\lambda > 1$.

Aiming to better exploit group structure, Boumal [34] proposes² to iterate $v \leftarrow \text{sgn}(Yv)$, where sgn rounds each entry to ± 1 , thus projecting to the group. This method is highly efficient, and is moreover observed to produce a better estimate than PCA once the signal-to-noise parameter λ is sufficiently large. However, this method does not appear to produce a meaningful estimate until λ is somewhat larger³ than 1. This behavior poses two challenges: can we devise an efficient iterative method combining the best features of PCA and the projected power method, which outperforms both statistically? As studied by [56], iterative methods based on *approximate message passing* are very effective in this setting, achieving

¹Gaussian orthogonal ensemble: a random symmetric matrix with off-diagonal entries $\mathcal{N}(0, 1)$, diagonal entries $\mathcal{N}(0, 2)$, and all entries independent (up to symmetry).

²Boumal’s paper targets the close variant of $U(1)$ synchronization. Projected power methods have appeared earlier, e.g. [110]. Their application to synchronization problems also appears in [41].

³A heuristic analysis similar to Section 2.6.3 suggests that $\lambda > \sqrt{\pi/2} \approx 1.253$ is required.

an optimal estimation error that beats both methods discussed above (Figure 2-1). Moving forward, can we find analogous iterative methods for groups other than $\mathbb{Z}/2$, and for more complicated observation models?




Figure 2-1: Comparison of iterative algorithms for Gaussian $\mathbb{Z}/2$ synchronization; log-error $\ln(1 - |\langle x, \hat{x} \rangle|/n)$ versus SNR λ . The three curves are projected power iteration (green), PCA (blue), and the AMP algorithm of the present thesis (black). (The special case of our AMP algorithm for $\mathbb{Z}/2$ appears in [56].) Each data point is an average of 200 trials with $n = 2000$ vertices.

Angular synchronization and Gaussian $U(1)$ synchronization. Singer [141] introduced a PCA approach for angular synchronization, a 2D analogue of the problem above with symmetry over $U(1)$ (unit-norm complex numbers: $z \in \mathbb{C}$ with $|z| = 1$), in which one estimates the orientations of noisy, randomly rotated copies of an unknown image. Again, a Gaussian simplification has been studied in [34]: one observes a matrix

$$Y = \frac{\lambda}{n} x x^* + \frac{1}{\sqrt{n}} W,$$

where $x \in U(1)^n$ is the signal to be recovered, a vector of unit complex numbers, and W is a GUE matrix. Boumal's algorithm now iterates $v \leftarrow f(Yv)$, where f divides each entry by its

norm, thus projecting to the unit circle. Again this method beats PCA only in a sufficiently high SNR regime, leaving open the question: can we achieve optimal estimation through an almost-linear-time, iterative algorithm? Results are known for slower convex programs [16].

The Gaussian model above is a rather drastic simplification of the problem of synchronizing entire images, however. A more elaborate Gaussian model (which the author and others first introduced in [120]) is the following: instead of observing only the matrix Y as above, suppose we are given matrices corresponding to different Fourier modes:

$$\begin{aligned} Y_1 &= \frac{\lambda_1}{n} x x^* + \frac{1}{\sqrt{n}} W_1, \\ Y_2 &= \frac{\lambda_2}{n} x^2 (x^2)^* + \frac{1}{\sqrt{n}} W_2, \\ &\vdots \\ Y_K &= \frac{\lambda_K}{n} x^K (x^K)^* + \frac{1}{\sqrt{n}} W_K, \end{aligned}$$

where x^k denotes entrywise power, and W_k are independent GUE matrices. Due to Fourier theory, a very large class of measurement models for $U(1)$ synchronization decomposes into matrix-based observations on different frequencies, in a manner resembling the model above.

With a PCA-based approach, it is not clear how to effectively couple the information from these matrices to give a substantially better estimate than could be derived from only one. Indeed, many spectral approaches to this synchronization problem and others apply PCA to only the first Fourier mode, discarding a great deal of useful data on other modes. Bandeira et al. [18] introduced a very general semidefinite relaxation for synchronization problems such as above, but its performance remains unclear even empirically: while this convex program can be solved in polynomial time, it is large enough to make experiments or application difficult. Can we hope for some efficient iterative algorithm to strongly leverage data from multiple ‘frequencies’ or ‘channels’ such as this?

Cryo-electron microscopy. Perhaps the biggest concrete goal in the study of synchronization is the orientation problem in *cryo-electron microscopy* (cryo-EM), a modern alter-

native to x-ray crystallography for imaging large biomolecules. One is given many noisy 2D projections (microscopy images) of an unknown molecule, each taken from a different unknown 3D orientation. The goal is to estimate the orientations, in order to assemble the images into an estimate of the molecule structure [142]. Thus, one is tasked with learning elements g_u of $SO(3)$, one for each image u , based on some loss function derived from the observed images. This loss function depends only on the relative alignments $g_u g_v^{-1}$, as there is no *a priori* reference frame. One previous approach to this problem, due to [43, 142], produces a matrix of pairwise image comparisons, and then attempts to extract the rotations g_u from the top eigenvectors of this matrix. However, it is reasonable to imagine that this approach could be significantly sub-optimal, for the reasons seen above: PCA does not exploit the significant group structure of the signal, and by linearizing into a single matrix, PCA only exploits the first “Fourier mode” of the observations, as in the previous problem.

Another method used in practice for cryo-EM and related problems is *alternating minimization*, which alternates between estimating the rotations by aligning the images with a previous guess of the molecule structure, and then estimating the molecule structure from the images using these rotations. This method only appears to succeed given a strong initial guess of the molecule structure, and then it is unclear whether the final estimate mainly reflects the observations or simply the initial guess, leading to a problem of model bias; see e.g. [42]. In this thesis we are interested in *de novo* estimation without a substantial initial guess, steering clear of this pitfall.

The complexity of the observation model and the group present a host of challenges, but centrally: can an improved iterative algorithm for the previous synchronization problems generalize to the noncommutative setting of groups such as $SO(3)$?

In this chapter we present an iterative algorithm to meet all of the challenges above. Our algorithm aims to solve a general formulation of the synchronization problem: it can apply to multiple-frequency problems for a large class of observation models, with symmetry over any compact group. Our approach is statistically powerful, empirically providing a better estimate than both PCA and the projected power method on $U(1)$ synchronization,

and leveraging multiple frequencies to give several orders of magnitude improvement in estimation error in experiments (see Figures 2-4 and 2-5). Indeed, we conjecture based on ideas from statistical physics that in many regimes our algorithm is statistically optimal, providing a minimum mean square error (MMSE) estimator asymptotically as the matrix dimensions become infinite (see Section 2.8). Finally, our approach is highly efficient, with each iteration taking time linear in the (matrix) input, and with roughly 15 iterations sufficing for convergence in experiments.

Our algorithm follows the framework of *approximate message passing* (AMP), based on belief propagation on graphical models [117] and the related *cavity method* in statistical physics [106]. Following a general blueprint, AMP algorithms have previously been derived and analyzed for compressed sensing [63, 64, 23, 81], sparse PCA [57], non-negative PCA [110], cone-constrained PCA [59], planted clique [58] and general structured PCA [125]. In fact, AMP has already been derived for $\mathbb{Z}/2$ synchronization under a Gaussian observation model [56], and our algorithm will generalize this one to all compact groups. A striking feature of AMP is that its asymptotic performance can be captured exactly by a particular fixed-point equation called *state evolution*, which has enabled the rigorous understanding of its performance on some problems [23, 81]. AMP is provably statistically optimal in many cases, including Gaussian $\mathbb{Z}/2$ synchronization (modulo a technicality whereby the proof supposes a small warm-start) [56].

AMP algorithms frequently take a form similar to the projected power method of Boumal described above, alternating between a matrix–vector product with the observations and an entrywise nonlinear transformation, together with an extra ‘Onsager’ correction term. In the case of $\mathbb{Z}/2$ or $U(1)$ synchronization, we will see that the AMP derivation reproduces Boumal’s algorithm, except with the projection onto the unit circle replaced by a soft, sigmoid-shaped projection function to the unit disk (see Figure 2-2), with the magnitude maintaining a quantitative measure of confidence. A “low-temperature limit” of AMP then recovers exactly Boumal’s algorithm, while the “high-temperature limit” is ordinary PCA; we will see that belief propagation suggests an optimal intermediate temperature based on the

signal-to-noise ratio. Integrating the usual AMP blueprint with the representation theory of compact groups, we obtain a broad generalization of this method, to synchronization problems with multiple frequencies and noncommutative groups such as $SO(3)$. In full generality, the nonlinear transformation has a simple interpretation through representation theory and the exponential function.

The rest of this chapter is organized as follows. We begin in Section 2.2 with an outline of our methods in the simplified cases of synchronization over $\mathbb{Z}/2$ and $U(1)$, motivating our approach from a detailed discussion of prior work and its shortfalls. In Section 2.3 we provide our general algorithm and the general model for which it is designed. Several experiments on this Gaussian model and other models are presented in Section 2.4, demonstrating strong empirical performance. We then offer two separate derivations of our AMP algorithm: in Section 2.5, we derive our algorithm as a simplification of belief propagation, and then in Section 2.6 we give an alternative self-contained derivation of the nonlinear update step and use this to provide a non-rigorous analysis of AMP (based on standard assumptions from statistical physics). In particular, we derive the state evolution equations that govern the behavior of AMP, and use these to identify the threshold above which AMP achieves non-trivial reconstruction. Namely, we see that AMP has the same threshold as PCA (requiring the SNR λ to exceed 1 on at least one frequency), but AMP achieves considerably better recovery error above the threshold. In Section 2.7 we argue for the correctness of the above non-rigorous analysis, providing both numerical and mathematical evidence. It is known that inefficient estimators can beat the $\lambda = 1$ threshold (see Chapter 3) but we conjecture that no efficient algorithm is able to break this barrier, thus concluding in Section 2.8 with an exploration of statistical-to-computational gaps that we expect to exist in synchronization problems, driven by ideas from statistical physics.

2.2 Intuition: iterative methods for $\mathbb{Z}/2$ and $U(1)$ synchronization

We begin with a discussion of synchronization methods over the cyclic group $\mathbb{Z}/2 = \{\pm 1\}$ and the group of unit-norm complex numbers (or 2D rotations) $U(1)$. These examples will suffice to establish intuition and describe much of the novelty of our approach, while avoiding the conceptual and notational complication of representation theory present in the general case. Sections 2.2.1, 2.2.2, and some of 2.2.3 discuss prior work on these problems in more depth, while Sections 2.2.3 and 2.2.4 develop a special case of our algorithm.

2.2.1 $\mathbb{Z}/2$ synchronization

The problem of *Gaussian $\mathbb{Z}/2$ synchronization* is to estimate a uniformly drawn signal $x \in \{\pm 1\}^n$ given the matrix

$$Y = \frac{\lambda}{n}xx^\top + \frac{1}{\sqrt{n}}W,$$

where W is a symmetric matrix whose off-diagonal⁴ entries are distributed independently (up to symmetry) as $\mathcal{N}(0, 1)$, and $\lambda > 0$ is a signal-to-noise parameter. With this scaling, the signal and noise are of comparable size in spectral norm; we can not hope to recover x exactly, but we can hope to produce an estimate $\hat{x} \in \{\pm 1\}$ that is correlated nontrivially with x , i.e. there exists $\varepsilon > 0$ (not depending on n) such that $\frac{1}{n^2}\langle x, \hat{x} \rangle^2 > \varepsilon$ with probability $1 - o(1)$ as $n \rightarrow \infty$. As $xx^\top = (-x)(-x)^\top$, we can only hope to estimate x up to sign; thus we aim to achieve a large value of $\langle x, \hat{x} \rangle^2$. We now review three algorithmic methods for this problem.

Spectral methods. With the scaling above, the spectral norm of the signal $\frac{\lambda}{n}xx^\top$ is λ , while that of the noise $\frac{1}{\sqrt{n}}W$ is 2. By taking the top eigenvector of Y , x may be estimated

⁴The diagonal entries are irrelevant because the diagonal entries of Y contain no information about x . Various conventions for the diagonal entries can be taken, such as $Y_{ii} = 0$ or $W_{ii} \sim \mathcal{N}(0, 2)$. Any such reasonable choice of diagonal entries will have negligible effect on the algorithms discussed here, e.g. the diagonal component of Y is $o(1)$ in spectral norm.

with significant correlation provided that λ is a large enough constant.

Specifically, the generative model for Y above is a special case of the *spiked Wigner model*, and the eigenvalues and eigenvectors of such spiked models are among the main objects of study in random matrix theory. When $\lambda > 1$, the (unit norm) top eigenvector $v_{\max}(Y)$ correlates nontrivially with x ; more specifically, as $n \rightarrow \infty$, we have $\frac{1}{n} \langle x, v_{\max}(Y) \rangle^2 \rightarrow 1 - 1/\lambda^2$ in probability [69, 25]. When $\lambda \leq 1$, this squared correlation tends to zero; in fact, this is known to be true of all estimators [56, 120], reflecting a sharp statistical phase transition.

Note that a top eigenvector may be computed through *power iteration* as follows: an initial guess $v^{(0)}$ is drawn randomly, and then we iteratively compute $v^{(t)} = Yv^{(t-1)}$, rescaling the result as appropriate. Thus each entry is computed as $v_u^{(t)} = \sum_w Y_{uw}v_w^{(t-1)}$; we can imagine that each entry w sends a ‘message’ $Y_{uw}v_w^{(t-1)}$ to each entry u – the ‘vote’ of entry w as to the identity of entry u – and then each entry sums the incoming votes to determine its new value. The result has both a sign, reflecting the weighted majority opinion as to whether that entry should ultimately be $+1$ or -1 , and also a magnitude, reflecting a confidence and serving as the weight in the next iteration. Thus we can envision the spectral method as a basic “message-passing algorithm.”

While this approach is effective as quantified above, it would seem to suffer from two drawbacks:

- the spectral method is rotation-invariant, and thus cannot exploit the entrywise ± 1 structure of the signal;
- the vertex weights can grow without bound, potentially causing a few vertices to exert undue influence.

Indeed, these drawbacks cause major issues in the *stochastic block model*, a variant of the model above with the Gaussian observations replaced by low-probability Bernoulli observations, usually envisioned as the adjacency matrix of a random graph. Here a few sporadically high-degree vertices can dominate the spectral method, causing asymptotically significant losses in the statistical power of this approach.

Projected power iteration. Our next stepping-stone toward AMP is the projected power method studied by [34, 41], a variant of power iteration that exploits entrywise structure. Here each iteration takes the form $v^{(t)} = \text{sgn}(Yv^{(t-1)})$, where the sign function $\text{sgn} : \mathbb{R} \rightarrow \{\pm 1\}$ applies entrywise. Thus each iteration is a majority vote that is weighted only by the magnitudes of the entries of Y ; the weights do not become more unbalanced with further iterations. Further, this algorithm is basis-dependent in a way that plausibly exploits the ± 1 structure of the entries.

Empirically, this algorithm obtains better correlation with the truth, on average, when $\lambda > 2.4$ approximately; see Figure 2-3. However, for very noisy models with $1 < \lambda < 2.4$, this method appears weaker than the spectral method. The natural explanation for this weakness is that this projected power method forgets the distinction between a 51% vote and a 99% vote, and thus is over-influenced by weak entries. This is particularly problematic for low signal-to-noise ratios λ , for which 51% votes are common. In fact, a heuristic analysis similar to Section 2.6.3 suggests that this method does not achieve the correct threshold for λ , failing to produce nontrivial correlation with the truth whenever $\lambda \leq \sqrt{\pi/2} \approx 1.253$.

Soft-threshold power iteration. A natural next step is to consider iterative algorithms of the form $v^{(t)} = f(Yv^{(t-1)})$, where f applies some function $\mathbb{R} \rightarrow [-1, 1]$ entrywise (by abuse of notation, we will also denote the entrywise function by f). Instead of the identity function, as in the spectral method, or the sign function, as in the projected power method, we might imagine that some continuous, sigmoid-shaped function performs best, retaining some sense of the confidence of the vote without allowing the resulting weights to grow without bound. It is natural to ask what the optimal function for this purpose is, and whether the resulting weights have any precise meaning.

Given the restriction to the interval $[-1, 1]$, one can imagine treating each entry as a sign with confidence in a more precise way, as the expectation of a distribution over $\{\pm 1\}$. At each iteration, each entry u might then obtain the messages $Y_{uw}v_w^{(t-1)}$ from all others, and compute the posterior distribution, summarized as an expectation $v_u^{(t)}$. As one can compute, this corresponds to the choice of transformation $f(t) = \tanh(\lambda t)$ where λ is the

signal-to-noise parameter from above (see Figure 2-2). The resulting algorithm is similar to [56].




Figure 2-2: Soft threshold functions used by AMP. The solid line is $f(t) = \tanh(t)$, used for $\mathbb{Z}/2$ synchronization. The dashed line is $f(t) = I_1(2t)/I_0(2t)$, used for $U(1)$ synchronization with one frequency. (I_k denotes modified Bessel functions of the first kind.)

2.2.2 Belief propagation and approximate message passing

The soft-projection algorithm above may remind the reader of belief propagation, due to [117] in the context of inference, and to [106] as the *cavity method* in the context of statistical physics. We may envision the problem of estimating x as probabilistic inference over a graphical model. The vertices of the model represent the unknown entries of x , and every pair of vertices u, w participates in an edge interaction based on the matrix entry $Y_{uw} = Y_{wu}$. Specifically, it may be computed that the posterior distribution for $x \in \{\pm 1\}^n$ after observing Y is given by

$$\Pr(x) \propto \prod_{u < w} \exp(\lambda Y_{uw} x_u x_w),$$

which is precisely the factorization property that a graphical model captures.

Given such a model, belief propagation proceeds in a fashion reminiscent of the previous algorithm: each vertex w sends a message to each neighbor u encoding the posterior distribution of x_u based on the previous distribution of x_w and the direct interaction λY_{uw} .

Each vertex u then consolidates all incoming messages into a new ‘posterior’ distribution on x_u given these messages, computed as if the messages were independent. However, belief propagation introduces a correction to this approach: rather than letting information passed from w to u propagate back to w on the next iteration, belief propagation is designed to pass information along only those paths that do not immediately backtrack. Specifically, at each iteration, the message from w to u is based on only the synthesis of messages (to w) from all vertices except u from the previous iteration.

This algorithm differs from the iterative methods presented above, both in this non-backtracking behavior, and in the fact that the transformation from the distribution at w to the message $w \rightarrow u$ is not necessarily linear (as in the multiplication $Y_{uw}v_w$ above). Both of these differences can be reduced by passing to the framework of *approximate message passing* [63], which simplifies belief propagation in dense models with weak interactions, through the following two observations (inspired by [145] in the theory of spin glasses):

- As the interaction λY_{uw} is small, scaling as $O(1/\sqrt{n})$ as $n \rightarrow \infty$, we may pass to an expansion in small Y_{uw} when computing the message $w \rightarrow u$ from the mean w . In this example, we find that the message $w \rightarrow u$ should be Rademacher with mean $\lambda Y_{uw} v_w^{(t-1)} + O(Y_{uw}^2)$, where $v_w^{(t-1)}$ is the mean of the distribution for x_w in the previous iteration. This linear expansion ensures that the main message-passing step can be expressed as a matrix–vector product.
- Rather than explicitly computing non-backtracking messages, which is computationally more involved, we may propagate the more naïve backtracking messages and then subtract the bias due to this simplification, which concentrates well. This correction term is called an *Onsager correction*. If vertex w passes messages to all neighbors based on its belief at iteration $t - 2$, and then all of these neighbors send return messages based on their new beliefs at time $t - 1$, then when updating the belief for vertex w at time t , one can explicitly subtract off the ‘reflected’ influence of the previous belief at time $t - 2$. It turns out that this is the only correction necessary: all other error contributions (e.g. 3-cycles) are $o(1)$.

Following these simplifications, one can arrive at an *approximate message passing* (AMP) algorithm for $\mathbb{Z}/2$ synchronization:

Algorithm 2.2.1 (AMP for $\mathbb{Z}/2$ synchronization [56]).

- Initialize $v^{(0)}$ to small (close to zero) random values in $[-1, 1]$. Initialize $v^{(-1)} = 0$.
- Iterate for $1 \leq t \leq T$:
 - Set $c^{(t)} = \lambda Y v^{(t-1)} - \lambda^2(1 - \langle (v^{(t-1)})^2 \rangle) v^{(t-2)}$, the Onsager-corrected sum of incoming messages.
 - Set $v_u^{(t)} = \tanh(c_u^{(t)})$ for each vertex u , the new estimated posterior mean.
- Return $\hat{x} = v^{(T)}$ (or the approximate MAP estimate $\hat{x} = \text{sgn}(v^{(T)})$ if a proper estimate in $\{\pm 1\}^n$ is desired).

Here $\langle (v^{(t-1)})^2 \rangle$ denotes the average of the squared entries of $v^{(t-1)}$. Detailed derivations of this algorithm appear in Sections 2.5 and 2.6 in much higher generality.

In the setting of $\mathbb{Z}/2$ synchronization, an algorithm equivalent to the above approach appears in [56], where a statistical optimality property is proven: if AMP is warm-started with a state $v^{(0)}$ with nontrivial correlation with the truth, then it converges to an estimate of x that achieves minimum mean-squared error (MMSE) asymptotically as $n \rightarrow \infty$. The warm-start requirement is technical and likely removable: if AMP is initialized to small randomness, with trivial correlation $O(1/\sqrt{n})$ with the truth, then its early iterations resemble PCA and should produce nontrivial correlation in $O(\log n)$ iterations. The statistical strength of AMP is confirmed empirically, as it appears to produce a better estimate than either PCA or the projected power method, for every $\lambda > 1$; see Figure 2-3.

2.2.3 AMP for Gaussian $U(1)$ synchronization with one frequency

As a first step toward higher generality, consider the following Gaussian synchronization model over the unit complex numbers $U(1)$. The goal is to estimate a uniformly drawn

signal $x \in U(1)^n$ given the matrix

$$Y = \frac{\lambda}{n}xx^* + \frac{1}{\sqrt{n}}W,$$

where W is a Hermitian matrix whose entries are distributed independently (up to Hermitian symmetry) as $\mathbb{C}\mathcal{N}(0, 1)$, the complex normal distribution given by $\mathcal{N}(0, 1/2) + \mathcal{N}(0, 1/2)i$, and where $\lambda > 0$ is a signal-to-noise parameter. As xx^* is invariant under a global phase shift of x , we can only hope to estimate up to the same ambiguity, and so we would like an estimator \hat{x} that maximizes $|\langle x, \hat{x} \rangle|^2$, where the inner product is conjugated in the second variable. Many of the previously discussed iterative techniques adapt to this new case.

Spectral methods. The same analysis of the spectral method holds in this case; thus when $\lambda > 1$, the top eigenvector achieves nontrivial correlation with x , while for $\lambda < 1$, the spectral method fails and nontrivial estimation is provably impossible [120].

Projected power method. After each matrix–vector product, we can project $v^{(t)}$ entrywise onto the unit circle, preserving the phase of each entry while setting the magnitude to 1. This algorithm is analyzed in [34] in a lower-noise regime, where it is shown to converge to the maximum likelihood estimator. A heuristic analysis similar to Section 2.6.3 suggests that this method does not achieve nontrivial correlation with the truth unless $\lambda > 2/\sqrt{\pi} \approx 1.128$.

Soft-threshold power method. One might imagine applying some entrywise function after each matrix–vector product, which preserves the phase of each entry while mapping the magnitude to $[0, 1]$. Thus the vector entries v_u live in the unit disk, the convex hull of the unit circle; these might be envisioned as estimates of the posterior expectation of x_u .

Belief propagation & AMP. Belief propagation is somewhat problematic in this setting: all messages should express a distribution over $U(1)$, and it is not *a priori* clear how this should be expressed in finite space. However, under the simplifications of approximate message passing, the linearity of the message-passing stage enables a small summary of this

distribution to suffice: we need only store the expectation of each distribution, a single value in the unit disk. Approximate message passing takes the following form:

Algorithm 2.2.2 (AMP for $U(1)$ synchronization with one frequency).

- Initialize $v^{(0)}$ to small random values in the unit disk $\text{conv}(U(1))$. Initialize $v^{(-1)} = 0$.
- Iterate for $1 \leq t \leq T$:
 - Set $c^{(t)} = \lambda Y v^{(t-1)} - \lambda^2(1 - \langle |v^{(t-1)}|^2 \rangle) v^{(t-2)}$, the Onsager-corrected sum of incoming messages.
 - Set $v_u^{(t)} = f(c_u^{(t)})$ for each vertex u , the new estimated posterior mean. Here f applies the function $f(t) = I_1(2t)/I_0(2t)$ to the magnitude, leaving the phase unchanged.
- Return $\hat{x} = v^{(T)}$ (or the approximate MAP estimate $\hat{x} = \text{phase}(v^{(T)})$ if a proper estimate in $U(1)^n$ is desired).

Here I_k denotes the modified Bessel functions of the first kind. The function f is depicted in Figure 2-2. Detailed derivations of this algorithm appear in Sections 2.5 and 2.6 in much higher generality.

2.2.4 AMP for Gaussian $U(1)$ synchronization with multiple frequencies

Consider now the following more elaborate synchronization problem. The goal is to estimate a spike $x \in U(1)^n$ from the observations

$$\begin{aligned}
Y_1 &= \frac{\lambda_1}{n} x x^* + \frac{1}{\sqrt{n}} W_1, \\
Y_2 &= \frac{\lambda_2}{n} x^2 (x^2)^* + \frac{1}{\sqrt{n}} W_2, \\
&\vdots
\end{aligned}$$

$$Y_K = \frac{\lambda_K}{n} x^K (x^K)^* + \frac{1}{\sqrt{n}} W_K,$$

where the W_k are independent Hermitian matrices whose entries are distributed independently (up to Hermitian symmetry) as $\mathbb{CN}(0, 1)$, the $\lambda_k > 0$ are signal-to-noise parameters, and x^k denotes the entrywise k th power of x . This multifrequency Gaussian model was first introduced by the author and others in [120]. The associated MAP estimation problem is an instance of the *non-unique games* framework, introduced by [18].

Thus we are given K independent noisy matrix-valued observations of x ; we can imagine these observations as targeting different *frequencies* or Fourier modes. Given two independent draws of $\lambda x x^* / n + W / \sqrt{n}$ as in the previous section, the spectral method applied to their average will produce a nontrivial estimate of x as soon as $\lambda > 1/\sqrt{2}$. However, under the multiple frequencies model above, with $K = 2$ and $\lambda_1 = \lambda_2 = \lambda$, nontrivial estimation is provably impossible for $\lambda < 0.937$ (see Chapter 3); we present non-rigorous evidence in Section 2.8 that the true statistical threshold should in fact remain $\lambda = 1$. Thus the multiple frequencies model would seem to confound attempts to exploit the multiple observations together. However, we will discuss how AMP enables us to obtain a much better estimate when $\lambda > 1$ than is possible with one matrix alone.

Let us return to the issue of belief propagation over $U(1)$, and of how to represent distributions. One crude approach might be to discretize $U(1)$ and express the density on a finite subset of points; however, this is somewhat messy (e.g. the discretization may not be preserved under rotation) and only becomes worse for more elaborate groups such as $SO(3)$ (here one can not even find arbitrarily fine discretizations on which the group acts transitively).

Instead, we could exploit the rich structure of Fourier theory, and express a distribution on $U(1)$ by the Fourier series of its density⁵. Thus, if $\mu_w^{(t)}$ is the belief distribution at vertex

⁵A dense subset of distributions satisfies appropriate continuity assumptions to discuss their densities with respect to uniform measure, a Fourier series, etc., and we will not address these analytic technicalities further.

w and time t , we can express:

$$\frac{d\mu_w^{(t)}}{d\theta/2\pi} = \sum_{k \in \mathbb{Z}} v_{w,k} e^{ik\theta},$$

with $v_{w,0} = 1$ and $v_{w,-k} = \overline{v_{w,k}}$. Computing the distributional message $m_{w \rightarrow u}$ from w to u , we obtain

$$\frac{dm_{w \rightarrow u}^{t+1}}{d\theta/2\pi} = 1 + \sum_{1 \leq |k| \leq K} \lambda_k(Y_k)_{uw} v_{w,k} e^{ik\theta} + O((Y_\bullet)_{uw}^2),$$

where we take $Y_{-k} = \overline{Y_k}$. As $(Y_k)_{uw}$ is order $1/\sqrt{n}$ in probability, this approximation will be asymptotically accurate. Thus it suffices to represent distributions by the coefficients $v_{w,k}$ with $|k| \leq K$. By conjugate symmetry, the coefficients with $1 \leq k \leq K$ suffice. The sufficiency of this finite description of each belief distribution is a key insight to our approach.

The other crucial observation concerns the remaining BP step of consolidating all incoming messages into a new belief distribution. As each incoming message is a small perturbation of the uniform distribution, the approximation $\log(1+x) \approx x$ allows us to express the log-density of the message distribution:

$$\log \frac{dm_{w \rightarrow u}^{t+1}}{d\theta/2\pi} = \sum_{1 \leq |k| \leq K} \lambda_k(Y_k)_{uw} v_{w,k} e^{ik\theta} + O((Y_\bullet)_{uw}^2).$$

We now add these log-densities to obtain the log-density of the new belief distribution, up to normalization:

$$\log \frac{d\mu_u^{t+1}}{d\theta/2\pi} + \text{const.} = \sum_{1 \leq |k| \leq K} \left(\sum_{w \neq u} \lambda_k(Y_k)_{uw} v_{w,k} \right) e^{ik\theta} + O((Y_\bullet)_{uw}^2).$$

We thus obtain the Fourier coefficients of the *log-density* of the new belief from the Fourier coefficients of the *density* of the old belief, by matrix-vector products. Remarkably, this tells us that the correct per-vertex nonlinear transformation to apply at each iteration is the transformation from Fourier coefficients of the log-density to those of the density! In other words, the transformation acts on Fourier series as composition with \exp , followed by normalization. (In section 2.6 we will see an alternative interpretation of this nonlinear

transformation as an MMSE estimator.)

The only constraints on a valid log-density are those of conjugate symmetry on Fourier coefficients; thus log-densities form an entire linear space. By contrast, densities are subject to non-negativity constraints, and form a nontrivial convex body in \mathbb{R}^K . The latter space is the analogue of the unit disk or the interval $[-1, 1]$ in the preceding examples, and this transformation from the Fourier series of a function to those of its exponential (together with normalization) forms the analogue of the preceding soft-projection functions.

We thus arrive at an AMP algorithm for the multiple-frequency problem:

Algorithm 2.2.3 (AMP for $U(1)$ synchronization with multiple frequencies).

- For each $1 \leq k \leq K$ and each vertex u , initialize $v_{u,k}^{(0)}$ to small random values in \mathbb{C} and initialize $v_{u,k}^{(-1)} = 0$.
- Iterate for $1 \leq t \leq T$:
 1. For each $1 \leq k \leq K$, set $c_k^{(t)} = \lambda_k Y_k v_k^{(t-1)} - \lambda_k^2 (1 - \langle (v_k^{(t-1)})^2 \rangle) v_k^{(t-2)}$, the vector of k th Fourier components of the estimated posterior log-densities, with Onsager correction.
 2. Compute $v_k^{(t)}$, the vector of k th Fourier components of the estimated posterior densities.
- Return $\hat{x} = v_1^{(T)}$ (or some rounding if a proper estimate in $U(1)^n$ is desired, or even the entire per-vertex posteriors represented by $c^{(T)}$).

Again, a more detailed derivation can be found in Sections 2.5 and 2.6.

It is worth emphasizing that only the expansion

$$\log \frac{d\mu_u^{(t)}}{d\theta/2\pi} + \text{const.} = 2\text{Re} \sum_{1 \leq k \leq K} c_{k,w}^{(t)} e^{ik\theta}$$

is an accurate expansion of the estimated vertex posteriors. While this log-density is band-limited, this still allows for the density to be a very spiked, concentrated function, without

suffering effects such as the Gibbs phenomenon. By contrast, the finitely many v coefficients that this algorithm computes do not suffice to express the Fourier expansion of the density, and a truncated expansion based on only the computed coefficients might even become negative.

We conclude this section by noting that nothing in our derivation depended crucially on the Gaussian observation model. The choice of model tells us how to propagate beliefs along an edge according to a matrix–vector product, but we could carry this out for a larger class of graphical models. The essential properties of a model, that enables this approach to adapt, are:

- The model can be expressed as a graphical model with only pairwise interactions:

$$\Pr(x) \propto \prod_{u < w} \mathcal{L}_{uw}(x_u, x_w).$$

- The interaction graph is dense (at least a constant fraction of pairs interact), with all pair potentials individually weak ($1 + O(1/\sqrt{n})$).
- The pair potentials $\mathcal{L}_{uw}(x_u, x_w)$ depend only on the group ratio $x_u x_w^{-1}$. (This is the core property of a *synchronization* problem.)
- The pair potentials \mathcal{L}_{uw} are band-limited as a function of $x_u x_w^{-1}$, to some finite collection of Fourier modes. This assumption (or approximation) allows the potentials to be expressed by a finite number of parameters.
- The noise is independent across edges, and the noise on different Fourier modes is uncorrelated. This is satisfied by the Gaussian model and is used in our derivation of the Onsager correction.

The formulation of AMP for general models of this form is discussed in the next section.

2.3 AMP over general compact groups

The approach discussed above for $U(1)$ synchronization with multiple frequencies readily generalizes to the setting of an arbitrary compact group G , with Fourier theory generalized to the representation theory of G . Just as the Fourier characters are precisely the irreducible representations of $U(1)$, we will represent distributions over G by an expansion in terms of irreducible representations, as described by the Peter–Weyl theorem. Under the assumption (or approximation) of band-limited pairwise observations, it will suffice to store a finite number of coefficients of this expansion. (Note that finite groups have a finite number of irreducible representations and so the band-limited requirement poses no restriction in this case. For infinite groups, the observations may not be band-limited, but we can approximate them arbitrarily-well as such by taking sufficiently many of the most important irreducible representations and discarding the rest.)

A geometric view on this is as follows. Belief propagation ideally sends messages in the space of distributions on G ; this is a form of formal convex hull on G , and is illustrated in the case of $\mathbb{Z}/2$ synchronization by sending messages valued in $[-1, 1]$. When G is infinite, however, this space is infinite-dimensional and thus intractable. We could instead ask whether the convex hull of G taken in some finite-dimensional embedding is a sufficient domain for messages. The key to our approach is the observation that, when observations are band-limited, it suffices to take an embedding of G described by a sum of irreducible representations.

This section will be devoted to presenting our AMP algorithm in full generality, along with the synchronization model that it applies to. In particular, the algorithm can run on the general graphical model formulation of Section 2.3.2, but when we analyze its performance we will restrict to the Gaussian observation model of Section 2.3.4.

2.3.1 Representation theory preliminaries

Haar measure

A crucial property of compact groups is the existence of a (normalized) Haar measure, a positive measure μ on the group that is invariant under left and right translation by any group element, normalized such that $\mu(G) = 1$ [40]. This measure amounts to a concept of ‘uniform distribution’ on such a group, and specializes to the ordinary uniform distribution on a finite group. Throughout this chapter, integrals of the form

$$\int_G f(g) \, dg,$$

are understood to be taken with respect to Haar measure.

Peter–Weyl decomposition

Fix a compact group G . We will be working with the density functions of distributions over G . In order to succinctly describe these, we use the representation-theoretic analogue of Fourier series: the Peter–Weyl decomposition. The Peter–Weyl theorem asserts that $L^2(G)$ (the space of square-integrable, complex scalar functions on G) is the closure of the span (with coefficients from \mathbb{C}) of the following basis, which is furthermore orthonormal with respect to the Hermitian inner product on $L^2(G)$ [40]:

$$R_{\rho ab}(g) = \sqrt{d_\rho} \rho(g)_{ab},$$

indexed over all complex irreducible representations ρ of G , and all $1 \leq a \leq d_\rho$, $1 \leq b \leq d_\rho$ where $d_\rho = \dim \rho$. The representations are assumed unitary (which is without loss of generality because any representation is isomorphic to a unitary one). The inner product is taken to be conjugate-linear in the second input.

Since we want our algorithm to be able to store the description of a function using finite space, we fix a finite list \mathcal{P} of irreducible representations to use. From now on, all Peter–Weyl

decompositions will be assumed to only use representations from \mathcal{P} ; we describe functions of this form as *band-limited*. We exclude the trivial representation from this list because we will only need to describe functions up to an additive constant. Given a real-valued function $f : G \rightarrow \mathbb{R}$, we will often write its Peter–Weyl expansion in the form

$$f(g) = \sum_{\rho} \langle \hat{f}_{\rho}, R_{\rho}(g) \rangle,$$

where \hat{f}_{ρ} and $R_{\rho}(g) = \sqrt{d_{\rho}} \rho(g)$ are $d_{\rho} \times d_{\rho}$ complex matrices. Here ρ ranges over the irreducibles in \mathcal{P} ; we assume that the functions f we are working with can be expanded in terms of only these representations. The matrix inner product used here is defined by $\langle A, B \rangle = \text{Tr}(AB^*)$. The Peter–Weyl coefficients of a function can be extracted by integration against the appropriate basis functions:

$$\hat{f}_{\rho} = \int_G R_{\rho}(g) f(g) \, dg.$$

By analogy to Fourier theory, we will sometimes refer to the coefficients \hat{f}_{ρ} as *Fourier coefficients*, and refer to the irreducible representations as *frequencies*.

Representations of real, complex, and quaternionic type

Every irreducible complex representation of a compact group G over \mathbb{C} is of one of three types: real type, complex type, or quaternionic type [40]. We will need to deal with each of these slightly differently. In particular, for each type we are interested in the properties of the Peter–Weyl coefficients that correspond to the represented function being real-valued.

A complex representation ρ is of *real type* if it can be defined over the reals, i.e. it is isomorphic to a real-valued representation. Thus in this case we assume without loss of generality that we are working with a real-valued ρ . In this case it is clear that if f is a real-valued function then (by integrating against R_{ρ}) \hat{f}_{ρ} is real. Conversely, if \hat{f}_{ρ} is real then the term $\langle \hat{f}_{\rho}, R_{\rho}(g) \rangle$ (from the Peter–Weyl expansion) is real.

A representation ρ is of *complex type* if ρ is not isomorphic to its conjugate representation

$\bar{\rho}$, which is the irreducible representation defined by $\bar{\rho}(g) = \overline{\rho(g)}$. We will assume that the complex-type representations in our list \mathcal{P} come in pairs, i.e. if ρ is on the list then so is $\bar{\rho}$. If f is real-valued, we see (by integrating against R_ρ and \bar{R}_ρ) that $\widehat{f}_\rho = \widehat{f}_{\bar{\rho}}$. Conversely, if $\widehat{f}_\rho = \widehat{f}_{\bar{\rho}}$ holds then $\langle \widehat{f}_\rho, R_\rho(g) \rangle + \langle \widehat{f}_{\bar{\rho}}, R_{\bar{\rho}}(g) \rangle$ is real.

Finally, a representation ρ is of *quaternionic type* if it can be defined over the quaternions in the following sense: d_ρ is even and $\rho(g)$ is comprised of 2×2 blocks, each of which encodes a quaternion by the following relation:

$$a + bi + cj + dk \quad \leftrightarrow \quad \begin{pmatrix} a + bi & c + di \\ -c + di & a - bi \end{pmatrix}.$$

Note that this relation respects quaternion addition and multiplication. Furthermore, quaternion conjugation (negate b, c, d) corresponds to matrix conjugate transpose. If a matrix is comprised of 2×2 blocks of this form, we will call it *block-quaternion*. Now let ρ be of quaternionic type (and assume without loss of generality that ρ takes the above block-quaternion form), and let f be a real function $G \rightarrow \mathbb{R}$. By integrating against R_ρ we see that \widehat{f}_ρ must also be block-quaternion. Conversely, if \widehat{f}_ρ is block-quaternion then $\langle \widehat{f}_\rho, R_\rho(g) \rangle$ is real; to see this, write $\langle \widehat{f}_\rho, R_\rho(g) \rangle = \text{Tr}(\widehat{f}_\rho R_\rho(g)^*)$, note that $\widehat{f}_\rho R_\rho(g)^*$ is block-quaternion, and note that the trace of any quaternion block is real.

2.3.2 Graphical model formulation

As in Section 2.2.2, we take the standpoint of probabilistic inference over a graphical model. Thus we consider the task of estimating $g \in G^n$ from observations that induce a posterior probability factoring into pairwise likelihoods:

$$\Pr(g) \propto \prod_{u < w} \mathcal{L}_{uw}(g_u, g_w). \quad (2.1)$$

We assume that the pair interactions $\mathcal{L}_{uw}(g_u, g_w)$ are in fact a function of $g_u g_w^{-1} \in G$, depending only on the relative orientation of the group elements. This factorization property

amounts to a graphical model for g , with each entry $g_u \in G$ corresponding to a vertex u , and each pair interaction \mathcal{L}_{uw} represented by an edge of the model.

Taking a Peter–Weyl decomposition of $\log(\mathcal{L}_{uw})$ as a function of $g_u g_w^{-1}$ allows us to write:

$$\mathcal{L}_{uw}(g_u, g_w) = \exp \sum_{\rho} \langle Y_{uw}^{\rho}, \rho(g_u g_w^{-1}) \rangle,$$

where ρ runs over all irreducible representations of G . We require coefficients $Y_{uw}^{\rho} \in \mathbb{C}^{d_{\rho} \times d_{\rho}}$ for which this expansion is real-valued (for all $g_u g_w^{-1}$). We also require the symmetry $\mathcal{L}_{uw}(g_u, g_w) = \mathcal{L}_{wu}(g_w, g_u)$, which means $Y_{uw}^{\rho} = (Y_{wu}^{\rho})^*$. Let Y_{ρ} be the $nd_{\rho} \times nd_{\rho}$ matrix with blocks Y_{uw}^{ρ} . For all u , define $\mathcal{L}_{uu} = 1$ and $Y_{uu}^{\rho} = 0$.

The input to our synchronization problem will simply be the coefficients Y_{ρ} . These define a posterior distribution μ on the latent vector g of group elements, and our goal is to approximately recover g up to a global right-multiplication by some group element.

We suppose that the observations are *band-limited*: $Y_{\rho} = 0$ except on a finite set \mathcal{P} of irreducible representations. This will allow us to reduce all Peter–Weyl decompositions to a finite amount of relevant information. We will always exclude the trivial representation from \mathcal{P} : this representation can only contribute a constant factor to each pair likelihood, which then disappears in the normalization, so that without loss of generality we can assume the coefficient of the trivial representation to always be zero.

Many synchronization problems (for instance, sensor localization) have noise on each pairwise measurement, and fit this graphical model formulation well. Other synchronization problems (for instance, cryo-EM) are based on per-vertex measurements with independent randomness; one can derive pairwise information by comparing these measurements, but these pairwise measurements do not have independent noise and do not strictly fit the model described above. Prior work has run into the same issue and achieved strong results nonetheless. Specifically, the *non-unique games* (NUG) model of [18] suggests the optimization

problem of minimizing an objective

$$\sum_{u < w} f_{uw}(g_u g_w^{-1}).$$

Such problems can be artificially placed into our framework by viewing them as the maximum likelihood estimation problem corresponding to the graphical model above, with

$$\mathcal{L}_{uw}(g_u, g_w) = \exp(-\beta f_{uw}(g_u g_w^{-1})) \quad (2.2)$$

for an arbitrary positive ‘inverse temperature’ β . For true probabilistic models, our approach attempts Bayes-optimal inference (minimizing mean squared error), while the NUG approach attempts maximum likelihood estimation which may have higher expected error. In our approach, the maximum likelihood problem can be recovered by scaling up all potentials, i.e. taking the low-temperature limit $\beta \rightarrow \infty$ in (2.2).

One might also formulate a version of our model that allows node potentials, as seen for instance in image segmentation [72] and some community detection problems [152]:

$$\Pr(g) \propto \left(\prod_{u < w} \mathcal{L}_{uw}(g_u, g_w) \right) \left(\prod_u \mathcal{L}_u(g_u) \right),$$

expressing a nontrivial prior or observation on each group element. Although this model is compatible with our methods (so long as the node potentials are also band-limited), we suppress this generality for the sake of readability.

2.3.3 AMP algorithm

We now state our AMP algorithm. The algorithm takes as input the log-likelihood coefficients $Y_\rho \in \mathbb{C}^{nd_\rho \times nd_\rho}$, for each ρ in a finite list \mathcal{P} of irreducibles (which must not contain the trivial representation; also for each complex-type representation ρ in the list, $\bar{\rho}$ must also appear in the list). The algorithm’s state at time t is comprised of the Fourier coefficients $C_\rho^{(t)} \in \mathbb{C}^{nd_\rho \times d_\rho}$, which are updated as follows.

Algorithm 2.3.1 (AMP for synchronization over compact groups).

- For each $\rho \in \mathcal{P}$ and each vertex u , initialize $V_{u,\rho}^{(0)}$ to small random values in \mathbb{C} and initialize $V_{u,\rho}^{(-1)} = 0$.
- Iterate for $1 \leq t \leq T$:
 - For each $\rho \in \mathcal{P}$, set

$$C_{u,\rho}^{(t)} = d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho V_{w,\rho}^{(t-1)} - d_\rho^{-2} |Y_{\text{typ}}^\rho|^2 V_{u,\rho}^{(t-2)} \sum_w (d_\rho I - (V_{w,\rho}^{(t-1)})^* V_{w,\rho}^{(t-1)}),$$

the Fourier coefficients of the estimated posterior log-densities, with Onsager correction. Here $|Y_{\text{typ}}^\rho|^2$ denotes the average squared-norm of the entries of Y_ρ .

- For each u and each $\rho \in \mathcal{P}$, set $V_{u,\rho}^{(t)} = \mathcal{E}_\rho(C_u^{(t)})$, where as in Section 2.2.4, the nonlinear transformation

$$\mathcal{E}_\rho(C) = \int_G R_\rho(g) \exp \left(\sum_{\rho'} \langle C_{\rho'}, R_{\rho'}(g) \rangle \right) dg / \int_G \exp \left(\sum_{\rho'} \langle C_{\rho'}, R_{\rho'}(g) \rangle \right) dg \quad (2.3)$$

takes the Fourier coefficients for a function f on G and returns those of $\exp \circ f$, re-normalized to have integral 1. These $V_{u,\rho}^{(t)}$ are the Fourier coefficients of the estimated posterior densities, truncated to the contribution from irreducibles \mathcal{P} , which suffice for the next iteration.

- Return the posteriors represented by $C_{u,\rho}^{(T)}$, or some rounding of these (e.g. the per-vertex MAP estimate).

This algorithm follows the intuition of Section 2.2, and derivations can be found in Sections 2.5 and 2.6.

Note that each iteration runs in time $O(n^2)$, which is linear in the input matrices. This runtime is due to the matrix–vector products; the rest of the iteration takes $O(n)$ time. We expect $O(\log n)$ iterations to suffice, resulting in a nearly-linear-time algorithm with respect

to the matrix inputs. Some applications may derive from per-vertex observations that are pairwise compared to produce edge observations, hacked into this framework by an abuse of probability; our algorithm then takes nearly-quadratic time with respect to the vertex observations. However, some such per-vertex applications produce matrices with a low-rank factorization $Y_\rho = U_\rho U_\rho^\top$, for which the matrix–vector product can be performed in $O(n)$ time.

2.3.4 Gaussian observation model

Our AMP algorithm handles the general graphical model formulation above, but we will be able to analyze its performance in more detail when restricted to the following concrete Gaussian observation model (which we first introduced in [120]), generalizing the Gaussian models of Section 2.2. First, latent group elements g_u are drawn independently and uniformly from G (from Haar measure). Then for each representation ρ in \mathcal{P} , we observe the $nd_\rho \times nd_\rho$ matrix

$$M_\rho = \frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho.$$

Here X_ρ is the $nd_\rho \times d_\rho$ matrix formed by vertically stacking the $d_\rho \times d_\rho$ matrices $\rho(g_u)$ for all vertices u . λ_ρ is a signal-to-noise parameter for the frequency ρ . The noise W_ρ is a Gaussian random matrix drawn from the GOE, GUE, or GSE, depending on whether ρ is of real type, complex type, or quaternionic type, respectively. In any case, W_ρ is normalized so that each off-diagonal⁶ entry has expected squared-norm 1. To be concrete, in the real case the entries are $\mathcal{N}(0, 1)$ and in the complex case, the real and imaginary parts of each entry are $\mathcal{N}(0, 1/2)$. For the quaternionic case, each 2×2 block encodes a quaternion value $a + bi + cj + dk$ in the usual way (see Section 2.3.1) where a, b, c, d are $\mathcal{N}(0, 1/2)$. The noise matrices W_ρ are independent across representations except when we have a conjugate pair of complex-type representations we draw M_ρ randomly as above and define $M_{\bar{\rho}} = \overline{M_\rho}$ and $\lambda_{\bar{\rho}} = \lambda_\rho$. Note that the normalization is such that the signal term has spectral norm λ_ρ and

⁶The diagonal entries (or diagonal 2×2 quaternion blocks) are unimportant, as discussed previously. One natural choice is to zero out those entries of M_ρ . Another is to take W_ρ with diagonal entries (or blocks) $\mathcal{N}(0, 2)$ (real case), $\mathcal{N}(0, 1)$ (complex case), or $\mathcal{N}(0, 1)I_2$ (quaternion case).

the noise term has spectral norm 2.

The author and others first introduced the above Gaussian synchronization model in [120]. Special cases of this model have been studied previously for synchronization over $\mathbb{Z}/2$ or $U(1)$ with a single frequency [16, 56, 82, 34] (previously implicit in [141]). In fact, [56] derives AMP for the $\mathbb{Z}/2$ case and proves that it is information-theoretically optimal. The idea of optimizing objective functions that have information on multiple frequencies comes from [18].

In Appendix A.1, we show how the Gaussian observation model fits into the graphical model formulation by deriving the corresponding coefficient matrices Y_ρ . In particular, we show that $Y_\rho = d_\rho \lambda_\rho M_\rho$, a scalar multiple of the observed Gaussian matrix.

2.3.5 Representation theory of some common examples

In this section we discuss the representation theory of a few central examples, namely \mathbb{Z}/L , $U(1)$, and $SO(3)$, and connect the general formalism back to the examples of Section 2.2.

Representations of \mathbb{Z}/L and $U(1)$. The irreducible representations of these groups are one-dimensional, described by the discrete Fourier transform and the Fourier series, respectively. $U(1)$ has frequencies indexed by $k \in \mathbb{Z}$, given by $\rho(g) = g^k$ where $g \in U(1)$ (i.e. a unit-norm complex number). All of these representations are of complex type. We will say “ $U(1)$ with K frequencies” to refer to the frequencies $1, \dots, K$ along with their conjugates, the frequencies $-1, \dots, -K$. Similarly, if we identify \mathbb{Z}/L with the complex L th roots of unity, we have frequencies defined the same way as above, except to avoid redundancy we restrict the range of k as follows. If L is odd, we allow $k \in \{1, 2, \dots, (L-1)/2\}$ along with their conjugates (negations). If L is even, we have complex-type representations $k \in \{1, 2, \dots, L/2 - 1\}$ (along with their conjugates), plus an additional real-type representation $k = L/2$. Again, “ \mathbb{Z}/L with K frequencies” means we take frequencies $1, \dots, K$ along with their conjugates (when applicable).

For the case of $\mathbb{Z}/2$ we can now see how the tanh function from the AMP algorithm of

Section 2.2.2 arises as a special case of the nonlinear transformation \mathcal{E} occurring in AMP. The only nontrivial representation of $\mathbb{Z}/2$ is the ‘parity’ representation in which -1 acts as -1 . In this context, \mathcal{E} will first input the Fourier series of a log-density with respect to the uniform measure on $\{\pm 1\}$:

$$\log \frac{d\mu_u}{dx} + \text{const.} = cx$$

and evaluate this at ± 1 to obtain the values $\pm c$. We then compute the exponential of this at each point to obtain the un-normalized density of e^{-c} at -1 and e^c at 1 . Normalizing, the density values are $e^{-c}/(e^{-c} + e^c)$ and $e^c/(e^{-c} + e^c)$, so that the new parity coefficient is

$$\mathcal{E}_{\text{parity}}(c) = \frac{e^c - e^{-c}}{e^c + e^{-c}} = \tanh c.$$

Representations of $SO(3)$. This group has one irreducible representation ρ_k of each odd dimension $d_k = 2k + 1$; the $k = 0$ representation is the trivial representation $\rho_0(g) = 1$, and the $k = 1$ representation is the standard representation of $SO(3)$ as rotations of three-dimensional space. All of these representations are of real type, and may be described as the action of rotations on the $2k + 1$ -dimensional space of homogeneous degree k spherical harmonics. Frequently in the literature (for instance in molecular chemistry), a complex basis for the spherical harmonics is given, and the representation matrices are the complex-valued Wigner D-matrices; however, the representation can be defined over the reals, as is demonstrated by any real orthogonal basis for the spherical harmonics. See e.g. Section II.5 of [40] for a more detailed account. As in the cases above, we will often refer to synchronization problems over “ $SO(3)$ with K frequencies”, in which the observations are assumed to be band-limited to the first K nontrivial irreducibles with $1 \leq k \leq K$.

2.4 Experimental results

We present a brief empirical exploration of the statistical performance of AMP in various settings, and as compared to other algorithms.

In Figure 2-3 we compare the performance of the spectral method, projected power




Figure 2-3: Comparison of iterative algorithms for Gaussian $\mathbb{Z}/2$ synchronization. The horizontal axis represents the signal-to-noise ratio λ , and the vertical axis depicts the log-error $\ln(1 - |\langle x, \hat{x} \rangle|/n)$ where $x \in \{\pm 1\}^n$ is the ground truth and $\hat{x} \in \{\pm 1\}^n$ is the (rounded) output of the algorithm. The four curves are projected power iteration (green), soft-threshold power iteration (red), spectral method (blue), and AMP (black). Each data point is an average of 200 trials with $n = 2000$ vertices.

iteration, soft-threshold power iteration without an Onsager correction, and full AMP (see Sections 2.2.1 and 2.2.2) for Gaussian $\mathbb{Z}/2$ synchronization. The spectral method achieves the optimal threshold of $\lambda = 1$ as to when nontrivial recovery is possible, but does not achieve the optimal correlation afterwards. The projected power method appears to asymptotically achieve the optimal correlation as $\lambda \rightarrow \infty$, but performs worse than the spectral method for small λ . Soft-thresholding offers a reasonable improvement on this, but the full AMP algorithm strictly outperforms all other methods. This reflects the optimality result of [56] and highlights the necessity for the Onsager term. The gains are fairly modest in this setting, but increase with more complicated synchronization problems.

Figures 2-4 and 2-5 compare the performance of AMP on Gaussian $U(1)$ synchronization with multiple frequencies; see Section 2.2.4 for the model. In sharp contrast to spectral methods, which offer no reasonable way to couple the frequencies together, AMP produces an estimate that is orders of magnitude more accurate than what is possible with a single




Figure 2-4: Gaussian $U(1)$ synchronization with K frequencies; from bottom to top, $K = 1, \dots, 6$. The signal-to-noise ratios λ_k are all equal, with the common value given by the horizontal axis. Each curve depicts the correlation $|\langle x, \hat{x} \rangle|/n$ between the ground truth and the AMP estimate. Each data point is an average of 50 trials with $n = 1000$ vertices.




Figure 2-5: Here the vertical axis depicts the log-error $\ln(1 - |\langle x, \hat{x} \rangle|/n)$. From top to bottom: $K = 1, \dots, 6$.

frequency.

In Figures 2-6 and 2-7, we see similar results over $SO(3)$, under the Gaussian model of Section 2.3.4. This also demonstrates that AMP is an effective synchronization algorithm for more complicated, non-abelian Lie groups.

This ability to exploit multiple frequencies represents a promising step toward improved algorithms for cryo-electron microscopy, which may be viewed as a synchronization problem over $SO(3)$. Some previous approaches to this problem effectively band-limit the observations to a single frequency and then apply a spectral method [142], and the experiments in Figures 2-4–2-7 demonstrate that our algorithm stands a compelling chance of achieving a higher-quality reconstruction.

We remark that some numerical issues arise when computing the nonlinear transformation \mathcal{E} in our AMP algorithm, which involves integration over the group. Our implementation of \mathcal{E} for $U(1)$ and $SO(3)$ is based on evaluating each log-density on a discretization of the group,




Figure 2-6: Gaussian $SO(3)$ synchronization with K frequencies; from bottom to top, $K = 1, 2, 3$. The signal-to-noise ratios λ_k are all equal, with the common value given by the horizontal axis. Each curve depicts the squared correlation $\|X^\top \hat{X}\|_F / (n\sqrt{3})$ between the ground truth and the AMP estimate. Here X and \hat{X} are $3n \times n$ matrices where each 3×3 block encodes an element of $SO(3)$ via the standard representation (rotation matrices). Each data point is an average of 5 trials with $n = 100$ vertices.




Figure 2-7: Here the vertical axis depicts the log-error $\ln(1 - \|X^\top \hat{X}\|_F / (n\sqrt{3}))$. From top to bottom: $K = 1, 2, 3$.

taking the pointwise exponential, and thus approximating each integral by a discrete sum. This approach is somewhat crude but appears to work adequately in our experiments; there is undoubtedly room for this numerical procedure to be improved. More sophisticated methods may be necessary to obtain adequate results on any higher-dimensional Lie groups. Note also that if the vertex posterior in question is extremely concentrated near a point, the numerical value of each integral will depend significantly on whether this spike lies near a discretization point; however, this should affect both the numerator and denominator integrals in (2.3) by approximately equal factors, so as to have a minimal effect on the normalized value of \mathcal{E}_ρ .

2.5 Derivation of AMP from belief propagation

In this section we derive the general AMP algorithm of Section 2.3 starting from belief propagation, similarly to [64]. We begin with the belief propagation update step (see e.g. [105]), writing messages $\mu_{u \rightarrow v}^{(t)}$ as densities with respect to Haar measure:

$$\frac{d\mu_{u \rightarrow v}^{(t)}}{dg_u} = \frac{1}{Z_{u \rightarrow v}^{(t)}} \prod_{w \neq u, v} \int_G \mathcal{L}_{uw}(g_u, g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w.$$

Here t denotes the timestep and $Z_{u \rightarrow v}^{(t)}$ is the appropriate normalization constant. Expand this (positive) probability density as the exponential of an L^2 function, expressed as a Peter–Weyl expansion:

$$\frac{d\mu_{u \rightarrow v}^{(t)}}{dg_u} = \exp \sum_{\rho ab} C_{u \rightarrow v, \rho ab}^{(t)} \overline{R_{\rho ab}(g_u)}.$$

We can extract these Fourier coefficients $C_{u \rightarrow v, \rho ab}^{(t)}$ by integrating against the basis functions above. Assume that ρ is not the trivial representation; then:

$$\begin{aligned} C_{u \rightarrow v, \rho ab}^{(t)} &= \int_G R_{\rho ab}(g_u) \log \frac{d\mu_{u \rightarrow v}^{(t-1)}}{dg_u} dg_u \\ &= \int_G R_{\rho ab}(g_u) \log \left(\frac{1}{Z_{u \rightarrow v}^{(t-1)}} \prod_{w \neq u, v} \int_G \mathcal{L}_{uw}(g_u, g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \right) dg_u \\ &= \int_G R_{\rho ab}(g_u) \sum_{w \neq u, v} \log \left(\int_G \exp \left(\sum_{\rho'} \langle Y_{uw}^{\rho'}, \rho'(g_u g_w^{-1}) \rangle \right) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \right) dg_u. \end{aligned}$$

As the $Y_{uw}^{\rho'}$ are small, we can pass to a linear expansion about these, incurring $o(1)$ error as $n \rightarrow \infty$:

$$\begin{aligned} &\approx \int_G R_{\rho ab}(g_u) \sum_{w \neq u, v} \left(\log \int_G \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w + \sum_{\rho'} \int_G \langle Y_{uw}^{\rho'}, \rho'(g_u g_w^{-1}) \rangle \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \bigg/ \int_G \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \right) dg_u \\ &= \sum_{w \neq u, v} \sum_{\rho'} \int_G \int_G R_{\rho ab}(g_u) \langle Y_{uw}^{\rho'}, \rho'(g_u g_w^{-1}) \rangle \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w dg_u. \end{aligned}$$

To progress further, we will expand the middle factor of the integrand:

$$\begin{aligned}\left\langle Y_{uw}^{\rho'}, \rho'(g_u g_w^{-1}) \right\rangle &= \left\langle Y_{uw}^{\rho'}, \rho'(g_u) \rho'(g_w)^* \right\rangle \\ &= \sum_{a'b'c'} Y_{uw,a'b'}^{\rho'} \overline{\rho'(g_u)_{a'c'}} \rho'(g_w)_{b'c'}.\end{aligned}$$

Returning to the previous derivation:

$$\begin{aligned}C_{u \rightarrow v, \rho ab}^{(t)} &= \sum_{w \neq u, v} \sum_{\rho'} \sum_{a'b'c'} \int_G \int_G R_{\rho ab}(g_u) Y_{uw,a'b'}^{\rho'} \overline{\rho'(g_u)_{a'c'}} \rho'(g_w)_{b'c'} \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w dg_u \\ &= \sum_{w \neq u, v} \sum_{\rho'} \sum_{a'b'c'} Y_{uw,a'b'}^{\rho'} \int_G R_{\rho ab}(g_u) \overline{\rho'(g_u)_{a'c'}} dg_u \cdot \int_G \rho'(g_w)_{b'c'} \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \\ &= d_{\rho'}^{-1} \sum_{w \neq u, v} \sum_{\rho'} \sum_{a'b'c'} Y_{uw,a'b'}^{\rho'} \int_G R_{\rho ab}(g_u) \overline{\rho'(g_u)_{a'c'}} dg_u \cdot \int_G R_{\rho' b' c'}(g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \\ &= d_{\rho'}^{-1} \sum_{w \neq u, v} \sum_{\rho'} \sum_{a'b'c'} Y_{uw,a'b'}^{\rho'} \delta_{\rho, \rho'} \delta_{a, a'} \delta_{b, b'} \int_G R_{\rho' b' c'}(g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w \\ &= d_{\rho}^{-1} \sum_{w \neq u, v} \sum_{b'} Y_{uw, ab'}^{\rho} \int_G R_{\rho b' b}(g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w.\end{aligned}$$

In matrix form,

$$C_{u \rightarrow v, \rho}^{(t)} = d_{\rho}^{-1} \sum_{w \neq u, v} Y_{uw}^{\rho} \int_G R_{\rho}(g_w) \frac{d\mu_{w \rightarrow u}^{(t-1)}}{dg_w} dg_w.$$

Let $\mathcal{E} : \bigoplus_{\rho} \mathbb{C}^{d_{\rho} \times d_{\rho}} \rightarrow \bigoplus_{\rho} \mathbb{C}^{d_{\rho} \times d_{\rho}}$ denote the transformation from the nontrivial Fourier coefficients $C_{u \rightarrow v, \rho}$ of $\log \frac{d\mu_{u \rightarrow v}^{(t)}}{dg_u}$ to the Fourier coefficients of $\frac{d\mu_{u \rightarrow v}^{(t)}}{dg_u}$. Then we have

$$C_{u \rightarrow v, \rho}^{(t)} = d_{\rho}^{-1} \sum_{w \neq u, v} Y_{uw}^{\rho} \mathcal{E}_{\rho}(C_{w \rightarrow u}^{(t-1)}).$$

The map \mathcal{E} amounts to exponentiation in the evaluation basis, except that the trivial Fourier coefficient is missing from the input, causing an unknown additive shift. This corresponds to an unknown multiplicative shift in the output, which we correct for by noting that $\frac{d\mu_{u \rightarrow v}^{(t)}}{dg_u}$ should normalize to 1. Thus \mathcal{E} amounts to exponentiation followed by normalization.

Explicitly, we can let

$$I_{\rho ab}(C) = \int_G R_{\rho ab}(g) \exp \left(\sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} \right) dg.$$

Then $\mathcal{E}_{\rho ab}(C) = I_{\rho ab}(C)/I_{\text{triv}}(C)$ where triv denotes the trivial representation $R_{\text{triv}}(g) = 1$.

2.5.1 Onsager correction

In this section we complete the derivation of AMP by replacing the non-backtracking nature by an *Onsager correction* term, reducing the number of messages from n^2 to n . This is similar to the derivation in Appendix A of [23].

In order to remove the non-backtracking nature of the AMP recurrence, let us define

$$\begin{aligned} C_{u,\rho}^{(t)} &= d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho \mathcal{E}_\rho(C_{w \rightarrow u}^{(t-1)}) \\ &= C_{u \rightarrow v, \rho}^{(t)} + \delta_{u \rightarrow v, \rho}^{(t)}, \end{aligned}$$

where $\delta_{u \rightarrow v, \rho}^{(t)} = d_\rho^{-1} Y_{uv}^\rho \mathcal{E}_\rho(C_{v \rightarrow u}^{(t-1)})$. Then, substituting $C_{w \rightarrow u}^{(t-1)} = C_w^{(t-1)} - \delta_{w \rightarrow u}^{(t-1)}$, we have

$$\begin{aligned} C_{u,\rho}^{(t)} &= d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho \mathcal{E}_\rho \left(\{C_w^{(t-1)} - d_\rho^{-1} Y_{wu}^{\rho'} \mathcal{E}_{\rho'}(C_{u \rightarrow w}^{(t-2)})\}_{\rho'} \right) \\ &\approx d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho \mathcal{E}_\rho(C_w^{(t-1)}) - d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho D\mathcal{E}_\rho|_{C_w^{(t-1)}} \left[\{d_\rho^{-1} Y_{wu}^{\rho'} \mathcal{E}_{\rho'}(C_{u \rightarrow w}^{(t-2)})\}_{\rho'} \right] \end{aligned}$$

where D denotes the total derivative

$$\approx d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho \mathcal{E}_\rho(C_w^{(t-1)}) - d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho D\mathcal{E}_\rho|_{C_w^{(t-1)}} \left[\{d_\rho^{-1} Y_{wu}^{\rho'} \mathcal{E}_{\rho'}(C_{u \rightarrow w}^{(t-2)})\}_{\rho'} \right].$$

Under the assumption that Y consists of per-edge $O(n^{-1/2})$ noise and $O(n^{-1})$ signal, the error incurred in these two steps should be $o(1)$. We thus reach an entirely non-backtracking

recurrence where the first term is a message-passing step and the second term is the so-called *Onsager correction*. It remains to simplify this. We focus on a single matrix coefficient of the correction:

$$\begin{aligned}
\text{Ons}_{u,\rho ab}^{(t)} &= d_\rho^{-1} \sum_{w \neq u} \sum_c Y_{uw,ac}^\rho D\mathcal{E}_\rho|_{C_w^{(t-1)}} \left[\{d_{\rho'}^{-1} Y_{wu}^{\rho'} \mathcal{E}_{\rho'}(C_u^{(t-2)})\} \right]_{cb} \\
&= d_\rho^{-1} \sum_{w \neq u} \sum_{\rho' c e f} Y_{uw,ac}^\rho \frac{\partial \mathcal{E}_{\rho cb}}{\partial C_{\rho' ef}} \Big|_{C_w^{(t-1)}} (d_{\rho'}^{-1} Y_{wu}^{\rho'} \mathcal{E}_{\rho'}(C_u^{(t-2)}))_{ef} \\
&= d_\rho^{-1} \sum_{w \neq u} \sum_{\rho' c e f h} Y_{uw,ac}^\rho \frac{\partial \mathcal{E}_{\rho cb}}{\partial C_{\rho' ef}} \Big|_{C_w^{(t-1)}} d_{\rho'}^{-1} Y_{wu,eh}^{\rho'} \mathcal{E}_{\rho' hf}(C_u^{(t-2)}).
\end{aligned}$$

As in the derivation of [23], we now make a few simplifying approximations which we expect to be correct in the large- n limit. We expect sufficiently little correlation between the quantity $Y_{uw,ac}^\rho Y_{wu,eh}^{\rho'}$ and the other factors that, by the law of large numbers (since there are many terms in the sum), we can replace this quantity by its expectation. We assume, as in the Gaussian model, that the noise component of Y is independent across edges, across frequencies, and across matrix entries (other than explicit dependencies such as symmetry). It follows that the only terms in which $Y_{uw,ac}^\rho Y_{wu,eh}^{\rho'}$ has significantly nonzero mean is when $\rho = \rho'$, $a = h$, and $c = e$. In this case we have $Y_{uw,ac}^\rho Y_{wu,eh}^{\rho'} = |Y_{uw,ac}^\rho|^2$. Further replacing this by its expected value (which we assume depends only on ρ) yields

$$\text{Ons}_{u,\rho ab}^{(t)} = d_\rho^{-2} |Y_{\text{typ}}^\rho|^2 \sum_f \mathcal{E}_{\rho a f}(C_u^{(t-2)}) \sum_{w \neq u} \sum_c \frac{\partial \mathcal{E}_{\rho cb}}{\partial C_{\rho c f}} \Big|_{C_w^{(t-1)}}$$

where $|Y_{\text{typ}}^\rho|^2$ denotes the average squared-norm of the entries of Y_ρ .

An interlude, understanding derivatives of \mathcal{E} :

$$\frac{\partial I_{\rho ab}}{\partial C_{\rho' cd}} = \int_G R_{\rho ab}(g) \overline{R_{\rho' cd}(g)} \exp \sum_{\rho'' a' b'} C_{\rho'' a' b'} \overline{R_{\rho'' a' b'}(g)} dg.$$

In particular,

$$\begin{aligned}\frac{\partial I_{\text{triv}}}{\partial C_{\rho ab}} &= \int_G \overline{R_{\rho ab}(g)} \exp \sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} dg \\ &= \overline{I_{\rho ab}}.\end{aligned}$$

Note the following convenient identity:

$$\begin{aligned}\sum_c \frac{\partial I_{\rho cb}}{\partial C_{\rho cf}} &= \int_G \left(\sum_c R_{\rho cb}(g) \overline{R_{\rho cf}(g)} \right) \exp \sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} dg \\ &= d_\rho \int_G \left(\sum_c \rho(g)_{cb} \rho(g^{-1})_{fc} \right) \exp \sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} dg \\ &= d_\rho \int_G \rho(g^{-1}g)_{fb} \exp \sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} dg \\ &= d_\rho \delta_{bf} \int_G \exp \sum_{\rho' a' b'} C_{\rho' a' b'} \overline{R_{\rho' a' b'}(g)} dg \\ &= d_\rho \delta_{bf} I_{\text{triv}}(C).\end{aligned}$$

Recalling that $\mathcal{E}_{\rho ab}(C) = I_{\rho ab}(C)/I_{\text{triv}}(C)$, we have

$$\begin{aligned}\sum_c \frac{\partial \mathcal{E}_{\rho cb}}{\partial C_{\rho cf}} &= \frac{I_{\text{triv}} \sum_c \frac{\partial I_{\rho cb}}{\partial C_{\rho cf}} - \sum_c I_{\rho cb} \overline{I_{\rho cf}}}{I_{\text{triv}}^2} \\ &= d_\rho \delta_{bf} - \sum_c \mathcal{E}_{\rho cb}(C) \overline{\mathcal{E}_{\rho cf}(C)} \\ &= (d_\rho I - \mathcal{E}_\rho(C)^* \mathcal{E}_\rho(C))_{fb}.\end{aligned}$$

Thus we obtain the following form for the Onsager correction:

$$\text{Ons}_{u,\rho}^{(t)} = d_\rho^{-2} |Y_{\text{typ}}^\rho|^2 \mathcal{E}_\rho(C_u^{(t-2)}) M_\rho^{(t)}, \quad M_\rho^{(t)} = \sum_w d_\rho I - \mathcal{E}_\rho(C_w^{(t-1)})^* \mathcal{E}_\rho(C_w^{(t-1)}),$$

with each AMP iteration reading as

$$C_{u,\rho}^{(t)} = d_\rho^{-1} \sum_{w \neq u} Y_{uw}^\rho \mathcal{E}_\rho(C_w^{(t-1)}) - \text{Ons}_{u,\rho}^{(t)}.$$

2.6 MMSE derivation and state evolution

The goal of this section is to derive the state evolution equations that govern the behavior of AMP on the Gaussian synchronization model of Section 2.3.4 (in the large n limit). Along the way, we will give an alternative derivation of the algorithm (excluding the Onsager term) which shows that the nonlinear function \mathcal{E} has an interpretation as an MMSE (minimum mean squared error) estimator. This derivation is similar to [56] and based on ideas first introduced by [63]. We do not give a proof that the state evolution equations derived here are correct (i.e. that AMP obeys them) but we will argue for their correctness in Section 2.7.

2.6.1 MMSE estimator

We begin by defining a ‘scalar’ problem: a simplification of the Gaussian synchronization model where we attempt to recover a single group element from noisy measurements. We will be able to analyze the Gaussian synchronization model by connection to this simpler model. (This is the idea of *single letterization* from information theory.) Suppose there is an unknown group element g drawn uniformly from G (Haar measure) and for each irreducible representation ρ in our list \mathcal{P} we are given a measurement $u_\rho = \mu_\rho \rho(g) + \sigma_\rho z_\rho$ (for some constants μ_ρ, σ_ρ). Here z_ρ is a $d_\rho \times d_\rho$ non-symmetric matrix of Gaussian entries (real, complex, or block-quaternion, depending on the type of ρ) with all entries (or blocks) independent and each entry normalized to have expected squared-norm 1. (Note that z_ρ is the same as an off-diagonal block of the matrix W_ρ from Section 2.3.4.) For ρ of complex type, we only get a measurement u_ρ for one representation in each conjugate pair, and define $u_{\bar{\rho}} = \overline{u_\rho}$. The MMSE estimator for $\rho(g)$ (minimizing the matrix mean squared error $\mathbb{E} \|\widehat{\rho(g)} - \rho(g)\|_F^2$) is simply the conditional expectation

$$\begin{aligned}
\mathbb{E} \left[\rho(g) \middle| \{u_q\}_q \right] &= \int_{h \in G} \rho(h) \exp \left(- \sum_q \frac{1}{2\sigma_q^2} \|u_q - \mu_q q(h)\|_F^2 \right) / \int_{h \in G} \exp \left(- \sum_q \frac{1}{2\sigma_q^2} \|u_q - \mu_q q(h)\|_F^2 \right) \\
&= \int_{h \in G} \rho(h) \exp \left(\sum_q \frac{\mu_q}{\sigma_q^2} \langle u_q, q(h) \rangle \right) / \int_{h \in G} \exp \left(\sum_q \frac{\mu_q}{\sigma_q^2} \langle u_q, q(h) \rangle \right) \\
&\equiv \mathcal{F}_\rho \left(\left\{ \frac{\mu_q}{\sigma_q^2} u_q \right\}_q \right)
\end{aligned}$$

where

$$\mathcal{F}_\rho(\{w_q\}_q) = \int_{h \in G} \rho(h) \exp \left(\sum_q \langle w_q, q(h) \rangle \right) / \int_{h \in G} \exp \left(\sum_q \langle w_q, q(h) \rangle \right).$$

Here q ranges over irreducible representations in our list \mathcal{P} (which includes both q and \bar{q} for representations of complex type). The likelihoods used in the above computation are derived similarly to those in Appendix A.1. We recognize \mathcal{F} as a rescaling of the function \mathcal{E} from the AMP update step.

2.6.2 AMP update step

Consider the Gaussian observation model $M_\rho = \frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho$ from Section 2.3.4. Similarly to [56], the MMSE-AMP update step (without Onsager term) is

$$U_\rho^{t+1} = M_\rho \mathcal{F}_\rho \left(\left\{ \frac{\mu_q^t}{(\sigma_q^t)^2} U_q^t \right\}_q \right)$$

where t indicates the timestep and $\mu_\rho^t, \sigma_\rho^t$ will be defined based on state evolution below. Here the AMP state U_ρ^t is $nd_\rho \times d_\rho$ with a $d_\rho \times d_\rho$ block for each vertex. \mathcal{F}_ρ is applied to each of these blocks separately. We will motivate this AMP update step below, but notice its similarity to the MMSE estimator above.

2.6.3 State evolution

The idea of state evolution is that the AMP iterates can be approximately modeled as ‘signal’ plus ‘noise’ [63]. Namely, we postulate that $U_\rho^t = \mu_\rho^t X_\rho + \sigma_\rho^t Z_\rho$ for some constants μ_t, σ_t , where Z_ρ is a $nd_\rho \times d_\rho$ Gaussian noise matrix with each $d_\rho \times d_\rho$ block independently distributed like z_ρ (from the scalar model) with $Z_{\bar{\rho}} = \overline{Z_\rho}$ for conjugate pairs. Recall X_ρ has blocks $\rho(g_u)$, the ground truth. Note that this sheds light on the AMP update step above: at each iteration we are given U_q^t , a noisy copy of the ground truth; the first thing we do is to apply the MMSE estimator entrywise.

We will derive a recurrence for how the parameters μ_ρ and σ_ρ change after one iteration. To do this, we assume that the noise W_ρ is independent from Z_ρ at each timestep. This assumption is far from true; however, it turns out that AMP’s Onsager term corrects for this (e.g. [23]). In other words, we derive state evolution by omitting the Onsager term and assuming independent noise at each timestep. Then if we run AMP (with the Onsager term and the same noise at each timestep), it behaves according to state evolution. We now derive state evolution:

$$\begin{aligned} U_\rho^{t+1} &= M_\rho \mathcal{F}_\rho \left(\left\{ \frac{\mu_q^t}{(\sigma_q^t)^2} U_q^t \right\}_q \right) \\ &= \left(\frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho \right) \mathcal{F}_\rho \left(\left\{ \frac{\mu_q^t}{(\sigma_q^t)^2} (\mu_q^t X_q + \sigma_q^t Z_q) \right\}_q \right) \\ &= \left(\frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho \right) \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right) \end{aligned}$$

where $\gamma_q^t = \left(\frac{\mu_q^t}{\sigma_q^t} \right)^2$

$$= \frac{\lambda_\rho}{n} X_\rho X_\rho^* \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right) + \frac{1}{\sqrt{nd_\rho}} W_\rho \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right).$$

First focus on the signal term:

$$\frac{\lambda_\rho}{n} X_\rho X_\rho^* \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right) \approx \lambda_\rho X_\rho \mathbb{E}_{g, z_q} \left[\rho(g)^* \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q \right\}_q \right) \right]$$

where g is drawn from Haar measure on G , and z_q is a non-symmetric Gaussian matrix of the appropriate type (as in Section 2.6.1). Define $A_\rho^t \in \mathbb{C}^{d_\rho \times d_\rho}$ to be the second matrix in the expression above:

$$A_\rho^t \equiv \mathbb{E}_{g, z_q} \left[\rho(g)^* \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q \right\}_q \right) \right].$$

We will see shortly that A_ρ^t is a multiple $a_\rho^t \in \mathbb{R}$ of the identity and so we can now write the signal term as $\lambda_\rho a_\rho^t X_\rho$. Therefore our new signal parameter is $\mu_\rho^{t+1} = \lambda_\rho a_\rho^t$.

We take a short detour to state some properties of A_ρ^t , which we prove in Appendix A.2.

Lemma 2.6.1. *A_ρ^t is a real multiple of the identity: $A_\rho^t = a_\rho^t I_{d_\rho}$ for some $a_\rho^t \in \mathbb{R}$. Furthermore, we have the following equivalent formulas for a_ρ^t :*

$$(i) \quad \mathbb{E}_{g, z_q} \left[\rho(g)^* \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q \right\}_q \right) \right]$$

$$(ii) \quad \mathbb{E}_{g, z_q} [\mathcal{F}_\rho(\cdots)^* \mathcal{F}_\rho(\cdots)]$$

$$(iii) \quad \mathbb{E}_{z_q} \left[\mathcal{F}_\rho \left(\left\{ \gamma_q^t I_{d_q} + \sqrt{\gamma_q^t} z_q \right\}_q \right) \right]$$

$$(iv) \quad \mathbb{E}_{z_q} [\mathcal{F}_\rho(\cdots)^* \mathcal{F}_\rho(\cdots)]$$

where \cdots denotes the argument to \mathcal{F}_ρ from the previous line.

Returning to state evolution, we now focus on the noise term:

$$\frac{1}{\sqrt{nd_\rho}} W_\rho \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right).$$

Each entry of this $nd_\rho \times d_\rho$ matrix is Gaussian. The variance (expected squared-norm) of

entry (i, j) is (approximately)

$$\begin{aligned}
\frac{1}{nd_\rho} \sum_{k=1}^{nd_\rho} \left| \mathcal{F}_\rho \left(\left\{ \gamma_q^t X_q + \sqrt{\gamma_q^t} Z_q \right\}_q \right) \right|_{k,j}^2 &\approx \frac{1}{d_\rho} \mathbb{E}_{g, z_q} \sum_{k=1}^{d_\rho} \left| \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q \right\}_q \right) \right|_{k,j}^2 \\
&= \frac{1}{d_\rho} \mathbb{E} [\mathcal{F}_\rho(\cdots)^* \mathcal{F}_\rho(\cdots)]_{jj} \\
&= \frac{1}{d_\rho} (A_\rho^t)_{jj} \\
&= \frac{1}{d_\rho} a_\rho^t.
\end{aligned}$$

We therefore have the new noise parameter $(\sigma_\rho^{t+1})^2 = \frac{a_\rho^t}{d_\rho}$.

To summarize, we now have the state evolution recurrence $\mu_\rho^{t+1} = \lambda_\rho a_\rho^t$ and $(\sigma_\rho^{t+1})^2 = \frac{a_\rho^t}{d_\rho}$.

2.6.4 Simplified AMP update step

Note that the state evolution recurrence implies the relation

$$\frac{\mu_\rho^{t+1}}{(\sigma_\rho^{t+1})^2} = d_\rho \lambda_\rho.$$

Provided our initial values of μ_ρ, σ_ρ satisfy this relation (which can always be arranged by scaling the initial U_ρ appropriately), our AMP update step (without Onsager term) becomes

$$U_\rho^{t+1} = M_\rho \mathcal{F}_\rho \left(\{d_\rho \lambda_\rho U_q^t\}_q \right).$$

This is convenient because we can implement AMP without keeping track of the state evolution parameters $\mu_\rho^t, \sigma_\rho^t$. Also note that this variant of AMP matches the original derivation after the rescaling $C_\rho^t = \sqrt{d_\rho \lambda_\rho} U_\rho^t$ (and excluding the Onsager term).

2.6.5 Reduction to single parameter (per frequency)

We will rewrite the state evolution recurrence in terms of a single parameter per frequency. This parameter will be γ_ρ^t , which was introduced earlier: $\gamma_\rho^t = \left(\frac{\mu_\rho^t}{\sigma_\rho^t} \right)^2$. Recall the state

evolution recurrence $\mu_\rho^{t+1} = \lambda_\rho a_\rho^t$ and $(\sigma_\rho^{t+1})^2 = \frac{a_\rho^t}{d_\rho}$. We therefore have the update step

$$\gamma_\rho^{t+1} = \left(\frac{\mu_\rho^{t+1}}{\sigma_\rho^{t+1}} \right)^2 = \frac{(\lambda_\rho a_\rho^t)^2}{a_\rho^t/d_\rho} = d_\rho \lambda_\rho^2 a_\rho^t.$$

Using part (iii) of Lemma 2.6.1 we can write this as:

$$\gamma_\rho^{t+1} = \lambda_\rho^2 \mathbb{E}_{z_q} \text{Tr } \mathcal{F}_\rho \left(\left\{ \gamma_q^t I_{d_q} + \sqrt{\gamma_q^t z_q} \right\}_q \right). \quad (2.4)$$

This is the final form of our state evolution recurrence. The relation between $\mu_\rho, \sigma_\rho, \gamma_\rho$ can be summarized as $\gamma_\rho = d_\rho \lambda_\rho \mu_\rho = d_\rho^2 \lambda_\rho^2 \sigma_\rho^2$.

We expect that the state evolution recurrence (2.4) exactly governs the behavior of AMP in the large n limit. Although the derivation above was heuristic, we discuss its correctness in Section 2.7. There is a caveat regarding how it should be initialized (see Section 2.7) but in practice we can imagine the initial γ value is a small random vector. (Note that the initialization $\gamma = \vec{0}$ is problematic because state evolution will never leave zero.) We expect that state evolution converges to some fixed point of the recurrence. Some complications arise if there are multiple fixed points (see Section 2.8) but we expect there to be a unique fixed point that is reached from any small initialization. This fixed point γ^* describes the output of AMP in the sense that (following the postulate of state evolution) the final AMP iterate is approximately distributed as $U_\rho \approx \mu_\rho^* X_\rho + \sigma_\rho^* Z_\rho$, which in terms of γ^* is (up to scaling) $U_\rho \approx \gamma_\rho^* X_\rho + \sqrt{\gamma_\rho^*} Z_\rho$. (See [23] for the precise sense in which we expect this to be true.) Note that one can use this to translate a γ^* value into any measure of performance, such as MSE. This gives an exact asymptotic characterization of the performance of AMP for any set of λ_ρ values. The most prominent feature of AMP's performance is the threshold at $\lambda = 1$, which we derive in the next section.

One can check that our state evolution recurrence matches the Bayes-optimal cavity and replica predictions of [82] for $\mathbb{Z}/2$ and $U(1)$ with one frequency. Indeed, we expect AMP to be statistically optimal in these settings (and many others too; see Section 2.8), and this has been proven rigorously for $\mathbb{Z}/2$ [56].

2.6.6 Threshold at $\lambda = 1$

In this section we use the state evolution occurrence to derive the threshold above which AMP achieves nontrivial recovery. In particular, if $\lambda_\rho < 1$ for all frequencies ρ then the AMP fixed point γ^* is equal to the zero vector and so AMP gives trivial performance (random guessing) in the large n limit. On the other hand, if $\lambda_\rho > 1$ for at least one frequency ρ then γ^* is nonzero and AMP achieves nontrivial recovery.

The zero vector is always a fixed point of state evolution. Whether or not AMP achieves nontrivial performance depends on whether the zero vector is a stable or unstable fixed point. Therefore we consider the regime where γ_ρ is small for all ρ . When the input $\{w_q\}_q$ to \mathcal{F}_ρ is small, we can approximate \mathcal{F}_ρ by its linearization.

$$\mathcal{F}_\rho(\{w_q\}_q) \approx \int_h \rho(h) \left[1 + \sum_q \langle w_q, q(h) \rangle \right] = \int_h \rho(h) \sum_q \langle w_q, q(h) \rangle$$

and so

$$\begin{aligned} \mathcal{F}_\rho(\{w_q\}_q)_{ab} &\approx \int_h \rho(h)_{ab} \sum_{qcd} w_{qcd} \overline{q(h)_{cd}} \\ &= \sum_{qcd} w_{qcd} \int_h \rho(h)_{ab} \overline{q(h)_{cd}} \\ &= \sum_{qcd} w_{qcd} \frac{1}{d_\rho} \delta_{\rho ab, qcd} \\ &= \frac{w_{\rho ab}}{d_\rho} \end{aligned}$$

which means $\mathcal{F}_\rho(\{w_q\}_q) \approx \frac{w_\rho}{d_\rho}$. Now the state evolution update step becomes

$$\begin{aligned} \gamma_\rho^{t+1} &= \lambda_\rho^2 \mathbb{E}_{z_q} \text{Tr} \mathcal{F}_\rho \left(\left\{ \gamma_q^t I_{d_q} + \sqrt{\gamma_q^t} z_q \right\}_q \right) \\ &\approx \lambda_\rho^2 \mathbb{E}_{z_q} \text{Tr} \frac{1}{d_\rho} \left(\gamma_\rho^t I_{d_\rho} + \sqrt{\gamma_\rho^t} z_q \right) \\ &= \lambda_\rho^2 \gamma_\rho^t. \end{aligned}$$

This means that when γ is small (but nonzero), γ_ρ shrinks towards zero if $\lambda_\rho < 1$ and grows in magnitude if $\lambda_\rho > 1$. We conclude the threshold at $\lambda = 1$.

2.7 Correctness of state evolution?

In this section we justify the heuristic derivation of state evolution in the previous section and argue for its correctness. We first discuss prior work that provides a rigorous foundation for the methods we used, in related settings. We then show numerically that our AMP algorithm obeys the state evolution equations.

2.7.1 Rigorous work on state evolution

State evolution was introduced along with AMP by [63], based on *density evolution* in the sparse setting of LDPC codes [126]. It was later proven rigorously that AMP obeys state evolution in the large n limit (in a particular formal sense) for certain forms of the AMP iteration [23, 81]. In particular, $\mathbb{Z}/2$ synchronization with Gaussian noise (a special case of our model) falls into this framework and thus admits a rigorous analysis [56]. Although the proofs of [23, 81] only consider the case of real-valued AMP, it has been stated [102] that the proof extends to the complex-valued case. This covers our synchronization model over $U(1)$ with one frequency. In order to cover our general formulation of AMP over any group with any number of frequencies, one needs to replace the complex numbers by a different real algebra (namely a product of matrix algebras). We expect that this generalization should follow from the existing methods.

There is, however, an additional caveat involving the initialization of state evolution. In practice, we initialize AMP to small random values. Recall that we only need to recover the group elements up to a global right-multiplication and so there exists a favorable global right-multiplication so that our random initialization has some correlation with the truth. However, this correlation is $o(1)$ and corresponds to $\gamma = \vec{0}$ in the large n limit. This means that technically, the formal proof of state evolution (say for $\mathbb{Z}/2$) tells us that for any fixed

t , AMP achieves $\gamma = \vec{0}$ after t iterations in the large n limit. Instead we would like to show that after $\omega(1)$ iterations we achieve a nonzero γ . It appears that proving this would require a non-asymptotic analysis of AMP, such as [131]. It may appear that this initialization issue can be fixed by initializing AMP with a spectral method, which achieves $\Omega(1)$ correlation with the truth; however this does not appear to easily work due to subtle issue about correlation between the noise and iterates. In practice, the initialization issue is actually not an issue at all: with a small random initialization, AMP consistently escapes from the trivial fixed point (provided some λ exceeds 1). One way to explain this is that when the AMP messages are small, the nonlinear function \mathcal{F} is essentially the identity (see Section 2.6.6) and so AMP is essentially just the power method; this roughly means that AMP automatically initializes itself to the output of the spectral method.

2.7.2 Experiments on state evolution

We now present experimental evidence that AMP obeys the state evolution equations. In Figure 2-8 we show two experiments, one with $U(1)$ and one with $SO(3)$. In both cases we see that the performance of AMP closely matches the state evolution prediction. We see some discrepancy near the $\lambda = 1$ threshold, which can be attributed to the fact that here we are running AMP with finite n whereas state evolution describes the $n \rightarrow \infty$ behavior.

2.8 Statistical-to-computational gaps

In various settings it has been shown, using standard but non-rigorous methods from statistical physics, that the analysis of AMP and state evolution yields a complete picture of the various “phase transitions” that occur in a computational problem (e.g. [97, 96]). In some settings, certain features of these predictions have been confirmed rigorously (e.g. [89, 22]). In this section we will use these methods to give non-rigorous predictions about statistical-to-computational gaps in the Gaussian synchronization model.

In Section 2.6.6 we have seen that (in the large n limit) AMP achieves nontrivial recovery




Figure 2-8: AMP compared to the state evolution equations experimentally. **Left:** $U(1)$ with K frequencies, for $K = 1, 2, 3, 4$ (from top to bottom) with $n = 100$. The solid line is AMP and the dotted line is the state evolution prediction. The horizontal axis is the signal-to-noise ratio λ , which we take to be equal on all frequencies. The vertical axis is the natural logarithm of error, which is defined as $\text{error} = 1 - |\langle x, \hat{x} \rangle|/n \in [0, 1]$ where $x \in U(1)^n$ is the truth and $\hat{x} \in U(1)^n$ is the (rounded) output of AMP. In particular, a log error value of zero (top of the figure) indicates trivial recovery (random guessing), and lower values are better. **Right:** $SO(3)$ with K frequencies, for $K = 1, 2$ (from top to bottom), with $n = 50$. Now error is measured as $\text{error} = 1 - \frac{1}{\sqrt{3n}} \|X^\top \hat{X}\|_F \in [0, 1]$ where X, \hat{X} are $3n \times n$ matrices whose 3×3 blocks encode elements of $SO(3)$ via the standard representation (3D rotation matrices).

if and only if $\lambda > 1$ on at least one frequency. In this section, we will see that it is sometimes statistically possible to succeed below this threshold, although no known efficient algorithm achieves this. A rigorous analysis of an inefficient estimator has indeed confirmed that the $\lambda = 1$ threshold can be beaten in some cases (see Chapter 3); the non-rigorous computations in this section give sharp predictions for exactly when this is possible.

2.8.1 Free energy

Recall the parameter $\gamma = \{\gamma_\rho\}_\rho$ from the state evolution recurrence (2.4); γ captures the amount of information that AMP's current state has about each frequency, with $\gamma_\rho = 0$ indicating no information and $\gamma_\rho \rightarrow \infty$ indicating complete knowledge.

An important quantity is the *Bethe free energy* per variable (also called the *replica symmetric potential function*) of a state γ , which for the Gaussian synchronization model is given

(up to constants) by

$$f(\gamma) = -\frac{1}{4} \sum_{\rho} d_{\rho}^2 \lambda_{\rho}^2 + \frac{1}{2} \sum_{\rho} d_{\rho} \gamma_{\rho} + \frac{1}{4} \sum_{\rho} \frac{\gamma_{\rho}^2}{\lambda_{\rho}^2} - \mathbb{E}_z \log \mathbb{E}_g \exp \left(\sum_{\rho} \langle \rho(g), \gamma_{\rho} I_{d_{\rho}} + \sqrt{\gamma_{\rho}} z_{\rho} \rangle \right)$$

where z_{ρ} is a $d_{\rho} \times d_{\rho}$ matrix of i.i.d. standard Gaussians (of the appropriate type: real, complex, or quaternionic, depending on ρ), and g is drawn from Haar measure on the group. We do not include the derivation of this expression, but it can be computed from belief propagation (as in [96]) or from the replica calculation (as in [82]).

Roughly speaking, the interpretation of the Bethe free energy is that it is the objective value that AMP is trying to minimize. AMP can be thought of as starting from the origin $\gamma = 0$ and performing naïve gradient descent in the free energy landscape until it reaches a local minimum; the value of γ at this minimum describes the final state of AMP. (It can be shown that the fixed points of the state evolution recurrence (2.4) are precisely the stationary points of the Bethe free energy.) As is standard for these types of problems, we conjecture that AMP is optimal among all polynomial-time algorithms. However, with no restriction on efficiency, the information-theoretically optimal estimator is given by the global minimum of the free energy. (This has been shown rigorously for the related problem of rank-one matrix estimation [22].) The intuition here is that the optimal estimator should use exhaustive search to enumerate all fixed points of AMP and return the one of lowest Bethe free energy. Note that just because we can compute the γ value that minimizes the Bethe free energy it does not mean we can achieve this γ with an efficient algorithm; γ represents correlation between the AMP iterates and the ground truth, and since the truth is unknown it is hard to find iterates that have a prescribed γ .

2.8.2 Examples

We now examine the Bethe free energy landscapes of some specific synchronization problems at various values of λ , and discuss the implications. Our primary examples will be $U(1)$ and \mathbb{Z}/L with various numbers of frequencies, as discussed in Section 2.3.5. Recall that

references to $U(1)$ or \mathbb{Z}/L “with K frequencies” means that observations are band-limited to the Fourier modes $e^{ik\theta}$ with $|k| \leq K$.

Our first example is $U(1)$ with a single frequency, shown in Figure 2-9. Here we see that the problem transitions from (statistically) ‘impossible’ to ‘easy’ (AMP achieves non-trivial recovery) at $\lambda = 1$, with no (computationally) ‘hard’ regime. In particular, AMP is statistically optimal for every value of λ .






Figure 2-9: Free energy landscape for $U(1)$ with 1 frequency. **Left:** $\lambda < 1$. The global minimum of free energy occurs at $\gamma = 0$, indicating that AMP or any other estimator achieves zero correlation with the truth. **Right:** $\lambda > 1$. Now the global minimum occurs at nonzero γ , and this is achieved by AMP. Therefore AMP achieves the statistically optimal MSE (mean squared error). This MSE departs continuously from zero at the $\lambda = 1$ threshold.

Our next example is a single-frequency problem that exhibits a computational gap (a ‘hard’ phase). In Figure 2-10 we take the alternating group A_4 with its irreducible 3-dimensional representation as the rotational symmetries of a tetrahedron. When $\lambda > 1$, AMP achieves statistically optimal performance but when λ is below 1 but sufficiently large, AMP gives trivial performance while the statistically optimal estimator gives nontrivial performance. This means we have a computational gap, i.e. there are values of λ below the AMP threshold ($\lambda = 1$) where nontrivial recovery is statistically possible.


Next we move on to some 2-frequency problems, where γ is now a 2-dimensional vector. In Figure 2-11 we see an example with no computational gap, and an example with a computational gap. Note that the free energy landscape at the AMP threshold $\lambda = (1, \dots, 1)$




(a) $\lambda = 0.8$, impossible



(b) $\lambda = 0.91$, impossible



(c) $\lambda = 0.92$, hard



(d) $\lambda = 1.1$, easy

Figure 2-10: Free energy landscape for A_4 with 1 frequency: the standard 3-dimensional representation (rigid motions of a tetrahedron). **(a)** $\lambda = 0.8$. The global minimizer is $\gamma = 0$ so no estimator achieves nontrivial recovery. **(b)** $\lambda = 0.91$. A new local minimum in the free energy has appeared, but the global minimum is still at $\gamma = 0$ and so nontrivial recovery remains impossible. **(c)** $\lambda = 0.92$. AMP is stuck at $\gamma = 0$ but the (inefficient) statistically optimal estimator achieves a nontrivial γ (the global minimum). AMP is not statistically optimal. This computational gap appears at $\lambda \approx 0.913$, at which point the global minimizer transitions discontinuously from $\gamma = 0$ to some positive value. **(d)** $\lambda = 1.1$. AMP achieves optimal recovery. The AMP γ value transitions discontinuously from zero to optimal at $\lambda = 1$.

reveals whether or not a computational gap exists: there is a gap if and only if the global minimum of free energy does not occur at the origin.

We now state some experimental results regarding which synchronization problems have




Figure 2-11: Free energy landscape for 2-frequency problems at the critical value $\lambda = (1, 1)$. Darker colors indicate lower free energy. **Left:** $\mathbb{Z}/6$ with 2 frequencies. Here the origin is the global minimizer of free energy and so there is no computational gap, i.e. nontrivial recovery is statistically impossible when both λ_1 and λ_2 are below 1. **Right:** $\mathbb{Z}/5$ with 2 frequencies. Here the global minimizer (marked with an X) does not lie at the origin and so there is a computational gap, i.e. there is a regime where nontrivial recovery is statistically possible yet AMP fails.

computational gaps. For $U(1)$ with (the first) K frequencies, there is a gap iff $K \geq 3$. For \mathbb{Z}/L with K frequencies, there is a gap for $K \geq 3$ and no gap for $K = 1$; when $K = 2$ there is only a gap for $L = 5$. For $SO(3)$ with K frequencies, there is a gap for all $K \geq 1$.

In Chapter 3 we will give some rigorous lower bounds for Gaussian synchronization problems, showing for instance that $U(1)$ with one frequency is statistically impossible below $\lambda = 1$. The non-rigorous results above predict further results that we were unable to show rigorously, e.g. $U(1)$ with two frequencies and $\mathbb{Z}/3$ (with one frequency) are statistically impossible below the $\lambda = 1$ threshold.

In the examples above we saw that when every λ is below 1, AMP gives trivial performance, and when some λ exceeds 1, AMP gives statistically optimal performance. However, the behavior can be more complicated, namely AMP can exhibit nontrivial but sub-optimal performance. In Figure 2-12 we show such an example: $\mathbb{Z}/25$ with 9 frequencies.




Figure 2-12: An example where AMP gives nontrivial but sub-optimal performance. Here we take $\mathbb{Z}/25$ with 9 frequencies. Set $\lambda_k = 0.8$ for $k = 1, \dots, 8$ and $\lambda_9 = 1.1$. Since we cannot visualize the free energy landscape in 9 dimensions, we instead plot the state evolution recurrence as it evolves over time (number of iterations t) from two different starting points. The bottom two curves correspond to AMP’s performance, where we initialize γ to be small: $\gamma = (0.05, 0.05)$. The solid line is γ_9 and the dashed line is γ_1 (which is representative of $\gamma_2, \dots, \gamma_8$). The top two curves correspond to a “warm start” $\gamma = (0.7, 0.7)$. We see that with the warm start, state evolution converges to a different fixed point with larger γ values, and thus better correlation with the truth. Furthermore, this fixed point has lower free energy (not shown) than the lower one, indicating that the information-theoretically optimal estimator outperforms AMP.

Chapter 3

Synchronization: contiguity and rigorous bounds

This chapter is adapted from joint work with Amelia Perry, Afonso Bandeira, and Ankur Moitra [120]. The journal version [122] is comprised of the first part of the preprint [120]; this chapter is based primarily on the second part but shares some content with [122].

3.1 Introduction

In the previous chapter we gave a sharp but non-rigorous analysis of the Gaussian synchronization model, determining the optimal mean squared error achievable for any given value of the signal-to-noise parameter(s) λ . In particular, we identified a threshold at $\lambda = 1$, above which AMP (or simply PCA) achieves nontrivial correlation with the truth. The goal of this chapter is to rigorously investigate whether the $\lambda = 1$ threshold is optimal, i.e. whether there is *any* estimator that can succeed when $\lambda \leq 1$.

Take for example, the simple case of $\mathbb{Z}/2$ synchronization: we observe

$$Y = \frac{\lambda}{n}xx^\top + \frac{1}{\sqrt{n}}W, \tag{3.1}$$

where $x \in \{\pm 1\}^n$ is the signal to be recovered, W is a GOE matrix¹, and $\lambda > 0$ is the signal-to-noise parameter.

We are interested in the following statistical questions:

- *Detection*: For what values of λ is it possible to consistently distinguish (with probability $1 - o(1)$ as $n \rightarrow \infty$) between a random matrix Y drawn from the $\mathbb{Z}/2$ model and a pure noise matrix ($\lambda = 0$)?
- *Recovery*: For what values of λ does there exist an estimator that achieves non-vanishing correlation with x as $n \rightarrow \infty$?

Random matrix theory gives a precise analysis of PCA (top eigenvector) for the $\mathbb{Z}/2$ model:

Theorem 3.1.1 ([69, 25]). *Let Y be drawn from the $\mathbb{Z}/2$ model (3.1).*

- *If $\lambda \leq 1$, the top eigenvalue of Y converges almost surely to 2 as $n \rightarrow \infty$, and the top (unit-norm) eigenvector v has trivial correlation with the spike: $\langle v, x \rangle^2 \rightarrow 0$ almost surely.*
- *If $\lambda > 1$, the top eigenvalue converges almost surely to $\lambda + 1/\lambda > 2$ and v has nontrivial correlation with the spike: $\langle v, x \rangle^2 \rightarrow 1 - 1/\lambda^2$ almost surely.*

Therefore PCA solves the detection and recovery problems precisely when $\lambda > 1$. Our goal is now to investigate whether *any* method (perhaps having nothing to do with eigenvalues or eigenvectors) can beat this threshold. (Recall that although AMP outperforms PCA in terms of mean squared error once above the threshold, it does not achieve a better threshold. In this chapter, we will only be concerned with the threshold.)

Our focus in this chapter will be on proving non-detection lower bounds, i.e. proving that the detection problem is statistically impossible when λ is below a certain value. To do this, we will use a second moment method associated with the classical notion of *contiguity*

¹Gaussian orthogonal ensemble: a random symmetric matrix with off-diagonal entries $\mathcal{N}(0, 1)$, diagonal entries $\mathcal{N}(0, 2)$, and all entries independent (up to symmetry).

[92]. If we can compute a particular second moment and show that it is finite, non-detection follows immediately. This method has been used recently to study detection thresholds in the stochastic block model [114, 19] and other planted models [109, 20].

Curiously, there is no formal relation between detection and recovery in general; there are simple (but pathological) example of problems for which one is possible but not the other. However, for all the problems considered in this chapter, we expect that the detection and recovery thresholds are identical. (This is known to be true, for instance, for the $\mathbb{Z}/2$ model: both the detection and recovery thresholds occur at $\lambda = 1$ [56, 120].) For models with Gaussian noise, [20] give a general method to transfer non-detection lower bounds to non-recovery lower bounds.

To give a broader perspective, we remark that the $\mathbb{Z}/2$ model above is a special case of the *spiked Wigner* model from random matrix theory, in which the signal (“spike”) x is any vector of norm \sqrt{n} (not necessarily entrywise ± 1). Theorem 3.1.1 holds in this more general setting and so PCA exhibits a threshold at $\lambda = 1$. An analogous spectral threshold (the celebrated BBP transition [12]) occurs in the related *spiked Wishart* (covariance) model. Work by the author and collaborators [120, 119] uses the second moment method to investigate the statistical detection threshold for spiked Wigner, spiked Wishart, and spiked tensor models; various assumptions on the structure of the signal x are considered.

The second moment method does not always give a sharp lower bound on the detection threshold. In cases where it is loose, it can sometimes be strengthened by conditioning away from certain “bad” events that are extremely rare but cause the second moment to blow up. In this chapter, we will make use of a variant of this idea due to [19]. More involved conditioning methods (due to the author and others) can give even tighter bounds [119, 20]. More generally, modified second moment methods of a similar nature have appeared in contexts such as branching Brownian motion [39], branching random walks [8, 36], the Gaussian free field [33, 38, 37], cover times for random walks [54], community detection in random networks [11, 150], and thresholds for random satisfiability problems (e.g. k -colorability, k -sat) [48, 49, 47, 45, 46].

For the related question of determining the recovery threshold, sharp bounds have been achieved in some settings using methods based on the I-MMSE formula (e.g. [56, 57]), and interpolation methods (e.g. [22, 95]). Some of these methods also give sharp bounds on the detection threshold [44, 9].

In this chapter we study the Gaussian synchronization model defined in the previous chapter, as well as a simpler “truth-or-Haar” model for synchronization over finite groups. As discussed above, we give lower bounds on the detection threshold using the notion of contiguity and the associated second moment method. We furthermore give upper bounds on the detection threshold by analyzing inefficient algorithms based on exhaustive search.

The rest of this chapter is organized as follows. In Section 3.2 we define contiguity and present the second moment method which will be the core of our proofs. In Section 3.3 we define the truth-or-Haar model and give both lower and upper bounds. In Section 3.4 we recall the Gaussian synchronization model and give both lower and upper bounds.

3.2 Contiguity and the second moment method

Contiguity and related ideas will play a crucial role in this chapter. To give some background, contiguity was first introduced by [92] and since then has found many applications throughout probability and statistics. This notion and related tools such as the *small sub-graph conditioning method* have been used to establish many fundamental results about random graphs (e.g. [128, 80, 108]; see [153] for a survey). It has also been used to show the impossibility of detecting community structure in certain regimes of the stochastic block model [114, 19]. We will take inspiration from many of these works, in how we go about establishing contiguity. It is formally defined as follows:

Definition 3.2.1 ([92]). Let distributions P_n, Q_n be defined on the measurable space $(\Omega_n, \mathcal{F}_n)$. We say that the sequence P_n is *contiguous* to Q_n , and write $P_n \triangleleft Q_n$, if for any sequence of events A_n , $Q_n(A_n) \rightarrow 0 \implies P_n(A_n) \rightarrow 0$ as $n \rightarrow \infty$.

Contiguity implies that the distributions P_n and Q_n cannot be reliably distinguished in the

following sense:

Claim 3.2.2. If $P_n \triangleleft Q_n$ then there is no a statistical test \mathcal{D} that takes a sample from either P_n or Q_n (say each is chosen with probability $\frac{1}{2}$) and correctly outputs which of the two distributions it came from with error probability $o(1)$ as $n \rightarrow \infty$.

Proof. Suppose that such a test \mathcal{D} exists. Let A_n be the event that \mathcal{D} outputs ‘ P_n .’ Since \mathcal{D} succeeds reliably when the sample comes from Q_n , we have $Q_n(A_n) \rightarrow 0$ (as $n \rightarrow \infty$). By contiguity this means $P_n(A_n) \rightarrow 0$. But this contradicts the fact that \mathcal{D} succeeds reliably when the sample comes from P_n . \square

Note that $P_n \triangleleft Q_n$ and $Q_n \triangleleft P_n$ are not the same. Nevertheless either of them implies non-distinguishability. Also, showing that two distributions are contiguous does not rule out the existence of a test that distinguishes between them with constant probability. In fact, for many pairs of contiguous random graph models, such tests do exist.

Our goal in this chapter is to show thresholds below which planted and pure noise models are contiguous. We will do this through computing a particular second moment, related to the χ^2 -divergence as $1 + \chi^2(P_n || Q_n)$, through a form of the second moment method:

Lemma 3.2.3 (explicit in [109], implicit in earlier work). *Let $\{P_n\}$ and $\{Q_n\}$ be two sequences of probability measures on $(\Omega_n, \mathcal{F}_n)$. If the second moment*

$$\mathbb{E}_{Q_n} \left[\left(\frac{dP_n}{dQ_n} \right)^2 \right]$$

exists and remains bounded as $n \rightarrow \infty$, then $P_n \triangleleft Q_n$.

All of the contiguity results in this chapter will follow through Lemma 3.2.3. The roles of P_n and Q_n are not symmetric, and we will always take P_n to be the spiked distribution (where a planted signal is present) and take Q_n to be the unspiked (pure noise) distribution, as the second moment is more tractable to compute in this direction. We include the proof of Lemma 3.2.3 here for completeness:

Proof. Let $\{A_n\}$ be a sequence of events. Using Cauchy–Schwarz,

$$\begin{aligned} P_n(A_n) &= \int_{A_n} dP_n = \int_{A_n} \frac{dP_n}{dQ_n} dQ_n \leq \sqrt{\int_{A_n} \left(\frac{dP_n}{dQ_n}\right)^2 dQ_n} \cdot \sqrt{\int_{A_n} dQ_n} \\ &\leq \sqrt{\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n}\right)^2} \cdot \sqrt{Q_n(A_n)}. \end{aligned}$$

The first factor is bounded; so if $Q_n(A_n) \rightarrow 0$ as $n \rightarrow \infty$, we must also have $P_n(A_n) \rightarrow 0$, as desired. \square

There will be times when the above second moment is infinite but we are still able to prove contiguity using a modified second moment that conditions on ‘good’ events. This idea is based on [19].

Lemma 3.2.4. *Let ω_n be a ‘good’ event that occurs with probability $1 - o(1)$ under P_n . Let \tilde{P} be the conditional distribution of P given ω_n . If*

$$\mathbb{E}_{Q_n} \left[\left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 \right]$$

remains bounded as $n \rightarrow \infty$, then $P_n \triangleleft Q_n$.

Proof. By Lemma 3.2.3 we have $\tilde{P}_n \triangleleft Q_n$. This implies $P_n \triangleleft Q_n$ because $\tilde{P}_n(A_n) \rightarrow 0$ implies $P_n(A_n) \rightarrow 0$ (since $P_n(\omega_n) \rightarrow 1$). \square

3.3 The truth-or-Haar model

3.3.1 Main results

In this section we study a very simple model for synchronization over finite groups: for each pair of group elements we either observe the true relative group element, or a uniformly random one.

Definition 3.3.1. Let G be a finite group and let $\tilde{p} \geq 0$. In the *truth-or-Haar* model $\text{ToH}(\tilde{p}, G)$ we first draw a vector $g \in G^n$ where each coordinate g_u is chosen independently from uniform (Haar) measure on G . For each unordered pair $\{u, v\}$ (with $u \neq v$), with probability $p = \frac{\tilde{p}}{\sqrt{n}}$ let $Y_{uv} = g_u g_v^{-1}$, and otherwise let Y_{uv} be drawn uniformly from G . Define $Y_{vu} = (Y_{uv})^{-1}$ and $Y_{uu} = 1$ (the identity element of G). We reveal the matrix $Y \in G^{n \times n}$.

The truth-or-Haar model is not interesting for infinite groups G . This is because if G is infinite, the detection problem can be solved for any $\tilde{p} > 0$ by checking whether there is a consistent triangle, i.e. three vertices u, v, w such that $Y_{uv}Y_{vw}Y_{wu} = 1$.

This problem has been studied previously by [141] for the case where the group G is the cyclic group \mathbb{Z}/L . It is important to note that since we only have pairwise measurements, we can only hope to recover the group elements up to a global right-multiplication by some group element.

[141] shows that for $G = \mathbb{Z}/L$ there is a spectral approach that succeeds at detection and recovery above the threshold $\tilde{p} > 1$. Specifically, the spectral method identifies each group element with a complex L th root of unity and takes the top eigenvalue (and eigenvector) of the complex-valued observed matrix Y . We expect that an efficient algorithm for detection exists for any finite group above this $\tilde{p} = 1$ threshold: for instance, if the group has a \mathbb{Z}/L quotient (for any L) we can apply the \mathbb{Z}/L spectral algorithm.

Using the second moment method, we will prove the following lower bound for the truth-or-Haar model:

Theorem 3.3.2. *Let G be a finite group of order L and let $\tilde{p} \geq 0$. If*

$$\tilde{p} < \tilde{p}_L^* \triangleq \sqrt{\frac{2(L-1)\log(L-1)}{L(L-2)}}$$

then $\text{ToH}(G, \tilde{p})$ is contiguous to $\text{ToH}(G, 0)$. For $L = 2$, $\tilde{p}_2^ = 1$ (the limit value of the 0/0 expression).*

The proof will span Sections 3.3.2 and 3.3.3. We provide some numerical values for the

critical value \tilde{p}^* .

L	2	3	4	5	6	10	100
\tilde{p}^*	1	0.961	0.908	0.860	0.819	0.703	0.305

Note that this lower bound matches the spectral threshold $\tilde{p} = 1$ when $L = 2$, but does not match it for $L \geq 3$. We also give an upper bound for the truth-or-Haar model using an inefficient algorithm:

Theorem 3.3.3. *Let G be a finite group of order $L \geq 2$. If*

$$\tilde{p} > \sqrt{\frac{4 \log L}{L-1}}$$

there is a computationally inefficient algorithm that can distinguish between the spiked and unspiked models.

The proof will be given in Section 3.3.4.

For small L , this theorem is not very interesting because the right-hand side exceeds the spectral threshold of 1. However, for $L \geq 11$, the right-hand side drops below 1, indicating that it is information-theoretically possible to go below the spectral threshold. However, we expect that no *efficient* algorithm can beat the spectral threshold.

As $L \rightarrow \infty$, this upper bound differs from the lower bound of Theorem 3.3.2 by a factor of $\sqrt{2}$. Here we expect that the upper bound is asymptotically tight and that the lower bound can be improved by a factor of $\sqrt{2}$ (asymptotically) using a more sophisticated conditioning method of the author and others [119]; here the event conditioned on depends not only on the signal but also on the noise.

3.3.2 Second moment computation

We will now establish contiguity results in the truth-or-Haar model. Let $p = \frac{\tilde{p}}{\sqrt{n}}$. Let P_n be the ‘spiked’ model $\text{ToH}_n(\tilde{p}, G)$ and let $Q_n = \text{ToH}_n(0, G)$ be the ‘unspiked’ model in which

the observations are completely random. We give an upper bound on the second moment:

$$\begin{aligned}
\frac{dP_n}{dQ_n} &= \mathbb{E} \prod_{g, u < v} \frac{p \mathbb{1}[Y_{uv} = g_u g_v^{-1}] + (1-p)/L}{1/L}, \\
\mathbb{E}_{Q_n} \left[\left(\frac{dP_n}{dQ_n} \right)^2 \right] &= \mathbb{E}_{g, g'} \prod_{u < v} \mathbb{E}_{Y_{uv} \sim Q_n} (pL \mathbb{1}[Y_{uv} = g_u g_v^{-1}] + 1-p)(pL \mathbb{1}[Y_{uv} = g'_u (g'_v)^{-1}] + 1-p) \\
&= \mathbb{E}_{g, g'} \prod_{u < v} \mathbb{E}_{Y_{uv} \sim Q_n} (p^2 L^2 \mathbb{1}[g_u g_v^{-1} = Y_{uv} = g_u g_v^{-1}] + p(1-p)L \mathbb{1}[Y_{uv} = g_u g_v^{-1}] \\
&\quad + p(1-p)L \mathbb{1}[Y_{uv} = g'_u (g'_v)^{-1}] + (1-p)^2) \\
&= \mathbb{E}_{g, g'} \prod_{u < v} (1 + p^2(L \mathbb{1}[g_u g_v^{-1} = g'_u (g'_v)^{-1}] - 1)) \\
&= \mathbb{E}_{g, g'} \prod_{u < v} (1 + p^2(L \mathbb{1}[g_u^{-1} g'_u = g_v^{-1} g'_v] - 1)) \\
&\leq \mathbb{E}_{g, g'} \prod_{u < v} \exp [p^2(L \mathbb{1}[g_u^{-1} g'_u = g_v^{-1} g'_v] - 1)] \\
&\leq \mathbb{E}_{g, g'} \prod_{u, v} \exp \left[\frac{p^2}{2} (L \mathbb{1}[g_u^{-1} g'_u = g_v^{-1} g'_v] - 1) \right] \\
&= e^{-n^2 p^2 / 2} \mathbb{E}_{g, g'} \exp \left[\frac{p^2 L}{2} \sum_{u, v} \mathbb{1}[g_u^{-1} g'_u = g_v^{-1} g'_v] \right].
\end{aligned}$$

3.3.3 The conditioning method

Our next step will be to make use of a result of [19] (Proposition 5) involving boundedness of a particular expectation involving multinomial random variables. We refer to this as the *conditioning method* because it involves conditioning away from bad events via Lemma 3.2.4. For convenience, we restate the setup and result of [19].

Let Δ_q denote the simplex $\{(\pi_1, \dots, \pi_q) : \pi_i \geq 0, \sum_i \pi_i = 1\}$. For $\pi \in \Delta_q$, let $\Delta_{q^2}(\pi)$ denote the set of $q \times q$ matrices whose row- and column-sums are given by π , namely:

$$\Delta_{q^2}(\pi) = \{\alpha \in \mathbb{R}^{q \times q} : \alpha_{ij} \geq 0 \ \forall i, j, \sum_{i=1}^q \alpha_{ij} = \pi_j \ \forall j, \sum_{j=1}^q \alpha_{ij} = \pi_i \ \forall i\}.$$

Fix a $q^2 \times q^2$ matrix A and some $\pi \in \Delta_q$. Let $\bar{\alpha} \in \Delta_{q^2}(\pi)$ be given by $\bar{\alpha}_{ij} = \pi_i \pi_j$, let $N \sim \text{Multinomial}(n, \bar{\alpha})$, and $X = (N - n\bar{\alpha})/\sqrt{n}$. Fix a sequence a_n such that $\sqrt{n} \ll a_n \ll n$ and define Ω_n to be the event that $|\sum_i N_{ij} - n\pi_j| \leq a_n \forall j$ and $|\sum_j N_{ij} - n\pi_i| \leq a_n \forall i$. Since $\sqrt{n} \ll a_n$, the probability of Ω_n converges to 1.

Proposition 3.3.4 ([19] Proposition 5). *Define*

$$m \triangleq \sup_{\alpha \in \Delta_{q^2}(\pi)} \frac{(\alpha - \bar{\alpha})^\top A(\alpha - \bar{\alpha})}{D(\alpha, \bar{\alpha})}$$

where D is the KL divergence: $D(\alpha, \bar{\alpha}) = \sum_{ij} \alpha_{ij} \log(\alpha_{ij}/\bar{\alpha}_{ij})$. If $m < 1$ then

$$\mathbb{E}[\mathbb{1}_{\Omega_n} \exp(X^\top AX)] \rightarrow \mathbb{E} \exp(Z^\top AZ) < \infty,$$

as $n \rightarrow \infty$, where $Z \sim \mathcal{N}(0, \text{diag}(\bar{\alpha}) - \bar{\alpha}\bar{\alpha}^\top)$. If instead $m > 1$ then

$$\mathbb{E}[\mathbb{1}_{\Omega_n} \exp(X^\top AX)] \rightarrow \infty$$

as $n \rightarrow \infty$.

The intuition behind the above result is the following. Think of $\alpha = N/n$ so that $X = \sqrt{n}(\alpha - \bar{\alpha})$. Thus we can write $\exp(X^\top AX) = \exp(n(\alpha - \bar{\alpha})^\top A(\alpha - \bar{\alpha}))$. The probability that a particular α occurs is asymptotically $\exp(-nD(\alpha, \bar{\alpha}))$. This means

$$\mathbb{E}[\mathbb{1}_{\Omega_n} \exp(X^\top AX)] \approx \int_{\alpha} \exp[n((\alpha - \bar{\alpha})^\top A(\alpha - \bar{\alpha}) - D(\alpha, \bar{\alpha}))] \quad (3.2)$$

where the integral is over α values for which Ω_n holds. Now apply the *saddle point method* (also called Laplace's method): as n becomes large, (3.2) is dominated by the value of α for which

$$(\alpha - \bar{\alpha})^\top A(\alpha - \bar{\alpha}) - D(\alpha, \bar{\alpha}) \quad (3.3)$$

is maximized. In particular, (3.2) is bounded as $n \rightarrow \infty$ if (3.3) is negative for every $\alpha \in \Delta_{q^2}(\pi)$. Rearranging this yields the condition $m < 1$ in Proposition 3.3.4. The fact that

we are restricting to the event Ω_n helps us here; if we did not include the indicator $\mathbb{1}_{\Omega_n}$, we would have to change $\Delta_{q^2}(\pi)$ to Δ_{q^2} (the simplex of dimension q^2) in the definition of m , which in some cases gives a larger value of m and thus a weaker result.

Let us now see how to use Proposition 3.3.4 to bound a conditional variant of the second moment from Section 3.3.2. Let $q = L$ and $\pi = (1/L, \dots, 1/L) \in \Delta_q$. For a vector $g \in G^n$, let $\omega_n(g)$ be the event

$$\left| |\{u : g_u = a\}| - n/L \right| \leq a_n.$$

We will compute the conditional second moment $\mathbb{E}_{Q_n} \left[\left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 \right]$ where \tilde{P} is the conditional distribution $P_n | \omega_n(g)$. Our goal is to show that this conditional second moment remains bounded as $n \rightarrow \infty$ so that contiguity $P_n \triangleleft Q_n$ follows from Lemma 3.2.4.

Similarly to Section 3.3.2, we compute

$$\frac{d\tilde{P}_n}{dQ_n} \leq (1 + o(1)) e^{-n^2 p^2 / 2} \mathbb{E}_{g, g'} \omega_n(g) \omega_n(g') \exp \left[\frac{p^2 L}{2} \sum_{u, v} \mathbb{1}[g_u^{-1} g'_u = g_v^{-1} g'_v] \right].$$

Let $N_{ab} = |\{u : g_u = a, g'_u = b\}|$ and note that $N \sim \text{Multinomial}(n, \bar{\alpha})$ where $\bar{\alpha} = \frac{1}{L^2} \mathbb{1}_{L^2}$. Define Ω_n as above (depending on N). Let $X = \frac{N - n\bar{\alpha}}{\sqrt{n}} \in \mathbb{R}^{L^2}$, and let A be the $L^2 \times L^2$ matrix $A_{ab, a'b'} = \frac{\tilde{p}^2 L}{2} \mathbb{1}\{a^{-1}b = a'^{-1}b'\}$ where (recall) $p = \frac{\tilde{p}}{\sqrt{n}}$. We can now write

$$\frac{d\tilde{P}_n}{dQ_n} \leq (1 + o(1)) \mathbb{E}[\mathbb{1}_{\Omega_n} \exp(X^\top A X)]$$

and so by Proposition 3.3.4 we have that $\mathbb{E}_{Q_n} \left[\left(\frac{d\tilde{P}_n}{dQ_n} \right)^2 \right]$ is bounded provided

$$\sup_{\alpha \in \Delta_{q^2}(\pi)} \frac{(\alpha - \bar{\alpha})^\top A (\alpha - \bar{\alpha})}{D(\alpha, \bar{\alpha})} < 1.$$

Rewrite the numerator:

$$\begin{aligned}
(\alpha - \bar{\alpha})^\top A(\alpha - \bar{\alpha}) &= \alpha^\top A\alpha - 2\alpha^\top A\bar{\alpha} + \bar{\alpha}^\top A\bar{\alpha} \\
&= \frac{\tilde{p}^2 L}{2} \left(\sum_{aba'b'} \alpha_{ab} \alpha_{a'b'} \mathbb{1}\{a^{-1}b = a'^{-1}b'\} - \frac{2}{L} + \frac{1}{L} \right) \\
&= \frac{\tilde{p}^2 L}{2} \left(\sum_{h \in G} \alpha_h^2 - \frac{1}{L} \right)
\end{aligned}$$

where $\alpha_h = \sum_{(a,b) \in S_h} \alpha_{ab}$ and $S_h = \{(a,b) \mid a^{-1}b = h\}$.

In Appendix B.1 we prove the following result which provides the solution to the optimization problem above and thus completes the proof of Theorem 3.3.2.

Proposition 3.3.5. *For $L \geq 2$,*

$$\sup_{\alpha} \frac{L}{2} \frac{(\sum_{h \in G} \alpha_h^2 - \frac{1}{L})}{D(\alpha, \bar{\alpha})} = \frac{L(L-2)}{2(L-1)\log(L-1)}$$

where α ranges over (vectorized) nonnegative $L \times L$ matrices with row- and column-sums equal to $\frac{1}{L}$. When $L = 2$, the right-hand side is taken to equal 1 (the limit value of the 0/0 expression).

3.3.4 Upper bound via exhaustive search

In this subsection, we show that exhaustive search outperforms spectral methods in the Truth-or-Haar Model when L is large enough. Specifically, we prove Theorem 3.3.3.

We will use an inefficient algorithm based on exhaustive search over all candidate solutions $g \in G^n$. Given an observed matrix Y valued in G , let $T(g)$ be the number of edges satisfied by g , i.e. the number of unordered pairs $\{u, v\}$ (with $u \neq v$) such that $Y_{uv} = g_u g_v^{-1}$. The algorithm will distinguish between $P_n = \text{ToH}(G, \tilde{p})$ and $Q_n = \text{ToH}(G, 0)$ by thresholding $T = \max_{g \in G^n} T(g)$ (at some cutoff to be determined later).

Suppose Y is drawn from P_n and let $g^* \in G^n$ be the true spike. Then $T(g) \sim \text{Binom}(N, p')$ where $N = \binom{n}{2}$ and $p' = p + \frac{1-p}{L}$. By Hoeffding's inequality,

$$P_n(T(g^*) \leq Np' - k) \leq \exp\left(-\frac{2k^2}{N}\right)$$

which in turn implies $P_n(T(g^*) \leq Np' - n \log n) = o(1)$.

Now suppose Y is drawn from Q_n and fix any $g \in G^n$. Then $T(g) \sim \text{Binom}(N, 1/L)$. By the Chernoff bound,

$$Q_n(T(g) \geq k) \leq \exp(-ND(k/N \parallel 1/L))$$

where $D(a \parallel b) = a \log(a/b) + (1-a) \log((1-a)/(1-b))$. By a union bound over all L^n choices for g ,

$$\begin{aligned} Q_n(T \geq Np' - n \log n) &\leq L^n \exp\left(-ND\left(\frac{Np' - n \log n}{N} \parallel \frac{1}{L}\right)\right) \\ &= \exp\left(n \log L - ND\left(p + \frac{1-p}{L} - \mathcal{O}\left(\frac{\log n}{n}\right) \parallel \frac{1}{L}\right)\right) \\ &= \exp(n \log L - ND(1/L + \Delta \parallel 1/L)) \end{aligned}$$

where $\Delta = p(1 - 1/L) - \mathcal{O}\left(\frac{\log n}{n}\right) = \frac{\tilde{p}(L-1)}{\sqrt{n}L} - o\left(\frac{1}{\sqrt{n}}\right)$

$$\begin{aligned}
&= \exp \left[n \log L - N \left((1/L + \Delta) \log \left(\frac{1/L + \Delta}{1/L} \right) + (1 - 1/L - \Delta) \log \left(\frac{1 - 1/L - \Delta}{1 - 1/L} \right) \right) \right] \\
&= \exp \left[n \log L - N \left((1/L + \Delta) \log (1 + L\Delta) + (1 - 1/L - \Delta) \log \left(1 - \frac{L\Delta}{L-1} \right) \right) \right] \\
&= \exp \left[n \log L - N \left((1/L + \Delta) \left(L\Delta - \frac{1}{2} L^2 \Delta^2 \right) \right. \right. \\
&\quad \left. \left. + \left(\frac{L-1}{L} - \Delta \right) \left(-\frac{L\Delta}{L-1} - \frac{L^2 \Delta^2}{2(L-1)^2} \right) + o(1/n) \right) \right] \\
&= \exp \left[n \log L - N \left(\Delta + L\Delta^2 - \frac{1}{2} L\Delta^2 - \Delta + \frac{L}{L-1} \Delta^2 - \frac{L\Delta^2}{2(L-1)} + o(1/n) \right) \right] \\
&= \exp \left[n \log L - \frac{n^2}{2} \Delta^2 \frac{1}{2} \left(L + \frac{L}{L-1} \right) + o(n) \right] \\
&= \exp \left[n \log L - \frac{n}{4} \tilde{p}^2 \left(\frac{L-1}{L} \right)^2 \left(\frac{L^2}{L-1} \right) + o(n) \right] \\
&= \exp \left[n \log L - \frac{n}{4} \tilde{p}^2 (L-1) + o(n) \right] \\
&= o(1)
\end{aligned}$$

provided $\log L < \tilde{p}^2(L-1)/4$, i.e.

$$\tilde{p} > \sqrt{\frac{4 \log L}{L-1}}.$$

Therefore, it is possible to reliably distinguish P_n and Q_n by thresholding T at $Np' - n \log n$.

3.4 The Gaussian synchronization model

3.4.1 The model

Recall that the truth-or-Haar model is only meaningful for finite groups. Thus, to study synchronization problems over infinite groups such as $U(1)$ (unit-norm complex numbers) we need the noise to be continuous in nature. This motivates our Gaussian synchronization model in which we add Gaussian noise to the true relative group elements $g_u g_v^{-1}$. This model

was defined in Section 2.3.4), and we gave a sharp non-rigorous analysis of the statistical and computational limits of the model in Chapter 2. We repeat the definition of the model here for the reader's convenience. (We will also use a slightly different convention for quaternionic-type representations in this chapter; see Remark 3.4.4 below.) The model is very general, allowing for any compact group, and for observations on different 'frequencies' (irreducible representations of the group).

We now define the model. In order to have a sensible notion of adding Gaussian noise to a group element, we need to introduce some representation theory. We will assume the reader is familiar with the basics of representation theory. See Section 1.4 or e.g. [40] for an introduction.

Since we will be discussing representations of quaternionic type, we need to recall basic facts about quaternions. (Quaternions and quaternionic-type representations can be skipped on a first reading.) Quaternions take the form $q = a + bi + cj + dk$ where $a, b, c, d \in \mathbb{R}$ and (non-commutative) multiplication follows the rules $i^2 = j^2 = k^2 = ijk = -1$. Like complex numbers, quaternions support the operations norm $|q| = \sqrt{a^2 + b^2 + c^2 + d^2}$, real part $\Re(q) = a$, and conjugate $\bar{q} = a - bi - cj - dk$ satisfying $\overline{q_1 q_2} = \bar{q}_2 \bar{q}_1$ and $q\bar{q} = \bar{q}q = |q|^2$. These allow for the natural notions of unitarity and conjugate transpose A^* for quaternion-valued matrices A . The algebra of quaternions is denoted by \mathbb{H} .

Let G be a compact group. The irreducible representations of G over \mathbb{C} are finite dimensional. Every irreducible representation of G over \mathbb{C} has one of three types: real, complex, or quaternionic. Representations of real type can be defined over the reals (i.e. each group element is assigned a matrix with real-valued entries). Representations of complex type are (unlike the other types) not isomorphic to their complex conjugate representation $\bar{\rho}$. Representations of quaternionic type can be assumed to take the following form: each 2×2 block of complex numbers encodes a quaternion value using the correspondence

$$a + bi + cj + dk \quad \leftrightarrow \quad \begin{pmatrix} a + bi & c + di \\ -c + di & a - bi \end{pmatrix}.$$

Alternatively, we can think of quaternionic-type representations as being defined over the quaternions (i.e. each group element is assigned a quaternion-valued matrix) with dimension half as large. We will assume that our irreducible representations (over \mathbb{C}) are defined over \mathbb{R} , \mathbb{C} , or \mathbb{H} , depending on whether their type is real, complex, or quaternionic (respectively). Representations of complex type come in conjugate pairs. Without loss of generality, all these representations can be taken to be unitary.

Let d_ρ be the dimension of representation ρ . For quaternionic-type representations we let d_ρ be the quaternionic dimension, which is half the complex dimension (this differs from our convention in Chapter 2; see Remark 3.4.4 below). (For real-type representations, the real and complex dimensions are the same.) In defining our Gaussian model we need to fix a finite list of representations (‘frequencies’) to work with.

Definition 3.4.1. Let G be a compact group. A *list of frequencies* Ψ is a finite set of non-isomorphic irreducible (over \mathbb{C}) representations of G . We do not allow the trivial representation to be included in this list. For representations of complex type, we do not allow ρ and its conjugate $\bar{\rho}$ to both appear in the list.

We need to introduce Gaussian noise of various types. The type of noise used will correspond to the type of the representation in question.

Definition 3.4.2. A *standard Gaussian* of real, complex, or quaternionic type is defined to be

- for real type, $\mathcal{N}(0, 1)$
- for complex type: $\mathcal{N}(0, 1/2) + \mathcal{N}(0, 1/2) i$
- for quaternionic type: $\mathcal{N}(0, 1/4) + \mathcal{N}(0, 1/4) i + \mathcal{N}(0, 1/4) j + \mathcal{N}(0, 1/4) k$

where each component is independent.

Note that the normalization ensures that the expected squared norm is 1.

Definition 3.4.3. Let a *GOE*, *GUE*, or *GSE* (respectively) matrix be a random Hermitian matrix where the off-diagonals are standard Gaussians of real, complex, or quaternionic type (respectively), and the diagonal entries are real Gaussians $\mathcal{N}(0, 2/\beta)$ where $\beta = 1, 2, 4$ (respectively) depending on the type. All entries are independent except for the Hermitian constraint.

These matrices are the well-known Gaussian orthogonal (resp. unitary, symplectic) ensembles from random matrix theory.

Remark 3.4.4. We point out that we have used slightly different conventions from Chapter 2. Namely, for quaternionic-type representations we have decreased (by a factor of 2) both the definition of d_ρ and the variance of quaternionic Gaussian noise. These conventions will be more convenient in this chapter. Roughly speaking, in Chapter 2 we wanted to think of quaternionic representations as defined over \mathbb{C} , whereas here we think of them defined over the quaternions \mathbb{H} .

We can now formally state the Gaussian synchronization model over any compact group.

Definition 3.4.5. Let G be a compact group and let Ψ be a list of frequencies. For each $\rho \in \Psi$, let $\lambda_\rho \geq 0$. The *Gaussian synchronization model* $\text{GSynch}(\{\lambda_\rho\}, G, \Psi)$ is defined as follows. To sample from the n th distribution, draw a vector $g \in G^n$ by sampling each coordinate independently from Haar (uniform) measure on G . Let X_ρ be the $nd_\rho \times d_\rho$ matrix formed by stacking the matrices $\rho(g_u)$ for all u . For each frequency $\rho \in \Psi$, reveal the $nd_\rho \times nd_\rho$ matrix

$$Y_\rho = \frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho$$

where W_ρ is an $nd_\rho \times nd_\rho$ Hermitian Gaussian matrix (GOE, GUE, or GSE depending on whether ρ has real, complex, or quaternionic type, respectively). If we write a scalar λ in place of $\{\lambda_\rho\}$ we mean that $\lambda_\rho = \lambda$ for all ρ .

When $\lambda_\rho > 1$ for at least one ρ , we can use PCA (top eigenvalue) to reliably distinguish between $P_n = \text{GSynch}_n(\{\lambda_\rho\}, G, \Psi)$ and $Q_n = \text{GSynch}_n(0, G, \Psi)$; this follows from

Theorem 3.1.1. If given K frequencies, all with the same λ , it may appear that one should be able to combine the frequencies in order to achieve the threshold $\lambda > 1/\sqrt{K}$; after all, this would be possible if given K independent observations of a single frequency. However, our contiguity results will show that $\lambda > 1/\sqrt{K}$ is not sufficient. In fact, we conjecture that $\lambda > 1$ is required for any efficient algorithm to succeed at detection (see Chapter 2), although there are inefficient algorithms that succeed below this (as we will show in Section 3.4.6).

3.4.2 Second moment computation

Let P_n be $\text{GSynch}(\{\lambda_\rho\}, G, \Psi)$ and let Q_n be $\text{GSynch}(0, G, \Psi)$. Let $\beta_\rho = 1, 2, 4$ for real-, complex-, or quaternionic-type (respectively). We will use the standard Hermitian inner product for matrices: $\langle A, B \rangle = \text{Tr}(AB^*)$ where B^* denotes the conjugate transpose of B .

$$\begin{aligned} \frac{dP_n}{dQ_n} &= \mathbb{E}_X \prod_{\rho \in \Psi} \frac{\exp\left(-\frac{\beta_\rho n d_\rho}{4} \left\| Y_\rho - \frac{\lambda_\rho}{n} X_\rho X_\rho^* \right\|_F^2\right)}{\exp\left(-\frac{\beta_\rho n d_\rho}{4} \|Y\|_F^2\right)} \\ &= \mathbb{E}_X \prod_{\rho} \exp\left(\frac{\beta_\rho \lambda_\rho d_\rho}{2} \Re \langle Y_\rho, X_\rho X_\rho^* \rangle - \frac{\beta_\rho \lambda_\rho^2 d_\rho}{4n} \|X_\rho X_\rho^*\|_F^2\right). \end{aligned}$$

$$\begin{aligned} &\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2 \\ &= \mathbb{E}_{Y \sim Q_n} \mathbb{E}_{X, X'} \prod_{\rho} \exp\left(\frac{\beta_\rho \lambda_\rho d_\rho}{2} \Re \langle Y_\rho, X_\rho X_\rho^* + X'_\rho X'^*_{\rho} \rangle - \frac{\beta_\rho \lambda_\rho^2 d_\rho}{4n} \|X_\rho X_\rho^*\|_F^2 - \frac{\beta_\rho \lambda_\rho^2 d_\rho}{4n} \|X'_\rho X'^*_{\rho}\|_F^2\right) \\ &= \mathbb{E}_{X, X'} \prod_{\rho} \mathbb{E}_{Y_\rho} \exp\left(\frac{\beta_\rho \lambda_\rho d_\rho}{2} \Re \langle Y_\rho, X_\rho X_\rho^* + X'_\rho X'^*_{\rho} \rangle - \frac{\beta_\rho \lambda_\rho^2 d_\rho}{4n} \|X_\rho X_\rho^*\|_F^2 - \frac{\beta_\rho \lambda_\rho^2 d_\rho}{4n} \|X'_\rho X'^*_{\rho}\|_F^2\right). \end{aligned}$$

Use the Gaussian moment-generating function to eliminate Y : if z is a scalar (from \mathbb{R} , \mathbb{C} , or \mathbb{H}) and y is a standard Gaussian of the same type, then $\mathbb{E} \exp(\Re(yz)) = \exp(\frac{1}{2\beta}|z|^2)$. Recall that Y_ρ (drawn from Q_n) is Hermitian with each off-diagonal entry $\frac{1}{\sqrt{nd_\rho}}$ times a standard Gaussian (of the appropriate type), and each diagonal entry real Gaussian $\mathcal{N}(0, \beta/2)$. Continuing from above,

$$\begin{aligned}
&= \mathbb{E}_{X, X'} \prod_{\rho} \exp \left(\frac{1}{2\beta_{\rho}} \frac{1}{nd_{\rho}} \beta_{\rho}^2 \lambda_{\rho}^2 d_{\rho}^2 \frac{1}{2} \|X_{\rho} X_{\rho}^* + X'_{\rho} X'_{\rho}^*\|_F^2 - \frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} \|X_{\rho} X_{\rho}^*\|_F^2 - \frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} \|X'_{\rho} X'_{\rho}^*\|_F^2 \right) \\
&= \mathbb{E}_{X, X'} \prod_{\rho} \exp \left(\frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} \|X_{\rho} X_{\rho}^* + X'_{\rho} X'_{\rho}^*\|_F^2 - \frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} \|X_{\rho} X_{\rho}^*\|_F^2 - \frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} \|X'_{\rho} X'_{\rho}^*\|_F^2 \right) \\
&= \mathbb{E}_{X, X'} \prod_{\rho} \exp \left(\frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{4n} 2 \Re \langle X_{\rho} X_{\rho}^*, X'_{\rho} X'_{\rho}^* \rangle \right) \\
&= \mathbb{E}_{X, X'} \prod_{\rho} \exp \left(\frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{2n} \|X_{\rho}^* X'_{\rho}\|_F^2 \right).
\end{aligned}$$

3.4.3 The sub-Gaussian method

We will aim to show contiguity at a point where all λ 's are equal: $\lambda_{\rho} = \lambda$ for all ρ . (Note however that if we show contiguity at some λ and we then decrease some of the individual λ_{ρ} 's, we still have contiguity because the second moment above will only decrease.) Ideally we want contiguity for all $\lambda < 1$, matching the spectral threshold.

For each $u \in [n]$ let Z_u be a vector in \mathbb{R}^D where $D = \sum_{\rho \in \Psi} \beta_{\rho} d_{\rho}^2$, formed as follows. First draw h_u independently from Haar measure on G . For each ρ , vectorize the matrix $\sqrt{\beta_{\rho} d_{\rho}} \rho(h_u)$ into a real-valued vector of length $\beta_{\rho} d_{\rho}^2$ by separating the β_{ρ} components of each of the d_{ρ}^2 entries. Finally, concatenate all these vectors together to form Z_u . Let $Z^{(G, \Psi)}$ denote the distribution that each Z_u follows.

We can rewrite the second moment as

$$\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2 = \mathbb{E}_Z \exp \left(\frac{\lambda^2}{2n} \left\| \sum_u Z_u \right\|^2 \right).$$

(Here $h_u = g_u^{-1}g'_u$.)

We will use the following definition of sub-Gaussian for vector-valued random variables.

Definition 3.4.6. We say $z \in \mathbb{R}^m$ is *sub-Gaussian with covariance proxy* $\sigma^2 I$ if $\mathbb{E}[z] = 0$ and for all vectors $v \in \mathbb{R}^m$,

$$\mathbb{E} \exp(\langle z, v \rangle) \leq \exp\left(\frac{1}{2} \sigma^2 \|v\|^2\right).$$

More generally we can allow for a covariance proxy Σ that is not a multiple of the identity by replacing $\sigma^2 \|v\|^2$ by $v^\top \Sigma v$, but we will not need this here. Standard methods in the theory of large deviations give the following multivariate sub-Gaussian tail bound.

Lemma 3.4.7. Suppose $z \in \mathbb{R}^m$ is sub-Gaussian with covariance proxy $\sigma^2 I$. Let $\varepsilon > 0$. For all $a \geq 0$,

$$\mathbb{P}[\|z\|^2 \geq a] \leq C \exp\left(-\frac{a(1-\varepsilon)}{2\sigma^2}\right)$$

where $C = C(\varepsilon, m)$ is a constant depending only on ε and the dimension m .

Proof. Let $v_1, \dots, v_C \in \mathbb{R}^m$ be a collection of unit vectors such that for every unit vector $\hat{z} \in \mathbb{R}^m$, there exists i satisfying $\langle \hat{z}, v_i \rangle \geq \sqrt{1-\varepsilon}$. If $\|z\|^2 \geq a$ then there must exist i such that $\langle z, v_i \rangle \geq \sqrt{a(1-\varepsilon)}$. For a fixed i and for any $t > 0$ we have

$$\begin{aligned} \mathbb{P}[\langle z, v_i \rangle \geq \sqrt{a(1-\varepsilon)}] &= \mathbb{P}[\exp(t\langle z, v_i \rangle) \geq \exp(t\sqrt{a(1-\varepsilon)})] \\ &\leq \mathbb{E}[\exp(\langle z, tv_i \rangle)] \exp(-t\sqrt{a(1-\varepsilon)}) \\ &\leq \exp\left(\frac{1}{2}\sigma^2 t^2\right) \exp(-t\sqrt{a(1-\varepsilon)}) \end{aligned}$$

setting $t = \sqrt{a(1-\varepsilon)}/\sigma^2$,

$$= \exp\left(-\frac{a(1-\varepsilon)}{2\sigma^2}\right).$$

The result now follows by a union bound over all i . □

The following theorem gives a sufficient condition for contiguity in terms of the sub-Gaussian property.

Theorem 3.4.8 (sub-Gaussian method). *Let G be a compact group and let Ψ be a list of frequencies. Suppose $Z^{(G, \Psi)}$ (defined above) is sub-Gaussian with covariance proxy $\sigma^2 I$. If $\lambda < 1/\sigma$ then $\text{GSynch}(\lambda, G, \Psi)$ is contiguous to $\text{GSynch}(0, G, \Psi)$.*

Proof. Note that $\sum_u Z_u$ is sub-Gaussian with covariance proxy $n\sigma^2 I$. From above we have

$$\begin{aligned}
\mathbb{E}_{Q_n} \left(\frac{dP_n}{dQ_n} \right)^2 &= \mathbb{E} \exp \left(\frac{\lambda^2}{2n} \left\| \sum_u Z_u \right\|^2 \right) \\
&= \int_0^\infty \mathbb{P} \left[\exp \left(\frac{\lambda^2}{2n} \left\| \sum_u Z_u \right\|^2 \right) \geq M \right] dM \\
&= \int_0^\infty \mathbb{P} \left[\|Z\|^2 \geq \frac{2n \log M}{\lambda^2} \right] dM \\
&\leq 1 + \int_1^\infty \mathbb{P} \left[\|Z\|^2 \geq \frac{2n \log M}{\lambda^2} \right] dM \\
&\leq 1 + \int_1^\infty C \exp \left(-\frac{(1-\varepsilon) 2n \log M}{2n\sigma^2 \lambda^2} \right) dM \\
&= 1 + \int_1^\infty C \exp \left(-\frac{(1-\varepsilon) \log M}{\sigma^2 \lambda^2} \right) dM \\
&= 1 + \int_1^\infty C M^{-(1-\varepsilon)/(\sigma^2 \lambda^2)} dM
\end{aligned}$$

which is finite provided that $(1-\varepsilon)/(\sigma^2 \lambda^2) > 1$. The second inequality uses Lemma 3.4.7. Since ε was arbitrary, this completes the proof. \square

Note that $\mathbb{E}[Z^{(G, \Psi)}] = 0$ (which is a requirement for sub-Gaussianity) is automatically satisfied; this follows from the Peter–Weyl theorem on orthogonality of matrix entries, which we will discuss in more detail in Section 3.4.5 (see also Section 2.3.1).

3.4.4 Applications of the sub-Gaussian method

In this section we use Theorem 3.4.8 to prove contiguity for some specific synchronization problems.

First we consider $U(1)$ with a single frequency. It was predicted in [82] that the statistical threshold for this problem should be the spectral threshold $\lambda = 1$; we now confirm this.

Theorem 3.4.9 ($U(1)$ with one frequency). *Consider the group $U(1)$ of unit-norm complex numbers under multiplication. Identify each element $e^{i\theta}$ of $U(1)$ with its angle θ . Let Ψ_1 be the list containing the single frequency $\rho : \theta \mapsto e^{i\theta}$. For any $\lambda < 1$, $\text{GSynch}(\lambda, U(1), \Psi_1)$ is contiguous to $\text{GSynch}(0, U(1), \Psi_1)$.*

Proof. We have $Z^{(U(1), \Psi_1)} = \sqrt{2}(\cos \theta, \sin \theta)$ where θ is drawn uniformly from $[0, 2\pi]$. Towards showing sub-Gaussianity we have, for any $v \in \mathbb{R}^2$,

$$\mathbb{E} \exp(\langle Z^{(U(1), \Psi_1)}, v \rangle) = \mathbb{E}_\theta \exp(\sqrt{2} v_1 \cos \theta + \sqrt{2} v_2 \sin \theta) = \mathbb{E}_\theta \exp(\sqrt{2} \|v\| \cos \theta).$$

Letting $w = \|v\|$, it is sufficient to show for all $w \geq 0$,

$$\mathbb{E}_\theta \exp(\sqrt{2} w \cos \theta) \leq \exp\left(\frac{1}{2} w^2\right).$$

This can be verified numerically but we also provide a rigorous proof. Using the Taylor expansion of \exp and the identity

$$\mathbb{E}_\theta [\cos^k \theta] = \begin{cases} \frac{(k-1)!!}{k!!} & k \text{ even} \\ 0 & k \text{ odd} \end{cases}$$

we have

$$\begin{aligned}
\mathbb{E}_\theta \exp \left(\sqrt{2} w \cos \theta \right) &= \mathbb{E}_\theta \sum_{k \geq 0} \frac{2^{k/2} w^k \cos^k \theta}{k!} = \sum_{k \geq 0} \frac{2^k w^{2k} \mathbb{E}_\theta \cos^{2k} \theta}{(2k)!} \\
&= \sum_{k \geq 0} \frac{2^k w^{2k} (2k-1)!!}{(2k)!(2k)!!} = \sum_{k \geq 0} \frac{2^k w^{2k}}{(2k)!!(2k)!!} \\
&\leq \sum_{k \geq 0} \frac{w^{2k}}{(2k)!!} = \sum_{k \geq 0} \frac{w^{2k}}{2^k k!} = \exp \left(\frac{1}{2} w^2 \right).
\end{aligned}$$

The exchange of expectation and infinite sum is justified by the Fubini–Tonelli theorem, provided we can show absolute convergence:

$$\sum_{k \geq 0} \mathbb{E}_\theta \left| \frac{2^{k/2} w^k \cos^k \theta}{k!} \right| \leq \sum_{k \geq 0} \left| \frac{2^{k/2} w^k}{k!} \right|$$

which converges by the ratio test. □

We now add a second frequency.

Example 3.4.10 ($U(1)$ with two frequencies). Consider again $U(1)$ but now let Ψ_2 be the list of two frequencies: $\rho_1 : \theta \mapsto e^{i\theta}$ and $\rho_2 : \theta \mapsto e^{2i\theta}$. For any $\lambda < \lambda^* \approx 0.9371$ (numerically computed), $\text{GSynch}(\lambda, U(1), \Psi_2)$ is contiguous to $\text{GSynch}(0, U(1), \Psi_2)$.

(We use “example” rather than “theorem” to indicate results that rely on numerical computations.) Although we are unable to show that the spectral threshold is optimal, note that this rules out the possibility that the threshold for two frequencies drops to $1/\sqrt{2}$ (which is what we would have if one could perfectly synthesize the frequencies). We expect that the true statistical threshold for this problem is $\lambda = 1$ and that our results are not tight here. We now move on to the case of \mathbb{Z}/L .

Details. We have $Z^{(U(1), \Psi_2)} = \sqrt{2}(\cos \theta, \sin \theta, \cos(2\theta), \sin(2\theta))$. Our threshold is $\lambda^* = 1/\sigma^*$ where

$$(\sigma^*)^2 = \sup_v \frac{2}{\|v\|} \log \mathbb{E} \left(\langle Z^{(U(1), \Psi_2)}, v \rangle \right) = \sup_v \frac{2}{\|v\|} \log \mathbb{E}_\theta \left(\sqrt{2}(v_1 \cos \theta + v_2 \sin \theta + v_3 \cos(2\theta) + v_4 \sin(2\theta)) \right).$$

By the change of variables $\theta \mapsto \theta - \theta_0$ (for some θ_0) we can rotate (v_1, v_2) arbitrarily, and so we can take $v_2 = 0$ and $v_1 \geq 0$ without loss of generality. By grid search over v_1, v_3, v_4 , we see numerically that the maximizer is $v^* = (0.720, 0, 0.559, 0)$ which yields contiguity for all $\lambda < \lambda^* \approx 0.937$. \square

Example 3.4.11 (\mathbb{Z}/L with one frequency). Now consider $\mathbb{Z}/L = \{0, 1, \dots, L-1\} \pmod{L}$ with $L \geq 2$ and Ψ_1 the list of one frequency: $j \mapsto \exp(2\pi i j/L)$. For $L = 3$, we have contiguity $\text{GSynch}(\lambda, \mathbb{Z}/L, \Psi_1) \triangleleft \text{GSynch}(0, \mathbb{Z}/L, \Psi_1)$ for all $\lambda < \lambda_3^* \approx 0.961$. For $L = 2$ and all $L \geq 4$, we have contiguity for all $\lambda < 1$.

Details. This is shown numerically in a manner similar to the examples above. Of course we cannot test this for all values of L , but we conjecture that the $\lambda^* = 1$ trend continues indefinitely. \square

We have that the spectral threshold is optimal for all L except 3. It is surprising that $L = 3$ is an exception here, but we expect that this is a weakness of our techniques and that the true threshold for $L = 3$ is also $\lambda = 1$.

Finally we give a coarse but general result for any group with any number of frequencies.

Theorem 3.4.12 (any group, any frequencies). *Let G be any group and let Ψ be any list of frequencies, with $D = \sum_{\rho \in \Psi} \beta_\rho d_\rho^2$. If $\lambda < 1/\sqrt{D}$ then $\text{GSynch}(\lambda, G, \Psi)$ is contiguous to $\text{GSynch}(0, G, \Psi)$.*

Proof. Since our representations ρ are unitary, we have $\|\rho(g)\|_F^2 = d_\rho$ for any $g \in G$, and so $\|Z^{(G, \Psi)}\|^2 = D$. This means for any vector v we have $|\langle Z^{(G, \Psi)}, v \rangle| \leq \|Z^{(G, \Psi)}\| \|v\| = \sqrt{D} \|v\|$. By Hoeffding's Lemma this implies the sub-Gaussian condition $\mathbb{E} \exp(\langle Z^{(G, \Psi)}, v \rangle) \leq \exp(\frac{1}{2} D \|v\|^2)$. \square

3.4.5 The conditioning method for finite groups

Here we give an alternative method to show contiguity for finite groups, based on the conditioning method of [19] (see Section 3.3.3). Let G be a finite group with $|G| = L$. Again

take all the λ 's to be equal: $\lambda_\rho = \lambda$ for all ρ . For $a, b \in G$, let $N_{ab} = |\{u \mid g_u = a, g'_u = b\}|$. Rewrite the second moment in terms of N_{ab} :

$$\mathbb{E}_{X, X'} \prod_{\rho} \exp \left(\frac{\beta_{\rho} \lambda_{\rho}^2 d_{\rho}}{2n} \|X_{\rho}^* X'_{\rho}\|_F^2 \right) = \mathbb{E}_{X, X'} \exp \left(\frac{\lambda^2}{2n} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c (X_{\rho}^* X'_{\rho})_c^2 \right)$$

where c ranges over all (real-valued) coordinates of entries of $\rho(g)$ (e.g. imaginary part of top right entry)

$$\begin{aligned} &= \mathbb{E}_{g, g'} \exp \left(\frac{\lambda^2}{2n} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c \left(\sum_u \rho(g_u^{-1} g'_u)_c \right)^2 \right) \\ &= \mathbb{E}_N \exp \left(\frac{\lambda^2}{2n} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c \left(\sum_{ab} N_{ab} \rho(a^{-1}b)_c \right)^2 \right) \\ &= \mathbb{E}_N \exp \left(\frac{1}{n} N^{\top} A N \right) \\ &= \mathbb{E}_N \exp (Y^{\top} A Y) \end{aligned}$$

where $Y = \frac{\vec{N} - n\bar{\alpha}}{\sqrt{n}}$, $\bar{\alpha} = \frac{1}{L^2} \mathbb{1}_{L^2}$, and A is the $L^2 \times L^2$ matrix

$$A_{ab, a'b'} = \frac{\lambda^2}{2} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c \rho(a^{-1}b)_c \rho(a'^{-1}b')_c.$$

To justify the last step, note that $\bar{\alpha}$ is in the kernel of A because all row- and column-sums of A are zero. This follows from the Peter–Weyl theorem on orthogonality of matrix coefficients, which we will discuss in more detail shortly. By Proposition 5 in [19] (Proposition 3.3.4 in this thesis) we have contiguity provided that

$$\sup_{\alpha} \frac{\alpha^{\top} A \alpha}{D(\alpha, \bar{\alpha})} < 1$$

where α ranges over (vectorized) $L \times L$ matrices with all row- and column-sums equal to $\frac{1}{L}$.

Theorem 3.4.13 (conditioning method). *Let G be a finite group of order L and let Ψ be a list of frequencies. Let \tilde{A} be the $L^2 \times L^2$ matrix $\tilde{A}_{ab,a'b'} = \frac{1}{2} \sum_{\rho \in \Psi} \beta_\rho d_\rho \sum_c \rho(a^{-1}b)_c \rho(a'^{-1}b')_c$ where $a, b, a', b' \in G$ and c ranges over (real) coordinates of matrix entries. Let $D(u, v)$ denote the KL divergence between two vectors: $D(u, v) = \sum_i u_i \log(u_i/v_i)$. If*

$$\lambda < \left[\sup_{\alpha} \frac{\alpha^\top \tilde{A} \alpha}{D(\alpha, \bar{\alpha})} \right]^{-1/2}$$

then $\text{GSynch}(\lambda, G, \Psi)$ is contiguous to $\text{GSynch}(0, G, \Psi)$. Here α ranges over (vectorized) $L \times L$ matrices with all row- and column-sums equal to $\frac{1}{L}$.

A finite group has only a finite number of irreducible representations (over \mathbb{C}), so let us now specialize to the case where our list Ψ contains all of them (excluding the trivial representation, and only taking one representation per conjugate pair). Expand the numerator:

$$\alpha^\top A \alpha = \frac{\lambda^2}{2} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c \left(\sum_{ab} \alpha_{ab} \rho(a^{-1}b)_c \right)^2 = \frac{\lambda^2}{2} \sum_{\rho} \beta_{\rho} d_{\rho} \sum_c \left(\sum_h \alpha_h \rho(h)_c \right)^2$$

where $\alpha_h = \sum_{(a,b) \in S_h} \alpha_{ab}$ and $S_h = \{(a, b) \mid a^{-1}b = h\}$. We now appeal to the Peter–Weyl theorem on the orthogonality of matrix coefficients: the basis functions $\chi_{\rho ij}(h) = \sqrt{d_{\rho}^{\mathbb{C}}} \rho(h)_{ij}$ (for all irreducible ρ over \mathbb{C} and matrix entries i, j) form an orthonormal basis for \mathbb{C}^G under the Hermitian inner product $\langle f_1, f_2 \rangle \triangleq \frac{1}{L} \sum_{h \in G} f_1(h) \overline{f_2(h)}$. Here $d_{\rho}^{\mathbb{C}}$ is the dimension as a complex representation, which is the same as d_{ρ} for real- and complex-type but equal to $2d_{\rho}$ for quaternionic-type. This means the above can be thought of as projecting the vector $\{\alpha_h\}_{h \in G}$ onto these basis elements and then computing the ℓ^2 norm of the result. By the basis-invariance of the ℓ^2 norm, we can rewrite the above as

$$\frac{\lambda^2 L^2}{2} \left[\frac{1}{L} \sum_h \alpha_h^2 - \left(\frac{1}{L} \sum_h \alpha_h \right)^2 \right] = \frac{\lambda^2 L}{2} \left[\sum_h \alpha_h^2 - \frac{1}{L} \right].$$

The second term here corrects for the fact that the trivial representation did not appear in our original expression. Note that the factor of $\beta = 2$ for complex representations corrects

for the fact that we were only using one representation per conjugate pair. The factor of $\beta = 4$ for quaternionic representations corrects for the fact that we were thinking of these representations as being defined over \mathbb{H} rather than \mathbb{C} ; the corresponding complex representation has dimension twice as large and represents each quaternion value by the following 2×2 complex matrix:

$$a + bi + cj + dk \quad \leftrightarrow \quad \begin{pmatrix} a + bi & c + di \\ -c + di & a - bi \end{pmatrix}.$$

(One factor of 2 comes from the fact that $d_\rho^{\mathbb{C}} = 2d_\rho$ and the other factor of 2 comes from the fact that the squared-Frobenius norm of this 2×2 matrix is twice the squared-norm of the associated quaternion.)

Note that we now have exactly the same optimization problem that we arrived at for the truth-or-Haar model with λ in place of \tilde{p} , so we can apply Proposition 3.3.5 to immediately obtain the following.

Theorem 3.4.14. *Let G be a finite group of order $L \geq 2$ and let Ψ_{all} be the list of **all** frequencies (excluding the trivial one and only taking one from each conjugate pair). If for all $\rho \in \Psi_{\text{all}}$,*

$$\lambda_\rho < \lambda_L^* \triangleq \sqrt{\frac{2(L-1)\log(L-1)}{L(L-2)}}$$

then $\text{GSynch}(\{\lambda_\rho\}, G, \Psi_{\text{all}})$ is contiguous to $\text{GSynch}(0, G, \Psi_{\text{all}})$. For $L = 2$, $\lambda_2^ = 1$ (the limit value of the $0/0$ expression).*

Here we have used the monotonicity of the second moment: if we show contiguity when all the λ_ρ 's are equal to some λ , and we then decrease some of the individual λ_ρ 's, we will still have contiguity.

Interestingly, our critical value λ_L^* is the same as our critical value \tilde{p}_L^* from the truth-or-Haar model. As discussed previously, this matches the spectral threshold $\lambda = 1$ only when $L = 2$. However, for small values of L , our λ_L^* is quite close to 1 (see the table in Section 3.3).

Also note that when $L = 3$, Theorem 3.4.14 matches (and proves rigorously) the numerical

value $\lambda^* \approx 0.961$ of Example 3.4.11 (obtained via the sub-Gaussian method). Note that when $L = 3$, \mathbb{Z}/L only has one frequency, so these two results apply to the same problem. However, we see that the conditioning method gains no advantage over the sub-Gaussian method in this case. This seems to be true in general for synchronization problems because there are no particularly ‘bad’ values for the spike due to symmetry of the group.

3.4.6 Upper bound via exhaustive search

In this subsection, we analyze the performance of exhaustive search in the Gaussian Synchronization Model. Specifically, we show:

Theorem 3.4.15. *Let G be a finite group of order L and let Ψ be a list of frequencies. If*

$$\sum_{\rho \in \Psi} \lambda_\rho^2 \beta_\rho d_\rho^2 > 4 \log L$$

there is a computationally inefficient algorithm that can distinguish between the spiked and unspiked models.

See Corollary 3.4.17 below for a simplification in the case of all frequencies.

Let $P_n = \text{GSynch}_n(\{\lambda_\rho\}, G, \Psi)$ and let $Q_n = \text{GSynch}_n(0, G, \Psi)$. By the Neyman–Pearson lemma, the most powerful test statistic for distinguishing P_n from Q_n is the likelihood ratio $\frac{dP_n}{dQ_n}$. Similarly to [20] we use the following modified likelihood ratio. For $g \in G^n$, let $V_\rho(g)$ be the $nd_\rho \times d_\rho$ matrix formed by stacking the matrices $\rho(g_u)$. Given $Y = \{Y_\rho\}$ drawn from either P_n or Q_n , our test is to compute $T = \max_{g \in G^n} T(g)$ where

$$T(g) = \sum_{\rho \in \Psi} \lambda_\rho \beta_\rho d_\rho \text{Tr}(V_\rho(g)^* Y_\rho V_\rho(g)).$$

If $T \geq \sum_\rho n \lambda_\rho^2 \beta_\rho d_\rho^2 - \sqrt{n \log n}$ then we answer ‘ P_n ’; otherwise, ‘ Q_n .’ The definition of $T(g)$ is motivated by the computation of $\frac{dP_n}{dQ_n}$ in Section 3.4; in fact, $T(g)$ is equal (up to constants) to $\frac{dP_n(Y|g)}{dQ_n(Y)}$. Note that this test is not computationally-efficient because it involves testing all

possible solutions $g \in G^n$. The best computationally-efficient test that we know of is PCA (or AMP), which succeeds if and only if at least one λ_ρ exceeds 1.

The proof of Theorem 3.4.15 will require the following computation.

Lemma 3.4.16. *Let V be a fixed $nd \times d$ matrix where each $d \times d$ block is unitary of some type ($\mathbb{R}, \mathbb{C}, \mathbb{H}$). Let W be an $nd \times nd$ Hermitian Gaussian matrix of the corresponding type (GOE, GUE, GSE , respectively). Let β be 1, 2, 4 (respectively) depending on the type. Then $\text{Tr}(V^* W V) \sim \mathcal{N}(0, 2n^2 d / \beta)$.*

Proof. Let u, v index the $d \times d$ blocks, and let a, b, c index the entries within each block.

$$\begin{aligned}
\text{Tr}(V^* W V) &= \sum_{u,v} \text{Tr}(V_u^* W_{uv} V_v) \\
&= \sum_{u < v} 2 \text{Tr} \Re(V_u^* W_{uv} V_v) + \sum_u \text{Tr}(V_u^* W_{uu} V_u) \\
&= \sum_{u < v} \sum_{a,b,c} 2 \Re[(V_u^*)_{ab} (W_{uv})_{bc} (V_v)_{ca}] + \sum_u \sum_{a,b,c} (V_u^*)_{ab} (W_{uu})_{bc} (V_u)_{ca} \\
&= \sum_{u < v} \sum_{a,b,c} 2 \Re[(V_u^*)_{ab} (W_{uv})_{bc} (V_v)_{ca}] + \sum_u \sum_{a,b < c} 2 \Re[(V_u^*)_{ab} (W_{uu})_{bc} (V_u)_{ca}] \\
&\quad + \sum_u \sum_{a,b} (V_u^*)_{ab} (W_{uu})_{bb} (V_u)_{ba} \\
&= \sum_{u < v} \sum_{b,c} 2 \mathcal{N}(0, |\sum_a (V_u^*)_{ab} (V_v)_{ca}|^2 / \beta) + \sum_u \sum_{b < c} 2 \mathcal{N}(0, |\sum_a (V_u^*)_{ab} (V_u)_{ca}|^2 / \beta) \\
&\quad + \sum_u \sum_b \mathcal{N}(0, 2 |\sum_a (V_u^*)_{ab} (V_u)_{ba}|^2 / \beta) \\
&= \mathcal{N}(0, 2 \sum_{u,v} \sum_{b,c} |\sum_a (V_u^*)_{ab} (V_v)_{ca}|^2 / \beta) \\
&= \mathcal{N}(0, 2 \sum_{u,v} \sum_{b,c} \sum_{a,a'} (\overline{V_u})_{ba} (V_v)_{ca} (\overline{V_v})_{ca'} (V_u)_{ba'} / \beta) \\
&= \mathcal{N}(0, 2 \sum_{u,v} \sum_{a,a'} \delta_{aa'} / \beta) \\
&= \mathcal{N}(0, 2n^2 d / \beta).
\end{aligned}$$

□

Proof of Theorem 3.4.15. We will now prove Theorem 3.4.15 by showing that (given the condition in the theorem) the test $T = \max_g T(g)$ (defined above) succeeds with probability $1 - o(1)$. If Y_ρ is drawn from the unspiked model $Q_n : \frac{1}{\sqrt{nd_\rho}}W$ then for any $g \in G^n$ we have

$$T(g) = \sum_{\rho} \lambda_{\rho} \beta_{\rho} d_{\rho} \frac{1}{\sqrt{nd_{\rho}}} \text{Tr}(V_{\rho}(g)^* W V_{\rho}(g)) = \sum_{\rho} \lambda_{\rho} \beta_{\rho} d_{\rho} \frac{1}{\sqrt{nd_{\rho}}} \mathcal{N}\left(0, \frac{2n^2 d_{\rho}}{\beta_{\rho}}\right) = \mathcal{N}\left(0, \sum_{\rho} 2n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2\right).$$

If instead Y_ρ is drawn from the spiked model $P_n : Y_{\rho} = \frac{\lambda_{\rho}}{n} X_{\rho} X_{\rho}^* + \frac{1}{\sqrt{nd_{\rho}}}W$ and we take g to be the ground truth g^* (so that $V_{\rho}(g) = X_{\rho}$), we have

$$T(g^*) = \sum_{\rho} \lambda_{\rho} \beta_{\rho} d_{\rho} \text{Tr}\left(\frac{\lambda_{\rho}}{n} X_{\rho}^* X_{\rho} X_{\rho}^* X_{\rho} + \frac{1}{\sqrt{nd_{\rho}}} V_{\rho}^* W V_{\rho}\right) = \sum_{\rho} n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 + \mathcal{N}\left(0, \sum_{\rho} 2n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2\right).$$

Using the Gaussian tail bound $\mathbb{P}[\mathcal{N}(0, \sigma^2) \geq t] \leq \exp\left(\frac{-t^2}{2\sigma^2}\right)$, we have that under the spiked model,

$$P_n \left[T \leq \sum_{\rho} n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 - \sqrt{n \log n} \right] \leq \exp \left(\frac{-n \log n}{2} \left(\sum_{\rho} 2n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 \right)^{-1} \right) = o(1).$$

Taking a union bound over all L^n choices for $g \in G^n$, we have that under the unspiked model,

$$\begin{aligned} Q_n \left[T \geq \sum_{\rho} n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 - \sqrt{n \log n} \right] &\leq L^n \exp \left(-\frac{1}{2} \left(\sum_{\rho} n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 - \sqrt{n \log n} \right)^2 \left(\sum_{\rho} 2n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 \right)^{-1} \right) \\ &= \exp \left(n \log L - \frac{1}{4} \sum_{\rho} n \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 + \mathcal{O}(\sqrt{n \log n}) \right) \end{aligned}$$

which is $o(1)$ provided $\sum_{\rho} \lambda_{\rho}^2 \beta_{\rho} d_{\rho}^2 > 4 \log L$. □

We can simplify the statement of the theorem in the case where all frequencies are present.

We note that if Ψ_{all} is the list of all frequencies then

$$\sum_{\rho \in \Psi_{\text{all}}} \beta_{\rho} d_{\rho}^2 = L - 1.$$

This follows from the “sum-of-squares” formula from the representation theory of finite groups. (The extra 1 comes from the fact that we don’t use the trivial representation in our list. The factor of $\beta = 2$ for complex-type representations accounts for the fact that we only use one representation per conjugate pair. The factor of $\beta = 4$ for quaternionic-type representations accounts for the fact that the complex dimension is twice the quaternionic dimension.) We therefore have the following corollary.

Corollary 3.4.17. *Let G be a finite group of order $L \geq 2$ and let Ψ_{all} be the list of **all** frequencies (excluding the trivial one and only taking one from each conjugate pair). If*

$$\lambda > \sqrt{\frac{4 \log L}{L - 1}}$$

*then an inefficient algorithm can distinguish the spiked and unspiked models, and so $\text{GSynch}(\lambda, G, \Psi_{\text{all}})$ is **not** contiguous to $\text{GSynch}(0, G, \Psi_{\text{all}})$.*

Note that for large L this differs from the lower bound of Theorem 3.4.14 by a factor of $\sqrt{2}$. As for the truth-or-Haar model, we expect that the upper bound is asymptotically tight and that the lower bound can be improved by a factor of $\sqrt{2}$ (asymptotically) using a more sophisticated conditioning method of the author and others [119]; here the event conditioned on depends not only on the signal but also on the noise.

Also note that the right-hand side matches Theorem 3.3.3 (upper bound for the truth-or-Haar model); interestingly, both our lower and upper bounds indicate that the all-frequencies Gaussian model behaves like the truth-or-Haar model with λ in place of \tilde{p} . In particular, we again see that an inefficient algorithm can beat the spectral threshold once $L \geq 11$.

Chapter 4

Orbit recovery: statistical limits

This chapter is adapted from joint work with Afonso Bandeira, Ben Blum-Smith, Amelia Perry, and Jonathan Weed [15]. An upcoming expanded version of [15] will include Joe Kileel as an additional author.

4.1 Introduction

Many computational problems throughout the sciences exhibit rich symmetry and geometry, especially in fields such as signal and image processing, computer vision, and microscopy. This is exemplified in cryo-electron microscopy (cryo-EM) [5, 142, 115], an imaging technique in structural biology that was recently awarded the 2017 Nobel Prize in Chemistry. This technique seeks to estimate the structure of a large biological molecule, such as a protein, from many noisy tomographic projections (2-dimensional images) of the molecule from random unknown directions in 3-dimensional space.

In cryo-EM, our signal of interest is the density θ of the molecule, considered as an element of the vector space of functions on \mathbb{R}^3 . We have access to observations of the following form: our microscopy sample contains many rotated copies $R_i\theta$ of the molecule, where $R_i \in \text{SO}(3)$ are random, unknown 3D rotations, and we observe the noisy projections $\Pi(R_i\theta) + \xi_i$, where Π denotes tomographic projection (in a fixed direction) and ξ_i is a large noise contribution,

perhaps Gaussian. This specific problem motivates the following general abstraction.

Fix a compact group G acting (by orthogonal transformations) on a vector space V . Throughout, the vector space will be taken to be \mathbb{R}^p and the group can be thought of as a subgroup of $O(p)$, the orthogonal group¹. Let $\theta \in V$ be the signal we want to estimate. We receive noisy measurements of its orbit as follows: for $i = 1, \dots, n$ we observe a sample of the form

$$y_i = g_i \cdot \theta + \xi_i$$

where g_i is drawn randomly (in Haar measure²) from G and $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$ is noise. The goal is to recover the *orbit* of θ under the action of G . We refer to this task as the *orbit recovery* problem.

This abstraction, already a rich object of study, neglects the tomographic projection in cryo-EM; we will also study a generalization of the problem which allows such a projection. We will also consider the additional extension of *heterogeneity* [83, 93, 94, 35], where mixtures of signals are allowed: we have K signals $\theta_1, \dots, \theta_K$, and each sample $y_i = g_i \cdot \theta_{k_i} + \xi_i$ comes from a random choice $1 \leq k_i \leq K$ of which signal is observed. This extension is of paramount importance for cryo-EM in practice, since the laboratory samples often contain one protein in multiple conformations, and understanding the range of conformations is key to understanding the function of the protein.

4.1.1 Prior work

Several special cases of the orbit recovery problem have been studied for their theoretical and practical interest. Besides cryo-EM, another such problem is *multi-reference alignment* (MRA) [17, 13, 118], a problem from signal processing [155, 123] with further relevance to

¹We alert the reader to the fact that we will use $O(p)$ to refer to the group of orthogonal matrices in dimension p and $O(g(n))$ as the standard big-O notation: $f(n) = O(g(n))$ if and only if there exists a constant C such that $f(n) \leq Cg(n)$ for all n sufficiently large. It will be clear from context which one is meant.

²We note that any distribution of g_i can be reduced to Haar by left multiplying y_i by a Haar-distributed group element. However, as illustrated in [3], it is sometimes possible to exploit deviations from Haar measure.

structural biology [60, 144]. In this problem, one observes noisy copies of a signal $\theta \in \mathbb{R}^p$, each with its coordinates permuted by a random cyclic shift. This is an example of the orbit recovery problem when G is taken to be the cyclic group \mathbb{Z}/p acting by cyclic permutations of the coordinates. Since the cyclic group \mathbb{Z}/p is simpler than $\text{SO}(3)$, understanding MRA has been seen as a useful stepping stone towards a full statistical analysis of cryo-EM.

Many prior methods for orbit recovery problems employ the so-called *synchronization* approach where the unknown group elements g_i are estimated based on pairwise comparison of the samples y_i . If the samples were noiseless, one would have $g_i g_j^{-1} y_j = y_i$; thus noisy samples still give some weak information about $g_i g_j^{-1}$. Synchronization is the problem of using such pairwise information to recover all the group elements g_i (up to a global right-multiplication by some group element). Once the group elements g_i are known, the underlying signal can often be easily recovered.

The synchronization approach has proven to be effective both in theory and practice when the noise is sufficiently small. However, once the noise level is large, no consistent estimation of the group elements g_i is possible [7]. Moreover, it is the high-noise regime that is the practically relevant one for many applications, including cryo-EM, where the presence of large noise is a primary obstruction to current techniques [140]. As a result, recent work has focused on approaches to cryo-EM and MRA which provably succeed even in the large-noise limit. One striking finding of this line of work is that the sample complexity of the statistical estimation problem increases drastically as the noise level increases. For instance, for the multi-reference alignment problem with noise variance σ^2 , consistent estimation of typical signals requires $\Omega(\sigma^6)$ samples [13, 4], with significantly worse rates for atypical signals. By contrast, when σ^2 is smaller than some threshold, only $O(\sigma^2)$ samples are required. Moreover, in contrast with the $O(\sigma^2)$ rate—which would hold even in the absence of a group action—the $\Omega(\sigma^6)$ bound obtained in previous works depends on particular properties of the cyclic group. In this work, we significantly extend this prior work by determining the sample complexity of the estimation problem in the high-noise regime for *general* groups.

The leading theoretical framework for the high-noise regime is the *invariant features*

approach [13, 26, 118, 35, 98]. This approach has a long history in the signal processing literature [86, 133, 134] and is analogous to the well known “method of moments” in statistics [148]. In brief, the invariant features approach bypasses entirely the problem of estimating the group elements and focuses instead on estimating features of the signal which are preserved by the action of the group. So long as these invariant features uniquely specify the orbit of the original signal, the invariants are sufficient statistics for the problem of recovering the orbit of the original signal. This simple approach yields optimal dependence of the sample complexity on the noise level for the multi-reference alignment problem [13, 118].

The application of invariant features to cryo-EM dates back to 1980 with the work of Kam [86], who partially solved cryo-EM by means of degree-2 invariant features, reducing the unknown molecule structure to a collection of unknown orthogonal matrices. Subsequent work has explored methods to estimate these orthogonal matrices [29], including recent work showing how two noiseless tomographic projections suffice to recover these orthogonal matrices [98]. Our work can be viewed as a degree-3 extension of Kam’s method that fully solves cryo-EM while circumventing the orthogonal retrieval issue, and without requiring any noiseless projections. Our approach is *ab initio*, i.e. it does not require an initial guess of what the molecule looks like and thus cannot suffer from *model bias*, which is a documented phenomenon [42] where the initial guess can have a significant effect on the result. *Ab initio* estimates are particularly useful to serve as a model-free starting point for popular iterative refinement algorithms such as RELION [136].

Throughout, we focus on the case where the group elements are Haar-distributed. In the basic orbit recovery problem (projection), any distribution of g_i can be reduced to Haar by left-multiplying each sample y_i by a Haar-distributed group element. However, as illustrated in [3], it is sometimes possible to exploit deviations from Haar measure. The situation is different when we add projection to the problem setup, as is the case with Cryo-EM; if the viewing direction is not distributed uniformly then there may exist parts of the molecule that are systematically imaged less than others, which can cause serious difficulties in reconstruction.

The present paper connects the orbit recovery problem to the invariant theory of groups, a classical and well-developed branch of algebra (see for example [84, 61, 143, 55]). Invariant theory’s traditional goal is to describe the ring of all polynomial functions on a vector space that are invariant under the action of a group – the *invariant algebra*. Since the 19th century, culminating in the pioneering work of David Hilbert [76, 77], it has been known that the invariant algebra is finitely generated in many cases of interest, and so a fundamental problem has been to bound the degrees of the generators. In 2002, Derksen and Kemper [55] introduced the notion of a *separating algebra* – a subring of the invariant algebra that separates all orbits of the group action which are separated by the full invariant algebra. Our connection to orbit recovery motivates the question of bounding the degree required to generate a separating algebra (see Section 4.3.4), a problem which has been recently studied [88, 62]. Our work also motivates the question of bounding the degree at which the *field* of invariant rational functions is generated as a field (see Section 4.3.3), which does not appear to have been the focus of research attention before.

4.1.2 Our contributions

In this chapter we extend the results of [13] and show that the method of moments yields optimal sample complexity for orbit recovery problems over *any* compact group. Specifically, we show that optimal sample complexity is achieved by an algorithm that estimates the moments from the samples and then solves a polynomial system of equations in order to find a signal θ that would produce such moments. As the sample complexity depends on the number of moments used, this gives rise to the algebraic question of how many moments suffice to determine the orbit of θ . Using tools from invariant theory and algebraic geometry, we investigate this question for various success criteria and obtain sharp results in a number of settings. Our main focus is on the case where the signal is assumed to be generic and the goal is to output a finite list of signals, one of which is the truth. In this case we give a simple efficient algorithm for determining the number of moments required for any given orbit recovery problem. The main step of the algorithm is to compute the rank of a particular

Jacobian matrix.

We note that ours is an information-theoretic result rather than a computational one because even with knowledge of the number of moments required, estimating the original signal still requires solving a particular polynomial system of equations and we do not attempt to give a computationally-efficient method for this. There are fast non-convex heuristic methods to solve these systems in practice [35] but we leave for future work the question of analyzing such methods rigorously and exploring whether or not they reach the information-theoretic limits determined in this paper. For the case of finite groups, another efficient method for solving the polynomial system is via tensor decomposition, which has been analyzed for MRA [118].

Concrete results for problems such as MRA and cryo-EM are in Section 4.4.

4.1.3 Motivating examples

In addition to the examples of MRA and cryo-EM, it is helpful to have the following motivating examples in mind:

1. Learning a “bag of numbers”: let G be the symmetric group S_p , acting on $V = \mathbb{R}^p$ by permutation matrices. Thus we observe random rearrangements of the entries of a vector, plus noise.
2. Learning a rigid body: let G be the rotation group $\text{SO}(p)$, acting on the matrix space $V = \mathbb{R}^{p \times m}$ by left-multiplication. We imagine the columns of our matrix as vertices defining a rigid body; thus we observe random rotations of this rigid body (with vertices labeled) plus noise.
3. S^2 registration: Let $S^2 \subseteq \mathbb{R}^3$ be the unit sphere. Let V be the finite-dimensional vector space of functions on $S^2 \rightarrow \mathbb{R}$ that are *band-limited*, i.e. linear combinations of spherical harmonics up to some fixed degree (spherical harmonics are the appropriate “Fourier basis” for functions on the sphere); let $\theta \in V$ be such a function $S^2 \rightarrow \mathbb{R}$. Let $G = \text{SO}(3)$, acting on the sphere by 3-dimensional rotation; this induces an action on

V via $(g \cdot \theta)(x) = \theta(g^{-1} \cdot x)$. Thus we observe many noisy copies of a fixed function on the sphere, each rotated randomly.

4.1.4 Problem statement

Throughout, we consider a compact (topological) group G acting linearly, continuously, and orthogonally on a finite-dimensional real vector space $V = \mathbb{R}^p$. In other words, G acts on V via a linear representation $\rho : G \rightarrow O(V)$, and ρ itself is a continuous function. Here $O(V)$ denotes the space of real orthogonal $p \times p$ matrices. Let $\text{Haar}(G)$ denote Haar measure (i.e., the “uniform distribution”) on G . We define the *orbit recovery* problem as follows.

Problem 4.1.1 (orbit recovery). Let $V = \mathbb{R}^p$ and let $\theta \in V$ be the unknown signal. Let G be a compact group that acts linearly, continuously, and orthogonally on V . For $i \in [n] = \{1, 2, \dots, n\}$ we observe

$$y_i = g_i \cdot \theta + \xi_i$$

where $g_i \sim \text{Haar}(G)$ and $\xi_i \sim \mathcal{N}(0, \sigma^2 I_{p \times p})$, all independently. The goal is to estimate θ . Note that we can only hope to recover θ up to action by G ; thus we aim to recover the *orbit* $\{g \cdot \theta : g \in G\}$ of θ .

In practical applications, σ is often known in advance and, when it is not, it can generally be estimated accurately on the basis of the samples. We therefore assume throughout that σ is known and do not pursue the question of its estimation in this work.

Our primary goal is to study the sample complexity of the problem: how must the number of samples n scale with the noise level σ (as $\sigma \rightarrow \infty$ with G and V fixed) in order for orbit recovery to be statistically possible? All of our results will furthermore apply to a generalized orbit recovery problem (Problem 4.2.3) allowing for *projection* and *heterogeneity* (see Section 4.1.6).

Our work reveals that it is natural to consider several different settings in which to state the orbit recovery problem. We consider the following two decisions:

1. Do we assume that θ is a *generic* signal, or do we allow for a *worst-case* signal? (Here *generic* means that there is a measure-zero set of disallowed signals.)
2. Do we want to output a θ' such that θ' (approximately) lies in the orbit of θ (*unique recovery*), or simply a finite list $\theta_1, \dots, \theta_s$ of candidates such that one of them (approximately) lies in the orbit of θ (*list recovery*)?

The terminology “list recovery” is borrowed from the idea of *list decoding* in the theory of error-correcting codes [67]. By taking all combinations of the two options above, there are four different recovery criteria. Strikingly, these different recovery criteria can be very different in terms of sample complexity, as the following examples show (see Section 4.4 for more details):

1. *Multi-reference alignment (MRA)*: Recall that this is the case $G = \mathbb{Z}/p$ acting on $V = \mathbb{R}^p$ by cyclic shifts. It is known [118] that if θ is generic then unique recovery is possible with $O(\sigma^6)$ samples. However, for a worst-case θ , many more samples are required (even for list recovery); as shown in [13], there are some very particular infinite families of signals that cannot be distinguished without $\Omega(\sigma^{2p})$ samples. This illustrates a large gap in difficulty between the generic and worst-case problems.
2. *Learning a rigid body*: Let G be the rotation group $\text{SO}(p)$ acting on the matrix space $\mathbb{R}^{p \times m}$ by left multiplication. We imagine the columns of our matrix as vertices defining a rigid body; thus we observe random rotations of this rigid body (with vertices labeled) plus noise. With $O(\sigma^4)$ samples it is possible to recover the rigid body up to reflection, so that list recovery (with a list of size 2) is possible. However, unique recovery (even for a generic signal) requires drastically more samples: $\Omega(\sigma^{2p})$.

We will address all four recovery criteria but our main focus will be on the case of *generic list recovery*, as it is algebraically the most tractable to analyze. For the following reasons we also argue that it is perhaps the most practically relevant case. Clearly real-world signals are generic. Also, unique recovery is actually impossible in some practical applications; for

instance, in cryo-EM it is impossible to determine the chirality of the molecule. (However, we can hope for unique recovery if we work over the group $O(3)$ instead of $SO(3)$.) Furthermore, one could hope to use application-specific clues to pick the true signal out from a finite list; for instance, in cryo-EM we might hope that the spurious solutions in our finite list do not look like “reasonable” molecules and can be thrown out.

4.1.5 Method of moments

Our techniques rely on estimation of the following moments:

Definition 4.1.2 (moment tensor). The *order- d moment tensor* is $T_d(\theta) \triangleq \mathbb{E}_g[(g \cdot \theta)^{\otimes d}]$ where $g \sim \text{Haar}(G)$.

We can estimate $T_d(\theta)$ from the samples by computing $\frac{1}{n} \sum_{i=1}^n y_i^{\otimes d}$ plus a correction term to cancel bias from the noise terms (see the full paper [15] for details). The moments $T_d(\theta)$ are related to polynomials that are invariant under the group action, which brings us to the fundamental object in invariant theory:

Definition 4.1.3 (invariant ring). Let $\mathbf{x} = (x_1, \dots, x_p)$ be a set of coordinate functions on $V = \mathbb{R}^p$, i.e. a basis for the dual V^* , so that $\mathbb{R}[\mathbf{x}] \triangleq \mathbb{R}[x_1, \dots, x_p]$ is the ring of polynomial functions $V \rightarrow \mathbb{R}$. We have an action of G on $\mathbb{R}[\mathbf{x}]$ given by $(g \cdot f)(\cdot) = f(g^{-1}(\cdot))$. (If we fix a basis for V , we can think of \mathbf{x} as indeterminate variables corresponding to the entries of $\theta \in V$.) The *invariant ring* $\mathbb{R}[\mathbf{x}]^G \subseteq \mathbb{R}[\mathbf{x}]$ is the ring consisting of polynomials f that satisfy $g \cdot f = f$ for all $g \in G$. An element of the invariant ring is called an *invariant polynomial* (or simply an *invariant*). Invariant polynomials can be equivalently characterized as polynomials of the form $\mathbb{E}_g[g \cdot f]$ where $f \in \mathbb{R}[\mathbf{x}]$ is any polynomial and $g \sim \text{Haar}(G)$.

The two objects above are equivalent in the following sense. The moment tensor $T_d(\theta)$ contains the same information as the set of evaluations $f(\theta)$ for all $f \in \mathbb{R}[\mathbf{x}]^G$ that are homogeneous of degree d . In particular, for any such polynomial f , $f(\theta)$ is a linear combination of the entries of $T_d(\theta)$.

The following algebraic question will be of central importance: when do the values of invariant polynomials (of degree $\leq d$) of θ determine the orbit of θ (in the appropriate sense)? As we see below, the sample complexity of the statistical problem is completely characterized by the answer to this question.

Warm up: hypothesis testing

Consider for now the simple problem of distinguishing between two fixed hypotheses $\theta = \tau_1$ and $\theta = \tau_2$, where τ_1 and τ_2 are two fixed vectors in V . One method is to find an invariant polynomial f for which $f(\tau_1) \neq f(\tau_2)$ and to estimate $f(\theta)$ using the samples. The sample complexity of this procedure depends on the degree of f because if f has degree d , we need $O(\sigma^{2d})$ samples to accurately estimate $f(\theta)$. We have the following (see the full paper [15] for the proof).

Theorem 4.1.4 (distinguishing upper bound). *Fix $\tau_1, \tau_2 \in V$. If there exists a degree- d invariant polynomial $f \in \mathbb{R}[\mathbf{x}]^G$ with $f(\tau_1) \neq f(\tau_2)$ then, using $O(\sigma^{2d})$ samples, it is possible to distinguish between $\theta = \tau_1$ and $\theta = \tau_2$ with type-I and type-II error probabilities each at most $1/3$.*

Here, $O(\sigma^{2d})$ hides factors that depend on G (and its action on V), τ_1 , and τ_2 , but not σ ; we are most interested in how the sample complexity scales as σ becomes large (with everything else held fixed). The error probability $1/3$ is arbitrary and can be boosted by taking more samples.

Furthermore, we have a matching lower bound to show that the method of moments is optimal: the sample complexity is driven by the minimum degree of an invariant polynomial that separates τ_1 and τ_2 .

Theorem 4.1.5 (distinguishing lower bound). *Fix $\tau_1, \tau_2 \in V$. Let d^* be the smallest positive integer d for which $T_d(\tau_1) \neq T_d(\tau_2)$. Then $\Omega(\sigma^{2d^*})$ samples are required to distinguish between $\theta = \tau_1$ and $\theta = \tau_2$ with type-I and type-II error probabilities each at most $1/3$.*

See the full paper [15] for the proof.

Recovery

We now address the problem of recovering the signal θ from the samples. Our goal is to recover the orbit of θ , defined as follows.

Definition 4.1.6. For $\theta_1, \theta_2 \in V$, define an equivalence relation $\overset{G}{\sim}$ by letting $\theta_1 \overset{G}{\sim} \theta_2$ if there exists $g \in G$ such that $g \cdot \theta_1 = \theta_2$. The *orbit* of θ (under the action of G) is the equivalence class of θ under $\overset{G}{\sim}$, i.e. the set $\{g \cdot \theta : g \in G\}$. Denote by V/G the set of orbits of V , that is, the equivalence classes of V modulo the relation $\overset{G}{\sim}$.

We need the following definitions to capture the notion of *approximately* recovering the orbit of θ .

Definition 4.1.7. For $\theta_1, \theta_2 \in V$, let

$$d_G(\theta_1, \theta_2) = \min_{g \in G} \|\theta_1 - g \cdot \theta_2\|_2.$$

This pseudometric induces a metric on the quotient space V/G in the obvious way, so we can write $d_G(\mathfrak{o}_1, \mathfrak{o}_2)$ for $\mathfrak{o}_1, \mathfrak{o}_2 \in V/G$. By slight abuse of notation, we write $d_G(\theta_1, \mathfrak{o}_2)$ for $d_G(\mathfrak{o}_1, \mathfrak{o}_2)$, where \mathfrak{o}_1 is the orbit of θ_1 .

Theorem 4.1.5 already shows that if the orbit of θ is not determined by knowledge of the first $d - 1$ moment tensors, then at least $\Omega(\sigma^{2d})$ samples are required to recover the orbit of θ . We are now ready to (informally) state our main result on recovery (see the full paper [15] for the proof), which provides a matching upper bound.

Theorem 4.1.8 (recovery upper bound, informal). *If the moments $T_1(\theta), \dots, T_d(\theta)$ uniquely determine the orbit of θ , then using $O(\sigma^{2d})$ samples, we can produce an estimator $\hat{\theta}$ such that $d_G(\theta, \hat{\theta}) \leq \varepsilon$ with high probability.*

The recovery procedure is based on estimating the moments $T_1(\theta), \dots, T_d(\theta)$ and solving a system of polynomial equations to recover a θ that is (approximately) consistent with those moments. The analogous result holds for list recovery (see the full paper [15]): if the

moments determine a finite number s of possibilities for the orbit of θ then we can output a list of s estimators, one of which is close to the orbit of θ .

We note again that $O(\sigma^{2d})$ only captures the dependence on σ in the limit $\sigma \rightarrow \infty$ with other parameters (such as θ and ε) held fixed.

Thus, we have reduced to the algebraic question of determining how many moments are necessary to determine the orbit of θ (either uniquely or in the sense of list recovery). In Section 4.3 we will use tools from invariant theory and algebraic geometry in order to address these questions.

4.1.6 Extensions: projection and heterogeneity

We now consider some extensions to the basic orbit recovery problem (Problem 4.1.1), motivated by the application of cryo-EM:

1. **Projection:** In cryo-EM, we do not observe a noisy 3-dimensional model of the rotated molecule; we only observe a 2-dimensional projection of it. We will model this projection by a linear map $\Pi : \mathbb{R}^p \rightarrow \mathbb{R}^q$ that maps a 3-dimensional model to its 2-dimensional projection (from a fixed viewing direction). The samples are then given by $y_i = \Pi(g_i \cdot \theta) + \xi_i$ where $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$.
2. **Heterogeneity:** In cryo-EM we observe images of many different copies of the same molecule, each rotated differently. However, if our sample is not pure, we may have a mixture of different molecules and want to recover the structure of all of them. We will model this by taking K different unknown signals $\theta_1, \dots, \theta_K$ along with positive mixing weights w_1, \dots, w_K which sum to 1. Each sample takes the form $y_i = g_i \cdot \theta_{k_i} + \xi_i$ where k_i is chosen at random according to the mixing weights.

In Section 4.2 we will formally define a generalization of the orbit recovery problem that allows for either (or both) of the above extensions. All of our methods will apply to this general case.

4.1.7 Outline of remainder of chapter

In Section 4.2, we define a generalization of Problem 4.1.1 which encompasses projection and heterogeneity, and specify the basic algebraic objects which relate to our generalized problem. In Section 4.3, we establish our basic algebraic results and specify the algebraic criteria that correspond to the different recovery criteria defined in Section 4.1.4. We also give an efficient algorithm to decide the algebraic criterion corresponding to generic list recovery. Finally, in Section 4.4, we apply our work to several examples of the orbit recovery problem, including MRA and cryo-EM. We conclude in Section 4.5 with questions for future work.

Sections 4.6 and 4.7 contain proofs of results from preceding sections. Appendix C.1 contains an account of the invariant theory of $\text{SO}(3)$.

4.2 General problem statement

Our results will apply not only to the basic orbit recovery problem (Problem 4.1.1) but to a generalization (Problem 4.2.3 below) that captures the projection and heterogeneity extensions discussed in Section 4.1.6. We first define mixing weights for heterogeneous problems.

Definition 4.2.1 (mixing weights). Let $w = (w_1, \dots, w_K) \in \Delta_K \triangleq \{(z_1, \dots, z_K) : z_k \geq 0 \ \forall k, \sum_{k=1}^K z_k = 1\}$. Let $k \stackrel{w}{\sim} [K]$ indicate that k is sampled from $[K] = \{1, \dots, K\}$ such that $k = \ell$ with probability w_ℓ . We will sometimes instead parametrize the mixing weights by $\bar{w}_k = w_k - 1/K$ so that \bar{w} lies in the vector space $\bar{\Delta} \triangleq \{(z_1, \dots, z_K) : \sum_{k=1}^K z_k = 0\}$.

In a heterogeneous problem with K different signals, we can only hope to recover the signals up to permutation. To formalize this, our compound signal will lie in a larger vector space V and we will seek to recover its orbit under a larger group G .

Definition 4.2.2 (setup for heterogeneity). Let \tilde{G} be a compact group acting linearly, continuously, and orthogonally on $\tilde{V} = \mathbb{R}^p$. Let $V = \tilde{V}^{\oplus K} \oplus \bar{\Delta}_K$, so that $\theta \in V$ encodes K different signals along with mixing weights: $\theta = (\theta_1, \dots, \theta_K, \bar{w})$. We let an element

(g_1, \dots, g_K, π) of the Cartesian product set $\tilde{G}^K \times S_K$ act on V as follows: first, each g_k acts on the corresponding θ_k , and then π permutes the θ_k and the coordinates of \bar{w} . Note that this action is linear and orthogonal (where $\bar{\Delta}$ uses the usual inner product inherited from \mathbb{R}^K). There is a natural group structure G on the set $\tilde{G}^K \times S_K$ such that the action just described is actually a group action by G : the semidirect product $G = \tilde{G}^K \rtimes_{\varphi} S_K$, where φ denotes the action of S_K on \tilde{G}^K by permutations of the factors. This is also called the *wreath product* of \tilde{G} by S_K and written $\tilde{G} \wr S_K$. The product topology on $\tilde{G}^K \times S_K$ makes G a topological group; it is compact with respect to this topology since all the factors are compact, and the action described above is continuous.

Of course, by taking $K = 1$ we recover the basic setup (without heterogeneity) as a special case. We are now ready to give the general problem statement.

Problem 4.2.3 (generalized orbit recovery). Let $\tilde{V} = \mathbb{R}^p$ and $W = \mathbb{R}^q$. Let \tilde{G} be a compact group acting linearly, continuously, and orthogonally on \tilde{V} . Let $\Pi : \tilde{V} \rightarrow W$ be a linear map. Let $\theta = (\theta_1, \dots, \theta_K, \bar{w}) \in V \triangleq \tilde{V}^{\oplus K} \oplus \bar{\Delta}_K$ be an unknown collection of K signals with mixing weights $w \in \Delta_K$. For $i \in [n] = \{1, 2, \dots, n\}$ we observe

$$y_i = \Pi(g_i \cdot \theta_{k_i}) + \xi_i$$

where $g_i \sim \text{Haar}(\tilde{G})$, $k_i \stackrel{w}{\sim} [K]$, $\xi_i \sim \mathcal{N}(0, \sigma^2 I_{q \times q})$, all independently. The goal is to estimate the orbit of θ under $G \triangleq \tilde{G}^K \rtimes S_K$.

Note that this serves as a reduction from the heterogeneous setup to the basic setup in the sense that we are still only concerned with recovering the orbit of a vector θ under the action of some compact group.

As discussed previously, we apply the method of moments. The moments are now defined as follows.

Definition 4.2.4 (moment tensor). For the generalized orbit recovery problem (Problem 4.2.3), the *order- d moment tensor* is $T_d(\theta) \triangleq \mathbb{E}_{g,k}[(\Pi(g \cdot \theta_k))^{\otimes d}]$ where $g \sim \text{Haar}(\tilde{G})$ and $k \stackrel{w}{\sim} [K]$. Equivalently, $T_d(\theta) = \sum_{k=1}^K w_k \mathbb{E}_g[(\Pi(g \cdot \theta_k))^{\otimes d}]$.

The invariant ring is defined as in Definition 4.1.3 but now for the larger group G acting on the larger V :

Definition 4.2.5 (invariant ring). Note that $\dim(V) = Kp + K - 1$ and let $\mathbf{x} = (x_1, \dots, x_{\dim(V)})$ be a basis for V^* ; here the last $K - 1$ variables correspond to $\bar{\Delta}$, e.g. they can correspond to $\bar{w}_1, \dots, \bar{w}_{K-1}$. We then let $\mathbb{R}[\mathbf{x}]^G \subseteq \mathbb{R}[\mathbf{x}]$ be the polynomials in \mathbf{x} that are invariant under the action of G (as in Definition 4.1.3).

Recall that in the basic orbit recovery problem, $T_d(\theta)$ corresponds precisely to the homogeneous invariant polynomials of degree d ; now $T_d(\theta)$ corresponds to a subspace of the homogeneous invariant polynomials of degree d . Specifically, the method of moments gives us access to the following polynomials (evaluated at θ):

Definition 4.2.6. Let U_d^T be the subspace (over \mathbb{R}) of the invariant ring $\mathbb{R}[\mathbf{x}]^G$ consisting of all \mathbb{R} -linear combinations of entries of $T_d(\mathbf{x})$. Let $U_{\leq d}^T = U_1^T \oplus \dots \oplus U_d^T \subseteq \mathbb{R}[\mathbf{x}]^G$. Here we write $T_d(\mathbf{x})$ for the collection of polynomials (one for each entry of $T_d(\theta)$) that map θ to $T_d(\theta)$.

We will be interested in whether the subspace $U_{\leq d}^T$ contains enough information to uniquely determine the orbit of θ (or determine a finite list of possible orbits) in the following sense.

Definition 4.2.7. A subspace $U \subseteq \mathbb{R}[\mathbf{x}]^G$ *resolves* $\theta \in V$ if there exists a unique $\mathbf{o} \in V/G$ such that $f(\theta) = f(\mathbf{o})$ for all $f \in U$. Similarly, U *list-resolves* θ if there are only finitely many orbits $\mathbf{o}_1, \dots, \mathbf{o}_s$ such that $f(\theta) = f(\mathbf{o}_i)$ for all $f \in U$.

Here we have abused notation by writing $f(\mathbf{o})$ to denote the (constant) value that f takes on every $\theta \in \mathbf{o}$. The following question is of central importance.

Question 4.2.8. Fix $\theta \in V$. How large must d be in order for $U_{\leq d}^T$ to uniquely resolve θ ? How large must d be in order for $U_{\leq d}^T$ to list-resolve θ ?

The answer depends on G and V but also on whether θ is a generic or worst-case signal, and whether we ask for unique recovery or list recovery. As discussed previously (see Section 4.1.5), the sample complexity of the generalized orbit recovery problem is $\Theta(\sigma^{2d})$ where

d is the minimal d from Question 4.2.8. Our algebraic results in Section 4.3 will give general methods to answer Question 4.2.8 for any G and V .

4.3 Algebraic results

In this section, we will consider the four recovery criteria defined in Section 4.1.4, and give algebraic characterizations of each case. As discussed previously (Question 4.2.8) it suffices to focus our attention on deciding when a subspace U resolves (or list-resolves) a parameter θ . We show below how to answer this question by purely algebraic means. Moreover, for generic list recovery, we show how this question can be answered algorithmically in polynomial time. For generic and worst-case unique recovery, we also give algorithms to decide the corresponding algebraic condition; however, these algorithms are not efficient.

Throughout, we assume the setup defined in Section 4.2 for the generalized orbit recovery problem. In particular, G is a compact group acting linearly and continuously on a finite-dimensional real vector space V (although we do not require in this section that the action be orthogonal). We have the invariant ring $\mathbb{R}[\mathbf{x}]^G$ corresponding to the action of G on V , and a subspace $U \subseteq \mathbb{R}[\mathbf{x}]^G$ (e.g. $U_{\leq d}^T$) of invariants that we have access to. We are interested in whether the values $f(\theta)$ for $f \in U$ determine the orbit of $\theta \in V$ under G . The specific structure of G and $U_{\leq d}^T$ (as defined in Section 4.2) will be largely unimportant and can be abstracted away.

4.3.1 Invariant theory basics

We will often need the following basic operator that averages a polynomial over the group G .

Definition 4.3.1 (Reynolds operator). The *Reynolds operator* $\mathcal{R} : \mathbb{R}[\mathbf{x}] \rightarrow \mathbb{R}[\mathbf{x}]^G$ is defined by

$$\mathcal{R}(f) = \mathbb{E}_{g \sim \text{Haar}(G)} [g \cdot f].$$

Note that the Reynolds operator is a linear projection from $\mathbb{R}[\mathbf{x}]$ to $\mathbb{R}[\mathbf{x}]^G$ that preserves the degree of homogeneous polynomials (i.e. a homogeneous polynomial of degree d gets mapped either to a homogeneous polynomial of degree d or to zero).

Observation 4.3.2. *Let $\mathbb{R}[\mathbf{x}]_d^G$ denote the vector space consisting of homogeneous invariants of degree d . We can obtain a basis for $\mathbb{R}[\mathbf{x}]_d^G$ by applying \mathcal{R} to each monomial in $\mathbb{R}[\mathbf{x}]$ of degree d . (This yields a spanning set which can be pruned to a basis if desired.)*

In our setting, we have the following basic fact from invariant theory.

Theorem 4.3.3 (e.g. [84] Theorem 4.1-3). *The invariant ring $\mathbb{R}[\mathbf{x}]^G$ is finitely generated as an \mathbb{R} -algebra. In other words, there exist generators $f_1, \dots, f_m \in \mathbb{R}[\mathbf{x}]^G$ such that $\mathbb{R}[f_1, \dots, f_m] = \mathbb{R}[\mathbf{x}]^G$.*

Furthermore, there is an algorithm to find a generating set; see Section 4.6.1. Another basic fact from invariant theory implies that the entire invariant ring is sufficient to determine the orbit of θ . (This is not always true for non-compact groups; see Example 2.3.1 in [55].)

Theorem 4.3.4 ([84] Theorem 6-2.2). *The full invariant ring $\mathbb{R}[\mathbf{x}]^G$ resolves every $\theta \in V$.*

Proof. Let $\mathfrak{o}_1, \mathfrak{o}_2 \in V/G$ be distinct (and therefore disjoint) orbits. Since G is compact and acts continuously, \mathfrak{o}_1 and \mathfrak{o}_2 are compact subsets of V . Thus by Urysohn's lemma there exists a continuous function $\tilde{f} : V \rightarrow \mathbb{R}$ such that $\tilde{f}(\tau) = 0 \ \forall \tau \in \mathfrak{o}_1$ and $\tilde{f}(\tau) = 1 \ \forall \tau \in \mathfrak{o}_2$. The Stone–Weierstrass theorem states that a continuous function on a compact domain can be uniformly approximated to arbitrary accuracy by a polynomial. This means there is a polynomial $f \in \mathbb{R}[\mathbf{x}]$ with $f(\tau) \leq 1/3 \ \forall \tau \in \mathfrak{o}_1$ and $f(\tau) \geq 2/3 \ \forall \tau \in \mathfrak{o}_2$. It follows that $h = \mathcal{R}(f)$ is an invariant polynomial that separates the two orbits: $h(\mathfrak{o}_1) \leq 1/3$ and $h(\mathfrak{o}_2) \geq 2/3$. \square

Thus, in order to determine the orbit of θ it is sufficient to determine the values of all invariant polynomials. (This condition is clearly also necessary in the sense that if the orbit is uniquely determined then so are the values of all invariants.)

Remark 4.3.5. In what follows we will be discussing algorithms that take the problem setup as input (including \tilde{G} and its action on \tilde{V} , along with Π, K) and decide whether or not $U_{\leq d}^T$ (for some given d) is capable of a particular recovery task (e.g. list recovery of a generic $\theta \in V$). We will always assume that these algorithms have a procedure to compute a basis for U_d^T (for any d) in exact symbolic arithmetic. This is non-trivial in some cases because $T_d(\mathbf{x})$ (and thus U_d^T) involves integration over the group (and may involve irrational values), but we will not worry about these details here. For the important case of $\text{SO}(3)$, it is possible to write down a basis for the invariants in closed form (see Appendix C.1).

Remark 4.3.6. We will draw from various references for algorithmic aspects of invariant theory. The case of finite groups is treated by [143]. Although the invariant ring is sometimes taken to be $\mathbb{C}[\mathbf{x}]^G$ instead of $\mathbb{R}[\mathbf{x}]^G$, this is unimportant in our setting because the two are essentially the same: since our group action is real, a basis for $\mathbb{R}[\mathbf{x}]^G$ (over \mathbb{R}) is a basis for $\mathbb{C}[\mathbf{x}]^G$ (over \mathbb{C}). The case of infinite groups is covered by [55]. Here the group is assumed to be a *reductive* group over \mathbb{C} (or another algebraically-closed field). This means in particular that the group is a subset of complex-valued matrices that is defined by polynomial constraints. Although compact groups such as $\text{SO}(3)$ do not satisfy this, the key property of a reductive group is the existence of a Reynolds operator satisfying certain properties; since this exists for compact groups (Definition 4.3.1), some (but not all) results still hold in our setting.

4.3.2 Generic list recovery

We will see that the case of list recovery of a generic signal is governed by the notion of algebraic independence.

Definition 4.3.7. Polynomials $f_1, \dots, f_m \in \mathbb{R}[\mathbf{x}]$ are *algebraically dependent* if there exists a nonzero polynomial $P \in \mathbb{R}[y_1, \dots, y_m]$ such that $P(f_1, \dots, f_m) = 0$ (i.e. $P(f_1(\mathbf{x}), \dots, f_m(\mathbf{x}))$ is equal to the zero polynomial). Otherwise, they are *algebraically independent*.

Definition 4.3.8. The *transcendence degree* of a subspace $U \subseteq \mathbb{R}[\mathbf{x}]$, denoted $\text{trdeg}(U)$ is the maximum value of m for which there exist algebraically independent $f_1, \dots, f_m \in U$. A set of $\text{trdeg}(U)$ such polynomials is called a *transcendence basis* of U .

We now present our algebraic characterization of the generic list recovery problem.

Theorem 4.3.9 (generic list recovery). *Let $U \subseteq \mathbb{R}[\mathbf{x}]^G$ be a finite-dimensional subspace. If $\text{trdeg}(U) = \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ then there exists a set $S \subseteq V$ of full measure such that if $\theta \in S$ then U list-resolves θ . Conversely, if $\text{trdeg}(U) < \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ then there exists a set $S \subseteq V$ of full measure such that if $\theta \in S$ then U does not list-resolve θ .*

The proof is deferred to Sections 4.6.2 and 4.6.3. A set has *full measure* if its complement has measure zero. The intuition behind Theorem 4.3.9 is that $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ is the number of degrees of freedom that need to be pinned down in order to learn the orbit of θ , and so we need this many algebraically independent constraints (invariant polynomials). Note that we have not yet given any bound on how large the finite list might be; we will address this in Section 4.3.3.

In order for Theorem 4.3.9 to be useful, we need a way to compute the transcendence degree of both $\mathbb{R}[\mathbf{x}]^G$ and U . In what follows, we will discuss methods for both of these: in Section 4.3.2 we show how to compute $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ analytically, and in Section 4.3.2 we give an efficient algorithm to compute $\text{trdeg}(U)$ for a subspace U . By taking $U = U_{\leq d}^T$ this yields an efficient algorithm to determine the smallest degree d at which $U_{\leq d}^T$ list-resolves a generic θ (thereby answering Question 4.2.8 for the case of generic list recovery).

Computing the transcendence degree of $\mathbb{R}[\mathbf{x}]^G$.

Intuitively, the transcendence degree of $\mathbb{R}[\mathbf{x}]^G$ is the number of parameters required to describe an orbit of G . For finite groups, this is simply the dimension of V :

Proposition 4.3.10 ([143] Proposition 2.1.1). *If G is a finite group, $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = \dim(V)$.*

For infinite groups, the situation may be slightly different. For instance, if $\text{SO}(3)$ acts on $V = \mathbb{R}^3$ in the standard way (rotations in 3 dimensions), then a generic orbit is a sphere, with dimension two. This means there is only one parameter to learn, namely the 2-norm, and we expect $\mathbb{R}[\mathbf{x}]^G$ to have transcendence degree 1 accordingly. On the other hand, if $\text{SO}(3)$ acts on a rich class of functions $S^2 \rightarrow \mathbb{R}$ (as in the S^2 registration problem; see Section 4.4.4)

then each orbit resembles a copy of $\text{SO}(3)$ which has dimension 3. This is formalized in the following.

Proposition 4.3.11 ([61] Corollary 6.2). *If G is an algebraic group, then*

$$\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = \dim(V) - \dim(G) + \min_{v \in V} \dim(G_v),$$

where G_v is the stabilizer at v of the action of G (that is, the subgroup of all $g \in G$ fixing v).

An alternate approach to the transcendence degree of $\mathbb{R}[\mathbf{x}]^G$ uses a central object in invariant theory: the *Hilbert series* (see e.g. [55]).

Definition 4.3.12. Let $\mathbb{R}[\mathbf{x}]_d^G$ be the subspace (over \mathbb{R}) of $\mathbb{R}[\mathbf{x}]^G$ consisting of homogeneous invariants of degree d . The *Hilbert series* of $\mathbb{R}[\mathbf{x}]^G$ is the formal power series

$$H(t) \triangleq \sum_{d=0}^{\infty} \dim(\mathbb{R}[\mathbf{x}]_d^G) t^d.$$

For a given G acting on V , there is an explicit formula (*Molien's formula*) for the Hilbert series:

Proposition 4.3.13 ([84] Remark 3-1.8). *Let $\rho : G \rightarrow \text{GL}(V)$ be the representation by which G acts on V . Then for $|t| < 1$, $H(t)$ converges and we have*

$$H(t) = \mathbb{E}_{g \sim \text{Haar}(G)} \det(I - t \rho(g))^{-1}.$$

This formula is tractable to compute, even for complicated groups; see Section 4.4.4 for details in the case of $\text{SO}(3)$. Once we have the Hilbert series, it is easy to extract $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ as follows.

Proposition 4.3.14. *The order of the pole at $t = 1$ of $H(t)$ is equal to $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$.*

The proof comes from [55]; see Section 4.6.4 for more details.

For heterogeneous problems ($K > 1$), the transcendence degree can be computed easily from the transcendence degree of the corresponding homogeneous ($K = 1$) problem.

Proposition 4.3.15. *Let \tilde{G} be a compact group acting linearly and continuously on \tilde{V} , and let $G = \tilde{G}^K \rtimes S_K$ act on $V = \tilde{V}^{\oplus K} \oplus \bar{\Delta}_K$ as in Definition 4.2.2. Let $\mathbb{R}[\mathbf{x}]^G$ be the invariant ring corresponding to the action of G on V , and let $\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}}$ be the invariant ring corresponding to the action of \tilde{G} on \tilde{V} (i.e. the $K = 1$ problem). Then $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = K \cdot \text{trdeg}(\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}}) + K - 1$.*

The proof can be found in Section 4.6.5. Note, however, that the result is intuitively reasonable by counting parameters. We know $\text{trdeg}(\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}})$ is the number of parameters required to describe an orbit of \tilde{G} acting on \tilde{V} . Thus, in the heterogeneity problem we have $\text{trdeg}(\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}})$ parameters for each of the K signals, plus an additional $K - 1$ parameters for the K mixing weights (since they sum to 1).

Algorithm for transcendence basis of U .

In this section we prove the following.

Theorem 4.3.16. *There is an efficient algorithm to perform the following task. Given a basis $\{u_1, \dots, u_s\}$ for a finite-dimensional subspace $U \subseteq \mathbb{R}[\mathbf{x}]$, output a transcendence basis for U .*

Our first ingredient is the following simple classical test for algebraic independence (see, e.g., [66, 24] for a proof).

Definition 4.3.17 (Jacobian). Given polynomials $f_1, \dots, f_m \in \mathbb{R}[\mathbf{x}] = \mathbb{R}[x_1, \dots, x_p]$, we define the *Jacobian matrix* $J_{\mathbf{x}}(f_1, \dots, f_m) \in (\mathbb{R}[\mathbf{x}])^{m \times p}$ by $(J_{\mathbf{x}}(f_1, \dots, f_m))_{ij} = \partial_{x_j} f_i$ where ∂_{x_j} denotes formal partial derivative with respect to x_j .

Proposition 4.3.18 (Jacobian criterion for algebraic independence). *Polynomials $\mathbf{f} = (f_1, \dots, f_m)$ are algebraically independent if and only if the Jacobian matrix $J_{\mathbf{x}}(\mathbf{f})$ has full row rank (over the field $\mathbb{R}(\mathbf{x})$).*

It suffices to test the rank of the Jacobian at a generic point \mathbf{x} .

Corollary 4.3.19. *Fix $\mathbf{f} = (f_1, \dots, f_m)$. Let $z \sim \mathcal{N}(0, I_{p \times p})$. If \mathbf{f} is algebraically dependent then $J_{\mathbf{x}}(\mathbf{f})|_{\mathbf{x}=z}$ does not have full row rank. If \mathbf{f} is algebraically independent then $J_{\mathbf{x}}(\mathbf{f})|_{\mathbf{x}=z}$ has full row rank with probability 1.*

Proof. An $m \times p$ matrix has deficient row rank if and only if either $m > p$ or every maximal square submatrix has determinant zero. Every such determinant of $J_{\mathbf{x}}(\mathbf{f})$ is a polynomial in \mathbf{x} ; if this polynomial is not identically zero then plugging in generic values for \mathbf{x} will not cause it to vanish. \square

Remark 4.3.20. In practice we may choose to plug in random *rational* values for \mathbf{x} so that the rank computation can be done in exact symbolic arithmetic. The Jacobian test will still succeed with overwhelming probability (provided we use a fine enough mesh of rational numbers). Also note that if we find *any* value of \mathbf{x} for which the Jacobian has full row rank, this constitutes a proof of algebraic independence.

Remark 4.3.21. In some cases (e.g. if the polynomials involve irrational values) it may be slow to compute the Jacobian rank in exact symbolic arithmetic. We can alternatively compute the singular values numerically and count how many are reasonably far from zero. This method works reliably in practice (i.e., it is extremely clear how to separate the zero and nonzero singular values) but does not constitute a rigorous proof of algebraic independence.

Curiously, although the Jacobian criterion gives an efficient test for algebraic dependence, it is much harder ($\#P$ -hard) to actually find the algebraic dependence (i.e., the polynomial relation) when one exists [87].

The Jacobian criterion implies the well-known fact that the collection of algebraically independent subsets of $\mathbb{R}[\mathbf{x}]$ form a *matroid*; this is called an *algebraic matroid* (see e.g. [137]). In particular, we have the following exchange property:

Proposition 4.3.22. *Let I, J be finite subsets of $\mathbb{R}[\mathbf{x}]$, each algebraically independent. If $|I| < |J|$ then there exists $f \in J \setminus I$ such that $I \cup \{f\}$ is algebraically independent.*

We next note that in the task from Theorem 4.3.16, a transcendence basis can always be taken from the basis $\{u_1, \dots, u_s\}$ itself.

Lemma 4.3.23. *Let U be a finite-dimensional subspace of $\mathbb{R}[\mathbf{x}]$ with basis $B = \{u_1, \dots, u_s\}$. If U contains r algebraically independent elements, then so does B .*

Proof. Let $B' \subseteq B$ be a maximal set of algebraically independent elements of B . If $|B| < r$ then by the exchange property (Proposition 4.3.22) there exists $v \in U \setminus B'$ such that $B' \cup \{v\}$ is algebraically independent. Write $v = \sum_{i=1}^s \alpha_i u_i$. Since B' is maximal, we have from the Jacobian criterion (Proposition 4.3.18) that for all $1 \leq i \leq s$, the row vector $J_{\mathbf{x}}(u_i)$ lies in the $\mathbb{R}(\mathbf{x})$ -span of $\mathcal{B} \triangleq \{J_{\mathbf{x}}(b)\}_{b \in B'}$. But this means that $J_{\mathbf{x}}(v) = \sum_{i=1}^s \alpha_i J_{\mathbf{x}}(u_i)$ lies in the $\mathbb{R}(\mathbf{x})$ -span of \mathcal{B} . By the Jacobian criterion this contradicts the fact that $B' \cup \{v\}$ is algebraically independent. \square

Proof of Theorem 4.3.16.

Let $\{u_1, \dots, u_s\}$ be a basis (or spanning set) for U . From above we have that the transcendence degree of U is the row rank of the Jacobian $J_{\mathbf{x}}(u_1, \dots, u_s)$ evaluated at a generic point \mathbf{x} . A transcendence basis for U is the set of u_i corresponding to a maximal linearly independent set of rows

We can use the following simple greedy algorithm to construct a transcendence basis. As input, receive a list of polynomials $\{u_1, \dots, u_s\}$. Initialize $I = \emptyset$. For $i = 1, \dots, s$, add $\{u_i\}$ to I if $I \cup \{u_i\}$ is algebraically independent, and do nothing otherwise. (Note that this condition can be efficiently tested by Corollary 4.3.19.) Output the resulting set I .

We now show correctness. Let I_i be the set after item u_i has been considered (and possibly added), and set $I_0 = \emptyset$. It suffices to show that for each $i \in \{0, \dots, s\}$, I_i is a maximal independent subset of $\{u_1, \dots, u_i\}$. We proceed by induction. The claim is vacuously true when $i = 0$. Assume it holds for $i - 1$. If I_i is not a maximal independent subset of $\{u_1, \dots, u_i\}$, then there exists an independent set $J \subseteq \{u_1, \dots, u_i\}$ with $|J| > |I|$, so by the exchange property (Proposition 4.3.22) there exists a u_j with $j \leq i$ such that $u_j \notin I_i$ and $I_i \cup \{u_j\}$ is independent. In particular, the subset $I_{j-1} \cup \{u_j\}$ of $I_i \cup \{u_j\}$ is independent. But the fact that u_j was not added at the $(j-1)$ th step implies that $I_{j-1} \cup \{u_j\}$ is not independent, a contradiction. So I_i is indeed maximal.

We obtain that $I = I_s$ is a maximal independent subset of $\{u_1, \dots, u_s\}$, and hence by Lemma 4.3.23 a transcendence basis of U . \square

4.3.3 Generic unique recovery

For list recovery problems, the following gives an explicit upper bound on the size of the list.

Theorem 4.3.24. *Let U be a subspace of the invariant ring $\mathbb{R}[\mathbf{x}]^G$. Let F_G be the field of fractions of $\mathbb{R}[\mathbf{x}]^G$. If $[F_G : \mathbb{R}(U)] = D < \infty$ then there exists a set $S \subseteq V$ of full measure such that for any $\theta \in S$, U list-resolves θ with a list of size $\leq D$.*

The proof is deferred to Section 4.6.2. Here $\mathbb{R}(U)$ is the smallest subfield of F_G containing both \mathbb{R} and U , and $[F_G : \mathbb{R}(U)]$ denotes the degree of a field extension; see Section 4.6.2 for more details. Since $[F_G : \mathbb{R}(U)] = 1$ is equivalent to $\mathbb{R}(U) = F_G$, we have the following criterion for unique recovery.

Corollary 4.3.25 (generic unique recovery). *If $\mathbb{R}(U) = F_G$ then there exists a set $S \subseteq V$ of full measure such that if $\theta \in S$ then U resolves θ .*

The intuition here is that we want to be able to learn every invariant polynomial by adding, multiplying, and dividing polynomials from U (and scalars from \mathbb{R}). We need θ to be generic so that we never divide by zero in the process.

Theorem 4.3.26. *For a finite-dimensional subspace $U \subseteq \mathbb{R}[\mathbf{x}]^G$, there is an algorithm to compute the degree of the field extension from Theorem 4.3.24. As input, the algorithm requires a basis for U and the ability to compute the Reynolds operator (Definition 4.3.1).*

We give the algorithm and the proof in Section 4.6.6. The algorithm uses Gröbner bases and is unfortunately inefficient to run in practice.

4.3.4 Worst-case unique recovery

We give a sufficient algebraic condition for worst-case unique recovery:

Theorem 4.3.27 (worst-case unique recovery). *Let $U \subseteq \mathbb{R}[\mathbf{x}]^G$ be a finite-dimensional subspace with basis $\{f_1, \dots, f_m\}$. If U generates $\mathbb{R}[\mathbf{x}]^G$ as an \mathbb{R} -algebra (i.e. $\mathbb{R}[f_1, \dots, f_m] = \mathbb{R}[\mathbf{x}]^G$) then U resolves every $\theta \in V$.*

Proof. Every element of $\mathbb{R}[\mathbf{x}]^G$ can be written as a polynomial in the f_i (with coefficients in \mathbb{R}). This means the values $f_1(\theta), \dots, f_m(\theta)$ uniquely determine all the values $f(\theta)$ for $f \in \mathbb{R}[\mathbf{x}]^G$ and so the result follows because $\mathbb{R}[\mathbf{x}]^G$ resolves every $\theta \in V$ (Theorem 4.3.4). \square

Theorem 4.3.28. *There is an algorithm to test whether or not U generates $\mathbb{R}[\mathbf{x}]^G$ as an \mathbb{R} -algebra. As input, the algorithm requires a basis for U and the ability to compute the Reynolds operator (Definition 4.3.1).*

We give the algorithm and the proof in Section 4.6.6. The algorithm uses Gröbner bases and is unfortunately inefficient to run in practice.

If G is a finite group, it is known that $\mathbb{R}[\mathbf{x}]^G$ has a generating set for which all elements have degree at most $|G|$ (this is *Noether's degree bound*; see Theorem 2.1.4 in [143]). It follows that $\mathbb{R}[\mathbf{x}]_{\leq |G|}^G$ resolves every $\theta \in V$. Recall (from Section 4.1.4) that this is tight for MRA: degree $|G|$ is necessary for worst-case signals.

A precise characterization of when U resolves every $\theta \in V$ is (by definition) that U should be a *separating set* or (equivalently) should generate a *separating algebra* (see [55] Section 2.4). The notions of generating and separating sets do not always coincide, as illustrated by Example 2.4.2 in [55]. Furthermore, generating sets may require strictly higher maximum degree [62].

4.3.5 Worst-case list recovery

We give a sufficient algebraic condition for worst-case list recovery:

Theorem 4.3.29 (worst-case list recovery). *Let $U \subseteq \mathbb{R}[\mathbf{x}]^G$ be a subspace with finite basis $\{f_1, \dots, f_m\}$. If $\mathbb{R}[\mathbf{x}]^G$ is finitely generated as a $\mathbb{R}[f_1, \dots, f_m]$ -module, then U list-resolves every $\theta \in V$.*

In other words, this condition says that there exists a basis $g_1, \dots, g_s \in \mathbb{R}[\mathbf{x}]^G$ such that every element of $\mathbb{R}[\mathbf{x}]^G$ can be written as a linear combination of g_1, \dots, g_s with coefficients from $\mathbb{R}[f_1, \dots, f_m]$. It is sufficient to take U to be a set of *primary invariants* from a *Hironaka decomposition* (see Section 4.6.4).

Proof. Since $\mathbb{R}[\mathbf{x}]^G$ finitely generated as an \mathbb{R} -algebra (Theorem 4.3.3), if $\mathbb{R}[\mathbf{x}]^G$ is finitely generated as a $\mathbb{R}[f_1, \dots, f_m]$ -module then it follows that (see [139] Section 5.3) every $h \in \mathbb{R}[\mathbf{x}]^G$ satisfies a monic polynomial

$$h^k + c_{k-1}h^{k-1} + \dots + c_1h + c_0 = 0$$

with $c_i \in \mathbb{R}[f_1, \dots, f_m]$. Letting h_1, \dots, h_s be generators for $\mathbb{R}[\mathbf{x}]^G$ (as an \mathbb{R} -algebra), we have that the values $f_1(\theta), \dots, f_m(\theta)$ determine a finite set of possible values for $h_1(\theta), \dots, h_s(\theta)$, each of which determines (at most) one orbit for θ . \square

4.4 Examples

In this section we work out some specific examples, determining the degree at which generic list recovery is possible using the methods of Section 4.3.2. (We focus on generic list recovery because our algorithms for the other recovery criteria are unfortunately too slow even for quite small examples.) We obtain several recovery theorems for problems such as MRA and cryo-EM within finite ranges of parameters where we have verified the Jacobian criterion using a computer, and beyond these parameter ranges, we state conjectural patterns.

The following themes emerge in the examples studied in this section. First, we see that many problems are possible at degree 3, which is promising from a practical standpoint. Second, we do not encounter any unexpected algebraic dependencies, and so we are able to show that heuristic parameter-counting arguments are correct. In particular, we see that if there are enough linearly independent invariants, there are also enough algebraically independent invariants.

4.4.1 Learning a bag of numbers

Let G be the symmetric group S_p acting on $V = \mathbb{R}^p$ by permutation matrices. The invariant ring consists of the symmetric polynomials, which are generated by the elementary symmetric polynomials e_1, \dots, e_p where e_i has degree i . Worst-case unique recovery is possible at degree

p since $\mathbb{R}[\mathbf{x}]_{\leq p}^G$ generates the full invariant ring. Furthermore, degree p is actually required, even for generic list recovery. This is because any invariant of degree $\leq p-1$ can be expressed as a polynomial in e_1, \dots, e_{p-1} and thus $\text{trdeg}(\mathbb{R}[\mathbf{x}]_{\leq p-1}^G) = p-1$. So this problem has a steep sample complexity of order σ^{2p} .

4.4.2 Learning a rigid body

Let G be the rotation group $\text{SO}(p)$ acting on the matrix space $\mathbb{R}^{p \times m}$ by left multiplication. We imagine the columns of our matrix as vertices defining a rigid body; thus we observe random rotations of this rigid body (with vertices labeled) plus noise. Let $U \in \mathbb{R}^{p \times m}$ be such a matrix signal. With $O(\sigma^4)$ samples, we can estimate the degree-2 Gram matrix $U^\top U$; taking a Cholesky factorization, we recover U up to left action by an element of the larger group $\text{O}(p)$. Thus we recover the rigid body up to a reflection ambiguity, demonstrating list recovery (with a list of size 2). Surprisingly, assuming $m \geq p$, we do not uniquely resolve a generic signal until degree p , where with $O(\sigma^{2p})$ samples we can estimate a $p \times p$ minor of U , which is a degree- p invariant that changes sign under reflection.

The impossibility of unique recovery until degree p is a consequence of the “first fundamental theorem” for the special orthogonal group $\text{SO}(p)$, which asserts that the invariant ring is generated by the entries of the Gram matrix $U^\top U$ together with the $p \times p$ minors of U (see for instance [84]); thus the invariants of degree $3, \dots, p-1$ carry no information in addition to the degree-2 invariants.

4.4.3 Multi-reference alignment (MRA)

Recall that this is the case of $G = \mathbb{Z}/p$ acting on $V = \mathbb{R}^p$ by cyclic shifts. It is already known that for the basic MRA problem (without projection or heterogeneity), generic unique recovery is possible at degree 3 for any p [13]. The methods of Section 4.3.2 confirm the weaker result that generic *list* recovery is possible at degree 3 (at least for the values of p that we tested). Note the stark contrast in difficulty from the case of the full symmetric group $G = S_p$ above.

Remark 4.4.1. This result for MRA is actually a special case of a more general phenomenon. Let G be any finite group and let V be the *regular representation* i.e. the space of functions $f : G \rightarrow \mathbb{R}$ with the action $(g \cdot f)(h) = f(g^{-1}h)$. (Note that for $G = \mathbb{Z}/p$ this is precisely the MRA problem.) It is known [85] that for this setup, the degree-3 invariants (the *triple correlation*) are sufficient to resolve a generic signal, and thus generic unique recovery is possible at degree 3.

We can also verify that for MRA with $p \geq 3$, generic list recovery is impossible at degree 2. This follows from Theorem 4.3.9 because $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = p$ (since G is finite) but the number of algebraically independent invariants of degree ≤ 2 is at most $\lfloor p/2 \rfloor + 1$. We can see this as follows. A basis for the invariants of degree ≤ 2 is $\{\mathcal{R}(x_1), \mathcal{R}(x_1^2), \mathcal{R}(x_1x_2), \mathcal{R}(x_1x_3), \dots, \mathcal{R}(x_1x_s)\}$ with $s = \lfloor p/2 \rfloor + 1$. Here \mathcal{R} denotes the Reynolds operator, which averages over cyclic shifts of the variables. For instance, $\mathcal{R}(x_1x_2) = \frac{1}{p}(x_1x_2 + x_2x_3 + x_3x_4 + \dots + x_px_1)$. Note that the basis above has size $\lfloor p/2 \rfloor + 2$ but there is an algebraic dependence within it because $\mathcal{R}(x_1)^2$ can be written in terms of the other basis elements. The claim now follows.

Generic list recovery is possible at degree 1 for $p = 1$ and at degree 2 for $p = 2$. (This is true even for worst-case unique recovery; recall from Section 4.3.4 that degree $|G|$ is always sufficient for this.)

We now move on to variants of the MRA problem.

MRA with projection

We now consider MRA with a projection step. We imagine that the coordinates of the signal are arranged in a circle so that G acts by rotating the signal around the circle. We then observe a projection of the circle onto a line so that each observation is the sum of the two entries lying “above” it on the circle. This is intended to resemble the tomographic projection in cryo-EM. We formally define the setup as follows.

Problem 4.4.2 (MRA with projection). Let $p \geq 3$ be odd. Let $V = \mathbb{R}^p$ and $G = \mathbb{Z}/p$ acting

on V by cyclic shifts. Let $q = (p - 1)/2$ and $W = \mathbb{R}^q$. Let $\Pi : V \rightarrow W$ be defined by

$$\Pi(v_1, \dots, v_p) = (v_1 + v_p, v_2 + v_{p-1}, \dots, v_{(p-1)/2} + v_{(p+3)/2}).$$

We call the associated generalized orbit recovery problem (Problem 4.2.3) *MRA with projection*. (We consider the homogeneous case $K = 1$.)

Note that since p is odd, there is one entry $v_{(p+1)/2}$ which is discarded by Π . The reason we consider the odd- p case rather than the seemingly more elegant even- p case is because generic list recovery is actually impossible in the even- p case. This is because the signals θ and $\theta + (c, -c, c, -c, \dots)$ cannot be distinguished from the samples, even if there is no noise.

Restricting now to odd p , note that we cannot hope for generic unique recovery because it is impossible to tell whether the signal is wrapped clockwise or counterclockwise around the circle. In other words, reversing the signal via $(\theta_1, \dots, \theta_p) \mapsto (\theta_p, \dots, \theta_1)$ does not change the distribution of samples. We can still hope for generic list recovery, hopefully with a list of size exactly 2. This degeneracy is analogous to the chirality issue in cryo-EM: it is impossible to determine the chirality of the molecule (i.e. if the molecule is reflected about some 2-dimensional plane through the origin, this does not change the distribution of samples).

It appears that, as in the basic MRA problem, generic list recovery is possible at degree 3. We proved this for p up to 21 by checking the Jacobian criterion (see Section 4.3.2) on a computer, and we conjecture that this trend continues.

Conjecture 4.4.3. *For MRA with projection, for any odd $p \geq 3$, generic list recovery is possible at degree 3.*

Note that generic list recovery is impossible at degree 2 because the addition of the projection step to basic MRA can only make it harder for $U_{\leq d}^T$ to list-resolve θ .

Heterogeneous MRA

We now consider heterogeneous MRA, i.e. the generalized orbit recovery problem (Problem 4.2.3) with $\tilde{G} = \mathbb{Z}/p$ acting on $\tilde{V} = \mathbb{R}^p$ by cyclic shifts, $K \geq 2$ heterogeneous components, and no projection (i.e., Π is the identity).

We will see that generic list recovery is possible at degree 3 provided that p is large enough compared to K . First note that the number of degrees of freedom to be recovered is $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = Kp + K - 1$ (see Propositions 4.3.10 and 4.3.15). Let us now count the number of distinct entries of $T_d(\mathbf{x})$ for $d \leq 3$. Note that $T_d(\mathbf{x})$ is symmetric (under permutations of indices) but we also have additional symmetries given by cyclic shifts, e.g. $(T_3(\mathbf{x}))_{i,j,k} = (T_3(\mathbf{x}))_{i+c,j+c,k+c}$ where c is an integer and the sums $i+c, j+c, k+c$ are computed modulo p . One can compute that $T_1(\mathbf{x})$ has 1 distinct entry, $T_2(\mathbf{x})$ has $\lfloor p/2 \rfloor + 1$ distinct entries, and $T_3(\mathbf{x})$ has $p + \lceil (p-1)(p-2)/6 \rceil$ distinct entries. The total number of distinct entries is

$$\mathcal{U} \triangleq p + 2 + \lfloor p/2 \rfloor + \lceil (p-1)(p-2)/6 \rceil.$$

By Theorem 4.3.9, list recovery is impossible when $\mathcal{U} < Kp + K - 1$. By testing the Jacobian condition, we observe that the converse also appears to hold. We tested this up to $K = 15$ and up to the corresponding critical p value.

Conjecture 4.4.4. *For heterogeneous ($K \geq 2$) MRA, generic list recovery is possible at degree 3 precisely if $\mathcal{U} \geq Kp + K - 1$. This condition on \mathcal{U} can be stated more explicitly as follows:*

- $K = 2$ requires $p \geq 1$.
- $K = 3$ requires $p \geq 12$.
- $K = 4$ requires $p \geq 18$.
- Each $K \geq 5$ requires $p \geq 6K - 5$.

Recent work [35] also studies the heterogeneous MRA problem. Similarly to the present work, they apply the method of moments and solve a polynomial system of equations in

order to recover the signals. To solve the system they use an efficient heuristic method that has no provable guarantees but appears to work well in practice. Their experiments suggest that if the signals have i.i.d. Gaussian entries, this method succeeds only when (roughly) $K \leq \sqrt{p}$ instead of the condition (roughly) $K \leq p/6$ that we see above (and that [35] also identified based on parameter-counting). In Chapter 5, we prove that indeed polynomial-time recovery is possible when $K \leq \tilde{\Omega}(\sqrt{p})$. We expect that this is a statistical-to-computational gap whereby it becomes computationally hard to efficiently solve the polynomial system once K exceeds \sqrt{p} .

4.4.4 S^2 registration

Recall that this is the case where the signal θ is a real-valued function defined on the unit sphere S^2 in \mathbb{R}^3 . The formal setup is as follows.

Let $G = \text{SO}(3)$. For each $\ell = 0, 1, 2, \dots$ there is an irreducible representation V_ℓ of $\text{SO}(3)$ of dimension $2\ell + 1$. These representations are of real type, i.e. they can be defined over the real numbers so that $V_\ell = \mathbb{R}^{2\ell+1}$. Let \mathcal{F} be a finite subset of $\{0, 1, 2, \dots\}$ and consider the orbit recovery problem in which G acts on $V = \oplus_{\ell \in \mathcal{F}} V_\ell$.

As intuition for the above setup, V_ℓ is a basis for the degree- ℓ *spherical harmonic* functions $S^2 \rightarrow \mathbb{R}$ defined on the surface of the unit sphere $S^2 \subseteq \mathbb{R}^3$. The spherical harmonics are a complete set of orthogonal functions on the sphere and can be used (like a “Fourier series”) to represent a function $S^2 \rightarrow \mathbb{R}$. Thus the signal $\theta \in V$ can be thought of as a function on the sphere, with $\text{SO}(3)$ acting on it by rotating the sphere. See Appendix C.1 for details on spherical harmonics.

The primary case of interest is $\mathcal{F} = \{1, \dots, F\}$ for some F (the number of “frequencies”). We will see that generic list recovery is possible at degree 3 so long as $F \geq 10$. We will see that it is convenient to not include $0 \in \mathcal{F}$, but we now justify why this is without loss of generality. V_0 is the trivial representation, i.e. the 1-dimensional representation on which every group element acts as the identity. In the interpretation of spherical harmonics, the V_0 -component is the mean value of the function over the sphere. We claim that the S^2 registration problem

with $0 \in \mathcal{F}$ can be easily reduced to the problem with $\mathcal{F}' = \mathcal{F} \setminus \{0\}$. This is because the V_0 -component is itself a degree-1 invariant; given the value of this invariant, one can subtract it off and reduce to the case without a V_0 -component (i.e. the case where the function on the sphere is zero-mean). Thus we have that e.g. generic list recovery is possible (at a given degree) for \mathcal{F} if and only if it is possible for \mathcal{F}' .

Using Proposition 4.3.11 we compute that $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = p - p'$, where

$$p = \dim(V) = \sum_{\ell \in \mathcal{F}} (2\ell + 1)$$

and

$$p' = \begin{cases} 0 & \ell_{\max} = 0 \\ 2 & \ell_{\max} = 1 \\ 3 & \ell_{\max} \geq 2 \end{cases} \quad \text{where } \ell_{\max} = \max_{\ell \in \mathcal{F}} \ell.$$

After all, V_0 is the trivial representation on the 1-dimensional vector space, with 3-dimensional stabilizer $\text{SO}(3)$, and V_1 is the standard 3-dimension representation of $\text{SO}(3)$ on \mathbb{R}^3 by rotations, which yields a one-dimensional $\text{SO}(2)$ stabilizer at each nonzero point. When $\ell_{\max} \geq 2$, the representation V is known to have zero-dimensional stabilizer at some points (see e.g. [68]).

In the following we restrict to the case $0 \notin \mathcal{F}$ for simplicity (but recall that this is without loss of generality). There are therefore no degree-1 invariants, i.e. $\mathbb{R}[\mathbf{x}]_1^G$ is empty. By Theorem 4.3.9, if $\dim(\mathbb{R}[\mathbf{x}]_2^G) + \dim(\mathbb{R}[\mathbf{x}]_3^G) < \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ then generic list recovery is impossible at degree 3; this rules out generic list recovery for $\mathcal{F} = \{1, 2, \dots, F\}$ when $F \leq 9$. (We will see below how to compute $\dim(\mathbb{R}[\mathbf{x}]_d^G)$.) Beyond this threshold, the situation is more hopeful:

Theorem 4.4.5. *If $\mathcal{F} = \{1, 2, \dots, F\}$ and $10 \leq F \leq 16$ then the degree-3 method of moments achieves generic list recovery.*

This theorem is based on computer verification of the Jacobian criterion for $10 \leq F \leq 16$ using exact arithmetic in a finite extension of \mathbb{Q} . This result lends credence to the following

conjecture.

Conjecture 4.4.6. *Consider the S^2 registration problem with $0 \notin \mathcal{F}$. We conjecture the following.*

- *Generic list recovery is possible at degree 3 if and only if $\dim(\mathbb{R}[\mathbf{x}]_2^G) + \dim(\mathbb{R}[\mathbf{x}]_3^G) \geq \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ (where $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ is computed above and $\dim(\mathbb{R}[\mathbf{x}]_d^G)$ can be computed from Proposition 4.4.7 below).*
- *In particular, if $\mathcal{F} = \{1, 2, \dots, F\}$ then generic list recovery is possible at degree 3 if and only if $F \geq 10$.*

The reason it is convenient to exclude the trivial representation is because it simplifies the parameter-counting: if we use the trivial representation then we have a degree-1 invariant f and so there is an algebraic relation between the degree-2 invariant f^2 and the degree-3 invariant f^3 .

We now discuss how to compute $\dim(\mathbb{R}[\mathbf{x}]_d^G)$. Using the methods in Section 4.6 of [55], we can give a formula for the Hilbert series of $\mathbb{R}[\mathbf{x}]^G$; see Section 4.7.1. However, if one wants to extract a specific coefficient $\dim(\mathbb{R}[\mathbf{x}]_d^G)$ of the Hilbert series, we give an alternative (and somewhat simpler) formula:

Proposition 4.4.7. *Consider S^2 registration with frequencies \mathcal{F} . Let $\chi_d(\phi) : \mathbb{R} \rightarrow \mathbb{R}$ be defined recursively by*

$$\begin{aligned} \chi_0(\phi) &= 1, \\ \chi_1(\phi) &= \sum_{\ell \in \mathcal{F}} \left[1 + 2 \sum_{m=1}^{\ell} \cos(m\phi) \right], \text{ and} \\ \chi_d(\phi) &= \frac{1}{d} \sum_{i=1}^d \chi_1(i\phi) \chi_{d-i}(\phi). \end{aligned}$$

Then we have

$$\dim(\mathbb{R}[\mathbf{x}]_d^G) = \frac{1}{\pi} \int_0^\pi (1 - \cos \phi) \chi_d(\phi) \, d\phi.$$

We give the proof in Section 4.7.2. Additionally, in Appendix C.1.6 we give explicit formulas for the invariants (up to degree 3), which yields a combinatorial analogue of Proposition 4.4.7 (up to degree 3).

4.4.5 Cryo-EM

We adapt the following simple model for the cryo-EM reconstruction problem. We will use properties of the 3-dimensional Fourier transform, including the projection-slice theorem; see e.g. [116] for a reference.

The signal is a 3-dimensional molecule, which we can think of as encoded by a density function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$. The 3-dimensional Fourier transform of f is $\hat{f} : \mathbb{R}^3 \rightarrow \mathbb{C}$ given by

$$\hat{f}(k_x, k_y, k_z) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-2\pi i(xk_x + yk_y + zk_z)} f(x, y, z) \, dx \, dy \, dz. \quad (4.1)$$

It is sufficient to learn \hat{f} because we can then recover f using the inverse Fourier transform. $\text{SO}(3)$ acts on the molecule by rotating it in 3-dimensional space (keeping the origin fixed). When f is rotated in (x, y, z) coordinates, \hat{f} is also rotated in (k_x, k_y, k_z) -coordinates by the same rotation. Each observation is a 2-dimensional image obtained by first rotating f by a random element of $\text{SO}(3)$ and then projecting f parallel to the z axis. Specifically, the projection of f is $f_{\text{proj}} : \mathbb{R}^2 \rightarrow \mathbb{R}$ given by

$$f_{\text{proj}}(x, y) = \int_{-\infty}^{\infty} f(x, y, z) \, dz.$$

By the *projection-slice theorem*, the 2-dimensional Fourier transform of f_{proj} is equal to the slice $\hat{f}_{\text{slice}} : \mathbb{R}^2 \rightarrow \mathbb{C}$ given by

$$\hat{f}_{\text{slice}}(k_x, k_y) = \hat{f}(k_x, k_y, 0).$$

Thus we think of \hat{f} as our unknown signal with $\text{SO}(3)$ acting by rotation, and with post-projection which reveals only the slice of \hat{f} lying in the plane $k_z = 0$.

This does not yet conform to our definition of a (generalized) orbit recovery problem because the signal needs to lie in a finite-dimensional real vector space. Instead of thinking of \hat{f} as a function on \mathbb{R}^3 , we fix a finite number S of nested spherical shells in \mathbb{R}^3 , each of different radius and all centered at the origin. We consider only the restriction of \hat{f} to these shells. We fix a finite number F of frequencies and on each shell we expand \hat{f} (restricted to that shell) in the basis of spherical harmonics, truncated to $1 \leq \ell \leq F$. (As in S^2 registration, we can discard the trivial representation $\ell = 0$ without loss of generality, and it is convenient to do so.) Being the Fourier transform of a real-valued function, \hat{f} satisfies

$$\hat{f}(-k_x, -k_y, -k_z) = \overline{\hat{f}(k_x, k_y, k_z)} \quad (4.2)$$

(see (4.1)) and so we can use a particular basis $H_{\ell m}$ of spherical harmonics for which the expansion coefficients are real; see Appendix C.1. We have now parametrized our signal by a finite number of real values $\theta_{s\ell m}$ with $1 \leq s \leq S$, $1 \leq \ell \leq F$, and $-\ell \leq m \leq \ell$. In particular, the restriction of \hat{f} to shell s has expansion

$$\sum_{1 \leq \ell \leq F} \sum_{-\ell \leq m \leq \ell} \theta_{s\ell m} H_{\ell m}.$$

$\text{SO}(3)$ acts on each shell by 3-dimensional rotation; see Section C.1 for the details of how $\text{SO}(3)$ acts on spherical harmonics. The projection Π reveals only the values on the equator $z = 0$ (or in spherical coordinates, $\theta = \pi/2$) of each shell. Using again the property (4.2), the output of Π on each shell has an expansion with real coefficients in a particular finite basis h_m ; see Section C.1.4.

Remark 4.4.8. There are various other choices one could make for the basis in which to represent the (Fourier transform of the) molecule. Each of our basis functions is the product of a spherical harmonic and a radial delta function (i.e. a delta function applied to the radius, resulting in a spherical shell). Another common basis is the Fourier–Bessel basis (used in e.g. [98]) where each basis function is the product of a spherical harmonic and a radial Bessel function. More generally we can take the product of spherical harmonics with any set of

radial basis function. It turns out that the choice of radial basis is unimportant because the resulting problem will be isomorphic to our case (spherical shells) and so the same results hold.

We now present our results on the above cryo-EM model. We focus on identifying the regime of parameters for which generic list recovery is possible at degree 3. Again using Proposition 4.3.11, we have for $F \geq 2$:

$$\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = \dim(V) - 3 = S \sum_{\ell=1}^F (2\ell + 1) - 3 = S(F^2 + 2F) - 3 \quad (4.3)$$

where again we have a zero-dimensional stabilizer.

In Appendix C.1 we give an explicit construction of the invariant polynomials in $U_{\leq 3}^T$. By testing the Jacobian criterion in exact arithmetic on small examples, we arrive at the following theorem:

Theorem 4.4.9. *Consider the homogeneous ($K = 1$) cryo-EM problem with S shells and F frequencies.*

- *If $S = 1$ then for any $F \geq 2$, generic list recovery is impossible at degree 3.*
- *If $2 \leq S \leq 4$ and $2 \leq F \leq 6$, the degree-3 method of moments achieves generic list recovery.*

The first assertion results from a simple counting argument: there are fewer invariants at degree ≤ 3 than degrees of freedom. The second part is by confirming that the Jacobian of the invariants has rank equal to $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$, through computer-assisted exact arithmetic over an appropriate finite extension of \mathbb{Q} .

In floating-point arithmetic, we have further verified that the Jacobian appears to have appropriate rank for $2 \leq S \leq 10$ and $2 \leq F \leq 10$, leading us to conjecture the following:

Conjecture 4.4.10. *If $S \geq 2$ then the degree-3 method of moments achieves generic list recovery (regardless of F).*

Intuitively, when there is a single shell ($S = 1$) there are simply not enough invariants in $U_{\leq 3}^T$. However, when $S \geq 2$, the number of invariants increases dramatically due to cross-terms that involve multiple shells.

Heterogeneous cryo-EM

We now consider heterogeneous cryo-EM ($K \geq 2$). By combining (4.3) with Proposition 4.3.15 we can compute $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$. Based on testing the Jacobian criterion on small examples, we conjecture that the degree-3 method of moments achieves generic list recovery if and only if $\dim(U_2^T) + \dim(U_3^T) \geq \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$. In other words, we expect no unexpected algebraic dependencies among $U_{\leq 3}^T$. (Recall that there are no degree-1 invariants since we are not using the trivial representation $\ell = 0$).

In Section C.1.6 we give a conjectured formula for the exact value of $\dim(U_2^T) + \dim(U_3^T)$ for all $S \geq 1, F \geq 2$. As a result we can determine for any given $S \geq 1$ and $F \geq 2$, the exact condition on K for which we believe generic list recovery is possible. For S and F large, this condition is approximately $K \leq S^2/4$.

4.5 Open questions

We leave the following as directions for future work.

1. Our methods require testing the rank of the Jacobian on a computer for each problem size. It would be desirable to have analytic results for e.g. (variants of) MRA in any dimension p .
2. We have given an efficient test for whether generic *list* recovery is possible, but have not given a similarly efficient test for generic *unique* recovery. In cases where unique recovery is impossible, it would be nice to give a tight bound on the size of the list; for instance, for MRA with projection, we conjecture that the list has size exactly 2 (due to “chirality”), but we lack a proof for this fact. Our algorithms for testing generic unique recovery are based on Gröbner bases, the calculation of which is known to be

computationally hard in the worst case [79]. Unfortunately, the algorithms we have proposed are also extremely slow in practice, though a faster implementation may be possible.

3. Our procedure for recovering θ from the samples involves solving a polynomial system of equations. While solving polynomial systems is NP-hard in general, the fact that the polynomials used in the orbit recovery problem have special structure leaves open the possibility of finding an efficient (polynomial time) method with rigorous guarantees. This is especially promising under additional assumptions such as random θ . Possible methods include tensor decomposition [118] and non-convex optimization [35]. We discuss this further in Chapter 5.
4. We have addressed the statistical limits of orbit recovery problems. However, the previous chapters of this thesis have indicated the presence of statistical-to-computational gaps in related synchronization problems, and we expect such gaps to appear in orbit recovery problems too. As discussed in Section 4.4.3, the results of [35] suggest a possible gap of this kind for heterogeneous MRA. We discuss this further in Chapter 5.

4.6 Proofs for Section 4.3: algebraic results

4.6.1 Algorithm for generators of $\mathbb{R}[\mathbf{x}]^G$

We know that $\mathbb{R}[\mathbf{x}]^G$ is finitely generated as an \mathbb{R} -algebra (Theorem 4.3.3). There are various algorithms to compute a finite set of generators for $\mathbb{R}[\mathbf{x}]^G$ [143, 55]. However, some require the group to be finite or to be reductive over an algebraically-closed field. One algorithm that certainly works in our context (compact groups) is Algorithm 2.2.5 in [143]. As input it requires the Hilbert series of $\mathbb{R}[\mathbf{x}]^G$ (which can be computed by Proposition 4.3.13) and a procedure to compute a basis for $\mathbb{R}[\mathbf{x}]_d^G$ (which can be done with the Reynolds operator by Observation 4.3.2). The idea is as follows. We keep a set of proposed generators f_1, \dots, f_m . At each step we compare the Hilbert series of $\mathbb{R}[\mathbf{x}]^G$ with the Hilbert series of $\mathbb{R}[f_1, \dots, f_m]$

(which can be computed using Gröbner bases). If these series differ at the t^d term, this means we are missing an invariant at degree d . To remedy this, we create a new homogeneous invariant of degree d using the Reynolds operator, and add it to our set of proposed generators. We repeat until the Hilbert series match.

4.6.2 Bounding the list size for generic signals

In this section we prove Theorem 4.3.24 and the first part of Theorem 4.3.9 (see Section 4.6.3 for the second part). Recall the following basic definitions and facts from field theory.

Definition 4.6.1. If F_2 is a subfield of F_1 , we write F_1/F_2 and call this a *field extension*. The *degree* of the extension, denoted $[F_1 : F_2]$, is the dimension of F_1 as a vector space over F_2 .

Proposition 4.6.2. *Let $\mathbb{R} \subseteq F_2 \subseteq F_1$ with F_1 finitely generated (as a field) over \mathbb{R} . Let r be the transcendence degree of F_1 (over \mathbb{R}). The field extension F_1/F_2 has finite degree if and only if F_1 contains r algebraically independent elements.*

Proof. This is a basic fact of field theory. If F_1 contains r algebraically independent elements then the extension F_1/F_2 is algebraic and finitely generated, and therefore has finite degree. Otherwise, the extension is transcendental and has infinite degree. \square

In light of the above (and using the fact that $\mathbb{R}[\mathbf{x}]^G$ is finitely generated), Theorem 4.3.24 implies the first part of Theorem 4.3.9 and so it remains to prove Theorem 4.3.24 (i.e. list size is bounded by $D \triangleq [F_G : \mathbb{R}(U)]$).

Proof of Theorem 4.3.24.

Write $F_U \triangleq \mathbb{R}(U)$. In characteristic zero, every algebraic extension is separable, so by the primitive element theorem, $F_G = F_U(\alpha)$ for some $\alpha \in F_G$. Since α generates a degree- D extension, α is the root of a degree- D polynomial

$$\alpha^D + b_{D-1}\alpha^{D-1} + \cdots + b_1\alpha + b_0 \quad (4.4)$$

with coefficients $b_i \in F_U$. Furthermore, every element of F_G can be expressed as

$$c_0 + c_1\alpha + \cdots + c_{D-1}\alpha^{D-1}$$

with $c_i \in F_U$. In particular, let g_1, \dots, g_k be generators for $\mathbb{R}[\mathbf{x}]^G$ (as an \mathbb{R} -algebra) and write

$$g_i = c_0^{(i)} + c_1^{(i)}\alpha + \cdots + c_{D-1}^{(i)}\alpha^{D-1}. \quad (4.5)$$

Let $S \subseteq V$ be the subset for which α and all the (finitely-many) coefficients $b_i, c_j^{(i)}$ have nonzero denominators; S is a non-empty Zariski-open set and thus has full measure. Now fix $\theta \in S$. Given the values $f(\theta)$ for all $f \in U$, each b_i takes a well-defined value in \mathbb{R} and so from (4.4) there are at most D possible values that $\alpha(\theta)$ can take. From (4.5), each value of $\alpha(\theta)$ uniquely determines all the values $g_i(\theta)$ and thus uniquely determines all the values $f(\theta)$ for $f \in \mathbb{R}[\mathbf{x}]^G$. Since $\mathbb{R}[\mathbf{x}]^G$ resolves θ (Theorem 4.3.4), this completes the proof. \square

4.6.3 Generic list recovery converse

In this section we prove the second part of Theorem 4.3.9 (the converse).

Let $p = \dim(V)$, $\text{trdeg}(U) = q$, and $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = r$ so that $q < r \leq p$. Let $\mathbf{f} = \{f_1, \dots, f_m\}$ be a basis for U , and let $\mathbf{g} = \{g_1, \dots, g_r\}$ be a transcendence basis for $\mathbb{R}[\mathbf{x}]^G$. Let $S \subseteq V$ be the set of points θ for which the Jacobian $J_{\mathbf{x}}(\mathbf{f})|_{\mathbf{x}=\theta}$ has row rank q and the Jacobian $J_{\mathbf{x}}(\mathbf{g})|_{\mathbf{x}=\theta}$ has row rank r ; by the Jacobian criterion (see Corollary 4.3.19), S is a non-empty Zariski-open set and thus has full measure.

Fix $\theta \in S$. For a sufficiently small open neighborhood $X \subseteq S$ containing θ we have the following. The Jacobian criterion on \mathbf{f} implies that $\{\tau \in X : \mathbf{f}(\tau) = \mathbf{f}(\theta)\}$ has dimension $p - q$. The Jacobian criterion on \mathbf{g} implies that every $z \in \mathbf{g}(X)$ has a preimage $\mathbf{g}^{-1}(z) \triangleq \{\tau \in X : \mathbf{g}(\tau) = z\}$ of dimension $p - r$. Since $p - q > p - r$ it follows that there are infinitely many $\theta_1, \theta_2, \dots \in X$ such that $\mathbf{f}(\theta_i) = \mathbf{f}(\theta)$ but the values $\mathbf{g}(\theta_1), \mathbf{g}(\theta_2), \dots$ are all distinct (and thus the θ_i belong to distinct orbits). Therefore U does not list-resolve θ .

4.6.4 Hilbert series and Hironaka decomposition

In this section we prove Proposition 4.3.14 on extracting the transcendence degree from the Hilbert series (as the pole order at $t = 1$). While this is a general property of finitely generated algebras over a field, there is an easy proof for invariant rings stemming from a key structural property of such rings called the *Cohen-Macaulay property* or *Hironaka decomposition*.

Theorem 4.6.3 ([55] Section 2.6). *The invariant ring $\mathbb{R}[\mathbf{x}]^G$ has the following structure. There exist homogeneous primary invariants $f_1, \dots, f_r \in \mathbb{R}[\mathbf{x}]^G$ and homogeneous secondary invariants $g_1, \dots, g_s \in \mathbb{R}[\mathbf{x}]^G$ such that*

- $\{f_1, \dots, f_r\}$ are algebraically independent, and
- any element of $\mathbb{R}[\mathbf{x}]^G$ can be written uniquely as a linear combination of g_1, \dots, g_s with coefficients from $\mathbb{R}[f_1, \dots, f_r]$.

The proof can be found in Section 2.6 of [55]; note that the only property of the group that is used is the existence of a Reynolds operator (and so the proof is valid for compact groups).

Proof of Proposition 4.3.14.

The Hironaka decomposition above implies that the Hilbert series takes the form

$$\frac{\sum_{j=1}^s t^{\deg(g_j)}}{\prod_{i=1}^r (1 - t^{\deg(f_i)})}$$

(this is equation (2.7.3) in [55]). It is now clear that the order of the pole at $t = 1$ is precisely r . But we can see as follows that f_1, \dots, f_r is a transcendence basis for $\mathbb{R}[\mathbf{x}]^G$ and so $r = \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$. As in the proof of Theorem 4.3.29, since $\mathbb{R}[\mathbf{x}]^G$ is a finitely generated $\mathbb{R}[f_1, \dots, f_r]$ -module, every $h \in \mathbb{R}[\mathbf{x}]^G$ satisfies a polynomial with coefficients in $\mathbb{R}[f_1, \dots, f_r]$, which is an algebraic dependence among $\{f_1, \dots, f_r, h\}$. \square

4.6.5 Transcendence degree for heterogeneity

In this section we prove Proposition 4.3.15. To recall the setup, we have \tilde{G} acting on \tilde{V} with associated variables $\tilde{\mathbf{x}}$. We also have $G = \tilde{G}^K \rtimes S_K$ acting on $V = \tilde{V}^{\oplus K} \oplus \bar{\Delta}_K$ with associated variables \mathbf{x} . Let us also introduce an intermediate group: $G' = \tilde{G}^K$, acting on V (with associated variables \mathbf{x}).

Partition the variables \mathbf{x} as follows. For $k = 1, \dots, K$, let $\mathbf{x}^{(k)} = (x_1^{(k)}, \dots, x_p^{(k)})$ be the variables corresponding to signal k . Let $\mathbf{z} = (z_1, \dots, z_{K-1})$ be the variables corresponding to the mixing weights $\bar{w}_1, \dots, \bar{w}_{K-1}$. Whenever we refer to z_K , this is just shorthand for $-\sum_{k=1}^{K-1} z_k$.

We first prove a simpler version of the result without the action of S_K .

Lemma 4.6.4. *Let $\tilde{r} = \text{trdeg}(\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}})$ and let $r = K\tilde{r} + K - 1$. Then*

$$\text{trdeg}(\mathbb{R}[\mathbf{x}]^{G'}) = r.$$

Proof. To show ‘ \geq ’ we need to exhibit r algebraically independent elements of $\mathbb{R}[\mathbf{x}]^{G'}$. Letting $f_1, \dots, f_{\tilde{r}}$ be a transcendence basis for $\mathbb{R}[\tilde{\mathbf{x}}]^{\tilde{G}}$, it suffices to take

$$I \triangleq \{f_i(\mathbf{x}^{(k)})\}_{1 \leq i \leq \tilde{r}, 1 \leq k \leq K} \cup \{z_1, \dots, z_{K-1}\}.$$

To show ‘ \leq ’ we first recall that we can obtain a spanning set for the subspace $\mathbb{R}[\mathbf{x}]_d^{G'}$ by applying the Reynolds operator \mathcal{R} (for G') to each degree- d monomial (in the variables \mathbf{x}). Such a monomial takes the form

$$m(\mathbf{x}) = M(\mathbf{z}) \prod_{k=1}^K m_k(\mathbf{x}^{(k)})$$

where M, m_k are monomials. Applying the Reynolds operator yields

$$\mathcal{R}(m(\mathbf{x})) = \mathbb{E}_{g_1, \dots, g_K \sim \tilde{G}} M(\mathbf{z}) \prod_{k=1}^K m_k(g_k \cdot \mathbf{x}^{(k)}) = M(\mathbf{z}) \prod_{k=1}^K \mathbb{E}_{g_k \sim \tilde{G}} m_k(g_k \cdot \mathbf{x}^{(k)}).$$

Note that $\mathcal{R}(m(\mathbf{x}))$ is the product of *pure* invariants, i.e. invariants that only involve variables from a single one of the blocks $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(K)}, \mathbf{z}$. It is clear that I (from above) is a maximal set of algebraically independent pure invariants. It is now easy to show using the Jacobian criterion (Proposition 4.3.18) that if any $\mathcal{R}(m(\mathbf{x}))$ is added to I , it will no longer be algebraically independent. The result now follows using basic properties of algebraic independence (Proposition 4.3.22 and Lemma 4.3.23). \square

Proof of Proposition 4.3.15.

Since $\mathbb{R}[\mathbf{x}]^G \subseteq \mathbb{R}[\mathbf{x}]^{G'}$, it is clear (in light of the above) that $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) \leq r$. Thus we need only to show $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) \geq r$ by demonstrating r algebraically independent invariants. Let e_1, \dots, e_K be the elementary symmetric functions in K variables. With f_i as above, we take the invariants

$$\{e_k(f_i(\mathbf{x}^{(1)}), \dots, f_i(\mathbf{x}^{(K)}))\}_{1 \leq i \leq r, 1 \leq k \leq K} \cup \{e_2(z_1, \dots, z_K), \dots, e_K(z_1, \dots, z_K)\}.$$

Note that $e_1(z_1, \dots, z_K)$ is not included because it is equal to 0. The fact that $e_k(f_i(\mathbf{x}^{(1)}), \dots, f_i(\mathbf{x}^{(K)}))$ are algebraically independent can be seen because $\{e_1, \dots, e_K\}$ is algebraically independent and $\{f_i(\mathbf{x}^{(k)})\}_{i,k}$ is algebraically independent. We can see that $\{e_k(z_1, \dots, z_K)\}_{k \geq 2}$ are algebraically independent as follows. An algebraic dependence would be a polynomial P such that $P(e_2(z_2, \dots, z_K), \dots, e_K(z_1, \dots, z_K))$ (now treating z_K as a separate variable) has a root $z_K = -\sum_{k=1}^{K-1} z_k$ and thus has $e_1(z_1, \dots, z_K)$ as factor. But this contradicts the fact that any symmetric polynomial has a *unique* representation in terms of the elementary symmetric polynomials. \square

4.6.6 Gröbner bases

In this section we show how to use Gröbner bases to test various algebraic conditions. In particular, we prove Theorems 4.3.26 and 4.3.28. The ideas from this section are mostly standard in the theory of Gröbner bases; see e.g. [50] for a reference.

Definition 4.6.5. A *monomial order* on $\mathbb{R}[\mathbf{x}]$ is a well-ordering on the set \mathcal{M} of all (monic)

monomials, satisfying $M \leq N \Leftrightarrow MP \leq NP$ for all $M, N, P \in \mathcal{M}$. We will say that a monomial order *favors* a variable x_i if the monomial x_i is larger (with respect to the monomial order) than any monomial not involving x_i . We write $\text{LM}(f)$ to denote the leading monomial of a polynomial f , i.e. the monomial occurring in f that is largest (with respect to the monomial order); $\text{LM}(f)$ does not include the coefficient.

Definition 4.6.6. A *Gröbner basis* of an ideal $I \subseteq \mathbb{R}[\mathbf{x}]$ is a finite subset $B \subseteq I$ such that for every $f \in I$ there exists $b \in B$ such that $\text{LM}(f)$ is a multiple of $\text{LM}(b)$. We call B a *reduced Gröbner basis* if all its elements are monic and it has the additional property that for every pair of distinct $b, b' \in B$, no monomial occurring in b is a multiple of $\text{LM}(b')$.

The following basic facts about Gröbner bases are proved in [50]. A Gröbner basis is indeed a basis, in that it generates the ideal. Every ideal $I \subseteq \mathbb{R}[\mathbf{x}]$ has a Gröbner basis, and has a unique reduced Gröbner basis. *Buchberger's algorithm* computes the reduced Gröbner basis of an ideal $I = \langle f_1, \dots, f_m \rangle$, given a list of generators f_i . (It is not a polynomial-time algorithm, however.)

Suppose we are interested in the relations between polynomials $f_1, \dots, f_m \in \mathbb{R}[\mathbf{x}]$. Introduce additional variables $\mathbf{t} = (t_1, \dots, t_m)$ and consider the ideal $I \triangleq \langle f_1(\mathbf{x}) - t_1, \dots, f_m(\mathbf{x}) - t_m \rangle \subseteq \mathbb{R}[\mathbf{x}, \mathbf{t}]$. Given f_1, \dots, f_m there is an algorithm to compute a Gröbner basis for the *elimination ideal*

$$J \triangleq \langle f_1(\mathbf{x}) - t_1, \dots, f_m(\mathbf{x}) - t_m \rangle \cap \mathbb{R}[\mathbf{t}].$$

In fact, the algorithm is simply to compute a Gröbner basis for I using a particular monomial order and then keep only the elements that depend only on \mathbf{t} (see Chapter 3 of [50]). The elimination ideal consists precisely of the polynomial relations among f_1, \dots, f_m :

Lemma 4.6.7. *For any polynomial $P \in \mathbb{R}[\mathbf{t}]$ we have: $P \in J$ if and only if $P(f_1(\mathbf{x}), \dots, f_m(\mathbf{x})) \equiv 0$.*

Proof. The direction ' \Rightarrow ' is clear because if we let $t_i = f_i(\mathbf{x})$ for all i then the generators of I vanish and so every element of I vanishes. To show the converse, it suffices to show that

for any polynomial $P \in \mathbb{R}[\mathbf{t}]$, $P(f_1(\mathbf{x}), \dots, f_m(\mathbf{x})) - P(t_1, \dots, t_m) \in I$. This can be shown inductively using the following key idea:

$$x_1x_2 - t_1t_2 = \frac{1}{2}(x_1 - t_1)(x_2 + t_2) + \frac{1}{2}(x_2 - t_2)(x_1 + t_1)$$

and so $x_1x_2 - t_1t_2 \in \langle x_1 - t_1, x_2 - t_2 \rangle$. □

Generation as an \mathbb{R} -algebra. Suppose we want to know whether $f_m \in \mathbb{R}[f_1, \dots, f_{m-1}]$. This is equivalent to asking whether there exists $P \in J$ of the form

$$P(\mathbf{t}) = t_m - Q(t_1, \dots, t_{m-1}) \tag{4.6}$$

for some $Q \in \mathbb{R}[t_1, \dots, t_{m-1}]$. Suppose that J contains an element P of the form (4.6). Compute a Gröbner basis B for J with respect to a monomial order that favors t_m . The leading monomial of P is t_m so by the definition of a Gröbner basis there must be an element $b \in B$ whose leading monomial divides t_m . Since $1 \notin J$ (by Lemma 4.6.7), the leading monomial of b is exactly t_m and so b takes the form (4.6). Therefore, $f_m \in \mathbb{R}[f_1, \dots, f_{m-1}]$ if and only if B contains an element of the form (4.6).

We can now prove Theorem 4.3.28: to test whether $\mathbb{R}[f_1, \dots, f_m] = \mathbb{R}[\mathbf{x}]^G$, compute generators g_1, \dots, g_s for $\mathbb{R}[\mathbf{x}]^G$ (see Section 4.6.1) and use the above to test whether each g_i is in $\mathbb{R}[f_1, \dots, f_m]$.

Generation as a field. Suppose we want to know whether $f_m \in \mathbb{R}(f_1, \dots, f_{m-1})$. This is equivalent to asking whether f_m can be expressed as a rational function of f_1, \dots, f_{m-1} (with coefficients in \mathbb{R}), which is equivalent (by multiplying through by the denominator) to asking whether there exists $P \in J$ of the form

$$P(\mathbf{t}) = t_m Q_1(t_1, \dots, t_{m-1}) - Q_2(t_1, \dots, t_{m-1}) \quad \text{with } Q_1 \notin J. \tag{4.7}$$

Suppose that J contains an element P of the form (4.7). Compute a *reduced* Gröbner basis B for J with respect to a monomial order that favors t_m . It is a basic property of Gröbner bases that P can be written as

$$P(\mathbf{t}) = \sum_i p_i(\mathbf{t})b_i(\mathbf{t})$$

where $p_i \in \mathbb{R}[\mathbf{t}]$ and $b_i \in B$ with $\text{LM}(p_i) \leq \text{LM}(P)$ and $\text{LM}(b_i) \leq \text{LM}(P)$. If no b_i involves the variable t_m then $Q_1 \in J$, a contradiction. Therefore some b_j must have degree 1 in t_m . Since B is a reduced Gröbner basis it cannot contain any element of the form (4.7) with $Q_1 \in J$. This completes the proof that $f_m \in \mathbb{R}(f_1, \dots, f_{m-1})$ if and only if B contains an element of the form (4.7).

Degree of field extension. Consider the setup from Theorem 4.3.24: given a finite set $U = \{f_1, \dots, f_m\} \subseteq \mathbb{R}[\mathbf{x}]^G$, we want to compute $[F_G : F_U]$ where $F_U = \mathbb{R}(U)$ and F_G is the field of fractions of $\mathbb{R}[\mathbf{x}]^G$. We can assume $[F_G : F_U]$ is finite (since we can efficiently test whether this is the case using Proposition 4.6.2 and the methods of Section 4.3.2). Let d be such that $\mathbb{R}[\mathbf{x}]_{\leq d}^G$ generates F_G as a field (over \mathbb{R}). (It is sufficient for $\mathbb{R}[\mathbf{x}]_{\leq d}^G$ to generate $\mathbb{R}[\mathbf{x}]^G$ as an \mathbb{R} -algebra; such a d can be computed via Section 4.6.1. If G is finite then $d = |G|$ is sufficient; see Section 4.3.4.) A generic element of $\mathbb{R}[\mathbf{x}]_{\leq d}^G$ will generate the field extension:

Lemma 4.6.8. *For all but a measure-zero set of $\alpha \in \mathbb{R}[\mathbf{x}]_{\leq d}^G$, $F_G = F_U(\alpha)$.*

This fact is related to the *primitive element theorem*. We include a proof for completeness.

Proof. The field extension F_G/F_U is finite and separable (since we're in characteristic zero), so by the fundamental theorem of Galois theory, there are only finitely many intermediate fields. (Take the normal closure of F_G/F_U ; then the intermediate fields are in bijection with a finite group, and only some of them lie inside F_G .) Let \mathcal{L} be the collection of intermediate fields of F_G/F_U that are proper subfields of F_G . We know $\mathbb{R}[\mathbf{x}]_{\leq d}^G$ is a subspace of F_G that generates F_G and therefore is not contained by any field in \mathcal{L} . This means each field $L \in \mathcal{L}$ intersects $\mathbb{R}[\mathbf{x}]_{\leq d}^G$ at a proper subspace V_L of $\mathbb{R}[\mathbf{x}]_{\leq d}^G$. The finite union $\cup_{L \in \mathcal{L}} V_L$ is a measure-zero subset of $\mathbb{R}[\mathbf{x}]_{\leq d}^G$, and any α outside of it satisfies $F_G = F_U(\alpha)$. \square

Let α be a generic element of $\mathbb{R}[\mathbf{x}]_{\leq d}^G$. In light of the above, $[F_G : F_U]$ is equal to the smallest positive integer D for which there exists a relation

$$Q_D(f_1, \dots, f_m)\alpha^D + \dots + Q_1(f_1, \dots, f_m)\alpha + Q_0(f_1, \dots, f_m) \equiv 0$$

for polynomials Q_i with $Q_D(f_1, \dots, f_m) \neq 0$. This can be tested similarly to field generation. Compute a reduced Gröbner basis B for the elimination ideal $J \subseteq \mathbb{R}[t_1, \dots, t_m, \tau]$ consisting of the relations among f_1, \dots, f_m, α ; use a monomial order that favors τ . Then $[F_G : F_U]$ is equal to the smallest positive integer D for which B contains an element of degree D in τ (or ∞ if B contains no element that involves τ). This proves Theorem 4.3.26.

Remark 4.6.9. An alternative to using Gröbner bases for the above tasks is to solve a (very large) linear system in order to find the minimal relation among a set of polynomials. There are bounds on the maximum possible degree of such a relation (if one exists) [87].

4.7 Proofs for S^2 registration

4.7.1 Formula for Hilbert series of $\mathbb{R}[\mathbf{x}]^G$

We can derive the Hilbert series of $\mathbb{R}[\mathbf{x}]^G$ for S^2 registration using the methods in Section 4.6 of [55].

Proposition 4.7.1. *Consider S^2 registration with frequencies \mathcal{F} . For $|t| < 1$, the Hilbert series of $\mathbb{R}[\mathbf{x}]^G$ is given by*

$$H(t) = \sum_{z \in \mathcal{P}} \text{Res}(f, z)$$

where

$$f(z) = \frac{1 - \frac{1}{2}(z + 1/z)}{z \prod_{\ell \in \mathcal{F}} \prod_{m=-\ell}^{\ell} (1 - tz^m)} = \frac{-z^{N-2}(1-z)^2}{2 \prod_{\ell \in \mathcal{F}} \left[\prod_{m=1}^{\ell} (z^m - t) \prod_{m=0}^{\ell} (1 - tz^m) \right]}$$

with $N = \frac{1}{2} \sum_{\ell \in \mathcal{F}} \ell(\ell + 1)$. Here $\text{Res}(f, z)$ denotes the residue (from complex analysis)

of the function f at the point z , and \mathcal{P} is the set of poles of $f(z)$ inside the unit circle (in \mathbb{C}). Namely, \mathcal{P} contains $t^{1/m}e^{2\pi i k/m}$ for all $m \in \{1, 2, \dots, \max_{\ell \in \mathcal{F}} \ell\}$ and for all $k \in \{0, 1, \dots, m-1\}$. If $N \leq 1$, \mathcal{P} also contains 0.

Proof. Recall Molien's formula (Proposition 4.3.13):

$$H(t) = \mathbb{E}_{g \sim \text{Haar}(G)} \det(I - t\rho(g))^{-1}.$$

Note that $\det(I - t\rho(g))$ depends only on the conjugacy class of g . In $\text{SO}(3)$, two elements are conjugate if and only if they rotate by the same angle ϕ . When $g \sim \text{Haar}(\text{SO}(3))$, the angle $\phi = \phi(g)$ is distributed with density function $\frac{1}{\pi}(1 - \cos \phi)$ on $[0, \pi]$ (see e.g. [135]). If g has angle ϕ , the matrix $\rho_\ell(g)$ by which it acts on the irreducible representation V_ℓ has eigenvalues $e^{-i\ell\phi}, e^{-i(\ell-1)\phi}, \dots, e^{i\ell\phi}$ (see e.g. [151]). The matrix $\rho(g)$ by which g acts on $V = \oplus_{\ell \in \mathcal{F}} V_\ell$ is block diagonal with blocks $\rho_\ell(g)$. Using the above we write an expression for the Hilbert series:

$$H(t) = \frac{1}{\pi} \int_0^\pi \frac{1 - \cos \phi}{\prod_{\ell \in \mathcal{F}} \prod_{m=-\ell}^\ell (1 - te^{im\phi})} d\phi = \frac{1}{2\pi} \int_0^{2\pi} \frac{1 - \frac{1}{2}(e^{i\phi} + e^{-i\phi})}{\prod_{\ell \in \mathcal{F}} \prod_{m=-\ell}^\ell (1 - te^{im\phi})} d\phi.$$

Now write this as a complex contour integral around the unit circle in \mathbb{C} and apply the residue theorem from complex analysis to arrive at the result. \square

4.7.2 Formula for dimension of $\mathbb{R}[\mathbf{x}]_d^G$

The dimension of $\mathbb{R}[\mathbf{x}]^G$ can be extracted as the coefficient of t^d in the Hilbert series from the previous section, but here we give a different formula based on character theory from representation theory. The *character* of a representation $\rho : G \rightarrow \text{GL}(V)$ (where V is a finite-dimensional real vector space) is the function $\chi_V : G \rightarrow \mathbb{R}$ defined by $\chi_V(g) = \text{tr}(\rho(g))$.

In our case, using the eigenvalues of $\rho_\ell(g)$ from the previous section, we have

$$\chi_{V_\ell}(g) = 1 + 2 \sum_{m=1}^\ell \cos(m\phi(g))$$

where $\phi(g)$ is the angle of rotation of g . For $V = \oplus_{\ell \in \mathcal{F}} V_\ell$ we then have $\chi_V(g) = \sum_{\ell \in \mathcal{F}} \chi_{V_\ell}(g)$.

As a representation of $G = \text{SO}(3)$, $\mathbb{R}[\mathbf{x}]_d$ is (isomorphic to) the d th symmetric power of V , denoted $S^d(V)$. (This is using the fact that a real representation is isomorphic to its dual.) There is a recursive formula for the character of $S^d(V)$:

$$\chi_{S^d(V)}(g) = \frac{1}{d} \sum_{i=1}^d \chi_V(g^i) \chi_{S^{d-i}(V)}(g).$$

This comes from the Newton–Girard formula for expressing complete homogeneous symmetric polynomials in terms of power sum polynomials.

The representation $\mathbb{R}[\mathbf{x}]_d = S^d(V)$ decomposes as the direct sum of irreducible representations V_ℓ . The subspace of $\mathbb{R}[\mathbf{x}]_d$ consisting of all copies of the trivial representation V_0 (the 1-dimensional representation on which every group element acts as the identity) is precisely $\mathbb{R}[\mathbf{x}]_d^G$. Thus, $\dim(\mathbb{R}[\mathbf{x}]_d^G)$ is the number of copies of the trivial representation in the decomposition of $\mathbb{R}[\mathbf{x}]_d$. This can be computed using characters: $\dim(\mathbb{R}[\mathbf{x}]_d^G) = \langle \chi_{S^d(V)}, \chi_{V_0} \rangle = \langle \chi_{S^d(V)}, 1 \rangle$ where $\langle f_1, f_2 \rangle \triangleq \mathbb{E}_{g \sim \text{Haar}(G)}[f_1(g)f_2(g)]$. Since characters are *class functions* (i.e. they are constant on conjugacy classes), we can compute this inner product by integrating over the angle ϕ (as in the previous section). This yields the formula stated in Proposition 4.4.7.

Chapter 5

Orbit recovery: computational limits

This chapter is based on joint work with Ankur Moitra.

5.1 Introduction

We refer the reader to Chapter 4 for the definitions and basic notions for the orbit recovery problem. While Chapter 4 studied the *statistical* limits of orbit recovery, here we discuss the problem of giving an *efficient* (polynomial time) and provably-correct algorithm for orbit recovery. In Section 5.2 we survey prior work, which has focused on the special case of multi-reference alignment (MRA). This work solves MRA by casting it as a *tensor decomposition* problem. In Section 5.3 we give heuristic predictions for the conditions under which general orbit recovery is solvable efficiently. These heuristics are based on analogy to the *tensor completion* problem. Finally, in Section 5.4 we prove a result regarding efficient solvability of heterogeneous MRA in the case where the signals are random. In particular, we show that recovery from the third moment tensor is efficiently possible so long as $K \leq \sqrt{p}/\text{polylog}(p)$ (where K is the number of heterogeneous components and p is the dimension). Up to log factors, this matches the conjectured threshold of [35].

5.2 Prior work: MRA

Prior work [118] has shown how to provably and efficiently solve the multi-reference alignment (MRA) problem. In this section we restrict to the setting of MRA, although many of the ideas are applicable to other finite groups.

Recall that in the MRA problem we consider the cyclic group $G = \mathbb{Z}/p$ acting on \mathbb{R}^p via cyclic shifts. As discussed in Chapter 4, with $O(\sigma^{2d})$ samples we can accurately estimate the d th moment tensor

$$T_d = \mathbb{E}_g (g \cdot \theta)^{\otimes d} = \frac{1}{p} \sum_{g \in G} (g \cdot \theta)^{\otimes d}.$$

We will attempt reconstruction from just a single T_d (usually $d = 3$). (It appears that there is not much to be gained by combining moments of different orders, since e.g. the third moment contains a lot more information than the second moment.)

The above is a special case of the *tensor decomposition* problem, which can be stated as follows. We are given (exactly or approximately) a low-rank tensor $\sum_{i=1}^m a_i^{\otimes d}$ for some vectors $a_1, \dots, a_m \in \mathbb{R}^p$, and the goal is to recover the components $\{a_i\}$. (If d is even, we can only hope to recover the a_i up to sign, but this sign ambiguity can be resolved using lower-order moments.)

If the components $\{a_i\}$ are linearly independent then *Jennrich's algorithm* (see e.g. [107]) can decompose the third moment, i.e. take $\sum_i a_i^{\otimes 3}$ as input and output the list of $\{a_i\}$. This algorithm is also robust to a certain amount of noise in the input [73]. Building on these results, [118] give an algorithm for MRA (with a generic signal) that decomposes the third moment tensor T_3 , recovering the list of all shifts of the signal: $\{g \cdot \theta\}_{g \in G}$. Note, however, that this method is limited to finite groups with $|G| \leq p$, because otherwise the components $\{g \cdot \theta\}_{g \in G}$ cannot be linearly independent (since there are too many of them).

Recall now the *heterogeneous MRA* problem in which the samples come from a mixture of K different signals $\theta_1, \dots, \theta_K$. For simplicity suppose the mixing weights are all equal. In

this case, the d th moment tensor takes the form

$$T_d = \frac{1}{Kp} \sum_{k=1}^K \sum_{g \in G} (g \cdot \theta_k)^{\otimes d}.$$

It is shown in [118] that provided $K \leq p/2$, Jennrich’s algorithm can be used to decompose the *fifth order* moment tensor and recover the list of all shifts of all signals: $\{g \cdot \theta_k\}_{k \in [K], g \in G}$. However, since this method uses the fifth moment, its sample complexity is $O(\sigma^{10})$. Ideally we would use only the third moment, requiring only $O(\sigma^6)$ samples.

We saw in Chapter 4 that generic list recovery for heterogeneous MRA is statistically possible with $O(\sigma^6)$ samples (i.e. using only the third moment and lower) provided roughly $K \leq p/6$. However, based on numerical simulations with a non-convex solver, it is conjectured by [35] that (in the case where the signals are i.i.d. Gaussian) *efficient* recovery requires roughly $K \leq \sqrt{p}$. In Section 5.3 we will see an alternative method for arriving at this conjecture.

In Section 5.4 we prove one side of the above conjecture (up to log factors). Specifically, we show that when the signals are i.i.d. Gaussian, heterogeneous MRA can be solved efficiently from the third moment so long as $K \leq p/\text{polylog}(p)$. This result builds on prior work on *random overcomplete 3-tensor decomposition* [70, 78, 100], i.e. the problem of recovering $\{a_i\}$ from $\sum_{i=1}^m a_i^{\otimes 3}$ in the case where the a_i are drawn independently from $\mathcal{N}(0, I/p)$. It is known that random overcomplete 3-tensor decomposition can be solved in polynomial time (using the *sum-of-squares hierarchy*) when roughly $m \leq p^{3/2}$ [100]. Note that the MRA third moment has $m = Kp$ components (all p shifts of all K signals) and so the condition $m \leq p^{3/2}$ is equivalent to the conjectured threshold $K \leq \sqrt{p}$. Intuitively, our result in Section 5.4 shows that when the components $\{a_i\}$ of a 3-tensor are taken to be all p shifts of K different random signals, these are “random enough” to behave as if all the a_i ’s were independently drawn from $\mathcal{N}(0, I/p)$. Formally, this amounts to showing certain concentration results for matrices involving shifts of random signals (specifically, equation (5.4) and Proposition 5.4.10).

5.3 Conjectures based on tensor completion

In this section we point out an analogy between general orbit recovery problems and tensor completion. Due to known results on tensor completion, this allows us to predict the computational limits of orbit recovery in high generality.

Consider a generalized orbit recovery problem, allowing for both projection and heterogeneity:

$$y_i = \Pi(g \cdot \theta_{k_i}) + \xi_i$$

with noise $\xi_i \sim \mathcal{N}(0, \sigma^2 I)$, signals $\theta_1, \dots, \theta_K \in \mathbb{R}^p$, and $k_i \in [K]$ chosen randomly according to mixing weights w_1, \dots, w_K . We will think of p being large and will ask how K should scale in order for polynomial-time recovery to be possible. We will restrict our attention to the common case where we hope to perform recovery given (an accurate estimate of) the third moment tensor

$$T_3 = \sum_{k=1}^K w_k \mathbb{E}_g (\Pi(g \cdot \theta_k))^{\otimes 3}.$$

We can rewrite this as

$$T_3 = \Pi^{\otimes 3} \mathbb{E}_g [g^{\otimes 3}] \left(\sum_{k=1}^K w_k \theta_k^{\otimes 3} \right)$$

where we are thinking of Π as a $q \times p$ matrix (so $\Pi^{\otimes 3}$ is $q^3 \times p^3$), g as a $p \times p$ matrix (so $\mathbb{E}_g [g^{\otimes 3}]$ is $p^3 \times p^3$), $\sum_{k=1}^K w_k \theta_k^{\otimes 3}$ as a $p^3 \times 1$ column vector, and T_3 as a $q^3 \times 1$ column vector. In other words, we have a rank- K tensor $\sum_{k=1}^K w_k \theta_k^{\otimes 3}$ and we observe linear measurements of it. The number of “independent” measurements we get is the number of algebraically independent degree-3 invariant polynomials that we have access to: $\text{trdeg}(U_3^T)$ (in the notation of Definition 4.2.6).

The above setup is reminiscent of the *tensor completion* problem in which there is a ground-truth $p \times p \times p$ tensor of rank r , and we observe m randomly-chosen entries. The goal is to recover the remaining hidden entries (using the fact that the tensor is guaranteed to be low-rank). It is known that tensor completion is possible in polynomial time when $m \geq \tilde{O}(rp^{3/2})$ (where \tilde{O} hides log factors) [21, 124]. (Furthermore, there are matching

sum-of-squares lower bounds suggesting that this is optimal [21].)

There are a few differences between orbit recovery and tensor completion. First, in orbit recovery we observe linear measurements that are not simply individual elements (so this is actually an instance of *tensor sensing*). Furthermore, the observed entries in tensor completion are chosen at random, whereas the linear measurements in orbit recovery are fixed. However, if we assume that the orbit recovery signal(s) are sufficiently random, we might expect that the linear measurements act as if they were random. The analogy between orbit recovery and tensor completion thus motivates the following informal conjecture.

Conjecture 5.3.1. *Consider the generalized orbit recovery problem (with heterogeneity and projection) with random signals. List recovery from the third moment is possible in polynomial time provided that $\text{trdeg}(U_3^T) \geq \tilde{O}(Kp^{3/2})$.*

For comparison, recall from Chapter 4 that *statistically*, recovery from the third moment requires roughly $\text{trdeg}(U_3^T) \geq Kp$ (assuming $p \gg \dim(G)$). Note that list recovery is necessary in some cases; for instance, in cryo-EM we can only hope to recover the molecule up to chirality.

For example, in heterogeneous MRA we have $\text{trdeg}(U_3^T) = \Theta(p^2)$, leading us to predict the computational threshold $K \leq \tilde{\Omega}(\sqrt{p})$. This matches the conjecture of [35] discussed above. For cryo-EM with S shells and F frequencies we have $\text{trdeg}(U_3^T) \sim S^3 F^2 / 4$ and $p \sim SF^2$ (see Appendix C.1.6) and so we expect efficient recovery when $K \leq \tilde{\Omega}(S^{3/2}/F)$. In particular, homogeneous ($K = 1$) cryo-EM should require $S^{3/2} \gg F$.

5.4 Heterogeneous MRA

In this section we prove the main result of this chapter: heterogeneous MRA with random signals can be solved in polynomial time provided $K \leq \tilde{\Omega}(\sqrt{p})$. The algorithm uses the *sum-of-squares hierarchy* and builds on previous sum-of-squares algorithms for random overcomplete tensor decomposition [70, 100].

5.4.1 Preliminaries

Notation

We will consider the group $G = \mathbb{Z}/p$ acting on \mathbb{R}^p via cyclic shifts.

For finite sets $X, Y \subseteq \mathbb{R}^p$, define the *Hausdorff distance*

$$d_H(X, Y) = \max\{\max_{x \in X} \min_{y \in Y} \|x - y\|, \max_{y \in Y} \min_{x \in X} \|x - y\|\}.$$

This will be used to measure the error of the algorithm's output.

For a matrix A , $\|A\|$ denotes the spectral norm. For a $d_1 \times d_2 \times d_3$ tensor E , we write $\|E\|_{\{1\}, \{2,3\}}$ to denote the spectral norm of the matrix obtained by flattening E to a $d_1 \times (d_2 d_3)$ matrix.

If \mathcal{A} and \mathcal{B} are systems of polynomial inequalities, we write $\mathcal{A} \vdash \mathcal{B}$ to denote that \mathcal{B} can be derived from \mathcal{A} via a constant-degree sum-of-squares proof (see e.g. [100]).

In our asymptotic notation, $o(1)$ refers to the limit $p \rightarrow \infty$ and \tilde{O} hides factors of $\log(p)$. We say that an event occurs with *high probability* if it has probability $1 - o(1)$ (as $p \rightarrow \infty$). We say that an event occurs with *overwhelming probability* if it has probability $1 - p^{-\omega(1)}$. If we have a polynomial (in p) number of events that each occur with overwhelming probability, their intersection occurs with high probability (by the union bound).

Concentration inequalities

Recall the following version of the matrix Chernoff bound (see e.g. [146]).

Theorem 5.4.1. *Let $\{B_k\}$ be a finite sequence of fixed matrices of dimension $d_1 \times d_2$, and let $\{\xi_i\}$ be a sequence of either i.i.d. Rademacher (uniform ± 1) or i.i.d. standard normal random variables. Define the variance parameter*

$$\sigma^2 = \max\{\|\sum_k B_k B_k^\top\|, \|\sum_k B_k^\top B_k\|\}.$$

Then for all $t > 0$,

$$\Pr \left\{ \left\| \sum_k \xi_k B_k \right\| \geq t \right\} \leq (d_1 + d_2) \cdot \exp(-t^2/2\sigma^2).$$

We will often use the basic variance bound

$$\sigma^2 \leq \sum_k \|B_k\|^2. \quad (5.1)$$

We will need the following decoupling result of [52].

Theorem 5.4.2 ([52] Theorem 1). *Let $\{X_i\}$ be a sequence of independent random variables in a measurable space (\mathcal{S}, S) and let $\{X_i^{(j)}\}$, $j = 1, \dots, k$ be k independent copies of $\{X_i\}$. Let f_{i_1, \dots, i_k} be a family of functions of k variables taking $(S \times \dots \times S)$ into a Banach space $(B, \|\cdot\|)$. Then for all $n \geq k \geq 2$, $t > 0$, there exist numerical constants C_k depending only on k so that*

$$\begin{aligned} & \Pr \left(\left\| \sum_{1 \leq i_1 \neq \dots \neq i_k \leq n} f_{i_1, \dots, i_k}(X_{i_1}^{(1)}, X_{i_2}^{(1)}, \dots, X_{i_k}^{(1)}) \right\| \geq t \right) \\ & \leq C_k \Pr \left(\left\| \sum_{1 \leq i_1 \neq \dots \neq i_k \leq n} f_{i_1, \dots, i_k}(X_{i_1}^{(1)}, X_{i_2}^{(2)}, \dots, X_{i_k}^{(k)}) \right\| \geq t/C_k \right) \end{aligned}$$

where the notation $1 \leq i_1 \neq \dots \neq i_k \leq n$ denotes that i_1, \dots, i_k are distinct values in $\{1, 2, \dots, n\}$.

The following concentration bound is a consequence of *hypercontractivity* (see e.g. Theorem 1.10 of [138]).

Theorem 5.4.3. *Consider a degree- q polynomial $f(Y) = f(Y_1, \dots, Y_n)$ of independent centered Gaussian or Rademacher random variables Y_1, \dots, Y_n . Let σ^2 be the variance of $f(Y)$. There exists an absolute constant $R > 0$ such that*

$$\Pr [|f(Y) - \mathbb{E}[f(Y)]| \geq t] \leq e^2 \cdot e^{-\left(\frac{t^2}{R\sigma^2}\right)^{1/q}}.$$

Fourier basis

It will help to pass to the Fourier basis of \mathbb{C}^p . The basis vectors are $\{\widehat{e}_k\}$ for $k \in \mathbb{Z}/p$ with

$$(\widehat{e}_k)_j = \frac{1}{\sqrt{p}} \exp(-2\pi i j k / p)$$

where i is the imaginary unit. This basis is convenient because $g \in \mathbb{Z}/p$ acts diagonally:

$$g \cdot \widehat{e}_k = g_k \widehat{e}_k$$

where

$$g_k \triangleq \exp(2\pi i g k / p).$$

In the standard real basis, we have $\theta \sim \mathcal{N}(0, I/p)$. In the Fourier basis, this corresponds to $\theta = \sum_k \widehat{\theta}_k \widehat{e}_k$ with $\widehat{\theta}_0 \sim \mathcal{N}(0, 1/p)$, $\widehat{\theta}_k \sim \mathcal{N}(0, 1/2p) + i\mathcal{N}(0, 1/2p)$, all independent except for the symmetry $\widehat{\theta}_{-k} = \overline{\widehat{\theta}_k}$ (complex conjugate). If p is even, there is an additional special case: $\widehat{\theta}_{p/2} \sim \mathcal{N}(0, 1/p)$.

5.4.2 Basic facts

Here we give a few basic facts about randomly-chosen MRA vectors that we will use throughout. We consider vectors $\theta^1, \dots, \theta^K \in \mathbb{R}^p$ drawn independently from $\mathcal{N}(0, I/p)$, with $K \leq \sqrt{p}$. Let $m = Kp$ and let $\{a_1, \dots, a_m\} = \{g \cdot \theta^k\}_{k \in [K], g \in \mathbb{Z}/p}$. First we show “incoherence” of $\{a_i\}$.

Lemma 5.4.4. *Let a_1, \dots, a_m be drawn randomly as above. For each i ,*

$$|\|a_i\|^2 - 1| \leq \tilde{O}(1/\sqrt{p})$$

with overwhelming probability. For each pair $i \neq j$,

$$|\langle a_i, a_j \rangle| \leq \tilde{O}(1/\sqrt{p})$$

with overwhelming probability.

Proof. The first statement follows from Bernstein's inequality for subexponential random variables (see e.g. [127]). When a_i and a_j are not shifts of the same signal θ^k , the second statement follows from the fact that if a_i is fixed then over the randomness of a_j we have $\langle a_i, a_j \rangle \sim \mathcal{N}(0, \|a_i\|^2/p)$. To bound $|\langle a_i, a_j \rangle|$ when a_i and a_j are two shifts of the same signal θ^k , apply the decoupling theorem (Theorem 5.4.2) to reduce to the case where a_i, a_j are drawn independently from $\mathcal{N}(0, I/p)$. \square

Lemma 5.4.5. *Let a_1, \dots, a_m be drawn randomly as above and let A be the $p \times m$ matrix with columns $\{a_i\}$. With high probability, $\|A\| \leq \tilde{O}(\sqrt{K})$.*

Proof. Write $A = \sum_{k=1}^K \sum_{i=1}^p \theta_i^k A^{ki}$ for the appropriate choice of fixed matrices A^{ki} . Note that $\|A^{ki}\| = 1$ since A^{ki} has a single 1 per row (and all other entries zero). The result now follows from the matrix Chernoff bound (Theorem 5.4.1) using the basic variance bound (5.1). \square

5.4.3 Main result

The following is our main result.

Theorem 5.4.6. *Let $K \leq \sqrt{p}$ and let $\theta^1, \dots, \theta^K$ be vectors in \mathbb{R}^p drawn independently from $\mathcal{N}(0, I/p)$. Let $\Theta = \{g \cdot \theta^k\}_{k \in [K], g \in \mathbb{Z}/p}$. There exists a polynomial-time algorithm that takes as input K and $\varepsilon > 0$ (with ε smaller than some constant) along with the 3-tensor*

$$T = \sum_{k=1}^K \sum_{g \in \mathbb{Z}/p} (g \cdot \theta^k)^{\otimes 3} + E \quad \text{with } \|E\|_{\{1\}, \{2,3\}} \leq \varepsilon \quad (5.2)$$

and outputs $\hat{\Theta} = \{\hat{\theta}^1, \dots, \hat{\theta}^{Kp}\}$ such that $d_H(\hat{\Theta}, \Theta) \leq \tilde{O}((K/\sqrt{p})^{1/4}) + O(\varepsilon^{1/4})$ with high probability over the random choice of Θ .

Remark 5.4.7. The connection between ε and the number of MRA samples n can be derived from [15] (the full version of Chapter 4) as follows. By Proposition 7.6 in [15], there is an

estimator that, with probability $1 - \delta$, outputs a tensor T such that every entry of the error E is $O(\sigma^3 \sqrt{\log(p/\delta)/n})$. Using the Frobenius norm bound $\|E\|_{\{1\},\{2,3\}} \leq \|E\|_F$ we see that $\|E\|_{\{1\},\{2,3\}} \leq \varepsilon$ is achievable with $n = O(\sigma^6 p^3 \log(p/\delta)/\varepsilon^2)$.

We will use the tensor decomposition framework of [100], in particular their Theorem 5.2 which we restate here.

Theorem 5.4.8 ([100] Theorem 5.2). *For every $\ell \in \mathbb{N}$, there exists an $n^{O(\ell)}$ -time algorithm with the following property: Let $\gamma > 0$ be smaller than some constant. Let $p, p' \in \mathbb{N}$ be numbers. Let $P : \mathbb{R}^p \rightarrow \mathbb{R}^{p'}$ be a polynomial with $\deg P \leq \ell$. Let $\{a_1, \dots, a_m\} \subseteq \mathbb{R}^p$ be a set of vectors such that $b_1 = P(a_1), \dots, b_m = P(a_m) \in \mathbb{R}^{p'}$ all have norm at least $1 - \gamma$ and $\|\sum_{i=1}^m b_i b_i^\top\| \leq 1 + \gamma$. Let \mathcal{A} be a system of polynomial inequalities in variables $u = (u_1, \dots, u_p)$ such that the vectors a_1, \dots, a_m satisfy \mathcal{A} and*

$$\mathcal{A} \vdash \left\{ \sum_{i=1}^m \langle b_i, P(u) \rangle^4 \geq (1 - \gamma) \|P(u)\|^4 \right\}. \quad (5.3)$$

Then, the algorithm on input \mathcal{A} and P outputs a set of unit vectors $\{b'_1, \dots, b'_m\} \subseteq \mathbb{R}^{p'}$ such that

$$d_H(\{b_1^{\otimes 2}, \dots, b_m^{\otimes 2}\}, \{(b'_1)^{\otimes 2}, \dots, (b'_m)^{\otimes 2}\}) \leq O(\gamma^{1/2}).$$

In our to apply the above theorem to our setting, we need to check the following conditions; this is the analogue of Proposition 7.1 in [100].

Proposition 5.4.9. *Let $\theta^1, \dots, \theta^K$ be drawn randomly as in Theorem 5.4.6 with $K \leq \sqrt{p}$. Let $m = Kp$ and let $\{a_1, \dots, a_m\} = \{g \cdot \theta^k\}_{k \in [K], g \in \mathbb{Z}/p}$. Let T and ε be as in (5.2). Let $\mathcal{A} = \{\langle T, u^{\otimes 3} \rangle \geq 1 - \eta, \|u\|^2 = 1\}$ for a particular $\eta = \varepsilon + \tilde{O}(K/\sqrt{p})$. Let $P : \mathbb{R}^p \rightarrow \mathbb{R}^{p^2}$ be given by $P(a) = a^{\otimes 2} - \|a\|^2 p^{-1} \sum_{i=1}^p e_i^{\otimes 2}$. For $i = 1, \dots, m$, let $b_i = P(a_i)$. Then with probability $1 - o(1)$,*

$$\left\| \sum_{i=1}^m b_i b_i^\top \right\| \leq 1 + \tilde{O}(K/\sqrt{p}) \quad (5.4)$$

and

$$\mathcal{A} \vdash \sum_{i=1}^m \langle b_i, P(u) \rangle^4 \geq (1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p})) \|P(u)\|^4. \quad (5.5)$$

Note that P can be thought of as projecting a (vectorized) $p \times p$ matrix onto the subspace of trace-zero matrices. This makes b_i a mean-zero vector, without which (5.4) would not hold. The proof of Proposition 5.4.9 will span Sections 5.4.4 and 5.4.5. Let us now see how Proposition 5.4.9 implies Theorem 5.4.6.

Proof of Theorem 5.4.6. Let $\gamma > 0$, to be chosen later. Define \mathcal{A} and P as in Proposition 5.4.9. We need to verify that the conditions of Theorem 5.4.8 hold with high probability. Using Lemma 5.4.4, the norm bound $\|b_i\| \geq 1 - \gamma$ holds for all i provided γ exceeds $\tilde{O}(1/\sqrt{p})$. By Proposition 5.4.9, the spectral bound $\|\sum_{i=1}^m b_i b_i^\top\| \leq 1 + \gamma$ holds provided γ exceeds $\tilde{O}(K/\sqrt{p})$. To check that a_1, \dots, a_m satisfy \mathcal{A} , write

$$\begin{aligned} \langle T, a_i^{\otimes 3} \rangle &= \|a_i\|^6 + \sum_{j \neq i} \langle a_i, a_j \rangle^3 + \langle E, a_i^{\otimes 3} \rangle \\ &\geq 1 - \tilde{O}(1/\sqrt{p}) - Kp\tilde{O}(1/p^{3/2}) - \|E\|_{\{1\},\{2,3\}} \|a_i\|^3 \\ &\geq 1 - \tilde{O}(K/\sqrt{p}) - \varepsilon(1 + \tilde{O}(1/\sqrt{p})) \\ &\geq 1 - \tilde{O}(K/\sqrt{p}) - \varepsilon \end{aligned}$$

where we have used Lemma 5.4.4 and the fact that ε is smaller than some constant. Thus we can choose $\eta = \varepsilon + \tilde{O}(K/\sqrt{p})$ such that a_1, \dots, a_m satisfy \mathcal{A} . The sum-of-squares proof (5.3) follows from Proposition 5.4.9 provided γ exceeds $O(\varepsilon) + \tilde{O}(K/\sqrt{p})$. We have now satisfied the conditions of Theorem 5.4.8 with $\gamma = O(\varepsilon) + \tilde{O}(K/\sqrt{p})$.

It remains to show how to extract our estimate $\hat{\Theta}$ from the outputs b'_i of Theorem 5.4.8. The guarantee of Theorem 5.4.8 implies that under some re-ordering of $\{b'_i\}$ we have $\|b_i \pm b'_i\|^2 \leq O(\gamma)$. Let B_i and B'_i be (respectively) the flattenings of b_i and b'_i to $p \times p$ matrices. Since $\|B_i - a_i a_i^\top\| \leq \tilde{O}(1/\sqrt{p})$ and $\|B_i \pm B'_i\|_F \leq O(\sqrt{\gamma})$, we have $\|\pm B'_i - a_i a_i^\top\| \leq O(\sqrt{\gamma})$. This means $\langle a_i, a'_i \rangle^2 \geq 1 - \tilde{O}(\sqrt{\gamma})$ where a'_i is the leading (unit-norm) eigenvector of B'_i . Flip the sign of a'_i if necessary so that $\langle T, (a'_i)^{\otimes 3} \rangle > 0$. Now output $\{a'_i\}$ as the estimator for $\{a_i\}$, and we have $\|a_i - a'_i\| \leq \tilde{O}(\gamma^{1/4})$ as desired. \square

The following matrix concentration result will be the key ingredient in the sum-of-squares

proof (5.5). Its proof can be found in Section 5.4.4.

Proposition 5.4.10. *With high probability, the matrix*

$$\sum_{i \neq j} \langle a_i, a_j \rangle (a_i \otimes a_j)(a_i \otimes a_j)^\top$$

(with i, j ranging over $[m]$) has spectral norm at most $\tilde{O}(K/\sqrt{p})$.

5.4.4 Spectral bounds

Proof of (5.4)

In this section we prove the first part of Proposition 5.4.9, namely (5.4).

Note that $\sum_{i=1}^m b_i b_i^\top = BB^\top$ where the columns of B are the b_i 's. We would like to bound $\|BB^\top\|$, but this is difficult because BB^\top does not concentrate near its expectation. Since BB^\top and $B^\top B$ have the same nonzero singular values, it is sufficient to instead bound $\|B^\top B\|$. This turns out to be much easier because the Gram matrix $B^\top B$ concentrates near the identity. Index the rows and columns of $B^\top B$ by pairs (k, g) with $1 \leq k \leq K$ and $g \in \mathbb{Z}/p$. We have

$$\begin{aligned} (B^\top B)_{kg, \ell h} &= \langle P(g \cdot \theta^k), P(h \cdot \theta^\ell) \rangle \\ &= \langle g \cdot \theta^k, h \cdot \theta^\ell \rangle^2 - \|\theta^k\|^2 \|\theta^\ell\|^2 / p. \end{aligned}$$

Since $|\|\theta^k\|^2 - 1| \leq \tilde{O}(1/\sqrt{p})$ by Lemma 5.4.4, we can replace the term $\|\theta^k\|^2 \|\theta^\ell\|^2 / p$ by simply $1/p$ while incurring an error of size $\tilde{O}(K/\sqrt{p})$ in Frobenius norm (which is an upper bound on spectral norm). Thus it is sufficient to show $\|C\| \leq 1 + \tilde{O}(K/\sqrt{p})$ where

$$C_{kg, \ell h} = \langle g \cdot \theta^k, h \cdot \theta^\ell \rangle^2 - 1/p.$$

Fix k, ℓ and focus on the corresponding $p \times p$ block $C^{k\ell}$ of C : $C_{gh}^{k\ell} \triangleq C_{kg, \ell h}$. Expand in

the Fourier basis:

$$\begin{aligned} C_{gh}^{k\ell} &= \left(\sum_a \overline{g_a \widehat{\theta}_a^k} h_a \widehat{\theta}_a^\ell \right) \left(\sum_b g_b \widehat{\theta}_b^k \overline{h_b \widehat{\theta}_b^\ell} \right) - 1/p = \sum_{ab} \overline{\widehat{\theta}_a^k \widehat{\theta}_a^\ell \widehat{\theta}_b^k \widehat{\theta}_b^\ell} g_a h_a g_b \overline{h_b} - 1/p \\ &= \sum_i \alpha_i^{k\ell} \overline{g_i} h_i \end{aligned}$$

where

$$\alpha_i^{k\ell} \triangleq \sum_{a-b=i} \overline{\widehat{\theta}_a^k \widehat{\theta}_a^\ell \widehat{\theta}_b^k \widehat{\theta}_b^\ell} - \frac{1}{p} \mathbb{1}_{i=0}$$

with a, b ranging over \mathbb{Z}/p (and subtraction $a - b$ performed mod p). Consider the quadratic form of $C^{k\ell}$ against unit vectors u, v written in the form $u_g = \sum_i \beta_i g_i / \sqrt{p}$ and $v_g = \sum_i \gamma_i g_i / \sqrt{p}$ with $\beta_i, \gamma_i \in \mathbb{C}$ such that $\sum_i \beta_i^2 = \sum_i \gamma_i^2 = 1$:

$$\begin{aligned} u^* C^{k\ell} v &= p^2 \mathbb{E}_{gh} \overline{u_g} C_{gh}^{k\ell} v_h = p \mathbb{E}_{gh} \left(\sum_i \overline{\beta_i g_i} \right) \left(\sum_i \alpha_i^{k\ell} \overline{g_i} h_i \right) \left(\sum_i \gamma_i h_i \right) \\ &= p \sum_i \alpha_{-i}^{k\ell} \overline{\beta_i} \gamma_i \end{aligned}$$

where the expectation is over Haar (uniform) measure on \mathbb{Z}/p . Thus we have $\|C^{k\ell}\| = \max_i p \alpha_i^{k\ell}$. The following two lemmas show that these values are tightly concentrated.

Lemma 5.4.11. *With high probability we have that for all k , $\|C^{kk}\| \leq 1 + \tilde{O}(1/\sqrt{p})$.*

Proof. When $k = \ell$ we have

$$\alpha_i^{kk} = \sum_{a-b=i} |\widehat{\theta}_a^k|^2 |\widehat{\theta}_b^k|^2 - \frac{1}{p} \mathbb{1}_{i=0}.$$

Note that

$$\mathbb{E}[|\widehat{\theta}_a^k|^2 |\widehat{\theta}_b^k|^2] = \begin{cases} 3/p^2 & a = b = 0, \\ 3/p^2 & a = b = p/2, \\ 2/p^2 & a = \pm b \notin \{0, p/2\}, \\ 1/p^2 & \text{otherwise.} \end{cases}$$

As a result we have $\mathbb{E}[p \alpha_i^{kk}] = 1 \pm O(1/p)$. The variance is $\text{Var}(p \alpha_i^{kk}) \leq O(1/p)$. Note

that $p\alpha_i^{kk}$ is a degree-4 polynomial of independent centered Gaussians (namely the real and imaginary parts of the θ_i^k). By hypercontractivity (Theorem 5.4.3) we have $|p\alpha_i^{kk} - 1| \leq \tilde{O}(1/\sqrt{p})$ with overwhelming probability. \square

Lemma 5.4.12. *With high probability we have that for all pairs k, ℓ with $k \neq \ell$, $\|C^{k\ell}\| \leq \tilde{O}(1/\sqrt{p})$.*

Proof. When $k \neq \ell$ we have

$$\mathbb{E}[\widehat{\theta}_a^k \widehat{\theta}_a^\ell \widehat{\theta}_b^k \widehat{\theta}_b^\ell] = \begin{cases} 1/p^2 & a = b, \\ 0 & \text{otherwise,} \end{cases}$$

and so $\mathbb{E}[p\alpha_i^{k\ell}] = 0$. The real and imaginary parts of $p\alpha_i^{k\ell}$ each have variance $O(1/p)$ and are each a degree-4 polynomial of independent centered Gaussians, so as above we can apply hypercontractivity (Theorem 5.4.3) to conclude $|p\alpha_i^{k\ell}| \leq \tilde{O}(1/\sqrt{p})$ with overwhelming probability. \square

It now follows from Lemmas 5.4.11 and 5.4.12 that $\|C\| \leq 1 + \tilde{O}(K/\sqrt{p})$ as desired. For the off-diagonal blocks we have used the fact that if M is $Kp \times Kp$ with each $p \times p$ block having spectral norm $\leq s$, then $\|M\| \leq Ks$. This completes the proof of (5.4).

For the sum-of-squares proof (5.5), we will also need the following corollary of the above.

Corollary 5.4.13. *With high probability,*

$$\{\|u\|^2 = 1\} \vdash \sum_{i=1}^m \langle a_i, u \rangle^4 \leq O(1) + \tilde{O}(K/\sqrt{p}).$$

Proof. We have

$$\begin{aligned}
\sum_i \langle a_i, u \rangle^4 &= \sum_i \langle a_i^{\otimes 2}, u^{\otimes 2} \rangle^2 \\
&= \sum_i \langle a_i^{\otimes 2} - \|a_i\|^2 p^{-1} \sum_{j=1}^p e_j^{\otimes 2} + \|a_i\|^2 p^{-1} \sum_{j=1}^p e_j^{\otimes 2}, u^{\otimes 2} \rangle^2 \\
&= \sum_i (\langle b_i, u^{\otimes 2} \rangle + \|a_i\|^2 \|u\|^2 / p)^2 \\
&\leq 4 \sum_i \langle b_i, u^{\otimes 2} \rangle^2 + 4 \sum_i (\|a_i\|^2 / p)^2 \\
&\leq 4 \left\| \sum_i b_i b_i^\top \right\| + O(K/p) \\
&\leq O(1) + \tilde{O}(K/\sqrt{p}).
\end{aligned}$$

□

Proof of Proposition 5.4.10

Let $M = \sum_{i \neq j} \langle a_i, a_j \rangle (a_i \otimes a_j)(a_i \otimes a_j)^\top$, the matrix we want to bound. We can write $M = M' - M''$ where

$$M' = \sum_{k\ell} \sum_{gh} \langle g \cdot \theta^k, h \cdot \theta^\ell \rangle (g \cdot \theta^k \otimes h \cdot \theta^\ell)(g \cdot \theta^k \otimes h \cdot \theta^\ell)^\top$$

and

$$M'' = \sum_k \sum_g \langle g \cdot \theta^k, g \cdot \theta^k \rangle (g \cdot \theta^k \otimes g \cdot \theta^k)(g \cdot \theta^k \otimes g \cdot \theta^k)^\top$$

with $k, \ell \in [K]$ and $g, h \in \mathbb{Z}/p$. Change basis to the Fourier basis and write out the entries of M', M'' :

$$\begin{aligned}
M'_{ab,cd} &= \sum_{k\ell} p^2 \mathbb{E}_{gh} \left(\sum_{e=1}^p \overline{g_e \widehat{\theta}_e^k} h_e \widehat{\theta}_e^\ell \right) g_a \widehat{\theta}_a^k \overline{h_b \widehat{\theta}_b^\ell} g_c \widehat{\theta}_c^k h_d \widehat{\theta}_d^\ell \\
&= \sum_{k\ell} p^2 \mathbb{1}_{a-b=c-d \triangleq \delta} \overline{\widehat{\theta}_{a-c}^k} \widehat{\theta}_{a-c}^\ell \widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^\ell} \widehat{\theta}_c^k \widehat{\theta}_{c-\delta}^\ell,
\end{aligned}$$

$$\begin{aligned}
M''_{ab,cd} &= \sum_k p \mathbb{E}_g \|\theta^k\|^2 g_a \widehat{\theta}_a^k g_b \overline{\widehat{\theta}_b^k} g_c \widehat{\theta}_c^k g_d \overline{\widehat{\theta}_d^k} \\
&= \sum_k p \|\theta^k\|^2 \mathbb{1}_{a-b=c-d \triangleq \delta} \widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^k} \widehat{\theta}_c^k \overline{\widehat{\theta}_{c-\delta}^k}.
\end{aligned}$$

Note that (after permuting the rows/columns), M is block-diagonal with one $p \times p$ block for each value of $\delta \triangleq a - b = c - d$. Let M^δ denote this block, i.e. $M_{ac}^\delta = M_{(a)(a-\delta), (c)(c-\delta)}$.

Now focus on a particular block (i.e. fix δ) and write $M^\delta = P^\delta + Q^\delta$ where

$$P_{ac}^\delta = \sum_k (p^2 |\widehat{\theta}_{a-c}^k|^2 - p \|\theta^k\|^2) \widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^k} \widehat{\theta}_c^k \overline{\widehat{\theta}_{c-\delta}^k}$$

and

$$Q_{ac}^\delta = \sum_{k \neq \ell} p^2 \overline{\widehat{\theta}_{a-c}^k} \widehat{\theta}_{a-c}^\ell \widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^\ell} \widehat{\theta}_c^\ell \overline{\widehat{\theta}_{c-\delta}^\ell}.$$

To bound $\|P^\delta\|$, write $P^\delta = p \sum_k D_k C_k D_k^*$ where $D_k = \text{diag}(\widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^k})$ and $(C_k)_{ac} = p |\widehat{\theta}_{a-c}^k|^2 - \|\theta^k\|^2$. Each $|\widehat{\theta}_i^k|$ is at most $\tilde{O}(1/\sqrt{p})$ with overwhelming probability and so $\|D_k\| \leq \tilde{O}(1/p)$. C_k is a circulant matrix and thus its eigenvalues can be written in closed form:

$$\lambda_j = \sum_{i=0}^{p-1} (p |\widehat{\theta}_i^k|^2 - \|\theta^k\|^2) \omega_j^i \quad \text{for } j = 0, 1, \dots, p-1$$

where $\{\omega_j\}$ are the p th roots of unity $\omega_j = 2\pi i j/p$. Note that $\lambda_0 = 0$ and for $j \neq 0$, we have $\sum_{i=0}^{p-1} \omega_j^i = 0$ and so

$$\lambda_j = \sum_{i=0}^{p-1} p |\widehat{\theta}_i^k|^2 \omega_j^i = \sum_{i=0}^{p-1} (p |\widehat{\theta}_i^k|^2 - 1) \omega_j^i.$$

Note that $p |\widehat{\theta}_i^k|^2 - 1$ is distributed either as $\chi_1^2 - 1$ or $\frac{1}{2} \chi_2^2 - 1$ (depending on i) which in either case is a (mean-zero) subexponential random variable. We have $|\widehat{\theta}_i^k| = |\widehat{\theta}_{-i}^k|$ but otherwise these values are independent. By Bernstein's inequality for subexponential random variables (see [127]) we have $|\lambda_j| \leq \tilde{O}(\sqrt{p})$ for all j , and so $\|C_k\| \leq \tilde{O}(\sqrt{p})$. Now

$$\|P^\delta\| = \|p \sum_k D_k C_k D_k^*\| \leq p \sum_k \|D_k\|^2 \|C_k\| \leq \tilde{O}(K/\sqrt{p}).$$

With high probability, this holds for all δ .

To bound $\|Q^\delta\|$, first apply a symmetrization trick: θ^k has the same distribution as $\sigma_k \theta^k$ where σ_k is an independent Rademacher random variable. Now Q^δ is distributed as $p \sum_{k \neq \ell} \sigma_k \sigma_\ell D_{k\ell} C_{k\ell} D_{k\ell}^*$ with $D_{k\ell} = \text{diag}(\widehat{\theta}_a^k \overline{\widehat{\theta}_{a-\delta}^\ell})$ and $(C_{k\ell})_{ac} = p \overline{\widehat{\theta}_{a-c}^k} \widehat{\theta}_{a-c}^\ell$. As above, we have $\|D_{k\ell}\| \leq \tilde{O}(1/p)$. Again, $C_{k\ell}$ is circulant with eigenvalues

$$\lambda_j = \sum_{i=0}^{p-1} p \overline{\widehat{\theta}_i^k} \widehat{\theta}_i^\ell \omega_j^i \quad \text{for } j = 0, 1, \dots, p-1.$$

Combining terms i and $-i$ (modulo p),

$$\begin{aligned} p \overline{\widehat{\theta}_i^k} \widehat{\theta}_i^\ell \omega_j^i + p \overline{\widehat{\theta}_{-i}^k} \widehat{\theta}_{-i}^\ell \omega_j^{-i} &= 2p \Re(\overline{\widehat{\theta}_i^k} \widehat{\theta}_i^\ell \omega_j^i) \\ &= 2p[(\Re(\widehat{\theta}_i^k) \Re(\widehat{\theta}_i^\ell) + \Im(\widehat{\theta}_i^k) \Im(\widehat{\theta}_i^\ell)) \cos(2\pi i j / p) \\ &\quad + (-\Re(\widehat{\theta}_i^k) \Im(\widehat{\theta}_i^\ell) + \Im(\widehat{\theta}_i^k) \Re(\widehat{\theta}_i^\ell)) \sin(2\pi i j / p)]. \end{aligned}$$

Apply Bernstein separately to the cosine and sine terms, yielding $|\lambda_j| \leq \tilde{O}(\sqrt{p})$ and so $\|p D_{k\ell} C_{k\ell} D_{k\ell}^*\| \leq \tilde{O}(1/\sqrt{p})$. By decoupling (Theorem 5.4.2) we can replace $\sigma_k \sigma_\ell$ with independent copies $\sigma_k^{(1)} \sigma_\ell^{(2)}$. Now apply the matrix Chernoff bound twice (once over the randomness of $\sigma^{(2)}$ and then again over the randomness of $\sigma^{(1)}$) using the basic variance bound (5.1) to conclude:

$$\begin{aligned} \text{for fixed } k, \quad \left\| \sum_{\ell \neq k} \sigma_\ell^{(2)} (p D_{k\ell} C_{k\ell} D_{k\ell}^*) \right\| &\leq \tilde{O}(\sqrt{K/p}), \\ \left\| \sum_k \sigma_k^{(1)} \left(\sum_{\ell \neq k} \sigma_\ell^{(2)} p D_{k\ell} C_{k\ell} D_{k\ell}^* \right) \right\| &\leq \tilde{O}(K/\sqrt{p}). \end{aligned}$$

Therefore with high probability we have $\|Q^\delta\| \leq \tilde{O}(K/\sqrt{p})$ for all δ .

5.4.5 SoS proof

In this section we prove (5.5), completing the proof of Proposition 5.4.9. The arguments follow Proposition 7.1 of [100], which in turn is based on Theorem 4.2 of [70].

Remark 5.4.14. The prior works [70] and [100] are written for the case where $\|a_i\| = 1$, whereas this is only approximately true in our case. For this reason, we have some extra factors of $\|a_i\|^2$ below, which do not appear in prior work. However, since these are $1 \pm \tilde{O}(1/\sqrt{p})$ they do not affect the result.

Recall $\mathcal{A} = \{\langle T, u^{\otimes 3} \rangle \geq 1 - \eta, \|u\| = 1\}$ where $\eta = \varepsilon + \tilde{O}(K/\sqrt{p})$, $T = \tilde{T} + E$, $\tilde{T} = \sum_{i=1}^m a_i^{\otimes 3}$ and $\|E\|_{\{1\}, \{2,3\}} \leq \varepsilon$. We immediately have

$$\mathcal{A} \vdash \langle \tilde{T}, u^{\otimes 3} \rangle \geq 1 - \eta - \varepsilon = 1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p}).$$

As in Claim 1 of [70] we can apply Cauchy-Schwarz to obtain

$$\mathcal{A} \vdash \langle \tilde{T}, u^{\otimes 3} \rangle^2 \leq \sum_{i=1}^m \|a_i\|^2 \langle a_i, u \rangle^4 + \sum_{i \neq j} \langle a_i, a_j \rangle \langle a_i, u \rangle^2 \langle a_j, u \rangle^2.$$

Similarly to Lemma 3 of [70], the matrix concentration result of Proposition 5.4.10 implies that the last term above is $\tilde{O}(K/\sqrt{p})$ and thus we have

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^4 \geq 1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p}).$$

As in the proof of Lemma 2 of [70], apply Cauchy-Schwarz again to obtain

$$\mathcal{A} \vdash \left(\sum_{i=1}^m \langle a_i, u \rangle^4 \right)^2 \leq \sum_{i=1}^m \|a_i\|^2 \langle a_i, u \rangle^6 + \sum_{i \neq j} \langle a_i, a_j \rangle \langle a_i, u \rangle^3 \langle a_j, u \rangle^3.$$

Following [70] (proof of Lemma 2), the last term above is $\tilde{O}(K/\sqrt{p})$. This uses the bounds $\|A\| \leq \tilde{O}(\sqrt{K})$ (Lemma 5.4.5) and $\sum_{i=1}^m \langle a_i, u \rangle^4 \leq O(1) + \tilde{O}(K/\sqrt{p})$ (Corollary 5.4.13). We now have

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^6 \geq 1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p}).$$

Again following [70] (proof of Lemma 2) we have

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^6 \leq 1 + \tilde{O}(K/\sqrt{p}).$$

This only uses the incoherence of $\{a_i\}$ (Lemma 5.4.4). This implies

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^4 \leq 1 + \tilde{O}(K/\sqrt{p}) \quad (5.6)$$

which is analogous to Lemma 2 of [70].

Now following Lemma 7.2 of [100], apply Cauchy-Schwarz again:

$$\mathcal{A} \vdash \left(\sum_{i=1}^m \langle a_i, u \rangle^6 \right)^2 \leq \sum_{i=1}^m \|a_i\|^2 \langle a_i, u \rangle^{10} + \left(\sum_{i=1}^m \langle a_i, u \rangle^4 \right)^2 \max_{i \neq j} |\langle a_i, a_j \rangle|.$$

By (5.6) and Lemma 5.4.4, the last term above is $\tilde{O}(1/\sqrt{p})$ and so

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^{10} \geq 1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p}).$$

Since $\langle a_i, u \rangle^2 \leq 1$ we arrive at

$$\mathcal{A} \vdash \sum_{i=1}^m \langle a_i, u \rangle^8 \geq 1 - O(\varepsilon) - \tilde{O}(K/\sqrt{p})$$

which is analogous to Lemma 7.2 of [100]. Now the desired result (5.5) follows similarly to the proof of Proposition 7.1 in [100]. All the steps we have used are captured by the sum-of-squares proof system.

Appendix A

Additional proofs for Chapter 2

A.1 Log-likelihood expansion for the Gaussian observation model

In this section we show how the Gaussian observation model fits into the graphical model formulation by deriving the corresponding coefficient matrices Y_ρ . In particular, we show that $Y_\rho = d_\rho \lambda_\rho M_\rho$, a scalar multiple of the observed Gaussian matrix.

We can write $\log \mathcal{L}_{uv}(g_u, g_v) = \sum_\rho \log \mathcal{L}_{uv}^\rho(g_u, g_v)$ and consider each representation separately. There are three cases for the three types of representations (see Section 2.3.1).

For convenience we recall the Gaussian observation model:

$$M_\rho = \frac{\lambda_\rho}{n} X_\rho X_\rho^* + \frac{1}{\sqrt{nd_\rho}} W_\rho.$$

Restricting to the u, v submatrix:

$$M_{uv}^\rho = \frac{\lambda_\rho}{n} \rho(g_u g_v^{-1}) + \frac{1}{\sqrt{nd_\rho}} W_{uv}^\rho.$$

Real type. Let ρ be of real type. Recall that in this case, each entry of W_{uv}^ρ is $\mathcal{N}(0, 1)$.

We have

$$\begin{aligned}\log \mathcal{L}_{uv}^\rho(g_u, g_v) &= \frac{-nd_\rho}{2} \left\| M_{uv}^\rho - \frac{\lambda_\rho}{n} \rho(g_u g_v^{-1}) \right\|_F^2 \\ &= \langle d_\rho \lambda_\rho M_{uv}^\rho, \rho(g_u g_v^{-1}) \rangle + \text{const.}\end{aligned}$$

Here $\|\cdot\|_F$ denotes the Frobenius norm. The additive constant in the last step depends on M_{uv}^ρ but not on g_u, g_v . Thus the log-likelihood coefficients are $Y_{uv}^\rho = d_\rho \lambda_\rho M_{uv}^\rho$ and so $Y_\rho = d_\rho \lambda_\rho M_\rho$.

Complex type. Now consider a representation ρ of complex type, along with its conjugate $\bar{\rho}$. Recall that in this case, each entry of W_{uv}^ρ has independent real and imaginary parts drawn from $\mathcal{N}(0, 1/2)$. We have

$$\begin{aligned}\log \mathcal{L}_{uv}^\rho(g_u, g_v) &= -nd_\rho \left\| M_{uv}^\rho - \frac{\lambda_\rho}{n} \rho(g_u g_v^{-1}) \right\|_F^2 \\ &= \langle d_\rho \lambda_\rho M_{uv}^\rho, \rho(g_u g_v^{-1}) \rangle + \langle d_\rho \lambda_\rho \overline{M_{uv}^\rho}, \overline{\rho(g_u g_v^{-1})} \rangle + \text{const.}\end{aligned}$$

Therefore we have $Y_\rho = d_\rho \lambda_\rho M_\rho$ and $Y_{\bar{\rho}} = d_\rho \lambda_\rho \overline{M_\rho} = d_{\bar{\rho}} \lambda_{\bar{\rho}} M_{\bar{\rho}}$.

Quaternionic type. Now consider a representation ρ of quaternionic type. Recall that in this case, W_{uv}^ρ is block-quaternion where each 2×2 block encodes a quaternion value whose 4 entries are drawn independently from $\mathcal{N}(0, 1/2)$. Note the following relation between the norm of a quaternion and its corresponding 2×2 matrix:

$$\|a + bi + cj + dk\|^2 \equiv a^2 + b^2 + c^2 + d^2 = \frac{1}{2} \left\| \begin{pmatrix} a + bi & c + di \\ -c + di & a - bi \end{pmatrix} \right\|_F^2.$$

We have

$$\begin{aligned}
\log \mathcal{L}_{uv}^\rho(g_u, g_v) &= -nd_\rho \cdot \frac{1}{2} \left\| M_{uv}^\rho - \frac{\lambda_\rho}{n} \rho(g_u g_v^{-1}) \right\|_F^2 \\
&= d_\rho \lambda_\rho \Re \left(\langle M_{uv}^\rho, \rho(g_u g_v^{-1}) \rangle \right) + \text{const} \\
&= d_\rho \lambda_\rho \langle M_{uv}^\rho, \rho(g_u g_v^{-1}) \rangle + \text{const}
\end{aligned}$$

where \Re denotes real part. In the last step we used the fact that M_{uv}^ρ and $\rho(g_u g_v^{-1})$ are block-quaternion and so their inner product is real (see Section 2.3.1). Therefore $Y_\rho = d_\rho \lambda_\rho M_\rho$.

A.2 Proof of Lemma 2.6.1

To see that (i) and (ii) are equal, recall the interpretation of \mathcal{F}_ρ as a conditional expectation: $\mathcal{F}_\rho(\cdots) = \mathbb{E}[\rho(g)|\cdots]$ where \cdots stands for $\{\gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q\}_q$. (This is related to the *Nishimori identities* in statistical physics.)

We have the following symmetry properties of \mathcal{F}_ρ .

Lemma A.2.1. (1) For any $\gamma_q^t \in \mathbb{R}$, $z_q \in \mathbb{C}^{d_\rho \times d_\rho}$, and $g, h \in G$, we have

$$\mathcal{F}_\rho \left(\left\{ \gamma_q^t q(hg) + \sqrt{\gamma_q^t} z_q \right\}_q \right) = \rho(h) \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} q(h^{-1}) z_q \right\}_q \right)$$

and

$$\mathcal{F}_\rho \left(\left\{ \gamma_q^t q(gh) + \sqrt{\gamma_q^t} z_q \right\}_q \right) = \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q q(h^{-1}) \right\}_q \right) \rho(h).$$

(2) Therefore, if we define

$$f_\rho(g) \equiv \mathbb{E}_{z_q} \mathcal{F}_\rho \left(\left\{ \gamma_q^t q(g) + \sqrt{\gamma_q^t} z_q \right\}_q \right)$$

we have $f_\rho(hg) = \rho(h) f_\rho(g)$ and $f_\rho(gh) = f_\rho(g) \rho(h)$.

Proof. Part (1) is a straightforward computation using the definition of \mathcal{F}_ρ . Part (2) follows from part (1) because z_q has the same distribution as $q(h^{-1}) z_q$ and as $z_q q(h^{-1})$. \square

We now return to the proof of Lemma 2.6.1. The equality of (i) and (iii) follows from part (2) of Lemma A.2.1. The equality of (ii) and (iv) follows from part (1) of Lemma A.2.1. Combining this with the equality of (i) and (ii) from above, we have now shown equality of (i),(ii),(iii),(iv). It remains to show that A_ρ^t is a real multiple of the identity.

Letting $e \in G$ be the identity, we have

$$f_\rho(e)\rho(g) = f_\rho(eg) = f_\rho(ge) = \rho(g)f_\rho(e)$$

and so by Schur's lemma, this means $f_\rho(e)$ is a (possibly complex) multiple of the identity. But $f_\rho(e)$ is just (iii), so we are done. To see that the multiple a_ρ^t is real, note that the trace of (ii) is real.

Appendix B

Additional proofs for Chapter 3

B.1 Proof of Proposition 3.3.5

In this section we prove Proposition 3.3.5, which we restate here for convenience.

Proposition B.1.1. *For $L \geq 2$,*

$$\sup_{\alpha} \frac{L}{2} \frac{(\sum_{h \in G} \alpha_h^2 - \frac{1}{L})}{D(\alpha, \bar{\alpha})} = \frac{LC}{2}$$

where

$$C = \frac{L-2}{(L-1)\log(L-1)}.$$

Here α ranges over (vectorized) nonnegative $L \times L$ matrices with row- and column-sums equal to $\frac{1}{L}$. When $L = 2$, we define $C = 1$ (the limit value).

Recall G is a finite group of order L , $\bar{\alpha} = \frac{1}{L^2} \mathbb{1}_{L^2}$ and $\alpha_h = \sum_{(a,b) \in S_h} \alpha_{ab}$ where $S_h = \{(a,b) \mid a^{-1}b = h\}$. D denotes the KL divergence, which in this case is

$$D(\alpha, \bar{\alpha}) = \sum_{ab} \alpha_{ab} \log(L^2 \alpha_{ab}) = 2 \log L + \sum_{ab} \alpha_{ab} \log(\alpha_{ab}).$$

Although α belongs to a compact domain, we write sup rather than max in the optimization above. This is because when $\alpha = \bar{\alpha}$, the numerator and denominator are both zero,

so we are really optimizing over $\alpha \neq \bar{\alpha}$.

A high-level sketch of the proof is as follows. First we observe that the optimal α value should be constant on each S_h , allowing us to reduce the problem to only the variables α_h . By local optimality, we show further that the optimal α should take a particular form where $\alpha_h = x$ for k out of the L group elements h , and $\alpha_h = y$ for the remaining ones (where $y = \frac{1-kx}{L-k}$ so that $\sum_h \alpha_h = 1$ as required). This allows us to reduce the problem to only the variables k and x . We then show that for a fixed k , the optimum value is $\frac{LC_k}{2}$ where

$$C_k = \frac{L - 2k}{k(L - k) \log\left(\frac{L-k}{k}\right)}$$

(defined to equal its limit value $\frac{2}{L}$ when $k = L/2$). Finally, we show that C_k is largest when $k = 1$, in which case we have $C_1 = C$ and the proof is complete.

Now we begin the proof in full detail. Note that the numerator of the optimization problem depends only on the sums α_h and not the individual entries α_{ab} . Furthermore, once we have fixed the α_h 's, the denominator is minimized by setting all the α_{ab} values equal within each S_h . (Think of the fact that the uniform distribution maximizes entropy.) Therefore we only need to consider matrices α that are constant on each S_h . Note that any such matrix has row- and column-sums equal to $1/L$ (since each row or column contains exactly one entry in each S_h), so we can drop this constraint. (Interestingly, the fact that this constraint doesn't help means that we do not actually benefit from conditioning away from 'bad' events in this case.) The denominator becomes

$$D(\alpha, \bar{\alpha}) = 2 \log L + \sum_h L \cdot \frac{\alpha_h}{L} \log\left(\frac{\alpha_h}{L}\right) = \log L + \sum_{h \in G} \alpha_h \log(\alpha_h)$$

and so we have a new equivalent optimization problem:

$$\sup_{\alpha \neq \bar{\alpha}} M(\alpha)$$

where

$$M(\alpha) = \frac{L}{2} \cdot \frac{\sum_{h \in G} \alpha_h^2 - \frac{1}{L}}{\log L + \sum_h \alpha_h \log(\alpha_h)}.$$

Now α is simply a vector of α_h values, with the constraints $\alpha_h \geq 0$ and $\sum_h \alpha_h = 1$. Accordingly, $\bar{\alpha}_h = \frac{1}{L}$ for all h .

We will show that the optimum value is $\frac{LC}{2}$. We first focus on showing one direction: $\sup_{\alpha \neq \bar{\alpha}} M(\alpha) \leq \frac{LC}{2}$. By multiplying through by the denominator of $M(\alpha)$ (which is positive for all $\alpha \neq \bar{\alpha}$ since it is the divergence), this is equivalent to

$$\max_{\alpha} T(\alpha) \leq 0$$

where

$$T(\alpha) = \sum_h \alpha_h^2 - \frac{1}{L} - C \left[\log L + \sum_h \alpha_h \log(\alpha_h) \right].$$

Note that $\alpha \neq \bar{\alpha}$ is no longer required (since $T(\bar{\alpha}) = 0$) and so we now have a maximization problem over a compact domain. We will restrict to values of α that are locally optimal for $T(\alpha)$. Compute partial derivatives:

$$\frac{\partial T}{\partial \alpha_h} = 2\alpha_h - C [\log(\alpha_h) + 1]$$

$$\frac{\partial^2 T}{\partial \alpha_h^2} = 2 - \frac{C}{\alpha_h}.$$

Note that $\frac{\partial T}{\partial \alpha_h} \rightarrow \infty$ as $\alpha_h \rightarrow 0^+$ (and this is the only place in the interval $[0, 1]$ where the derivative blows up), and so a maximizer α for $T(\alpha)$ should have no coordinates set to zero. $\frac{\partial T}{\partial \alpha_h}$ is decreasing when $\alpha_h < \frac{C}{2}$, and increasing when $\alpha_h > \frac{C}{2}$. In particular, $\frac{\partial T}{\partial \alpha_h}(\alpha_h)$ is (at most) 2-to-1. If some coordinate of α is 1 then the rest would have to be 0, which we already ruled out. Therefore, all coordinates of a maximizer are strictly between 0 and 1, which means $\frac{\partial T}{\partial \alpha_h}$ must be equal for all coordinates. Since the derivative is 2-to-1, this means a maximizer can have at most two different α_h values.

We can therefore restrict to α for which k out of the L coordinates have the value x , and

the remaining $L - k$ coordinates have the value $y = \frac{1-kx}{L-k}$ (since the sum of coordinates must be 1). Therefore it is sufficient to show

$$\min_{1 \leq k \leq L/2} \min_{0 \leq x \leq 1/k} T_k(x) \geq 0$$

where

$$T_k(x) = C \left[\log L + kx \log x + (1 - kx) \log \left(\frac{1 - kx}{L - k} \right) \right] - \left[kx^2 + \frac{(1 - kx)^2}{L - k} - \frac{1}{L} \right].$$

Although it only makes sense for k to take integer values, we will show that the above is still true when k is allowed to be any real number in the interval $[0, L/2]$.

Define

$$t_k(x) = C_k \left[\log L + kx \log x + (1 - kx) \log \left(\frac{1 - kx}{L - k} \right) \right] - \left[kx^2 + \frac{(1 - kx)^2}{L - k} - \frac{1}{L} \right].$$

Note that this is the same as $T_k(x)$ but with C replaced by C_k (defined above). In the following two lemmas we will show $\min_{k,x} t_k(x) \geq 0$ and $C_k \leq C_1 = C$ for all k . It follows that $T_k(x) \geq t_k(x)$ (since the coefficient of C in $T_k(x)$ is the KL divergence, which is nonnegative). This completes the proof of the upper bound $\sup_{\alpha \neq \bar{\alpha}} M(\alpha) \leq \frac{LC}{2}$ because

$$\min_{k,x} T_k(x) \geq \min_{k,x} t_k(x) \geq 0.$$

Lemma B.1.2. *For any $k \in [1, L/2]$, we have*

$$\min_{x \in [0, 1/k]} t_k(x) \geq 0.$$

Proof. We relax k to be a real number in the interval $(0, L/2)$. The $k = L/2$ case will follow by continuity. Compute the fourth derivative:

$$\frac{d^4 t_k}{dx^4} = C_k \left[\frac{2k}{x^3} + \frac{2k^4}{(1 - kx)^3} \right] > 0.$$

Since the fourth derivative is strictly positive, the second derivative is convex. It follows that the first derivative $\frac{dt_k}{dx}$ has at most three zeros. One can check explicitly that these zeros are $\frac{1}{L} < \frac{1}{2k} < \frac{L-k}{kL}$. Using concavity of the second derivative, the middle zero $\frac{1}{2k}$ is a local maximum of $t_k(x)$ and the global minimum of $t_k(x)$ is achieved at either $\frac{1}{L}$ or $\frac{L-k}{kL}$. Both of these attain the value $t_k(x) = 0$, completing the proof. \square

Lemma B.1.3. *For all $k \in [1, L/2]$, $C_k \leq C_1 = C$.*

Proof. We will show that C_k is monotone decreasing in k on the interval $(0, L/2)$, by showing that its derivative is negative. It then follows that we should take the smallest allowable value for k , i.e. $k = 1$. Compute the derivative:

$$\frac{dC_k}{dk} = \frac{L(L-2k) - (k^2 + (L-k)^2) \log\left(\frac{L-k}{k}\right)}{k^2(k-L)^2 \log^2\left(\frac{L-k}{k}\right)}.$$

The denominator is positive, so it suffices to show that the numerator is negative. Applying the bound $\log(x) < 2\left(\frac{x-1}{x+1}\right)$, valid for all $x \geq 1$, we see that the numerator is at most $-\frac{(L-2k)^3}{L} < 0$. \square

This completes the proof of the upper bound $\sup_{\alpha \neq \bar{\alpha}} M(\alpha) \leq \frac{LC}{2}$. The matching lower bound is achieved by taking the α value corresponding to $k = 1$ and $x = \frac{L-1}{L}$. (For $L = 2$, this corresponds to the singularity $\bar{\alpha}$, but the optimum is achieved in the limit $x \rightarrow \frac{L-1}{L} = \frac{1}{2}$.)

Appendix C

Additional proofs for Chapter 4

C.1 Spherical harmonics and $\text{SO}(3)$ invariants

C.1.1 Spherical harmonics

We follow the conventions of [31]. Parametrize the unit sphere by angular spherical coordinates (θ, ϕ) with $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. (Here $\theta = 0$ is the north pole and $\theta = \pi$ is the south pole.) For integers $\ell \geq 0$ and $-\ell \leq m \leq \ell$, define the complex spherical harmonic

$$Y_{\ell m}(\theta, \phi) = (-1)^m N_{\ell m} P_{\ell}^m(\cos \theta) e^{im\phi}$$

with normalization factor

$$N_{\ell m} = \sqrt{\frac{(2\ell + 1)(\ell - m)!}{4\pi(\ell + m)!}}$$

where $P_{\ell}^m(x)$ are the associated Legendre polynomials

$$P_{\ell}^m(x) = \frac{1}{2^{\ell} \ell!} (1 - x^2)^{m/2} \frac{d^{\ell+m}}{dx^{\ell+m}} (x^2 - 1)^{\ell}.$$

In the S^2 registration problem we are interested in representing a real-valued function on the sphere, in which case we use an expansion (with real coefficients) in terms of the real

spherical harmonics:

$$S_{\ell m}(\theta, \phi) = \begin{cases} \frac{(-1)^m}{\sqrt{2}}(Y_{\ell m}(\theta, \phi) + \overline{Y_{\ell m}}(\theta, \phi)) = \sqrt{2}N_{\ell m}P_{\ell}^m(\cos \theta) \cos(m\phi) & m > 0, \\ Y_{\ell 0}(\theta, \phi) = N_{\ell 0}P_{\ell}^0(\cos \theta) & m = 0, \\ \frac{(-1)^m}{i\sqrt{2}}(Y_{\ell|m|}(\theta, \phi) - \overline{Y_{\ell|m|}}(\theta, \phi)) = \sqrt{2}N_{\ell|m|}P_{\ell}^{|m|}(\cos \theta) \sin(|m|\phi) & m < 0. \end{cases}$$

Here $\overline{Y_{\ell m}}$ is the complex conjugate of $Y_{\ell m}$, which satisfies the identity

$$\overline{Y_{\ell m}}(\theta, \phi) = (-1)^m Y_{\ell(-m)}(\theta, \phi). \quad (\text{C.1})$$

Above we have also used the identity $P_{\ell}^{-m} = (-1)^m \frac{(\ell-m)!}{(\ell+m)!} P_{\ell}^m$, which implies $N_{\ell(-m)} P_{\ell}^{-m} = (-1)^m N_{\ell m} P_{\ell}^m$.

In the cryo-EM problem we are instead interested in representing the Fourier transform of a real-valued function. Such a function f has the property that if \vec{r} and $-\vec{r}$ are antipodal points on the sphere, $f(-\vec{r}) = \overline{f(\vec{r})}$. For this type of function we use an expansion (with real coefficients) in terms of a new basis of spherical harmonics:

$$H_{\ell m}(\theta, \phi) = \begin{cases} \frac{1}{\sqrt{2}}(Y_{\ell m}(\theta, \phi) + (-1)^{\ell+m} Y_{\ell(-m)}(\theta, \phi)) & m > 0, \\ i^{\ell} Y_{\ell 0}(\theta, \phi) & m = 0, \\ \frac{i}{\sqrt{2}}(Y_{\ell|m|}(\theta, \phi) - (-1)^{\ell+m} Y_{\ell(-|m|)}(\theta, \phi)) & m < 0. \end{cases}$$

One can check that $H_{\ell m}(-\vec{r}) = \overline{H_{\ell m}(\vec{r})}$ using (C.1) along with the fact $Y_{\ell m}(-\vec{r}) = (-1)^{\ell} Y_{\ell m}(\vec{r})$ which comes from $P_{\ell}^m(-x) = (-1)^{\ell+m} P_{\ell}^m(x)$.

C.1.2 Wigner D-matrices

We will mostly work in the basis of complex spherical harmonics $Y_{\ell m}$ since the formulas are simpler. The analogous results for the other bases can be worked out by applying the appropriate change of basis.

Let $V_{\ell} \simeq \mathbb{C}^{2\ell+1}$ be the vector space consisting of degree- ℓ complex spherical harmonics represented in the basis $\{Y_{\ell m}\}_{-\ell \leq m \leq \ell}$, i.e. $v \in \mathbb{C}^{2\ell+1}$ encodes the spherical harmonic

$\sum_{m=-\ell}^{\ell} v_m Y_{\ell m}$. These V_{ℓ} (for $\ell = 0, 1, 2, \dots$) are the irreducible representations of $\text{SO}(3)$. Each can also be defined over the real numbers by changing basis to the real spherical harmonics $S_{\ell m}$.

A group element $g \in \text{SO}(3)$ acts on a (spherical harmonic) function $f : S^2 \rightarrow \mathbb{R}$ via $(g \cdot f)(x) = f(g^{-1}x)$. The action of g on V_{ℓ} is given by the *Wigner D-matrix* $D^{\ell}(g) \in \mathbb{C}^{(2\ell+1) \times (2\ell+1)}$ defined as in [31].

We will need the following orthogonality properties of the Wigner D-matrices. First, the standard Schur orthogonality relations from representation theory yield

$$\mathbb{E}_{g \sim \text{Haar}(\text{SO}(3))} \overline{D_{mk}^{\ell}(g)} D_{m'k'}^{\ell'}(g) = \frac{1}{2\ell+1} \mathbb{1}_{\ell=\ell'} \mathbb{1}_{m=m'} \mathbb{1}_{k=k'}.$$

We also have [129]

$$D_{mk}^{\ell}(g) D_{m'k'}^{\ell'}(g) = \sum_{L=|\ell-\ell'|}^{\ell+\ell'} \langle \ell m \ell' m' | L (m+m') \rangle \langle \ell k \ell' k' | L (k+k') \rangle D_{(m+m')(k+k')}^L(g)$$

where $\langle \ell_1 m_1 \ell_2 m_2 | \ell m \rangle$ is a *Clebsch-Gordan coefficient*. There is a closed-form expression for these coefficients [32]:

$$\begin{aligned} \langle \ell_1 m_1 \ell_2 m_2 | \ell m \rangle &= \mathbb{1}_{m=m_1+m_2} \sqrt{\frac{(2\ell+1)(\ell+\ell_1-\ell_2)!(\ell-\ell_1+\ell_2)!(\ell_1+\ell_2-\ell)!}{(\ell_1+\ell_2+\ell+1)!}} \times \\ &\quad \sqrt{(\ell+m)!(\ell-m)!(\ell_1-m_1)!(\ell_1+m_1)!(\ell_2-m_2)!(\ell_2+m_2)!} \times \\ &\quad \sum_k \frac{(-1)^k}{k!(\ell_1+\ell_2-\ell-k)!(\ell_1-m_1-k)!(\ell_2+m_2-k)!(\ell-\ell_2+m_1+k)!(\ell-\ell_1-m_2+k)!} \end{aligned}$$

where the sum is over all k for which the argument of every factorial is nonnegative.

C.1.3 Moment tensor

Let \mathcal{F} be a multi-set of frequencies from $\{1, 2, \dots\}$ and consider the action of $G = \text{SO}(3)$ on $V = \oplus_{\ell \in \mathcal{F}} V_{\ell}$. Recall that we want an explicit formula for $T_d(\mathbf{x}) = \mathbb{E}_g[(\Pi(g \cdot \mathbf{x}))^{\otimes d}]$ with

$g \sim \text{Haar}(G)$ (where Π can be the identity in the case of no projection). We have

$$\mathbb{E}_g[(\Pi(g \cdot \mathbf{x}))^{\otimes d}] = \Pi^{\otimes d} \mathbb{E}_g[\rho(g)^{\otimes d}] \mathbf{x}^{\otimes d}$$

(where $\mathbf{x}^{\otimes d}$ is a column vector of length $\dim(V)^d$) and so we need an explicit formula for the matrix $\mathbb{E}_g[\rho(g)^{\otimes d}]$. Here $\rho(g)$ is block diagonal with blocks $D^\ell(g)$ for $\ell \in \mathcal{F}$. There are no degree-1 invariants since we have excluded the trivial representation $\ell = 0$. For the degree-2 invariants $\mathbb{E}_g[\rho(g)^{\otimes 2}]$, consider a particular block $\mathbb{E}_g[D^{\ell_1}(g) \otimes D^{\ell_2}(g)]$ for some pair (ℓ_1, ℓ_2) . The entries in this block can be computed using the above orthogonality relations (and using $D_{00}^0(g) = 1$):

$$\begin{aligned} \mathbb{E}_g[(D^{\ell_1}(g))_{m_1 k_1} (D^{\ell_2}(g))_{m_2 k_2}] &= \mathbb{1}_{\ell_1=\ell_2} \mathbb{1}_{m_1=-m_2} \mathbb{1}_{k_1=-k_2} \langle \ell_1 m_1 \ell_2 m_2 | 0 0 \rangle \langle \ell_1 k_1 \ell_2 k_2 | 0 0 \rangle \\ &= \mathbb{1}_{\ell_1=\ell_2} \mathbb{1}_{m_1=-m_2} \mathbb{1}_{k_1=-k_2} \frac{(-1)^{m_1+k_1}}{2\ell_1+1} \end{aligned}$$

using the special case $\langle \ell_1 m_1 \ell_2 m_2 | 0 0 \rangle = \mathbb{1}_{\ell_1=\ell_2} \mathbb{1}_{m_1=-m_2} \frac{(-1)^{\ell_1+m_1}}{\sqrt{2\ell_1+1}}$.

Similarly, for degree 3 we have

$$\begin{aligned} \mathbb{E}_g[(D^{\ell_1}(g))_{m_1 k_1} (D^{\ell_2}(g))_{m_2 k_2} (D^{\ell_3}(g))_{m_3 k_3}] &= \\ \mathbb{1}_{|\ell_2-\ell_3| \leq \ell_1 \leq \ell_2+\ell_3} \mathbb{1}_{m_1+m_2+m_3=0} \mathbb{1}_{k_1+k_2+k_3=0} &\frac{(-1)^{m_1+k_1}}{2\ell_1+1} \langle \ell_2 m_2 \ell_3 m_3 | \ell_1 (m_2+m_3) \rangle \langle \ell_2 k_2 \ell_3 k_3 | \ell_1 (k_2+k_3) \rangle. \end{aligned}$$

C.1.4 Projection

Let $V = \oplus_{\ell \in \mathcal{F}} V_\ell$ with \mathcal{F} a subset of $\{1, 2, \dots\}$. Let $\Pi : V \rightarrow W$ be the projection that takes a complex spherical harmonic function and reveals only its values on the equator $\theta = \pi/2$. In cryo-EM this projection is applied separately to each shell (see Section 4.4.5). Letting $L = \max_{\ell \in \mathcal{F}} \ell$, the functions $b_{-L}, b_{-L+1}, \dots, b_L$ (from the circle S^1 to \mathbb{R}) form a basis for W , where $b_m(\phi) = e^{im\phi}$. The projection Π takes the form

$$\Pi(Y_{\ell m}) = (-1)^m N_{\ell m} P_\ell^m(0) b_m$$

extended by linearity. By taking a binomial expansion of $(x^2 - 1)^\ell$ it can be shown that

$$P_\ell^m(0) = \begin{cases} 0 & (\ell + m) \text{ odd}, \\ \frac{(-1)^{(\ell-m)/2}}{2^\ell \ell!} \binom{\ell}{(\ell+m)/2} (\ell + m)! & (\ell + m) \text{ even}. \end{cases} \quad (\text{C.2})$$

For cryo-EM, if we use the basis $H_{\ell m}$ so that the expansion coefficients are real, the output of the projection can be expressed (with real coefficients) in the basis

$$h_m(\phi) = \begin{cases} \frac{1}{\sqrt{2}}(e^{im\phi} + (-1)^m e^{-im\phi}) & m > 0, \\ 1 & m = 0, \\ \frac{i}{\sqrt{2}}(e^{i|m|\phi} - (-1)^m e^{-i|m|\phi}) & m < 0, \end{cases}$$

where the projection Π takes the form

$$\Pi(H_{\ell m}) = (-1)^m N_{\ell|m|} P_\ell^{|m|}(0) h_m$$

extended by linearity.

C.1.5 Explicit construction of invariants

Consider the cryo-EM setup (see Section 4.4.5) with S shells and F frequencies. We will cover S^2 registration as the special case $S = 1$ (without projection). Use the basis of complex spherical harmonics, with corresponding variables $x_{s\ell m}$ with $1 \leq s \leq S$, $1 \leq \ell \leq F$, and $-\ell \leq m \leq \ell$. One can change variables to $S_{\ell m}$ or $H_{\ell m}$ but for our purposes of testing the rank of the Jacobian it is sufficient to just work with $Y_{\ell m}$ (since the change of variables has no effect on the rank of the Jacobian).

Recall that in Section C.1.3 we computed expressions for the matrices $\mathbb{E}_g[D^{\ell_1}(g) \otimes D^{\ell_2}(g)]$ and $\mathbb{E}_g[D^{\ell_1}(g) \otimes D^{\ell_2}(g) \otimes D^{\ell_3}(g)]$, and in particular they are rank-1. Using this we can explicitly compute the entries of $T_d(\mathbf{x})$ and thus extract a basis for U_2^T and U_3^T . We present the results below.

Degree-2 invariants

Without projection, the degree-2 invariants are

$$\mathcal{I}_2(s_1, s_2, \ell) = \frac{1}{2\ell + 1} \sum_{|k| \leq \ell} (-1)^k x_{s_1 \ell k} x_{s_2 \ell(-k)}$$

for $s_1, s_2 \in \{1, \dots, S\}$ and $\ell \in \{1, \dots, F\}$. Swapping s_1 with s_2 results in the same invariant, so take $s_1 \leq s_2$ to remove redundancies.

With projection, the degree-2 invariants are

$$\mathcal{P}_2(s_1, s_2, m) = (-1)^m \sum_{\ell \geq |m|} N_{\ell m} N_{\ell(-m)} P_\ell^m(0) P_\ell^{-m}(0) \mathcal{I}_2(s_1, s_2, \ell)$$

with $s_1, s_2 \in \{1, \dots, S\}$ and $m \in \{-F, \dots, F\}$. Negating m or swapping s_1 with s_2 results in the same invariant (up to sign) so take $s_1 \leq s_2$ and $m \geq 0$ to remove redundancies. Recall the expression (C.2) for $P_\ell^m(0)$.

Degree-3 invariants

Let $\Delta(\ell_1, \ell_2, \ell_3)$ denote the predicate $|\ell_2 - \ell_3| \leq \ell_1 \leq \ell_2 + \ell_3$ (which captures whether ℓ_1, ℓ_2, ℓ_3 can be the side-lengths of a triangle). Without projection, the degree-3 invariants are

$$\mathcal{I}_3(s_1, \ell_1, s_2, \ell_2, s_3, \ell_3) = \frac{1}{2\ell_1 + 1} \sum_{\substack{k_1 + k_2 + k_3 = 0 \\ |k_i| \leq \ell_i}} (-1)^{k_1} \langle \ell_2 \ k_2 \ \ell_3 \ k_3 | \ell_1(-k_1) \rangle x_{s_1 \ell_1 k_1} x_{s_2 \ell_2 k_2} x_{s_3 \ell_3 k_3}$$

for $s_i \in \{1, \dots, S\}$ and $\ell_i \in \{1, \dots, F\}$ satisfying $\Delta(\ell_1, \ell_2, \ell_3)$. There are some redundancies here. First, permuting the three (s_i, ℓ_i) pairs (while keeping each pair in tact) results in the same invariant (up to scalar multiple). Also, some of the above invariants are actually zero; specifically, this occurs when $(s_1, \ell_1) = (s_2, \ell_2) = (s_3, \ell_3)$ with ℓ_1 odd, or when $(s_1, \ell_1) = (s_2, \ell_2) \neq (s_3, \ell_3)$ with ℓ_3 odd (or some permutation of this case).

With projection, the degree-3 invariants are

$$\mathcal{P}_3(s_1, m_1, s_2, m_2, s_3, m_3) = (-1)^{m_1} \sum_{\ell_1, \ell_2, \ell_3 : \Delta(\ell_1, \ell_2, \ell_3)} N_{\ell_1 m_1} N_{\ell_2 m_2} N_{\ell_3 m_3} P_{\ell_1}^{m_1}(0) P_{\ell_2}^{m_2}(0) P_{\ell_3}^{m_3}(0) \langle \ell_2 m_2 \ell_3 m_3 | \ell_1 (-m_1) \rangle \mathcal{I}_3(s_1, \ell_1, s_2, \ell_2, s_3, \ell_3)$$

for $s_i \in \{1, \dots, S\}$ and $m_i \in \{-F, \dots, F\}$ such that $m_1 + m_2 + m_3 = 0$. There are again redundancies under permutation: permuting the three (s_i, m_i) pairs results in the same invariant. Negating all three m 's also results in the same invariant. There are additional non-trivial linear relations (see Section C.1.6 below).

C.1.6 Counting the number of invariants

S^2 registration

For the case of S^2 registration ($S = 1$) the above degree-2 and degree-3 invariants without projection (with redundancies removed as discussed above) form a basis for $\mathbb{R}[\mathbf{x}]_2^G \oplus \mathbb{R}[\mathbf{x}]_3^G$ (although we have not made this rigorous). Thus, counting these invariants gives a combinatorial analogue of Proposition 4.4.7.

Cryo-EM

In this section we give a formula for $\text{trdeg}(U_{\leq 3}^T)$ for (heterogeneous) cryo-EM (with projection), valid for all $K \geq 1$, $S \geq 1$ and $F \geq 2$. The formula is conjectural but has been tested (via the Jacobian criterion) for various small values of K, S, F .

The number of algebraically independent degree-2 invariants turns out to be the number of distinct \mathcal{I}_2 invariants (i.e. without projection). The number of such invariants is $\frac{1}{2}S(S+1)F$.

For degree 3, things are more complicated because the projected invariants \mathcal{P}_3 have smaller dimension than the \mathcal{I}_3 . We start by counting the number of distinct (up to scalar multiple) \mathcal{P}_3 invariants. To this end, let $\mathcal{X}(S, F)$ be the set of equivalence classes of tuples

$(s_1, m_1, s_2, m_2, s_3, m_3)$ with $s_i \in \{1, \dots, S\}$ and $m_i \in \{-F, \dots, F\}$, modulo the relations

$$(s_1, m_1, s_2, m_2, s_3, m_3) \sim (s_2, m_2, s_1, m_1, s_3, m_3) \sim (s_1, m_1, s_3, m_3, s_2, m_2) \quad (\text{permutation})$$

$$(s_1, m_1, s_2, m_2, s_3, m_3) \sim (s_1, -m_1, s_2, -m_2, s_3, -m_3) \quad (\text{negation}).$$

There are some non-trivial linear relations among the distinct \mathcal{P}_3 invariants, which we must now account for. The number of such relations is

$$E(S) \triangleq 2S + 4S(S-1) + S(S-1)(S-2).$$

This can be broken down as follows. For each $k \in \{1, 2, 3\}$ there are $2k$ relations for each size-3 multi-subset $\{s_1, s_2, s_3\}$ of $\{1, \dots, S\}$ with exactly k distinct elements. (We do not currently have a thorough understanding of what exactly the linear relations *are*, but we have observed that the above pattern holds.)

We can now put it all together and state our conjecture. We will also use the formula (4.3) for $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$, extended to the heterogeneous case via Proposition 4.3.15.

Conjecture C.1.1. *Consider heterogeneous cryo-EM with $F \geq 2$ frequencies.*

- $\text{trdeg}(\mathbb{R}[\mathbf{x}]^G) = K[S(F^2 + 2F) - 3] + K - 1,$
- $\dim(U_2^T) = \frac{1}{2}S(S+1)F,$
- $\dim(U_3^T) = |\mathcal{X}(S, F)| - E(S),$
- *generic list recovery is possible at degree 3 if and only if $\dim(U_2^T) + \dim(U_3^T) \geq \text{trdeg}(\mathbb{R}[\mathbf{x}]^G).$*

When S and F are large, the dominant term in $\dim(U_2^T) + \dim(U_3^T)$ is $|\mathcal{X}(S, F)| \approx S^3 F^2 / 4$ and so generic list recovery is possible when (roughly) $K \leq S^2 / 4$.

Remark C.1.2. When S is large compared to F we have $\dim(U_2^T) > \text{trdeg}(\mathbb{R}[\mathbf{x}]^G)$ and so we might expect generic list recovery to be possible at degree 2. However, this appears to

not be the case because unexpected algebraic dependencies are encountered in this regime, i.e. $\text{trdeg}(U_2^T) < \text{trdeg}(\mathbb{R}[\mathbf{x}]^G) < \dim(U_2^T)$. We have not observed instances where such unexpected algebraic dependencies affect the feasibility of generic list recovery at degree 3.

Bibliography

- [1] Emmanuel Abbe, Afonso S. Bandeira, Annina Bracher, and Amit Singer. Decoding binary node labels from censored edge measurements: Phase transition and efficient recovery. *IEEE Transactions on Network Science and Engineering*, 1(1):10–22, 2014.
- [2] Emmanuel Abbe, Afonso S. Bandeira, and Georgina Hall. Exact recovery in the stochastic block model. *IEEE Transactions on Information Theory*, 62(1):471–487, 2016.
- [3] Emmanuel Abbe, Tamir Bendory, William Leeb, João Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *arXiv preprint arXiv:1710.02793*, 2017.
- [4] Emmanuel Abbe, Joao Pereira, and Amit Singer. Sample complexity of the boolean multireference alignment problem. *arXiv preprint arXiv:1701.07540*, 2017.
- [5] Marc Adrian, Jacques Dubochet, Jean Lepault, and Alasdair W McDowall. Cryo-electron microscopy of viruses. *Nature*, 308(5954):32–36, 1984.
- [6] Amit Agrawal, Ramesh Raskar, and Rama Chellappa. What is the range of surface reconstructions from a gradient field? In *European Conference on Computer Vision*, pages 578–591. Springer, 2006.
- [7] Cecilia Aguerrebere, Mauricio Delbracio, Alberto Bartesaghi, and Guillermo Sapiro. Fundamental limits in multi-image alignment. *IEEE Trans. Signal Process.*, 64(21):5707–5722, 2016.
- [8] Elie Aïdékon. Convergence in law of the minimum of a branching random walk. *The Annals of Probability*, 41(3A):1362–1426, 2013.
- [9] Ahmed El Alaoui, Florent Krzakala, and Michael I. Jordan. Finite size corrections and likelihood ratio fluctuations in the spiked wigner model. *arXiv preprint arXiv:1710.02903*, 2017.
- [10] Arash A. Amini and Martin J. Wainwright. High-dimensional analysis of semidefinite relaxations for sparse principal components. In *IEEE International Symposium on Information Theory*, pages 2454–2458. IEEE, 2008.

- [11] Ery Arias-Castro and Nicolas Verzelen. Community detection in dense random networks. *The Annals of Statistics*, 42(3):940–969, 2014.
- [12] Jinho Baik, Gérard Ben Arous, and Sandrine Péché. Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, pages 1643–1697, 2005.
- [13] Afonso Bandeira, Philippe Rigollet, and Jonathan Weed. Optimal rates of estimation for multi-reference alignment. *arXiv preprint arXiv:1702.08546*, 2017.
- [14] Afonso S. Bandeira. *Convex Relaxations for Certain Inverse Problems on Graphs*. PhD thesis, Princeton University, June 2015.
- [15] Afonso S. Bandeira, Ben Blum-Smith, Amelia Perry, Jonathan Weed, and Alexander S. Wein. Estimation under group actions: recovering orbits from invariants. *arXiv:1712.10163*, 2017.
- [16] Afonso S. Bandeira, Nicolas Boumal, and Amit Singer. Tightness of the maximum likelihood semidefinite relaxation for angular synchronization. *arXiv:1411.3272*, 2014.
- [17] Afonso S. Bandeira, Moses Charikar, Amit Singer, and Andy Zhu. Multireference alignment using semidefinite programming. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 459–470. ACM, 2014.
- [18] Afonso S. Bandeira, Yutong Chen, and Amit Singer. Non-unique games over compact groups and orientation estimation in cryo-EM. *arXiv:1505.03840*, May 2015.
- [19] Jess Banks, Cristopher Moore, Joe Neeman, and Praneeth Netrapalli. Information-theoretic thresholds for community detection in sparse networks. In *29th Annual Conference on Learning Theory*, pages 383–416, june 2016.
- [20] Jess Banks, Cristopher Moore, Nicolas Verzelen, Roman Vershynin, and Jiaming Xu. Information-theoretic bounds and phase transitions in clustering, sparse PCA, and submatrix localization. *arXiv:1607.05222*, 2017.
- [21] Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *Conference on Learning Theory*, pages 417–445, 2016.
- [22] Jean Barbier, Mohamad Dia, Nicolas Macris, Florent Krzakala, Thibault Lesieur, and Lenka Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Advances in Neural Information Processing Systems*, pages 424–432, 2016.
- [23] Mohsen Bayati and Andrea Montanari. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57(2):764–785, 2011.

- [24] Malte Beecken, Johannes Mittmann, and Nitin Saxena. Algebraic independence and blackbox identity testing. *Information and Computation*, 222:2–19, 2013.
- [25] Florent Benaych-Georges and Raj Rao Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [26] Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, and Amit Singer. Bispectrum inversion with application to multireference alignment. *arXiv preprint arXiv:1705.00641*, 2017.
- [27] Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *Conference on Learning Theory*, pages 1046–1066, 2013.
- [28] Quentin Berthet and Philippe Rigollet. Optimal detection of sparse principal components in high dimension. *The Annals of Statistics*, 41(4):1780–1815, 2013.
- [29] Tejal Bhamre, Teng Zhang, and Amit Singer. Orthogonal matrix retrieval in cryo-electron microscopy. In *Biomedical Imaging (ISBI), 2015 IEEE 12th International Symposium on*, pages 1048–1052. IEEE, 2015.
- [30] Christopher M. Bishop. Bayesian PCA. *Advances in neural information processing systems*, pages 382–388, 1999.
- [31] Miguel A Blanco, M Flórez, and M Bermejo. Evaluation of the rotation matrices in the basis of real spherical harmonics. *Journal of Molecular Structure: Theochem*, 419(1):19–27, 1997.
- [32] Arno Böhm. *Quantum mechanics: foundations and applications*. Springer Science & Business Media, 2013.
- [33] Erwin Bolthausen, Jean-Dominique Deuschel, and Giambattista Giacomin. Entropic repulsion and the maximum of the two-dimensional harmonic. *The Annals of Probability*, 29(4):1670–1692, 2001.
- [34] Nicolas Boumal. Nonconvex phase synchronization. *arXiv:1601.06114*, 2016.
- [35] Nicolas Boumal, Tamir Bendory, Roy R. Lederman, and Amit Singer. Heterogeneous multireference alignment: a single pass approach. *arXiv preprint arXiv:1710.02590*, 2017.
- [36] Maury Bramson, Jian Ding, and Ofer Zeitouni. Convergence in law of the maximum of nonlattice branching random walk. *arXiv preprint arXiv:1404.3423*, 2014.
- [37] Maury Bramson, Jian Ding, and Ofer Zeitouni. Convergence in law of the maximum of the two-dimensional discrete gaussian free field. *Communications on Pure and Applied Mathematics*, 69(1):62–123, 2016.

- [38] Maury Bramson and Ofer Zeitouni. Tightness of the recentered maximum of the two-dimensional discrete gaussian free field. *Communications on Pure and Applied Mathematics*, 65(1):1–20, 2012.
- [39] Maury D. Bramson. Maximal displacement of branching brownian motion. *Communications on Pure and Applied Mathematics*, 31(5):531–581, 1978.
- [40] Theodor Bröcker and Tammo tom Dieck. *Representations of compact Lie groups*, volume 98. Springer Science & Business Media, 2013.
- [41] Yuxin Chen and Emmanuel Candès. The projected power method: An efficient algorithm for joint alignment from pairwise differences. *arXiv:1609.05820*, 2016.
- [42] Jon Cohen. Is high-tech view of HIV too good to be true? *Science*, 341(6145):443–444, 2013.
- [43] Ronald R. Coifman, Yoel Shkolnisky, Fred J. Sigworth, and Amit Singer. Reference free structure determination through eigenvectors of center of mass operators. *Applied and Computational Harmonic Analysis*, 28(3):296–312, 2010.
- [44] Amin Coja-Oghlan, Florent Krzakala, Will Perkins, and Lenka Zdeborova. Information-theoretic thresholds from the cavity method. *arXiv preprint arXiv:1611.00814*, 2016.
- [45] Amin Coja-Oghlan and Konstantinos Panagiotou. Going after the k-SAT threshold. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 705–714. ACM, 2013.
- [46] Amin Coja-Oghlan and Konstantinos Panagiotou. The asymptotic k-SAT threshold. *Advances in Mathematics*, 288:985–1068, 2016.
- [47] Amin Coja-Oghlan and Dan Vilenchik. Chasing the k-colorability threshold. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 380–389. IEEE, 2013.
- [48] Amin Coja-Oghlan and Lenka Zdeborová. The condensation transition in random hypergraph 2-coloring. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 241–250. SIAM, 2012.
- [49] Amin Coja-Oghlan and Konstantinos Panagiotou. Catching the k-NAESAT threshold. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 899–908. ACM, 2012.
- [50] David Cox, John Little, and Don O’Shea. *Ideals, varieties, and algorithms: an introduction to computational algebraic geometry and commutative algebra*. Springer New York, 2007.

- [51] Mihai Cucuringu, Yaron Lipman, and Amit Singer. Sensor network localization by eigenvector synchronization over the euclidean group. *ACM Transactions on Sensor Networks (TOSN)*, 8(3):19, 2012.
- [52] Victor H de la Peña and Stephen J Montgomery-Smith. Decoupling inequalities for the tail probabilities of multivariate U-statistics. *The Annals of Probability*, pages 806–816, 1995.
- [53] Aurelien Decelle, Florent Krzakala, Cristopher Moore, and Lenka Zdeborová. Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Physical Review E*, 84(6):066106, 2011.
- [54] Amir Dembo, Yuval Peres, Jay Rosen, and Ofer Zeitouni. Cover times for brownian motion and random walks in two dimensions. *Annals of mathematics*, pages 433–464, 2004.
- [55] Harm Derksen and Gregor Kemper. *Computational invariant theory*. Springer, 2015.
- [56] Yash Deshpande, Emmanuel Abbe, and Andrea Montanari. Asymptotic mutual information for the binary stochastic block model. In *IEEE International Symposium on Information Theory (ISIT)*, pages 185–189. IEEE, 2016.
- [57] Yash Deshpande and Andrea Montanari. Information-theoretically optimal sparse PCA. In *IEEE International Symposium on Information Theory*, pages 2197–2201. IEEE, 2014.
- [58] Yash Deshpande and Andrea Montanari. Finding hidden cliques of size $\sqrt{N/e}$ in nearly linear time. *Foundations of Computational Mathematics*, 15(4):1069–1128, 2015.
- [59] Yash Deshpande, Andrea Montanari, and Emile Richard. Cone-constrained principal component analysis. In *Advances in Neural Information Processing Systems*, pages 2717–2725, 2014.
- [60] R. Diamond. On the multiple simultaneous superposition of molecular structures by rigid body transformations. *Protein Science*, 1(10):1279–1287, October 1992.
- [61] Igor Dolgachev. *Lectures on invariant theory*, volume 296. Cambridge University Press, 2003.
- [62] M. Domokos. Degree bounds for separating invariants of abelian groups. *arXiv preprint arXiv:1602.06597*, 2016.
- [63] David L. Donoho, Arian Maleki, and Andrea Montanari. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.

- [64] David L. Donoho, Arian Maleki, and Andrea Montanari. Message passing algorithms for compressed sensing: I. motivation and construction. *IEEE Information Theory Workshop (ITW)*, pages 115–144, 2010.
- [65] Daniel Egloff, Markus Leippold, and Liuren Wu. The term structure of variance swap rates and optimal variance swap investments. *Journal of Financial and Quantitative Analysis*, 45(5):1279, 2010.
- [66] Richard Ehrenborg and Gian-Carlo Rota. Apolarity and canonical forms for homogeneous polynomials. *European Journal of Combinatorics*, 14(3):157–181, 1993.
- [67] Peter Elias. List decoding for noisy channels. Technical Report 335, Research Laboratory of Electronics, Massachusetts Institute of Technology, 1957.
- [68] Gábor Etesi. Spontaneous symmetry breaking in the $SO(3)$ gauge theory to discrete subgroups. *Journal of Mathematical Physics*, 37(4):1596–1601, 1996.
- [69] Delphine Féral and Sandrine Péché. The largest eigenvalue of rank one deformation of large Wigner matrices. *Communications in Mathematical Physics*, 272(1):185–228, 2007.
- [70] Rong Ge and Tengyu Ma. Decomposing overcomplete 3rd order tensors using sum-of-squares algorithms. *arXiv preprint arXiv:1504.05287*, 2015.
- [71] Arvind Giridhar and P.R. Kumar. Distributed clock synchronization over wireless networks: Algorithms and analysis. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 4915–4920. IEEE, 2006.
- [72] Amir Globerson, Tim Roughgarden, David Sontag, and Cafer Yildirim. How hard is inference for structured prediction? In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2181–2190, 2015.
- [73] Navin Goyal, Santosh Vempala, and Ying Xiao. Fourier pca and robust tensor decomposition. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 584–593. ACM, 2014.
- [74] Olivier Guédon and Roman Vershynin. Community detection in sparse networks via Grothendieck’s inequality. *Probability Theory and Related Fields*, pages 1–25, 2015.
- [75] Bruce Hajek, Yihong Wu, and Jiaming Xu. Achieving exact cluster recovery threshold via semidefinite programming. *IEEE Transactions on Information Theory*, 62(5):2788–2797, 2016.
- [76] David Hilbert. Über die Theorie der algebraischen Formen. *Mathematische Annalen*, 36:473–531, 1890.

- [77] David Hilbert. Über die vollen Invariantensysteme. *Mathematische Annalen*, 42:313–370, 1893.
- [78] Samuel B Hopkins, Tselil Schramm, Jonathan Shi, and David Steurer. Fast spectral algorithms from sum-of-squares proofs: tensor decomposition and planted sparse vectors. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 178–191. ACM, 2016.
- [79] Dũng T. Huynh. A superexponential lower bound for Gröbner bases and Church-Rosser commutative Thue systems. *Inform. and Control*, 68(1-3):196–206, 1986.
- [80] Svante Janson. Random regular graphs: asymptotic distributions and contiguity. *Combinatorics, Probability and Computing*, 4(04):369–405, 1995.
- [81] Adel Javanmard and Andrea Montanari. State evolution for general approximate message passing algorithms, with applications to spatial coupling. *Information and Inference*, 2(2):115–144, 2013.
- [82] Adel Javanmard, Andrea Montanari, and Federico Ricci-Tersenghi. Phase transitions in semidefinite relaxations. *Proceedings of the National Academy of Sciences*, 113(16):E2218–E2223, 2016.
- [83] Slavica Jonić. Cryo-electron microscopy analysis of structurally heterogeneous macromolecular complexes. *Computational and structural biotechnology journal*, 14:385–390, 2016.
- [84] Victor Kač. *Invariant theory [lecture notes]*. <https://people.kth.se/~laksov/notes/invariant.pdf>, 1994.
- [85] Ramakrishna Kakarala. Completeness of bispectrum on compact groups. *arXiv preprint arXiv:0902.0196*, 1, 2009.
- [86] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [87] Neeraj Kayal. The complexity of the annihilating polynomial. In *24th Annual IEEE Conference on Computational Complexity (CCC’09)*, pages 184–193. IEEE, 2009.
- [88] Martin Kohls and Hanspeter Kraft. Degree bounds for separating invariants. *arXiv preprint arXiv:1001.5216*, 2010.
- [89] Florent Krzakala, Jiaming Xu, and Lenka Zdeborová. Mutual information in rank-one matrix estimation. *arXiv:1603.08447*, 2016.
- [90] Can M. Le, Elizaveta Levina, and Roman Vershynin. Sparse random graphs: regularization and concentration of the laplacian. *arXiv:1502.03049*, 2015.

- [91] Can M. Le and Roman Vershynin. Concentration and regularization of random graphs. *arXiv:1506.00669*, 2015.
- [92] Lucien Le Cam. *Locally Asymptotically Normal Families of Distributions. Certain Approximations to Families of Distributions and Their Use in the Theory of Estimation and Testing Hypotheses*. Berkeley & Los Angeles, 1960.
- [93] Roy R. Lederman and Amit Singer. A representation theory perspective on simultaneous alignment and classification. *arXiv preprint arXiv:1607.03464*, 2016.
- [94] Roy R Lederman and Amit Singer. Continuously heterogeneous hyper-objects in cryo-EM and 3-D movies of many temporal dimensions. *arXiv preprint arXiv:1704.02899*, 2017.
- [95] Marc Lelarge and Léo Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv:1611.03888*, 2016.
- [96] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborov. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 680–687. IEEE, 2015.
- [97] Thibault Lesieur, Florent Krzakala, and Lenka Zdeborová. Phase transitions in sparse PCA. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1635–1639. IEEE, 2015.
- [98] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. *arXiv preprint arXiv:1710.08076*, 2017.
- [99] Robert B. Litterman and Jose Scheinkman. Common factors affecting bond returns. *The Journal of Fixed Income*, 1(1):54–61, 1991.
- [100] Tengyu Ma, Jonathan Shi, and David Steurer. Polynomial-time tensor decompositions with sum-of-squares. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 438–446. IEEE, 2016.
- [101] Zongming Ma and Yihong Wu. Computational barriers in minimax submatrix detection. *The Annals of Statistics*, 43(3):1089–1116, 2015.
- [102] Arian Maleki, Laura Anitori, Zai Yang, and Richard G. Baraniuk. Asymptotic analysis of complex LASSO via complex approximate message passing (CAMP). *IEEE Transactions on Information Theory*, 59(7):4290–4308, 2013.
- [103] Laurent Massoulié. Community detection thresholds and the weak Ramanujan property. In *Proceedings of the 46th Annual Symposium on Theory of Computing*, pages 694–703. ACM, 2014.

- [104] Frank McSherry. Spectral partitioning of random graphs. In *Proceedings of the 42nd Symposium on Foundations of Computer Science*, pages 529–537. IEEE, 2001.
- [105] Marc Mezard and Andrea Montanari. *Information, physics, and computation*. Oxford University Press, 2009.
- [106] Marc Mézard, Giorgio Parisi, and M.A. Virasoro. SK model: The replica solution without replicas. *Europhys. Lett.*, 1(2):77–82, 1986.
- [107] Ankur Moitra. *Algorithmic aspects of machine learning*. Lecture notes (MIT), 2014.
- [108] Michael S. O. Molloy, Hanna Robalewska, Robert W. Robinson, and Nicholas C. Wormald. 1-factorizations of random regular graphs. *Random Structures and Algorithms*, 10(3):305–321, 1997.
- [109] Andrea Montanari, Daniel Reichman, and Ofer Zeitouni. On the limitation of spectral methods: From the gaussian hidden clique problem to rank-one perturbations of gaussian tensors. In *Advances in Neural Information Processing Systems*, pages 217–225, 2015.
- [110] Andrea Montanari and Emile Richard. Non-negative principal component analysis: Message passing algorithms and sharp asymptotics. *IEEE Transactions on Information Theory*, 62(3):1458–1484, 2016.
- [111] Andrea Montanari and Subhabrata Sen. Semidefinite programs on sparse random graphs and their application to community detection. In *Proceedings of the 48th Annual SIGACT Symposium on Theory of Computing*, pages 814–827. ACM, 2016.
- [112] Elchanan Mossel, Joe Neeman, and Allan Sly. A proof of the block model threshold conjecture. *arXiv:1311.4115*, 2013.
- [113] Elchanan Mossel, Joe Neeman, and Allan Sly. Belief propagation, robust reconstruction and optimal recovery of block models. In *Conference on Learning Theory*, volume 35, pages 356–370, 2014.
- [114] Elchanan Mossel, Joe Neeman, and Allan Sly. Reconstruction and estimation in the planted partition model. *Probability Theory and Related Fields*, 162(3-4):431–461, 2015.
- [115] Eva Nogales. The development of cryo-EM into a mainstream structural biology technique. *Nature methods*, 13(1):24–27, 2016.
- [116] Brad Osgood. *Chapter 8: n-dimensional Fourier Transform*. Lecture Notes for EE 261: The Fourier Transform and its Applications. <https://see.stanford.edu/materials/lsoftae261/book-fall-07.pdf>, 2007.

- [117] Judea Pearl. Fusion, propagation, and structuring in belief networks. *Artificial intelligence*, 29(3):241–288, 1986.
- [118] Amelia Perry, Jonathan Weed, Afonso Bandeira, Philippe Rigollet, and Amit Singer. The sample complexity of multi-reference alignment. *arXiv preprint arXiv:1707.00943*, 2017.
- [119] Amelia Perry, Alexander S. Wein, and Afonso S. Bandeira. Statistical limits of spiked tensor models. *arXiv:1612.07728*, 2016.
- [120] Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA for spiked random matrices and synchronization. *arXiv:1609.05573*, 2016.
- [121] Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Message-passing algorithms for synchronization problems over compact groups. *Communications on Pure and Applied Mathematics*, to appear. *arXiv:1610.04583*, 2016.
- [122] Amelia Perry, Alexander S. Wein, Afonso S. Bandeira, and Ankur Moitra. Optimality and sub-optimality of PCA I: Spiked random matrix models. *The Annals of Statistics*, to appear.
- [123] R. G. Pita, M. R. Zurera, P. J. Amores, and F. L. Ferreras. Using multilayer perceptrons to align high range resolution radar signals. In *Artificial Neural Networks: Formal Models and Their Applications - ICANN 2005*, pages 911–916. Springer Berlin Heidelberg, 2005.
- [124] Aaron Potechin and David Steurer. Exact tensor completion with sum-of-squares. *arXiv preprint arXiv:1702.06237*, 2017.
- [125] Sundeep Rangan and Alyson K. Fletcher. Iterative estimation of constrained rank-one matrices in noise. In *IEEE International Symposium on Information Theory (ISIT)*, pages 1246–1250. IEEE, 2012.
- [126] Thomas J. Richardson and Rüdiger L. Urbanke. The capacity of low-density parity-check codes under message-passing decoding. *IEEE transactions on information theory*, 47(2):599–618, 2001.
- [127] Philippe Rigollet and Jan-Christian Hütter. High-dimensional statistics. *Lecture notes*, 2018.
- [128] Robert W. Robinson and Nicholas C. Wormald. Almost all regular graphs are Hamiltonian. *Random Structures & Algorithms*, 5(2):363–374, 1994.
- [129] Morris E. Rose. Elementary theory of angular momentum. *Physics Today*, 10:30, 1957.

- [130] David M. Rosen, Luca Carlone, Afonso S. Bandeira, and John J. Leonard. A certifiably correct algorithm for synchronization over the special euclidean group. *arXiv:1611.00128*, 2016.
- [131] Cynthia Rush and Ramji Venkataramanan. Finite sample analysis of approximate message passing. *arXiv:1606.01800*, 2016.
- [132] Alaa Saade, Florent Krzakala, and Lenka Zdeborová. Spectral clustering of graphs with the Bethe Hessian. In *Advances in Neural Information Processing Systems*, pages 406–414, 2014.
- [133] B. M. Sadler. Shift and rotation invariant object reconstruction using the bispectrum. In *Workshop on Higher-Order Spectral Analysis*, pages 106–111, Jun 1989.
- [134] Brian M Sadler and Georgios B Giannakis. Shift-and rotation-invariant object reconstruction using the bispectrum. *JOSA A*, 9(1):57–69, 1992.
- [135] Eugene Salamin. Application of quaternions to computation with rotations. Technical report, Working Paper, 1979.
- [136] S. H. W. Scheres. RELION: implementation of a bayesian approach to cryo-EM structure determination. *J. Struct. Biol.*, 180(3):519–530, 2012.
- [137] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Springer Science & Business Media, 2003.
- [138] Warren Schudy and Maxim Sviridenko. Concentration and moment inequalities for polynomials of independent random variables. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 437–446. Society for Industrial and Applied Mathematics, 2012.
- [139] Igor R. Shafarevich. *Basic algebraic geometry*. Springer, 1994.
- [140] Fred J. Sigworth. Principles of cryo-EM single-particle image processing. *Microscopy*, 65(1):57–67, 2016.
- [141] Amit Singer. Angular synchronization by eigenvectors and semidefinite programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011.
- [142] Amit Singer and Yoel Shkolnisky. Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming. *SIAM Journal on Imaging Sciences*, 4(2):543–572, 2011.
- [143] Bernd Sturmfels. *Algorithms in invariant theory*. Springer Science & Business Media, 2008.
- [144] D. L. Theobald and P. A. Steindel. Optimal simultaneous superpositioning of multiple structures with missing data. *Bioinformatics*, 28(15):1972–1979, 2012.

- [145] David J. Thouless, Philip W. Anderson, and Robert G. Palmer. Solution of ‘Solvable model of a spin glass’. *Philosophical Magazine*, 35(3):593–601, 1977.
- [146] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [147] BK Vainshtein and AB Goncharov. Determination of the spatial orientation of arbitrarily arranged identical particles of unknown structure from their projections. In *Soviet Physics Doklady*, volume 31, page 278, 1986.
- [148] A. W. van der Vaart. *Asymptotic statistics*, volume 3 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, Cambridge, 1998.
- [149] Marin Van Heel. Angular reconstitution: a posteriori assignment of projection directions for 3D reconstruction. *Ultramicroscopy*, 21(2):111–123, 1987.
- [150] Nicolas Verzelen and Ery Arias-Castro. Community detection in sparse random networks. *The Annals of Applied Probability*, 25(6):3465–3510, 2015.
- [151] Dimitri Vvedensky. *Chapter 8: Irreducible Representations of $SO(2)$ and $SO(3)$* . Group theory course [lecture notes]. <http://www.cmth.ph.ic.ac.uk/people/d.vvedensky/groups/Chapter8.pdf>, 2001.
- [152] Haolei Weng and Yang Feng. Community detection with nodal information. *arXiv:1610.09735*, 2016.
- [153] Nicholas C. Wormald. Models of random regular graphs. *London Mathematical Society Lecture Note Series*, pages 239–298, 1999.
- [154] Lenka Zdeborová and Florent Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [155] J. P. Zwart, R. van der Heiden, S. Gelsema, and F. Groen. Fast translation invariant classification of HRR range profiles in a zero phase representation. *Radar, Sonar and Navigation, IEE Proceedings*, 150(6):411–418, 2003.