

ORBIT RECOVERY FOR SPHERICAL FUNCTIONS

TAMIR BENDORY, DAN EDIDIN, JOSH KATZ, AND SHAY KREYMER

ABSTRACT. Orbit recovery is a central problem in both mathematics and applied sciences, with important applications to structural biology. This paper focuses on recovering generic orbits of functions on \mathbb{R}^n and the sphere S^{n-1} under the rotation action of $\mathrm{SO}(n)$. Specifically, we demonstrate that invariants of degree three (called the bispectrum) suffice to recover generic orbits of functions in finite-dimensional approximations of $L^2(\mathbb{R}^n)$ obtained by band-limiting the spherical component and discretizing the radial direction. In particular, our main result explicitly bounds the number of samples in the radial direction required for recovery from the degree three invariants. From an application perspective, the most important case is $\mathrm{SO}(3)$, which arises in many scientific fields, and in particular, plays a central role in leading structural biology applications such as cryo-electron tomography and cryo-electron microscopy. Our result for $\mathrm{SO}(3)$ states that considering three spherical shells (i.e., samples in the radial direction) is sufficient to recover generic orbits, which verifies an implicit conjecture made in [5]. Our proof technique provides an explicit, computationally efficient algorithm to recover the signal by successively solving systems of linear equations. We implemented this algorithm and demonstrated its effectiveness on two protein structures.

1. INTRODUCTION

Orbit recovery problems lie at the intersection of invariant theory and practical scientific applications, in particular in structural biology [7]. A prototypical example is the multi-reference alignment (MRA) problem, where the goal is to recover a signal from noisy observations subject to random group actions [6, 5, 8]. Invariant polynomials have emerged as a powerful tool for signal estimation in this setting, as they bypass the need to estimate the unknown group elements, a task that becomes particularly difficult in high-noise regimes. Remarkably, it has been shown that, from a statistical standpoint, this approach is optimal for signal recovery in such regimes [2, 1]. Beyond MRA, the problem of constructing polynomials that separate all orbits in a representation has been studied for over a century, with significant advances made in recent years. See for example [19, 21, 20]. In applications, however, it is typically only necessary to separate orbits almost surely or generically. Recent work has demonstrated that the number and degrees of polynomials required to separate generic orbits is often much lower than those necessary to recover all orbits [5, 25, 16, 23].

The simplified, and most studied MRA model, is when the group of circular shifts \mathbb{Z}_L acts on \mathbb{R}^L . In this case, it was shown that a generic signal can be recovered from the third-degree invariant [8, 35]. Follow-up papers extended the analyses to more intricate groups and transformations, such as $\mathrm{SO}(2)$ [33, 22], the dihedral group [11, 23] and dilations [40]; see [5] for the presentation of a wide range of different cases.

Although recovering an orbit from invariants of degree three would seem to be an intractable computational problem due to the difficult nature of solving systems of polynomial

equations, several of the examples that have generic degree three separation also have efficient and provable algorithms. For example, in [8] a series of efficient algorithms, including frequency marching and semi-definite programming, were designed for the case of \mathbb{Z}_L acting on \mathbb{R}^L . A Jennrich-type algorithm for the same case was designed in [35] and this was extended to the regular representation of any finite group in [24]. Spectral algorithms were also designed in [18]; see [1, 22] for related algorithms.

In this work, we study representations of $\mathrm{SO}(n)$ of the form

$$(1) \quad V_L = (L^2(S^{n-1})_L)^R,$$

where $L^2(S^{n-1})_L$ refers to L -band-limited functions on the $(n-1)$ -sphere. Our goal is to determine conditions that ensure that invariants of degree three can recover the orbit of a generic signal in V_L . Invariants of degree three are also referred to as the bispectrum, and we use the two terms interchangeably. We follow a frequency marching strategy, first suggested by the authors of [5] and also used in [32]. To derive the theoretical results, we use a small subset of the degree-three invariants to recursively solve the components of higher frequency from previously solved invariants of lower degree. Precisely, the invariants we consider give rise to linear equations for the component of a signal at higher frequency with coefficients determined by the lower frequency components. The main results of this paper, Theorem 2.1, gives a bound on the number of shells R (i.e., samples in the radial direction) as a function of n and L needed to ensure that we can recover the generic orbit from invariants of degree at most three. The proof is based on showing that the set of linear equations we consider always has full rank. The main result is presented in Section 2 and proved in Sections 3 and 4.

Section 5 is devoted to the analysis of the $n = 3$ case, the main application of our work. The problem of identifying an unknown signal in \mathbb{R}^3 up to rotation by an element in $\mathrm{SO}(3)$ occurs in a number of contexts, including two related techniques in molecular imaging, cryo-electron microscopy (cryo-EM) [34, 29] and cryo-electron tomography (cryo-ET) [37]. For the group $\mathrm{SO}(3)$, our main result, Theorem 2.2 states that $R = 3$ shells (independent of L) are sufficient to recover the generic signal in $(L^2(S^2)_L)^3$, for any L , from invariants of degree at most three. This affirms an implicit conjecture made in [5] and the bound is known to be optimal.

In Section 6, we implement the proposed frequency marching algorithm for $\mathrm{SO}(3)$ to recover the orbits using degree-three invariants. Numerically, there is no reason to discard any of the available equations. Leveraging all of them allows us to implement a numerically stable, frequency-marching algorithm that provably recovers the signal by successively solving systems of linear equations. We demonstrate the effectiveness, efficiency, and robustness of our approach for two molecular structures in the Electron Microscopy Data Bank (EMDB).

2. PROBLEM FORMULATION AND STATEMENT OF RESULTS

2.1. Formulation of the problem. Recall that a polynomial $f \in \mathbb{R}[x_1, \dots, x_n]$ is *harmonic* if $\Delta f = 0$ where $\Delta = \sum_{i=1}^n \frac{\partial^2}{\partial x_i^2}$ is the Laplacian. Since the Laplacian is rotation invariant, the translation of a harmonic polynomial by a rotation is again a harmonic polynomial. We denote by H_ℓ the vector space of homogeneous harmonic polynomials of degree ℓ with the action of $\mathrm{SO}(n)$ by rotation. Since a homogeneous polynomial is uniquely determined by its values on the sphere, we view elements of H_ℓ as spherical harmonic polynomials.

The space of homogeneous polynomials of degree ℓ on S^{n-1} decomposes as an $\text{SO}(n)$ representation into irreducible components as

$$(2) \quad \text{Sym}^\ell(S^{n-1}) = \bigoplus_{i=0}^{\lfloor \ell/2 \rfloor} H_{\ell-2i}.$$

Equation (2) yields the following formula for the dimension of H_ℓ :

$$(3) \quad m_\ell := \dim H_\ell = \binom{n+\ell-1}{n-1} - \binom{n+\ell-3}{n-1}.$$

When $n = 3$, then $m_\ell = 2\ell + 1$ and for general n we have that $m_\ell \sim O(\ell^{n-2})$.

A basic result in harmonic analysis states that we have a decomposition of $\text{SO}(n)$ representations $L^2(S^{n-1}) = \bigoplus_{\ell=0}^\infty H_\ell$. When $n = 3$, the representations H_ℓ are the usual spaces of spherical harmonics studied in the physics literature; see for example [30, 4]. We denote the finite sum $\bigoplus_{\ell=0}^L H_\ell$ as $L^2(S^{n-1})_L$ and refer to this finite-dimensional representation of $\text{SO}(n)$ as the space of L -bandlimited functions on the $(n-1)$ -sphere. In this work, we consider the finite-dimensional $\text{SO}(n)$ representation of the form (1), which we view as a discrete approximation for $L^2(\mathbb{R}^n)$, as a representation of $\text{SO}(n)$. Our goal is to determine bounds on R —the number of spherical shells—which ensure that the generic signal $f \in V_L$ can be recovered from invariants of degree at most three.

2.2. Main results. We are now ready to present the main results of this paper. The proofs are provided in Sections 4 and Section 5, and a technical background is provided in Section 3.

Theorem 2.1. *Let $V_L = (L^2(S^{n-1})_L)^R$ be the finite-dimensional approximation of $L^2(\mathbb{R}^n)$ by functions which are defined on R spherical shells and L -bandlimited on each shell. If $R \geq \frac{\dim H_\ell + (\lfloor \ell/2 \rfloor - 1)}{(\ell-1)}$, for all $\ell > 1$, where $\dim H_\ell$ is given in (3), then the $\text{O}(n)$ orbit of a generic real valued $f \in V$ is determined by $\text{SO}(n)$ -invariants of degree at most three.*

In the case of $\text{SO}(3)$, the most important case from the scientific perspective, we can eliminate the reflection ambiguity, and we obtain the following result. Importantly, in this case, only three spherical shells are required, independent of L .

Theorem 2.2 (The $\text{SO}(3)$ case.). *If $R \geq 3$, then for any $L \geq 0$ the invariants of degree at most three separate generic orbits in $(L^2(S^2)_L)^R$.*

Theorem 2.2 improves the linear bound of $R \geq L + 2$ given in [25] and proves that the hypothesis of [5, Theorem 4.19] is always satisfied when $R \geq 3$, thereby proving an implicit conjecture made there. Moreover, the bound $R \geq 3$ is optimal since it is known that if $R < 3$ then invariants of degree at most three cannot recover signals with small band limit [5]. It is an open question as to whether or not there is a constant multiplicity (independent of the band limit) so that degree 3 invariants recover a generic orbit for $\text{SO}(n)$ with $n > 3$.

2.3. Applications to sample complexity analysis in multi-reference alignment. The multi-reference alignment (MRA) model entails estimating a signal f in a representation V of a compact group G from m observations of the form

$$(4) \quad y_i = g_i \cdot f + \varepsilon_i,$$

where $g_1, \dots, g_m \in G$ are random elements drawn from a uniform distribution over the group G , and $\varepsilon_i \sim \mathcal{N}(0, \sigma^2 I)$ are i.i.d. realization of a Gaussian noise with variance σ^2 .

The goal is to estimate the signal f from y_1, \dots, y_m , while the group elements are treated as nuisance variables. Because there is no way to distinguish f from $g \cdot f$, the MRA problem is an orbit recovery problem. This model was first suggested as an abstraction of the cryo-EM model [7], and the problem has been studied in more generality as a prototype of statistical models with intrinsic algebraic structures, e.g., [6, 5, 9].

Sample complexity refers to the number of observations required to accurately estimate a signal. It was shown that the sample complexity of the MRA model, in the high noise regime $\sigma \rightarrow \infty$, is determined by the lowest order moment that identifies the signal [2, 1]. If the distribution over the group is uniform, then the moments are equivalent to invariant polynomials. In particular, as $m, \sigma \rightarrow \infty$, a necessary condition for accurate signal recovery is that $m/\sigma^{2d} \rightarrow \infty$, where d is the lowest order moment of the observations that determines the orbit of f uniquely. For example, if the first and second moments do not determine the signal, but the third moment does, then m must scale faster than σ^6 for accurate recovery. Based on this result, beginning in [5], the problem of identifying representations of compact groups for which the third moment can determine generic orbits has been studied by a number of authors [16, 25, 23, 24, 15].

Immediate corollaries of Theorem 2.1 and Theorem 2.2 are the following:

Corollary 2.3 (The sample complexity of multi-reference alignment for $\text{SO}(n)$). *Consider the MRA model (4) with $\sigma \rightarrow \infty$, where $V = (\oplus_{\ell=0}^L H_\ell)^R$ and $G = \text{SO}(n)$. If $R \geq \frac{\dim H_\ell + (\lceil \ell/2 \rceil - 1)}{(\ell - 1)}$, then the minimal number of observations required for accurate recovery of the $\text{O}(n)$ orbit of f , regardless of any specific algorithm, is $m/\sigma^6 \rightarrow \infty$.*

Corollary 2.4 (The sample complexity of multi-reference alignment for $\text{SO}(3)$). *Consider the MRA model (4) with $\sigma \rightarrow \infty$, where $V = (\oplus_{\ell=0}^L H_\ell)^R$ and $G = \text{SO}(3)$. If $R \geq 3$, then the minimal number of observations required for accurate recovery of the $\text{SO}(3)$ orbit of f , regardless of any specific algorithm, is $m/\sigma^6 \rightarrow \infty$.*

An interesting extension of the MRA model is the heterogeneous MRA model, where each measurement corresponds to a translated and noisy observation of one of K underlying signals (the model in (4) corresponds to the special case ($K = 1$)). In this setting, it is not possible to recover the third-degree invariant of each individual signal; only their average can be estimated. Nevertheless, recent works [5, 17] provide strong evidence that, in various MRA settings, it is possible to estimate the orbits of multiple signals simultaneously ($K > 1$) from the average of their degree-three invariants. We conjecture that a similar phenomenon holds in the model studied in this paper. In particular, a compelling direction for future work is to establish bounds on orbit recovery from third-degree invariants in the heterogeneous case, as a function of the number of shells R and the number of signals K .

2.4. Implications to structural biology. The MRA model (4) with $G = \text{SO}(3)$ fits the problem of subtomogram averaging in cryo-ET, which is one of the main steps in recovering molecular structures in in-situ environments [38]. Since the noise level in cryo-ET data is extremely high, Corollary 2.4 implies that the number of observations must be larger than σ^6 for recovery.

The problem of single-particle reconstruction in cryo-EM follows a similar model, but we have access only to the tomographic projections of the rotated copies of the signal [7]. Namely, The cryo-EM model reads

$$(5) \quad y_i = P(g_i \cdot f) + \varepsilon_i,$$

where $Pf(x_1, x_2) = \int_{x_3} f(x_1, x_2, x_3)$ is a tomographic projection. Thus, the results of this paper cannot be directly applied to derive the sample complexity of cryo-EM. Yet, following the results of this paper, and [5, Conjecture 4.17], we conjecture that for enough spherical shells, a signal can be recovered from the degree-three invariants in the cryo-EM model.

Conjecture 2.5 (The sample complexity of the cryo-EM model). *Consider the cryo-EM model (5). Then, if $R \geq 3$ is large enough, then the minimal number of observations required for accurate recovery of the $O(n)$ orbit of f , regardless of any specific algorithm, is $m/\sigma^6 \rightarrow \infty$.*

2.5. Related work. In [5, Conjecture 4.11], the authors conjecture that if $L \geq 10$ then the generic orbit in $L^2(S^2)_L = \oplus_{\ell=0}^L H_\ell$ can be determined up to a finite list of orbits from invariants of degree at most three. They also showed that the bound $L \geq 10$ is sharp. This conjecture was recently proved in [26]. Although [5, Conjecture 4.11] is about finite list recovery, it is possible that for L sufficiently large, the generic orbit in V_L can also be recovered from the bispectrum; i.e. an orbit can be determined up to a list of size one. To that end, [32] produces a frequency marching algorithm, which can be used to determine, for ℓ sufficiently large, the ℓ -th frequency component of $f \in V_L$, denoted by f^ℓ , from a portion of the bispectrum and the prior knowledge of the components f^i with $i < \ell$. Each step in the frequency march requires the solution to a linear system of equations for coefficients of f^ℓ , and the condition number of this system can be bounded in terms of the parameters of the model. However, the algorithm of [32] cannot be used to directly recover an unknown function $f \in V_L$ from invariants of low degree, since it requires as prior input an unspecified number of components of f of low frequency.

We also mention that when the support of the signal on the sphere is small then it was demonstrated in [12] that the $SO(3)$ -orbit can also be recovered from the bispectrum. In addition, a recent series of works showed that recovery from the second moment is possible, if the signal is known to lie in a semi-algebraic set of low dimension; for example, if the signal is sparse or is in the image of a deep neural network [13, 9, 3].

3. REPRESENTATION AND INVARIANT THEORY OF HARMONIC POLYNOMIALS

The decomposition (2) implies that H_ℓ is the irreducible representation of $SO(n)$ with highest weight vector ℓL_1 in the notation of [27, Section 19]. As a consequence, we obtain the following proposition.

Proposition 3.1. *The tensor product $H_i \otimes H_{\ell-i}$ contains a unique copy H_ℓ .*

Proof. The highest weight vector contained in the tensor product $H_i \otimes H_{\ell-i}$ is, in the notation of [27], $iL_1 + (\ell-i)L_1 = \ell L_1$, which appears with multiplicity one. Hence, $H_i \otimes H_{\ell-i}$ contains a single copy H_ℓ as a highest weight representation. \square

3.1. Invariants of $L^2(S^{n-1})_L$. We now describe a collection of polynomial invariants of degree at most three in $V_L := L^2(S^{n-1})_L$. To produce these invariant polynomials we first identify invariant tensors in $V_L^{\otimes k}$ (for $1 \leq k \leq 3$) and then symmetrize* them to produce invariant polynomials in $\text{Sym}^k V_L$. To begin, choose an orthonormal basis $e_1^\ell, \dots, e_{m_\ell}^\ell$ for each

*Recall that if V is a vector space then the symmetrization of a tensor $t = v_1 \otimes \dots \otimes v_k \in V^{\otimes k}$ is the symmetric tensor $St = \frac{1}{k!} \sum_{\sigma \in S_k} (v_{\sigma(1)} \otimes \dots \otimes v_{\sigma(k)})$. This operation can be extended linearly to define a projection $S: V^{\otimes k} \rightarrow \text{Sym}^k V$. The symmetrization of an arbitrary element $x \in V^{\otimes k}$ is its image under this projection.

summand H_ℓ in V_L . An element $f \in V_L$ can be represented as an L -tuple $f = (f^0, \dots, f^L)$ and each f^ℓ can be expanded as $f^\ell = \sum_{k=1}^{m_\ell} f_k^\ell e_k^\ell$.

Invariants of degree one. If $f \in V_L$ then the invariants of degree one in f is just the projection of f to the trivial one dimensional summand H_0 ; i.e., the only invariant of degree one is f^0 .

Invariants of degree two. The representations H_ℓ are self-dual via the quadratic form defining $\mathrm{SO}(n)$. Hence $V_L \otimes V_L$ is identified with $\mathrm{Hom}(V_L, V_L)$. By Schur's lemma, $(V_L \otimes V_L)^{\mathrm{SO}(n)} = \bigoplus_{\ell=0}^L (H_\ell \otimes H_\ell)^G$ and $(H_\ell \otimes H_\ell)^G$ is generated by the symmetric tensor $\sum_{k=1}^{m_\ell} e_k^\ell \otimes e_k^\ell$. Thus, the invariants of degree 2 in V_L are generated by the quadratic norms; i.e., the invariant polynomials $\sum_{k=1}^{m_\ell} (f_k^\ell)^2$.

Invariants of degree three. Here, we only construct a subset of the degree three invariants that are sufficient to distinguish generic orbits. Our approach to constructing invariants uses representation theory. Specifically, we will use representation-theoretic methods to construct invariant tensors in $(V_L)^{\otimes 3}$ and then symmetrize them to obtain an invariant polynomial. The following lemma identifies, for each ℓ , a set of $\lfloor \ell/2 \rfloor$ distinct invariant polynomials of degree three, which are linear in the ℓ -th component vector of $f \in L^2(S^{n-1})_\ell$.

Lemma 3.2. *For fixed $\ell \leq L$ and for $1 \leq i \leq \lfloor \ell/2 \rfloor$, the tensor product $H_i \otimes H_{\ell-i} \otimes H_\ell$ contains a nontrivial invariant tensor whose symmetrization is not identically zero. Hence, we obtain an invariant polynomial of degree 3, which we label as $I_3(i, \ell - i, \ell) \in \mathrm{Sym}^3(H_i + H_{\ell-i} + H_\ell)^{\mathrm{SO}(n)}$.*

Proof. By Proposition 3.1, the tensor product $H_i \otimes H_{\ell-i}$ contains a single copy of the irreducible representation H_ℓ . Hence, by Schur's lemma, $(H_i \otimes H_{\ell-i} \otimes H_\ell)^G$ is one-dimensional. To produce an invariant polynomial, let $\phi_\ell: H_i \otimes H_{\ell-i} \rightarrow H_\ell$ be the projection onto the summand isomorphic to H_ℓ . If $T_1, \dots, T_{m_\ell} \in H_i \otimes H_{\ell-i}$ is an orthonormal basis for this summand, which is identified with the prechosen basis, then $\sum_{k=1}^{m_\ell} T_k \otimes e_k^\ell$ is an invariant in $H_i \otimes H_{\ell-i} \otimes H_\ell$. If $i, \ell - i, \ell$ are distinct, then it is immediate that its symmetrization is non-zero. In the case where $i = \ell - i$ we note that the summand in $H_i \otimes H_i$ isomorphic to H_{2i} , necessarily lies in $\mathrm{Sym}^2 H_i$. The reason is that the highest weight of H_{2i} equals the highest weight in the (reducible) representation $\mathrm{Sym}^2 H_i$ which is $(2i)L_1$. Thus, in this case, the T_k are already symmetric, so the symmetrization of $\sum_{k=1}^{m_\ell} T_k \otimes e_k^\ell$ also cannot vanish. \square

The invariants constructed in Lemma 3.2 can be made explicit. Let $f = (f^1 + \dots + f^L) \in V_L$ with $f^\ell = \sum_{k=1}^{m_\ell} f_k^\ell e_k^\ell$. We can write the projection of $f^i \otimes f^{\ell-i}$ to the summand H_ℓ as $\sum_{k=1}^{m_\ell} C(f^i, f^{\ell-i})_k T_k$, where $C(f^i, f^{\ell-i})_k$ are bilinear forms in the coefficients $\{f_s^i, f_t^{\ell-i}\}$. Precisely, we have

$$(6) \quad C(f^i, f^{\ell-i})_k = \sum_{s=1}^{m_i} \sum_{t=1}^{m_{\ell-i}} c_{s,t}^k f_s^i f_t^{\ell-i}.$$

Here, the coefficients $c_{s,t}^k$ are the Clebsch-Gordan coefficients for $\mathrm{SO}(n)$. Since the projection $H_i \otimes H_{\ell-i} \rightarrow H_\ell$ is surjective, we know that for generic choice of $f_i, f_{\ell-i}$ the coefficients $C(f_i, f_{\ell-i})_k$ are identically non-zero and that there can be no linear relations between the $C(f^i, f^{\ell-i})_k$ as bilinear forms in $f_s^i, f_t^{\ell-i}$. The coefficients of the invariant part of the symmetrization of $f^i \otimes f^{\ell-i} \otimes f^\ell$ are the invariant polynomials $\sum_{k=1}^{m_\ell} C(f^i, f^{\ell-i})_k f_k^\ell$, which we denote by $I(i, \ell - i, \ell)(f^i, f^{\ell-i}, f^\ell)$.

4. PROOF OF THEOREM 2.1

We view an element $f \in (L^2(S^{n-1})_L)^R$ as an R -tuple of functions $f[1], \dots, f[R]$, one for each of the R shells. Each function $f[r]$ is assumed to be band-limited to frequency L , so we can write $f[r] = \sum_{\ell=0}^L f^\ell[r]$. In turn, the function $f^\ell[r]$ can be expanded in terms of an orthonormal basis $\{e_k^\ell[r]\}_{k=1}^{m_\ell}$ as $f^\ell[r] = \sum_{k=1}^{m_\ell} f_k^\ell[r] e_k^\ell[r]$. Following similar notation used in [10], we let A^ℓ be the $m_\ell \times R$ matrix such that $A_{k,r}^\ell = f_k^\ell[r]$. Our goal is to recover, for a generic f , the matrices A^ℓ from invariants of degree at most three.

Given a set of vectors $\{f^j[m]\}$ with $1 \leq j \leq \lceil \ell/2 \rceil$ and $1 \leq m \leq R$, the invariants $I_3(i, \ell - i, \ell)(f^i[m], f^{\ell-i}[m], f^\ell[r])$ determine linear equations for the coefficients $f_k^\ell[r]$

$$(7) \quad \sum_{k=1}^{m_\ell} C(f^i[m], f^{\ell-i}[n])_k f_k^\ell[r] = I_3(i, \ell - i, \ell)(f^i[m], f^{\ell-i}[n], f^\ell[r]).$$

If $i \neq j$, let

$$S_{i,j}[R] = \{(1, 1), \dots, (1, R), (2, 1), \dots, (R, 1)\} \subset [1, R]^2,$$

and let

$$S_{i,i}[R] = \{(1, 1), \dots, (1, R)\}.$$

Consider the system of $R(\ell - 1) - (\lceil \ell/2 \rceil - 1)$ linear equations for the coefficients $f^\ell[r]_k$

$$(8) \quad \left\{ \sum_{k=1}^{m_\ell} C(f^i[m], f^{\ell-i}[n])_k f_k^\ell[r] = I_3(i, \ell - i, \ell)(f^i[m], f^{\ell-i}[n], f^\ell[r]) \right\}_{1 \leq i \leq \lceil \ell/2 \rceil, (m,n) \in S_{i,\ell-i}}$$

Proposition 4.1. *For a generic choice $f^j[m]$ with $j < \ell$ and $m = 1, \dots, R$, the linear system (8) has full rank.*

Given Proposition 4.1, the system (8) will have a unique solution for $f^\ell[r]$ once $R(\ell - 1) - (\lceil \ell/2 \rceil - 1) \geq \dim H_\ell$. Since the invariants of degree two determine the Gram matrices $(A^\ell)^T A^\ell$, we know $A^1 = (f^1[1], \dots, f^1[R])$ up to the action of $O(n)$. Once we choose a representative for A_1 in its $O(n)$ -orbit, we can determine the matrices A^2, \dots, A^L inductively using Proposition 4.1.

Proof of Proposition 4.1. For fixed i, m, n we have the linear equation

$$\sum_{k=1}^{m_\ell} C(f^i[m], f^{\ell-i}[n])_k f_k^\ell[r] = I_3(i, \ell - i, \ell)(f^i[m], f^{\ell-i}[n], f^\ell[r]),$$

where we view $C_k(f^i[m], f^{\ell-i}[n])$ as coefficients in variables $x_{s,t}^i[m, n] = f_s^i[m] f_t^{\ell-i}[n]$. For distinct pairs $(m, n), (m', n') \in S_{(i,\ell-i)}[R]$, there are no algebraic relations between the sets of $m_i m_{\ell-i}$ monomials $\{x_{s,t}^i[m, n] = f_s^i[m] f_t^{\ell-i}[n]\}_{s,t}$ and $\{x_{s,t}^i[m', n'] = f_s^i[m'] f_t^{\ell-i}[n']\}$, so we may view them as distinct sets of variables. Thus, we obtain an $|S_{i,\ell-i}[R]| \times m_\ell$ linear system whose coefficient matrix is

$$M_i = \begin{bmatrix} C(\{x_{s,t}^i[1, 1]\})_1 & C(\{x_{s,t}^i[1, 1]\})_2 & \dots & C(\{x_{s,t}^i[1, 1]\})_{m_\ell} \\ \vdots & \vdots & \vdots & \vdots \\ C(\{x_{s,t}^i[1, R]\})_1 & C(\{x_{s,t}^i[1, R]\})_2 & \dots & C(\{x_{s,t}^i[1, R]\})_{m_\ell} \\ C(\{x_{s,t}^i[2, 1]\})_1 & C(\{x_{s,t}^i[2, 1]\})_2 & \dots & C(\{x_{s,t}^i[2, 1]\})_{m_\ell} \\ \vdots & \vdots & \vdots & \vdots \\ C(\{x_{s,t}^i[R, 1]\})_1 & C(\{x_{s,t}^i[R, 1]\})_2 & \dots & C(\{x_{s,t}^i[R, 1]\})_{m_\ell} \end{bmatrix}$$

if $i \neq \ell - i$, and

$$M_i = \begin{bmatrix} C(\{x_{s,t}^i[1,1]\})_1 & C(\{x_{s,t}^i[1,1]\})_2 & \dots & C(\{x_{s,t}^i[1,1]\})_{m_\ell} \\ \vdots & \vdots & \ddots & \vdots \\ C(\{x_{s,t}^i[1,R]\})_1 & C(\{x_{s,t}^i[1,R]\})_2 & \dots & C(\{x_{s,t}^i[1,R]\})_{m_\ell} \end{bmatrix}$$

if $i = k - i$. Here, $C(\{x_{s,t}^i[m,n]\})_\ell$ is of the form $\sum_{s=1}^{m_i} \sum_{t=1}^{m_{\ell-i}} c_{s,t}^\ell x_{s,t}^i[m,n]$. For fixed (m,n) , the forms $C(\{x_{s,t}^i[m,n]\})_\ell$ are linearly independent and thus the matrix M_i must have full rank over the field of rational functions $\mathbb{R}(\{x_{s,t}^i[m,n]\}_{s,t,m,n})$. Likewise, as we vary i in the range $1 \leq i \leq \lfloor \ell/2 \rfloor$, the matrix

$$\begin{bmatrix} M_1 \\ M_2 \\ \vdots \\ M_{\lfloor \ell/2 \rfloor} \end{bmatrix}$$

has full rank over the field of rational functions $\mathbb{R}(\{x_{s,t}^i[m,n]\}_{i,s,t,m,n})$. This implies that when we substitute values for the coefficients, we obtain, for generic choices of the functions $f^i[n]$, for $0 \leq i \leq \lfloor \frac{\ell}{2} \rfloor$, a maximal rank matrix. \square

Remark 4.2. We can also consider the $R^2 \times \lfloor k/2 \rfloor$ linear system

$$(9) \quad \left\{ \sum_{k=1}^{m_\ell} C(f^i[m], f^{\ell-i}[n])_\ell f_k^\ell[r] = I_3(i, \ell - i, \ell)(f^i[m], f^{\ell-i}[n], f^\ell[r]) \right\}_{1 \leq i \leq \lfloor \ell/2 \rfloor, (m,n) \in [1,R]^2}$$

If we can show that (9) has full rank, then for generic f we can improve the bound in Theorem 2.1 to $R \geq \sqrt{\frac{\dim H_\ell}{\lfloor \ell/2 \rfloor}}$. To do this, it suffices to find a single set of values for the $f_s^j[m]$ for which the system has full rank. However, without explicit knowledge of the Clebsch-Gordan coefficients, this is very difficult.

5. ORBIT RECOVERY OVER $\text{SO}(3)$

The invariants we obtain for general $\text{SO}(n)$ are not given explicitly due to the difficulty of choosing natural bases for the irreducible representations H_ℓ and computing the Clebsch-Gordan coefficients. However, for $\text{SO}(3)$ there are explicit functional bases for the H_ℓ and closed formulas for the Clebsch-Gordan coefficients. For these reasons, it is possible to use a mix of computational and theoretical techniques to answer orbit separation questions. Before we prove Theorem 2.2, we recall some known results, which we use in the proof.

For $\text{SO}(3)$, we have an explicit decomposition of $H_i \otimes H_j$ into irreducibles for any i, j . Precisely, we know that

$$(10) \quad H_i \otimes H_j = \bigoplus_{\ell=|i-j|}^{i+j} H_\ell.$$

Hence, H_ℓ appears with multiplicity one in any tensor product $H_i \otimes H_j$ with $i+j \geq \ell \geq i-j$. Using a basis for H_ℓ of spherical harmonic functions $Y_\ell^k(\theta, \phi)$, we can obtain a large number of explicit invariants of degree three:

$$(11) \quad I_3(\ell_1, \ell_2, \ell_3)(f^{\ell_1}, f^{\ell_2}, f^{\ell_3}) = \sum_{k_1+k_2+k_3=0, |k_i| \leq \ell_i} (-1)^{k_1} \langle \ell_2 k_2 \ell_3 k_3 | \ell_1 (-k_1) \rangle f_{k_1}^{\ell_1} f_{k_2}^{\ell_2} f_{k_3}^{\ell_3},$$

where

$$\begin{aligned}
\langle \ell_1 m_1 \ell_2 m_2 | \ell m \rangle &= \delta_{m, m_1 + m_2} \sqrt{\frac{(2\ell + 1)(\ell + \ell_1 - \ell_2)!(\ell - \ell_1 + \ell_2)!(\ell_1 + \ell_2 - \ell)!}{(\ell_1 + \ell_2 + \ell + 1)!}} \\
&\quad \times \sqrt{(\ell + m)!(\ell - m)!(\ell_1 - m_1)!(\ell_1 + m_1)!(\ell_2 - m_2)!(\ell_2 + m_2)!} \\
(12) \quad &\times \sum_k \frac{(-1)^k}{k!(\ell_1 + \ell_2 - \ell - k)!(\ell_1 - m_1 - k)!(\ell_2 + m_2 - k)!(\ell - \ell_2 + m_1 + k)!(\ell - \ell_1 - m_2 + k)!}.
\end{aligned}$$

When $\ell_3 = \ell$ and $\ell_1, \ell_2 < \ell$ but $\ell_1 + \ell_2 \geq \ell$, then (12) gives a linear equation for the unknowns f_k^ℓ in terms of the coefficients $f_s^{\ell_1}, f_t^{\ell_2}$ in the expansions $f^{\ell_1} = \sum_{s=-\ell_1}^{\ell_1} f_s^{\ell_1} Y_{\ell_1}^s$ and $f^{\ell_2} = \sum_{t=-\ell_2}^{\ell_2} f_t^{\ell_2} Y_{\ell_2}^t$ of f^{ℓ_1} and f^{ℓ_2} in spherical harmonics. In [5, Section 4.6.1], the authors proved, using numerical techniques, that if $R \geq 3$ then for generic f the unknown vectors $f^\ell[r]$ for $2 \leq \ell \leq 16$ are determined from the vectors $f^1[m]$ for $m = 1, \dots, R$. In addition, if $R \geq 3$, then invariants of degree at most three determine the coefficients $f^1[m]$ up to the action of $\text{SO}(3)$ (as opposed to $\text{O}(3)$).

Proof of Theorem 2.2. When $R = 3$, the expression $3(\ell - 1) - (\lceil \ell/2 \rceil - 1) \geq 2\ell + 1$ when $\ell \geq 6$. In particular, this means by Proposition 4.1 we can determine $f^\ell[r]$ for $k \geq 6$ from the vectors $f^i[m]$ for $1 \leq i < \ell$ and $1 \leq m \leq 3$. By the results of [5], we know that invariants of degree at most three determine the $\text{SO}(3)$ orbit of the unknown vectors $f^i[m]$ for $1 \leq i \leq 6$. \square

Remark 5.1. (The optimality of Theorem 2.2) The result of Theorem 2.2 is as strong as possible in the following sense. If $R < 3$, then, as previously observed in [5], for L small enough invariants of degree at most three in $V = (L^2(S^2)_L)^R$ do not generate a field of transcendence degree equal to $\dim V / \text{SO}(3)$. In this case, a frequency marching algorithm cannot be used to recover a generic orbit from invariants of degree at most three.

6. NUMERICAL EXPERIMENTS

In this section, we are interested in numerically investigating the performance of the frequency marching algorithm and its performance in the presence of noise.

6.1. Setting. Let $f \in \mathbb{R}^3$ be a smooth, real-valued function. We assume that f is band-limited in the sense that it can be represented using a truncated 3-D spherical-Bessel expansion:

$$(13) \quad f(cr, \theta, \varphi) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \sum_{s=1}^{S(\ell)} x_{\ell, m, s} \mathcal{Y}_\ell^m(\theta, \varphi) j_{\ell, s}(r).$$

where c is the bandlimit, $S(\ell)$ is determined by the Nyquist criterion [14], and $j_{\ell, s}(r)$ is the normalized spherical Bessel function:

$$(14) \quad j_{\ell, s}(r) = \frac{4}{|j_{\ell+1}(u_{\ell, s})|} j_\ell(u_{\ell, s} r),$$

with j_ℓ the spherical Bessel function of order ℓ and $u_{\ell,s}$ the s -th positive zero of j_ℓ . The real spherical harmonics \mathcal{Y}_ℓ^m are defined using the complex spherical harmonics:

$$(15) \quad \mathcal{Y}_\ell^m(\theta, \varphi) = \begin{cases} \sqrt{2} \Re Y_\ell^m(\theta, \varphi), & m > 0, \\ Y_\ell^0(\theta, \varphi), & m = 0, \\ \sqrt{2} \Im Y_\ell^{|m|}(\theta, \varphi), & m < 0, \end{cases}$$

where the complex spherical harmonics Y_ℓ^m are defined as

$$(16) \quad Y_\ell^m(\theta, \varphi) = \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos \theta) e^{im\varphi},$$

where P_ℓ^m are the associated Legendre polynomials with the Condon–Shortley phase. We set $c = 1/2$ to match the Nyquist sampling rate [31].

The sum over s in (13) can be interpreted as summing over concentric radial “shells.” To investigate the effect of shell count, we limit the sum to a fixed number R :

$$(17) \quad f(cr, \theta, \varphi) = \sum_{\ell=0}^{\ell_{\max}} \sum_{m=-\ell}^{\ell} \sum_{s=1}^R x_{\ell,m,s} \mathcal{Y}_\ell^m(\theta, \varphi) j_{\ell,s}(r), \quad r \leq 1.$$

The bispectrum is then defined as

$$(18) \quad B(x)[\ell_1, \ell_2, \ell_3, s_1, s_2, s_3] = \sum_{\substack{m_1+m_2+m_3=0 \\ |m_i| \leq \ell_i}} (-1)^{m_1} \langle \ell_2 m_2 \ell_3 m_3 \mid \ell_1 (-m_1) \rangle x_{\ell_1, m_1, s_1} x_{\ell_2, m_2, s_2} x_{\ell_3, m_3, s_3},$$

where the Clebsch–Gordan coefficients $\langle \ell_1 m_1 \ell_2 m_2 \mid \ell m \rangle$ are defined in (12).

6.2. Algorithm. While recovering f from its bispectrum is a non-convex problem, recovering the coefficients corresponding to the ℓ -th frequency, given the coefficients of the $1, \dots, \ell-1$ coefficients, is a linear problem. Thus, we recover f by successively solving the linear system of equations (18); this is the frequency marching algorithm. The expansion coefficient for $\ell = 0$ is just the mean of the volume (first-degree invariant). The $\ell = 1$ component of the signal f , denoted $A^1 \in \mathbb{R}^{3 \times R}$, can be determined up to the action of $O(3)$ from the Gram matrix $(A^1)^T A^1$, which is a set of invariants of degree 2. Now let $U \in \mathbb{R}^{3 \times R}$ be a matrix with $U^T U = (A^1)^T A^1$. We have two distinct $SO(3)$ orbits represented by U and its reflection U' . Either U or U' is in the same $SO(3)$ -orbit of A^1 . In order to determine the correct orbit we compare the determinants of the 3×3 minors of A^1 with both of U and U' . These determinants can be written as $I_3(1, 1, 1)(f^1[r_1], f^1[r_2], f^3[r_3])$ for $1 \leq r_1, r_2, r_3 \leq R$ in our notation and hence they are invariants of degree three.

After solving for the $\ell = 1$ term up to translation by an element of $SO(3)$, we fix a representative for the orbit. Once the action of $SO(3)$ has been fixed, we sequentially recover the rest of the coefficients up to $\ell = \ell_{\max}$. Specifically, for $\ell = 2$, we use the coefficients for $\ell = 0, 1$, for $\ell = 3$ the coefficients of $\ell = 0, 1, 2$ and so on. The reconstruction is of the orbit of the vector of coefficients, i.e., the volume is estimated up to a rotation. Algorithm 1 provides a pseudocode of the algorithm. In our numerical experiments, for simplicity, we assume that we know the coefficients corresponding to frequencies $\ell = 0$ and $\ell = 1$, thereby removing the rotation ambiguity.

Algorithm 1 Volume reconstruction from its bispectrum using frequency marching

-
- 1: **Inputs:** The bispectrum (18), and the spherical-Bessel coefficients corresponding to $\ell = 0, 1$
 - 2: **Outputs:** The spherical-Bessel coefficients $x_{\ell,m,s}$ (17), up to frequency ℓ_{\max} and number of shells R
 - 3: **for** frequency $\ell' = 2$ to ℓ_{\max} **do**
 - 4: Compute $\{x_{\ell',m,s}\}_{m=-\ell',s=1}^{m=\ell',s=R}$ by solving the linear system of equations (18) given coefficients up to ℓ' .
 - 5: **end for**
-

6.3. Numerical experiments. We consider two volumes in our experiments: the TRPV1 structure [28], available from the Electron Microscopy Data Bank (EMDB) under accession code EMD-8117[†], and the Plasmodium falciparum 80S ribosome [39], available as EMD-2660. Both volumes were downsampled to 31^3 voxels and expanded using varying values of ℓ_{\max} and R . The experiments were conducted on a MacBook Pro (2023) equipped with an Apple M3 Pro chip (12-core CPU, 18-core GPU) and 36GB of unified memory, running macOS 15.5. The code used to reproduce all numerical experiments is publicly available at <https://github.com/krshay/orbit-recovery-for-spherical-functions>.

Recovery in the absence of noise. We begin by demonstrating successful volume reconstructions from clean bispectrum data; see Figure 1. These results use $\ell_{\max} = 10$ and $R = 8$ for both volumes. The code was executed natively without virtualization, and took approximately 17 seconds for the reconstruction. The molecular visualizations were produced using UCSF Chimera [36]. We emphasize that the algorithm accurately recovers the volume for $R \geq 3$, as expected from the theory, and we provide an example with additional shells for visual clarity.




FIGURE 1. Volume reconstructions from clean bispectrum, with $\ell_{\max} = 10$ and $R = 8$ for the TRPV1 structure and the Plasmodium falciparum 80S ribosome. Left (gray): ground truth after expanding to the corresponding maximal frequency ℓ_{\max} ; right (yellow): reconstruction.

[†]<https://www.ebi.ac.uk/emdb/>

Condition number analysis. In practice, it is important that the recursive system of linear equations has good condition numbers in order to ensure the stability of the solutions. Fortunately, our numerical analysis indicates that these systems are generally well-conditioned. Tables 1 and 2 report the condition number. Generally, the condition number improves with the frequency and number of shells.

TABLE 1. Condition numbers for TRPV1

ℓ	3 shells	4 shells	5 shells
2	8749	1391	609.9
3	247.9	115.0	25.12
4	27.32	15.11	15.65
5	17.52	11.98	11.65
6	19.36	11.68	10.96
7	21.64	16.44	13.85
8	15.38	10.92	9.707
9	14.52	11.29	10.56
10	15.24	11.35	10.63

TABLE 2. Condition numbers for *Plasmodium falciparum* 80S ribosome

ℓ	3 shells	4 shells	5 shells
2	3.037	3.209	3.384
3	7.301	3.213	3.138
4	3.309	2.593	2.614
5	3.835	2.792	2.339
6	2.829	2.105	1.705
7	2.665	2.163	1.875
8	2.811	1.948	1.771
9	2.774	2.210	2.138
10	2.217	2.004	1.995

Robustness to Noise. In practical settings, the bispectrum is not available in its exact form but rather estimated from noisy observations. To evaluate the algorithm’s robustness to noise, we consider the following experimental setup. Let f_{true} denote the ground-truth volume—in this case, the *Plasmodium falciparum* 80S ribosome. We simulate noisy measurements according to the model $V_i = f_{\text{true}} + \varepsilon_i$, where each ε_i is a Gaussian noise vector with zero mean and variance $\sigma^2 = 0.5$. A total of 500 noisy measurements V_i are generated. For each noisy measurement, we compute its bispectrum and average the results across all 500 instances. The coefficient vector is then estimated using Algorithm 1. The experiment is performed with a bandlimit of $\ell_{\text{max}} = 10$, and the number of radial shells R is varied from 3 to 8. The resulting recovery errors are presented in Figure 2. The recovery error is computed as

$$(19) \quad \text{Relative Error} = \frac{\|x - \hat{x}\|_F}{\|x\|_F},$$

where x denotes the ground-truth expansion coefficients, \hat{x} the recovered coefficients, and $\|\cdot\|_F$ the Frobenius norm. Interestingly, while theoretical guarantees suggest that 3 shells suffice for accurate reconstruction, the results indicate that using more than 3 shells yields improved robustness in the presence of noise.




FIGURE 2. Recovery error of the *Plasmodium falciparum* 80S ribosome as a function of the number of shells, under additive Gaussian noise.

ACKNOWLEDGMENT

T.B. and D.E. are supported in part by BSF under Grant 2020159. T.B. is also supported in part by NSF-BSF under Grant 2019752, in part by ISF under Grant 1924/21, and in part by a grant from The Center for AI and Data Science at Tel Aviv University (TAD). D.E was also partially supported by NSF grant DMS2205626.

REFERENCES

- [1] Emmanuel Abbe, Tamir Bendory, William Leeb, João M Pereira, Nir Sharon, and Amit Singer. Multireference alignment is easier with an aperiodic translation distribution. *IEEE Transactions on Information Theory*, 65(6):3565–3584, 2018.
- [2] Emmanuel Abbe, Joao M Pereira, and Amit Singer. Estimation in the group action channel. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 561–565. IEEE, 2018.
- [3] Tal Amir, Tamir Bendory, Nadav Dym, and Dan Edidin. The stability of generalized phase retrieval problem over compact groups. *arXiv preprint arXiv:2505.04190*, 2025.
- [4] Sheldon Axler, Paul Bourdon, and Wade Ramey. *Harmonic function theory*, volume 137 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2001.
- [5] Afonso S Bandeira, Ben Blum-Smith, Joe Kileel, Jonathan Niles-Weed, Amelia Perry, and Alexander S Wein. Estimation under group actions: recovering orbits from invariants. *Applied and Computational Harmonic Analysis*, 66:236–319, 2023.
- [6] Afonso S Bandeira, Moses Charikar, Amit Singer, and Andy Zhu. Multireference alignment using semi-definite programming. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pages 459–470, 2014.
- [7] Tamir Bendory, Alberto Bartesaghi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE signal processing magazine*, 37(2):58–76, 2020.
- [8] Tamir Bendory, Nicolas Boumal, Chao Ma, Zhizhen Zhao, and Amit Singer. Bispectrum inversion with application to multireference alignment. *IEEE Transactions on Signal Processing*, 66(4):1037–1050, 2017.
- [9] Tamir Bendory, Nadav Dym, Dan Edidin, and Arun Suresh. A transversality theorem for semi-algebraic sets with application to signal recovery from the second moment and cryo-EM. *arXiv preprint arXiv:2405.04354*, 2024.
- [10] Tamir Bendory and Dan Edidin. The sample complexity of sparse multireference alignment and single-particle cryo-electron microscopy. *SIAM J. Math. Data Sci.*, 6(2):254–282, 2024.
- [11] Tamir Bendory, Dan Edidin, William Leeb, and Nir Sharon. Dihedral multi-reference alignment. *IEEE Transactions on Information Theory*, 68(5):3489–3499, 2022.
- [12] Tamir Bendory, Ido Hadi, and Nir Sharon. Compactification of the rigid motions group in image processing. *SIAM Journal on Imaging Sciences*, 15(3):1041–1078, 2022.
- [13] Tamir Bendory, Oscar Michelin, and Amit Singer. Sparse multi-reference alignment: Sample complexity and computational hardness. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8977–8981. IEEE, 2022.
- [14] Tejal Bhamre, Teng Zhang, and Amit Singer. Anisotropic twicing for single particle reconstruction using autocorrelation analysis. *arXiv preprint arXiv:1704.07969*, 2017.
- [15] Ben Blum-Smith and Harm Derksen. Generic orbits, normal bases, and generation degree for fields of rational invariants. *arXiv:2506.05650*, 2025.
- [16] Ben Blum-Smith, Thays Garcia, Rawin Hidalgo, and Consuelo Rodriguez. Degree bounds for fields of rational invariants of $\mathbb{Z}/p\mathbb{Z}$ and other finite groups. *J. Pure Appl. Algebra*, 228(10):Paper No. 107693, 45, 2024.
- [17] Nicolas Boumal, Tamir Bendory, Roy R Lederman, and Amit Singer. Heterogeneous multireference alignment: A single pass approach. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2018.
- [18] Hua Chen, Mona Zehni, and Zhizhen Zhao. A spectral method for stable bispectrum inversion with application to multireference alignment. *IEEE Signal Processing Letters*, 25(7):911–915, 2018.
- [19] Harm Derksen and Gregor Kemper. *Computational invariant theory*, volume 130 of *Encyclopaedia of Mathematical Sciences*. Springer, Heidelberg, enlarged edition, 2015. With two appendices by Vladimir L. Popov, and an addendum by Norbert A’Campo and Popov, Invariant Theory and Algebraic Transformation Groups, VIII.
- [20] Matyas Domokos. Degree bound for separating invariants of abelian groups. *Proceedings of the American Mathematical Society*, 145:3695–3708, 2017.

- [21] Jan Draisma, Gregor Kemper, and David Wehlau. Polarization of separating invariants. *Canad. J. Math.*, 60(3):556–571, 2008.
- [22] Gil Drozatz, Tamir Bendory, and Nir Sharon. Provable algorithms for multi-reference alignment over $SO(2)$. *arXiv preprint arXiv:2504.19140*, 2025.
- [23] Dan Edidin and Josh Katz. The reflection invariant bispectrum: signal recovery in the dihedral model. *arXiv preprint arXiv:2408.09599*, 2024.
- [24] Dan Edidin and Josh Katz. Orbit recovery from invariants of low degree in representations of finite groups. In *Proceedings of the 2025 SampTA conference, Vienna*, 2025.
- [25] Dan Edidin and Matthew Satriano. Orbit recovery for band-limited functions. *arXiv:2306.00155*, 2023.
- [26] Zhou Fan, Roy R Lederman, Yi Sun, Tianhao Wang, and Sheng Xu. Maximum likelihood for high-noise group orbit estimation and single-particle cryo-em. *The Annals of Statistics*, 52(1):52–77, 2024.
- [27] William Fulton and Joe Harris. *Representation theory*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991. A first course, Readings in Mathematics.
- [28] Yuan Gao, Erhu Cao, David Julius, and Yifan Cheng. TRPV1 structures in nanodiscs reveal mechanisms of ligand and lipid action. *Nature*, 534(7607):347–351, 2016.
- [29] Zvi Kam. The reconstruction of structure from electron micrographs of randomly oriented particles. *Journal of Theoretical Biology*, 82(1):15–39, 1980.
- [30] Anthony W. Knap. *Lie groups beyond an introduction*, volume 140 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, second edition, 2002.
- [31] Eitan Levin, Tamir Bendory, Nicolas Boumal, Joe Kileel, and Amit Singer. 3D ab initio modeling in cryo-EM by autocorrelation analysis. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 1569–1573. IEEE, 2018.
- [32] Allen Liu and Ankur Moitra. Algorithms from invariants: Smoothed analysis of orbit recovery over $SO(3)$. *arXiv preprint arXiv:2106.02680*, 2021.
- [33] Chao Ma, Tamir Bendory, Nicolas Boumal, Fred Sigworth, and Amit Singer. Heterogeneous multireference alignment for images with application to 2D classification in single particle reconstruction. *IEEE Transactions on Image Processing*, 29:1699–1710, 2019.
- [34] Jacqueline LS Milne, Mario J Borgnia, Alberto Bartesaghi, Erin EH Tran, Lesley A Earl, David M Schauder, Jeffrey Lengyel, Jason Pierson, Ardan Patwardhan, and Sriram Subramaniam. Cryo-electron microscopy—a primer for the non-microscopist. *The FEBS journal*, 280(1):28–45, 2013.
- [35] Amelia Perry, Jonathan Weed, Afonso S Bandeira, Philippe Rigollet, and Amit Singer. The sample complexity of multireference alignment. *SIAM Journal on Mathematics of Data Science*, 1(3):497–517, 2019.
- [36] Eric F Pettersen, Thomas D Goddard, Conrad C Huang, Gregory S Couch, Daniel M Greenblatt, Elaine C Meng, and Thomas E Ferrin. UCSF Chimera—a visualization system for exploratory research and analysis. *Journal of Computational Chemistry*, 25(13):1605–1612, 2004.
- [37] Martin Turk and Wolfgang Baumeister. The promise and the challenges of cryo-electron tomography. *FEBS letters*, 594(20):3243–3261, 2020.
- [38] Abigail JI Watson and Alberto Bartesaghi. Advances in cryo-ET data processing: meeting the demands of visual proteomics. *Current Opinion in Structural Biology*, 87:102861, 2024.
- [39] Wilson Wong, Xiao-chen Bai, Alan Brown, Israel S Fernandez, Eric Hanssen, Melanie Condron, Yan Hong Tan, Jake Baum, and Sjors HW Scheres. Cryo-EM structure of the Plasmodium falciparum 80S ribosome bound to the anti-protozoan drug emetine. *Elife*, 3:e03080, 2014.
- [40] Liping Yin, Anna Little, and Matthew Hirn. Bispectrum unbiasing for dilation-invariant multi-reference alignment. *IEEE Transactions on Signal Processing*, 2024.