

Compactification of the Rigid Motions Group in Image Processing*

Tamir Bendory[†], Ido Hadi[‡], and Nir Sharon[‡]

Abstract. Image processing problems in general, and in particular in the field of single-particle cryo-electron microscopy, often require considering images up to their rotations and translations. Such problems were tackled successfully when considering images up to rotations only, using quantities which are invariant to the action of rotations on images. Extending these methods to cases where translations are involved is more complicated. Here we present a computationally feasible and theoretically sound approximate invariant to the action of rotations and translations on images. It allows one to approximately reduce image processing problems to similar problems over the sphere, a compact domain acted on by the group of three-dimensional rotations, a compact group. We show that this invariant is induced by a family of mappings deforming, and thereby compactifying, the group structure of rotations and translations of the plane, i.e., the group of rigid motions, into the group of three-dimensional rotations. Furthermore, we demonstrate its viability in two image processing tasks: multireference alignment and classification. To our knowledge, this is the first instance of a quantity that is either exactly or approximately invariant to rotations and translations of images that both rests on a sound theoretical foundation and is applicable in practice.

Key words. compactification, group-invariant classification, group invariants, multireference alignment

MSC codes. 00A69, 22E70, 65D18, 68U05

DOI. 10.1137/21M1429448

1. Introduction. In image processing applications, there is often a need to consider images up to rotation, up to translation, or up to both. Two examples of this are found in multi-reference alignment (MRA) and image classification problems. In one of its more well-studied forms, the aim in MRA is to estimate an image up to rotation from a dataset of noisy rotated copies of it [3, 4, 17, 19]. In classification problems of the type we consider in this paper, the aim is to estimate the similarity up to rotation between every pair of images in a noisy dataset. In both cases, rather than using rotations, the problems can be considered over translations or both rotations and translations.

Rotations and translations of the plane form a group, $SE(2)$, usually referred to as rigid motions or the special Euclidean group. This group is noncompact with an algebraic structure of a semidirect product of the translation and rotation groups. The group structure enables

*Received by the editors June 25, 2021; accepted for publication (in revised form) April 13, 2022; published electronically July 14, 2022.

<https://doi.org/10.1137/21M1429448>

Funding: The work of the authors was supported by NSF-BSF award 2019752. The work of the first author was also partially supported by the Zimin Institute for Engineering Solutions Advancing Better Lives, BSF grant 2020159, and ISF grant 1924/21. The work of the third author was also partially supported by BSF grant 2018230.

[†]School of Electrical Engineering, Faculty of Engineering, Tel Aviv University, Tel Aviv Israel (bendory@tauex.tau.ac.il).

[‡]School of Mathematical Sciences, Faculty of Exact Sciences, Tel Aviv University, Tel Aviv Israel (idohadi@mail.tau.ac.il, nsharon@tauex.tau.ac.il).

one to formalize the notion of considering an image up to rotation, translation, or both in group-theoretic terms. In particular, an orbit of an image consists of all shifted and rotated copies of the image. Thus, the group action induces an equivalence relation on any set of images where the orbits act as the equivalence classes. Therefore, we are often required to consider only the orbits' representatives.

Using this group-theoretic formalism, an MRA problem amounts to estimating an orbit, or more accurately a representative of an orbit, from noisy samples of members of that orbit. Classification problems amount to estimating the minimal distance between orbits from noisy samples of members of these orbits. Formalizing these problems in this way enables well-developed concepts from group theory to be brought to bear on them. In particular, it leads to a wider applicability of what we refer to here as the invariants approach to orbit-related problems. Central to this approach are quantities that are invariant to the action of a group on an image. Formally, these invariants are operators on images that are constant on every orbit. It is trivial to find such invariant operators. For example, every operator that is constant on images is also invariant to any action by any group on it. However, combining insights from group theory and harmonic analysis on groups, a wide variety of nontrivial invariants were derived which not only are constant on every orbit, but also uniquely determine an orbit [12, 6, 14, 17, 3, 4, 19]. These invariants can be thought of as injective functions on the set of orbits of a group action and therefore essentially characterize the set of orbits of the action of a group on an image.

Using these invariants yielded a computationally inexpensive solution to the MRA [17] and classification [30] problems with respect to the rotation group. In both cases, an $SO(2)$ -invariant was evaluated on every image in the dataset, resulting in an $SO(2)$ -invariant representation of every image in the dataset. In [17], the MRA problem was handled by using these representations to estimate the $SO(2)$ -invariant representation of the underlying image, and then inverted to yield an estimate of the image itself, up to rotation. In [30], the classification problem was tackled by taking the distance between pairs of images to be the distance between their $SO(2)$ -invariant representations.

In both [17] and [30], as well as in our work, the motivating real-world application is single-particle reconstruction in cryogenic electron microscopy (cryo-EM). In cryo-EM, the goal is to reconstruct the three-dimensional (3D) structure of biological macromolecules from images generated by an electron microscope. While [17] and [30] handled MRA and classification, respectively, up to rotation only, in cryo-EM, as in other applications, there is often a need to consider images up to both rotations and translations [5]. In particular, the images used to estimate the 3D structure can be considered up to rotations and translations. One would therefore like to generalize the approach of [17] and [30] by using an orbit-characterizing $SE(2)$ -invariant.

Unfortunately, applying the invariant approach to MRA and classification up to both rotation and translations is considerably harder. Deriving computationally feasible orbit-characterizing invariants for the rigid motions group $SE(2)$ is hard, since this group is hard to analyze algebraically for two reasons. First, it is noncompact and harmonic analysis over noncompact groups is considerably more complicated than over compact groups, such as $SO(n)$, the rotation group of $\mathbb{R}^n$. Second, as a semidirect product it has a rather complicated algebraic structure. Thus, so far as we know, no computationally feasible orbit-characterizing invariants with sound theory behind them are known for $SE(2)$.

Since working over compact groups is simpler, both theoretically and in practice, we seek to compactify orbit-related problems over the rigid motions group. Specifically, we wished to approximately reduce orbit-related image processing problems over $SE(2)$ to similar problems over compact groups. This approach has worked in the past. A heuristically derived approximate $SE(2)$ -invariant was applied by Kondor [15] to an image classification problem. The invariant was computed by projecting images from the plane onto the sphere and computing an $SO(3)$ -invariant on the sphere. Similarly, in [18] a synchronization problem over $SE(2)$ itself, a problem where the objective was to estimate elements of $SE(2)$, was approximately reduced to a problem over the compact group $SO(3)$ (for more information on synchronization, see [7, 20, 22]).

Here we extend the theoretical approach of [18] to derive and theoretically justify the use of the approximate invariant used by [15]. We show the projection used by [15] is induced by contraction maps, a family of mappings from $SE(2)$ to $SO(3)$ that rigorously formalize the notion of continuously deforming one group into another. We then prove this projection maps an orbit of the action of $SE(2)$ on images to approximately the same orbit of the action of $SO(3)$ on the space of functions on the sphere. We then apply our approximate $SE(2)$ -invariant to MRA and classification problems over $SE(2)$ and show they provide viable solutions on simulated data, thus demonstrating its potential. Our MATLAB code to reproduce all numerical experiments is available at <https://github.com/idohadi/CompactificationImageProcessing>. To our knowledge, we provide here the first instance of a quantity that is either exactly or approximately $SE(2)$ -invariant on images that both rests on a sound theoretical foundation and is computable in practice.

1.1. Structure of the paper. We begin in section 2 by constructing a projection of compactly supported functions on the plane onto functions on the sphere and use it to define an approximate $SE(2)$ invariant. After deriving the theory, we discuss in section 3 several computational aspects of projecting discrete images onto the sphere and of calculating their spherical bispectrum. In particular, we demonstrate numerically that our approximate $SE(2)$ invariant is deserved to be named so. In the rest of the paper, we apply our approximate $SE(2)$ invariant to two problems of interest. First, in section 4, we flesh out a modified invariants approach to MRA using our approximate $SE(2)$ invariant. We perform several numerical experiments demonstrating our approach is viable. Second, in section 5 we discuss image classification up to rotation and translation. We use our approximate $SE(2)$ invariant to derive an approximately $SE(2)$ -invariant measure of the distance between images. We then demonstrate by numerical experiments that our approach is able to identify similar images up to rotations and translations and can do so even for large translation sizes. For both MRA and classification, we show that our approach outperforms methods that take only rotations into account.

2. Compactification of functions on the plane: Theoretical aspects. Given a function on the plane, we wish to define a function on the sphere which preserves, in some sense, two key properties of the original function. First, the original function needs to be recoverable from its spherical counterpart. Second, the group action of $SE(2)$ on the original function needs to be reflected, in some sense, in the group action of $SO(3)$ on its spherical counterpart.

In this section, we construct a mapping between functions on the plane and functions on the sphere which has these two properties. Center to this construction is a special family of maps, called contraction maps, between the Lie algebras of $SE(2)$ and $SO(3)$. Because the subgroup structure of a Lie group is recapitulated in its Lie algebra, contraction maps provide a sense in which one can continuously deform one Lie group into another [8]. We use these maps to construct and study the properties of the mapping described above.

Subsection 2.1 introduces the notation used throughout the paper. In particular, we introduce explicitly the group-theoretic formalization of considering images up to rotations, translations, or both. **subsection 2.2** introduces contraction maps in general and the family of contraction maps between $SE(2)$ and $SO(3)$ we use in particular. In **subsection 2.3**, we use this family of contraction maps to induce a map between functions on the plane and functions on the sphere. We also provide an explicit formula for this map and prove that the original function on the plane can be recovered from its spherical counterpart. In **subsection 2.4**, we study the manner in which the group action on the sphere is related to the original group action on the plane. In particular, we show that a member of the orbit of the original image is mapped to a neighborhood of the orbit of its spherical counterpart. This neighborhood gets at most exponentially larger with the size of the translational part of the element of $SE(2)$ acting on the original function. In the final section, **subsection 2.5**, we introduce the spherical bispectrum of a function on the sphere, an invariant to the action of $SO(3)$ that uniquely determines functions on the sphere, up to an action of $SO(3)$. Taken together, we show that when it is evaluated on images projected onto the sphere, the spherical bispectrum is an approximate $SE(2)$ invariant.

2.1. Notation. Throughout the paper, lowercase letters are scalars, lowercase boldface letters are column vectors, and uppercase boldface letters are matrices. The standard Euclidean norm is $\|\cdot\| = \|\cdot\|_2$, where $\|\mathbf{x}\|_2 = \sqrt{\sum_{j=1}^n |x_j|^2}$ for $\mathbf{x} = (x_1, x_2, \dots, x_n)^\top$ in $\mathbb{R}^n$ or in $\mathbb{C}^n$. The open r -ball in $\mathbb{R}^2$ is $B_r := \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\| < r\}$. The closed r -ball is denoted by $\bar{B}_r := \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\| \leq r\}$. Finally, $\partial B_r := \{\|\mathbf{x}\| \in \mathbb{R}^2 \mid \|\mathbf{x}\| = r\}$.

The elementwise complex conjugate of an $n \times m$ matrix $\mathbf{M}$ is denoted by $\mathbf{M}^*$. Thus, for example, $\mathbf{x}^* = (x_1^*, x_2^*, \dots, x_n^*)^\top$ for every $\mathbf{x} \in \mathbb{C}^n$. The transpose of an $n \times m$ matrix $\mathbf{M}$ is denoted by $\mathbf{M}^\top$ and its adjoint is denoted by $\mathbf{M}^\dagger$. Note that $\mathbf{M}^\dagger = (\mathbf{M}^*)^\top$.

The sphere is the unit 2-sphere, $S^2 := \{\mathbf{x} \in \mathbb{R}^3 \mid \|\mathbf{x}\| = 1\}$. The north pole of the sphere is denoted by $\mathbf{n} := (0, 0, 1)^\top$. The $d_{S^2}(\mathbf{x}_1, \mathbf{x}_2) = \arccos(\mathbf{s}_2^\top \mathbf{s}_1)$ is the great-circle distance metric on the sphere. $SO(3)$ acts on a function on the sphere $f : S^2 \rightarrow \mathbb{C}$ by $R \bullet f(\mathbf{s}) = f(R^\top \mathbf{s})$.

$SO(2)$ is the group of rotations of $\mathbb{R}^2$ and $SO(3)$ is the group of rotations of $\mathbb{R}^3$ or equivalently of S^2 . The group of rigid motions, also referred to here as the group of rotations and translations of the plane, is denoted by $SE(2) = \mathbb{R}^2 \rtimes SO(2)$. If $g = (\mathbf{x}, R) \in SE(2)$, we refer to $\mathbf{x}$ as the translational part of g and to R as the rotational part.

As we mentioned in the introduction, rotations and translations of the plane form a group, enabling one to formalize the notion of considering an image up to rotation, translation, or both in group-theoretic terms. If G is a (multiplicative) group, we say it acts on a set X if every $g \in G$ induces a map $X \rightarrow X$ denoted by $x \mapsto g \bullet x$ such that $g_1 \bullet (g_2 \bullet x) = (g_1 g_2) \bullet x$ for all $g_1, g_2 \in G$ and $x \in X$. An orbit of the action of G on X is the set $G \bullet x = \{g \bullet x \mid g \in G\}$ of all elements of X that can be reached by applying G to x . It is easy to prove that $G \bullet x_1 = G \bullet x_2$

for $x_1, x_2 \in X$ if and only if $x_1 = g \bullet x_2$ for some $g \in G$. Therefore, the action of a group on a set induces an equivalence relation on the set with the orbits functioning as the equivalence classes. Considering an image up to rotation, translation, or both amounts to considering it as a representative of its orbit with respect to the action of the rotation group, the translation group, or the translation and rotation group, respectively, on the space of images.

In order to formalize the latter statement, we model images, as is conventional, as compactly supported, square-integrable real-valued functions on the plane and denote by $L_c^2(\mathbb{R}^2)$ the space of such functions. The rotation group of the plane $SO(2)$ acts on $L_c^2(\mathbb{R}^2)$ by $f \mapsto R \bullet f$, where $R \bullet f(\mathbf{x}) = f(R^\top \mathbf{x})$. The (additive) translation group of the plane $\mathbb{R}^2$ acts on $L_c^2(\mathbb{R}^2)$ by $f \mapsto \mathbf{b} \bullet f$, where $\mathbf{b} \bullet f(\mathbf{x}) = f(\mathbf{x} - \mathbf{b})$. The group of translations and rotations of the plane $SE(2)$, usually referred to as the group of rigid motions of the plane or the special Euclidean group of the plane, has a more complicated algebraic structure. It is the semidirect product of the translation group and the rotation group, $SE(2) = \mathbb{R}^2 \rtimes SO(2)$, and so its group operation is $(\mathbf{b}_1, R_1)(\mathbf{b}_2, R_2) = (\mathbf{b}_1 + R_1 \mathbf{b}_2, R_1 R_2)$. This group acts on $\mathbb{R}^2$ by $(\mathbf{b}, R) \bullet \mathbf{x} = R\mathbf{x} + \mathbf{b}$ and on $C_c(\mathbb{R}^2)$ by $f \mapsto (\mathbf{b}, R) \bullet f$, where $(\mathbf{b}, R) \bullet f(\mathbf{x}) = f((\mathbf{b}, R)^{-1} \mathbf{x})$. The inverse of $(\mathbf{b}, R) \in SE(2)$ can readily be shown to be $(\mathbf{b}, R)^{-1} = (-R^\top \mathbf{b}, R^\top)$.

2.2. Contraction maps. A popular construction in physics obtains a Lie algebra as the limit of a sequence of Lie algebras. This construction was first made rigorous in a landmark paper [11]. Dooley [8] gave a coordinate-free formulation of this construction using contraction maps. This section follows his approach.

Consider a semisimple compact Lie group G and its Lie algebra $\mathfrak{g}$ with Lie brackets $[\cdot, \cdot]$. Let $G_0 \leq G$ be a continuous subgroup of G and denote its Lie algebra by $\mathfrak{g}_0$. Since $G_0 \leq G$, Lie correspondence ensures that $\mathfrak{g}_0$ is a Lie subalgebra of $\mathfrak{g}$. If $\mathfrak{g} = \mathfrak{g}_0 \oplus V$, where

$$[\mathfrak{g}_0, V] = \{[g_0, v] \mid g_0 \in \mathfrak{g}_0, v \in V\} \subset V,$$

define for every $\lambda > 0$ the contraction map of G with respect to G_0 :

$$\Psi_\lambda : V \rtimes G_0 \rightarrow G \text{ such that } (v, g_0) \mapsto \exp_G\left(\frac{v}{\lambda}\right) g_0.$$

As noted in [8], treated as distinct groups, both $V \rtimes G_0$ and G have $\mathfrak{g}$ as the underlying vector space of their Lie algebra. Yet, one can obtain the Lie brackets of $V \rtimes G_0$ by transforming the Lie brackets of G using Ψ_λ and taking the limit $\lambda \rightarrow \infty$. In this sense, the family of the contraction maps $\{\Psi_\lambda \mid \lambda > 0\}$ enables one to continuously deform one Lie algebra to produce another in the limit.

In our case, $G = SO(3)$ and $G_0 = SO(2)$, where $\mathfrak{so}(3)$ is the space of skew-symmetric 3×3 real matrices equipped with the commutator as the Lie bracket, that is, the space of real 3×3 matrices satisfying $S^\top = -S^\top$ with the Lie bracket $[S_1, S_2] = S_1 S_2 - S_2 S_1$. Therefore, as a vector space $\mathfrak{so}(3) = \text{Span}\{S_{i,j} \mid 1 \leq i < j \leq 3\}$, where $S_{i,j}$ is a real skew-symmetric 3×3 matrix with 1 and -1 at the (i, j) and (j, i) elements, respectively, and zero otherwise. Denoting $S_1 = S_{1,2}$, $S_2 = S_{1,3}$, and $S_3 = S_{2,3}$, one can easily show that

$$[S_1, S_2] = -S_3, \quad [S_2, S_3] = -S_1, \quad \text{and} \quad [S_1, S_3] = S_2.$$

Thus, $\mathfrak{g}_0 = \mathfrak{so}(2) \cong \text{Span}\{S_1\}$ and $V \cong \text{Span}\{S_2, S_3\} \cong \mathbb{R}^2$. Therefore, the contraction map of $\text{SO}(3)$ with respect to $\text{SO}(2)$ is

$$(2.1) \quad \Psi_\lambda : \text{SE}(2) \rightarrow \text{SO}(3) \text{ such that } (\mathbf{x}, R) \mapsto \exp_{\text{SO}(3)}\left(\frac{\mathbf{x}}{\lambda}\right) R.$$

Crucially for what follows, this contraction map is smooth, because the exponential map and the group operation map $\text{SO}(3) \times \text{SO}(3) \rightarrow \text{SO}(3)$ are all smooth.

In [18], the contraction map (2.1) was used to reduce a synchronization problem over $\text{SE}(n)$ to a synchronization problem over $\text{SO}(n)$. In particular, [18, Prop. 3.4] shows that (2.1) can be treated as an approximate group homomorphism. While their result holds for a general category of groups, we state it here for $\text{SO}(3)$.

Theorem 2.1. *Let $g_j = (\mathbf{x}_j, R_j) \in \text{SE}(2)$ ($j = 1, 2$). Then $\Psi_\lambda(g^{-1}) = \Psi_\lambda(g)^{-1}$ for all $g \in \text{SE}(2)$ and if also $\|\mathbf{x}_1\| + \|\mathbf{x}_2\| \leq 0.59\lambda$, then for all $g_1, g_2 \in \text{SE}(2)$*

$$(2.2) \quad \|\Psi_\lambda(g_1 g_2) - \Psi_\lambda(g_1) \Psi_\lambda(g_2)\|_F \leq \frac{C}{\lambda^2}, \text{ as } \lambda \rightarrow \infty,$$

while C is a constant independent of λ , but dependent on g_1 and g_2 . Here $\|\cdot\|_F$ is the Frobenius norm and the codomain of Ψ_λ is explicitly the 3×3 real orthogonal matrices with determinant 1.

In the following sections, we extend this result to functions on homogeneous spaces of $\text{SE}(2)$ and $\text{SO}(3)$. In our extension, Theorem 2.9, the dependence on λ is no longer asymptotic and it is no longer required that $\|\mathbf{x}_1\| + \|\mathbf{x}_2\| \leq 0.59\lambda$.

2.3. A mapping between functions on the plane and functions on the sphere.

In this section, we prove that the contraction map (2.1) between $\text{SE}(2)$ and $\text{SO}(3)$ induces an invertible mapping between function spaces on homogeneous spaces of these groups. A homogeneous space is a space a group acts on transitively, i.e., where every two elements in the space are related by the action of an element of the group. Every homogeneous space H of a some group G is $H \cong G/G_0$ for some subgroup $G_0 \leq G$ (see, e.g., [10, Chap. II(4)]). In particular, $\mathbb{R}^2 \cong \text{SE}(2)/\text{SO}(2)$ and $S^2 \cong \text{SO}(3)/\text{SO}(2)$. Denote their quotient maps by $\eta_{\text{SE}(2)} : \text{SE}(2) \rightarrow \mathbb{R}^2$ and $\eta_{\text{SO}(3)} : \text{SO}(3) \rightarrow S^2$, respectively. It is important to note that in the latter case, we identify $\text{SO}(2)$ with the subgroup of $\text{SO}(3)$ preserving the north pole $\mathbf{n} = (0, 0, 1)^\top$. Therefore, we write $\eta_{\text{SE}(2)}(\mathbf{x}, R) = \mathbf{x}$ and $\eta_{\text{SO}(3)}(R) = R\mathbf{n}$.

Using this notation, we prove the existence of a projection of compactly supported functions on the plane onto functions on the sphere, or more informally, a projection of images from the plane onto the sphere. This projection is shown to be induced by the contraction map (2.1). Stated formally, we prove the following.

Theorem 2.2. *Fix $\lambda > 0$. If $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a smooth function compactly supported within $B_{\lambda\pi} = \{\mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x}\| < \lambda\pi\}$, then the following hold:*

- **EXISTENCE.** *There is a smooth function $\kappa_\lambda f : S^2 \rightarrow \mathbb{R}$ satisfying*

$$(2.3) \quad \kappa_\lambda f \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda = f \circ \eta_{\text{SE}(2)}$$

or, equivalently, that the diagram in Figure 1 commutes.

- **UNIQUENESS.** *This $\kappa_\lambda f$ is the only smooth function on the sphere satisfying (2.3).*




Figure 1. Diagram of Theorem 2.2. There are two equivalent paths from $\mathrm{SE}(2)$ to $\mathbb{R}$. The blue path passes through a function on the plane. We complete the green path by proving the existence of $\kappa_\lambda f$ (dashed).

Throughout the paper, we refer to $\kappa_\lambda f$ as the projection of f onto the sphere and to λ as the scaling parameter of the projection.

Figure 1 can be taken as a visual demonstration of Theorem 2.2. In it, we see two equivalent paths from $\mathrm{SE}(2)$ to $\mathbb{R}$. The blue path passes through the quotient map and a given function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$. In Theorem 2.2, we prove the existence and uniqueness of the green path, the one passing through $\mathrm{SO}(3)$ and S^2 . Most of this latter path is already given to us. We use the properties of the contraction map Ψ_λ and the quotient map $\mathrm{SO}(2) \rightarrow S^2$ to complete the green path.

Remark 2.3. In this paper, smooth functions are C^∞ functions and all smooth manifolds are also C^∞ . Theorem 2.2 and several other results in this paper are proved only for smooth functions. While this limits their generality, it is in fact consistent with our primary motivation. In cryo-EM, the potential density maps are all smooth. Therefore, restricting ourselves to smooth functions allows us to leverage the well-developed and well-known theory of smooth manifolds, while still ensuring the results retain their potential applicability.

Our proof of Theorem 2.2 relies on several lemmas. We state the lemmas here and leave their proofs to Appendix A. The first lemma provides an important property of the quotient map $\eta_{\mathrm{SE}(2)}$ and of $\eta_{\mathrm{SO}(3)} \circ \Psi_\lambda$. Recall that a function $f : M \rightarrow N$ between smooth manifolds M and N is a smooth submersion if it is a smooth function with surjective differential at every point in its domain; that is, $d_p f : T_p M \rightarrow T_{f(p)} N$ is surjective for every $p \in M$, where $T_p M$ and $T_{f(p)} N$ are the tangent spaces of M and N at p and $f(p)$, respectively.

Lemma 2.4. *The maps $\eta_{\mathrm{SE}(2)}$ and $\eta_{\mathrm{SO}(3)} \circ \Psi_\lambda$ are smooth submersions.*

Since $\eta_{\mathrm{SO}(3)}(R) = R\mathbf{n}$, the mapping $\eta_{\mathrm{SO}(3)} \circ \Psi_\lambda$ satisfies

$$(2.4) \quad \eta_{\mathrm{SO}(3)} \circ \Psi_\lambda(\mathbf{x}, R) = \exp_{\mathrm{SO}(3)}(\mathbf{x})R\mathbf{n} = \exp_{\mathrm{SO}(3)}(\mathbf{x})\mathbf{n},$$

where the last transition follows from our identification of $\mathrm{SO}(2) \leq \mathrm{SO}(3)$ with the subgroup of rotations preserving $\mathbf{n}$. Overall, it follows that $\eta_{\mathrm{SO}(3)} \circ \Psi_\lambda(\mathbf{x}, R)$ depends only on $\mathbf{x}$ and can be treated as a function mapping the plane onto the sphere. Figure 2 visually demonstrates how this mapping works, whereas Lemma 2.5 formally characterizes its fibers, the preimage of singletons.




Figure 2. The mapping $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ can be seen as mapping lines through the origin of the plane (magenta) onto great circles through the north pole on the sphere (magenta). Consequently, the origin and the red, dashed circles are mapped to the north pole and the blue, dotted circles are mapped to the south pole.

Lemma 2.5. Fix $\lambda > 0$. Let $\mathbf{s} \in S^2 \setminus \{\pm \mathbf{n}\}$. Denote $\tilde{\mathbf{x}} = (\phi \cos \theta, \phi \sin \theta)^\top$, where (θ, ϕ) are the spherical coordinates of $\mathbf{s}$. The fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ are

(2.5)

$$(\eta_{\text{SO}(3)} \circ \Psi_\lambda)^{-1}(\mathbf{n}) = \{(\mathbf{x}, R) \in \text{SE}(2) \mid \|\mathbf{x}\| = 2n\lambda\pi \text{ for } n \in \mathbb{N} \cup \{0\}, R \in \text{SO}(2)\},$$

(2.6)

$$(\eta_{\text{SO}(3)} \circ \Psi_\lambda)^{-1}(-\mathbf{n}) = \{(\mathbf{x}, R) \in \text{SE}(2) \mid \|\mathbf{x}\| = (2n+1)\lambda\pi \text{ for } n \in \mathbb{N} \cup \{0\}, R \in \text{SO}(2)\},$$

(2.7)

$$(\eta_{\text{SO}(3)} \circ \Psi_\lambda)^{-1}(\mathbf{s}) = \{(\mathbf{x}, R) \in \text{SE}(2) \mid \mathbf{x} = (2n\pi + 1)\lambda\tilde{\mathbf{x}} \text{ for } n \in \mathbb{Z}, R \in \text{SO}(2)\}.$$

The final lemma enables us to lift a smooth, compactly supported function on the plane onto $\text{SE}(2)$ while respecting the fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$.

Lemma 2.6. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a smooth function. If f is compactly supported within $B_{\lambda\pi}$, then there is a unique smooth function $\tilde{f} : \text{SE}(2) \rightarrow \mathbb{R}$ which satisfies the following conditions:

- (i) $\tilde{f}$ is constant on fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$.
- (ii) $f \circ \eta_{\text{SE}(2)} = \tilde{f}$ on $B_{\lambda\pi} \times \text{SO}(2)$.

We begin by proving the existence of the projection.

Proof of EXISTENCE in Theorem 2.2. Let $\tilde{f}$ be the extension of f described in Lemma 2.6. By Lemma 2.6(i), it is constant on fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$. Because $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ is a smooth submersion by Lemma 2.4, there is $\tilde{F} : S^2 \rightarrow \mathbb{R}$ satisfying $\tilde{f} = \tilde{F} \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda$ (see Theorem 4.30 in [16]). By Lemma 2.6(ii), it follows that $\tilde{F} \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda = f \circ \eta_{\text{SE}(2)}$. We therefore take $\kappa_\lambda f := \tilde{F}$. ■

The projection we constructed in the [existence proof](#) of [Theorem 2.2](#) has an explicit, easy-to-compute formula.

Proposition 2.7. *If $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a smooth compactly supported function supported within $B_{\lambda\pi}$, then*

$$(2.8) \quad \kappa_\lambda f(\theta, \phi) = f(\lambda\theta \cos \phi, \lambda\theta \sin \phi).$$

Formula (2.8) is virtually identical to the one used by Kondor [15]. While Kondor chose it based on a heuristic, we here showed that it is derived from the contraction map (2.1), a more fundamental mapping between $\text{SE}(2)$ and $\text{SO}(3)$ that relates their underlying group structure.

The proof of [Proposition 2.7](#) relies on the following lemma, which provides an explicit formula for $\eta_{\text{SO}(3)} \circ \Psi_\lambda$. It is proved in [Appendix A](#).

Lemma 2.8. $\eta_{\text{SO}(3)} \circ \Psi_\lambda(\mathbf{x}, R) = \left(\frac{\mathbf{x}^\top}{\|\mathbf{x}\|} \sin\left(\frac{\|\mathbf{x}\|}{\lambda}\right), \cos\left(\frac{\|\mathbf{x}\|}{\lambda}\right) \right)^\top$ for all $(\mathbf{x}, R) \in \text{SE}(2)$.

Proof of Proposition 2.7. Let $\mathbf{s} \in S^2 \setminus \{-\mathbf{n}\}$ with spherical coordinate representation (θ, ϕ) . Let $(\mathbf{x}, R) \in \text{SE}(2)$ be a vector satisfying $\eta_{\text{SO}(3)} \circ \Psi_\lambda(\mathbf{x}, R) = \mathbf{s}$. By [Lemma 2.5](#), we can choose $\mathbf{x} \in B_{\lambda\pi}$. By [Lemma 2.8](#),

$$(2.9) \quad \mathbf{s} = \left(\frac{\mathbf{x}^\top}{\|\mathbf{x}\|} \sin\left(\frac{\|\mathbf{x}\|}{\lambda}\right), \cos\left(\frac{\|\mathbf{x}\|}{\lambda}\right) \right)^\top$$

if and only if $\theta = \frac{\|\mathbf{x}\|}{\lambda}$ and ϕ is the angle of $\mathbf{x}$ in polar coordinates. Now, for every $R \in \text{SO}(2)$ we have

$$\begin{aligned} f(\lambda\theta \cos \phi, \lambda\theta \sin \phi) &= f(\|\mathbf{x}\| \cos \phi, \|\mathbf{x}\| \sin \phi) && \text{by (2.9)} \\ &= f(\mathbf{x}) \\ &= f \circ \eta_{\text{SE}(2)}(\mathbf{x}, R) \\ &= \kappa_\lambda f \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda(\mathbf{x}, R) && \text{by Theorem 2.2} \\ &= \kappa_\lambda f(\mathbf{s}) = \kappa_\lambda f(\theta, \phi). \end{aligned}$$

We use the explicit formula in [Proposition 2.7](#) to prove uniqueness in [Theorem 2.2](#).

Proof of UNIQUENESS in Theorem 2.2. Note that the only property of $\kappa_\lambda f$ we used in the proof of [Proposition 2.7](#) is the fact it satisfies (2.3). Therefore, if $g_1, g_2 : S^2 \rightarrow \mathbb{R}$ satisfy (2.3), they are identical on $S^2 \setminus \{-\mathbf{n}\}$. Finally, it is easy to see that on $-\mathbf{n}$ both g_1 and g_2 must equal zero, since the support of f does not intersect the fiber of $-\mathbf{n}$ (2.6). We conclude $g_1 = g_2$, and therefore there is exactly one function on the sphere satisfying (2.3). ■

2.4. The group action approximation theorem. As shown above, the contraction map (2.1) is a mapping between groups which induces a mapping between function spaces these groups act on. Therefore, it makes sense to try to capture the way the action of $\text{SE}(2)$ on its corresponding function space is approximated by the action of $\text{SO}(3)$ on its corresponding function space by how well the mapping between function spaces commutes with the group action.

Theorem 2.1 of [18] does exactly that for the approximation of $\text{SE}(2)$ by $\text{SO}(3)$. We model our approach similarly to theirs and prove that

$$\kappa_\lambda((\mathbf{x}, R) \bullet f) \approx \Psi_\lambda((\mathbf{x}, R)) \bullet \kappa_\lambda f \quad \text{for all } \mathbf{x} \in \mathbb{R}^2.$$

The following theorem makes this idea precise.

Theorem 2.9. *Let $\tilde{\lambda} \geq \lambda \geq 1$, let $(\mathbf{b}, R) \in \text{SE}(2)$, and let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a smooth function compactly supported within $B_{\lambda\pi}$. Then,*

(i) *if $(\mathbf{b}, R)$ is a purely rotational element of $\text{SE}(2)$, i.e., $\mathbf{b} = \mathbf{0}$, then*

$$(2.10) \quad \kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f) = \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f;$$

(ii) *assume $(\mathbf{b}, R)$ is such that $(\mathbf{b}, R) \bullet f$ is supported within $B_{\lambda\pi}$ and $\Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f$ is supported within $\eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}(B_{\lambda\pi})$. Then*

$$(2.11) \quad \|\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f) - \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f\|_2 \leq \frac{4L\pi e^{\|\mathbf{b}\| + \lambda\pi}}{\tilde{\lambda}^2} \left(1 - \cos\left(\frac{\lambda\pi}{\tilde{\lambda}}\right)\right)^{1/2},$$

where L is the Lipschitz constant of $\kappa_{\tilde{\lambda}} f$ in the sense defined in (2.20).

Remark 2.10. It is worth emphasizing that the bound in (2.11) depends on the function being projected through L , the Lipschitz constant of its projection onto the sphere.

Remark 2.11. We note an apparent difference between **Theorem 2.9** and **Theorem 2.1**. In (2.2) of **Theorem 2.1** from [18], the bound behaves as $O(\frac{1}{\lambda^2})$ when the scaling parameter satisfies $\lambda \rightarrow \infty$. In (2.11) of **Theorem 2.9**, when we hold λ fixed and take $\tilde{\lambda} \rightarrow \infty$, it appears at first glance that the bound behaves as $O(\frac{1}{\tilde{\lambda}^3})$. However, we expect that the Lipschitz constant L of the projection $\kappa_{\tilde{\lambda}} f$ is $O(\tilde{\lambda})$, and thus the bound behaves as $O(\tilde{\lambda}^2)$. To see why, note that our projection maps two fixed points on the plane to two points on the sphere, as demonstrated visually in **Figure 2**. Therefore, if we fix two points on the plane, the distance between the value of their corresponding points on the sphere does not depend on $\tilde{\lambda}$. However, as $\tilde{\lambda}$ increases, the spherical cap to which the support of f is mapped grows smaller and specifically the maximal distance between any two points in it grow smaller by $O(\frac{1}{\tilde{\lambda}})$. Thus, the Lipschitz constant of $\kappa_{\tilde{\lambda}} f$ should scale asymptotically like $L_\lambda O(\tilde{\lambda})$, where L_λ is the Lipschitz constant of $\kappa_\lambda f$. Overall, we expect the entire bound to be $O(\frac{1}{\tilde{\lambda}^2})$ as $\tilde{\lambda} \rightarrow \infty$ for fixed λ , as in **Theorem 2.1**.

Our proof of **Theorem 2.9** begins with a lemma providing a sense in which the matrix exponential approximately preserves the matrix group multiplication map.

Lemma 2.12. *If X and Y are real $n \times n$ matrices and $\lambda \geq 1$, then*

$$(2.12) \quad \left\| \exp\left(\frac{X}{\lambda} + \frac{Y}{\lambda}\right) - \exp\left(\frac{X}{\lambda}\right) \exp\left(\frac{Y}{\lambda}\right) \right\| \leq \frac{C}{\lambda^2},$$

where

$$(2.13) \quad C := 2e^{\|X\| + \|Y\|} - e^{\|X\|} - e^{\|Y\|} - \|X\| - \|Y\|.$$

We particularize [Lemma 2.12](#) to our contraction map between $\text{SE}(2)$ and $\text{SO}(3)$.

Corollary 2.13. *Let $\tilde{\lambda} \geq 1$. If $(\mathbf{x}_1, R_1), (\mathbf{x}_2, R_1) \in \text{SE}(2)$, then*

$$(2.14) \quad \|\Psi_{\tilde{\lambda}}((\mathbf{x}_1, R_1)(\mathbf{x}_2, R_2)) - \Psi_{\tilde{\lambda}}(\mathbf{x}_1, R_1)\Psi_{\tilde{\lambda}}(\mathbf{x}_2, R_2)\| \leq \frac{C(\|\mathbf{x}_1\|, \|\mathbf{x}_2\|)}{\tilde{\lambda}^2},$$

where

$$(2.15) \quad C(r_1, r_2) := 2e^{r_1+r_2} - e^{r_1} - e^{r_2} - r_1 - r_2.$$

Proof. Throughout this proof we assume that $\exp = \exp_{\text{SO}(3)}$. Since we identify $\mathbb{R}^2 \cong \text{Span}\{S_1, S_2\}$ (see [subsection 2.2](#)) and the matrix exponential commutes with conjunction, for every $R \in \text{SO}(2)$ and $\mathbf{x} \in \mathbb{R}^2$ we have

$$(2.16) \quad \exp(R^\top \mathbf{x}) = \bar{R}^\top \exp(\mathbf{x}) \bar{R} \text{ with } \bar{R} := \begin{bmatrix} R & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix}.$$

Now, denote the left-hand side of (2.14) by M . From (2.1) and the definition of $\bar{R}$ in (2.16), it follows that

$$\Psi_{\tilde{\lambda}}((\mathbf{x}_1, R_1)(\mathbf{x}_2, R_2)) = \Psi_{\tilde{\lambda}}(\mathbf{x}_1 + R_1 \mathbf{x}_2, R_1 R_2) = \exp\left(\frac{\mathbf{x}_1 + R_1 \mathbf{x}_2}{\tilde{\lambda}}\right) \bar{R}_1 \bar{R}_2$$

and

$$\Psi_{\tilde{\lambda}}(\mathbf{x}_1, R_1)\Psi_{\tilde{\lambda}}(\mathbf{x}_2, R_2) = \exp\left(\frac{\mathbf{x}_1}{\tilde{\lambda}}\right) \bar{R}_1 \exp\left(\frac{\mathbf{x}_2}{\tilde{\lambda}}\right) \bar{R}_2.$$

Since the matrix 2-norm is rotation-invariant, it follows from the computations above that

$$M = \left\| \exp\left(\frac{\mathbf{x}_1 + R_1 \mathbf{x}_2}{\tilde{\lambda}}\right) - \exp\left(\frac{\mathbf{x}_1}{\tilde{\lambda}}\right) \bar{R}_1 \exp\left(\frac{\mathbf{x}_2}{\tilde{\lambda}}\right) \bar{R}_1^\top \right\|.$$

From (2.16), it now follows that

$$M = \left\| \exp\left(\frac{\mathbf{x}_1 + R_1 \mathbf{x}_2}{\tilde{\lambda}}\right) - \exp\left(\frac{\mathbf{x}_1}{\tilde{\lambda}}\right) \exp\left(\frac{R_1 \mathbf{x}_2}{\tilde{\lambda}}\right) \right\|.$$

Taking $X = \mathbf{x}_1$, $Y = \mathbf{x}_2$, it follows from [Lemma 2.12](#) that $M \leq \frac{C}{\tilde{\lambda}^2}$ with

$$\begin{aligned} C &= 2e^{\|\mathbf{x}_1\| + \|R_1 \mathbf{x}_2\|} - e^{\|\mathbf{x}_1\|} - e^{\|R_1 \mathbf{x}_2\|} - \|\mathbf{x}_1\| - \|R_1 \mathbf{x}_2\| \\ &= 2e^{\|\mathbf{x}_1\| + \|\mathbf{x}_2\|} - e^{\|\mathbf{x}_1\|} - e^{\|\mathbf{x}_2\|} - \|\mathbf{x}_1\| - \|\mathbf{x}_2\|. \end{aligned}$$

■

[Corollary 2.13](#) is an analogue of [Theorem 2.1](#) of [18]. It differs from it in three key respects. First, the requirement that the translational parts satisfy an inequality constraint is dropped. Second, the dependence on the scaling parameters λ and $\tilde{\lambda}$ is no longer asymptotic. Third, the dependence of the bound on the elements of $\text{SE}(2)$ is explicit.

Before proving [Theorem 2.9](#), we introduce another auxiliary lemmas, which we prove in [Appendix A](#).

Lemma 2.14. $d_{S^2}(\mathbf{x}, \mathbf{y}) \leq \sqrt{\pi} \|\mathbf{x} - \mathbf{y}\|$.

We are now ready to prove our main result, [Theorem 2.9](#).

Proof of Theorem 2.9. Recall that $\eta_{\text{SO}(3)}(R) = R\mathbf{n}$. Denote

$$(2.17) \quad \begin{aligned} K^2 &:= \|\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f) - \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f\|_2^2 \\ &= \int_{S^2} |\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f)(\mathbf{x}) - \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f(\mathbf{x})|^2 dS^2(\mathbf{x}). \end{aligned}$$

Fix $\mathbf{x} \in S^2$ and let $\mathbf{b}_1 \in B_{\tilde{\lambda}\pi}$ such that $\eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}(\mathbf{b}_1, I) = \mathbf{x}$. We focus on two terms we subtract from one another in the integrand of (2.17). Unraveling the definitions, it follows that the first term satisfies

$$(2.18) \quad \begin{aligned} \kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f)(\mathbf{x}) &= (\mathbf{b}, R) \bullet f(\mathbf{b}_1) \\ &= f\left((\mathbf{b}, R)^{-1} \mathbf{b}_1\right) \\ &= f \circ \eta_{\text{SE}(2)}\left((\mathbf{b}, R)^{-1}(\mathbf{b}_1, I)\right) \\ &= \kappa_{\tilde{\lambda}} f \circ \eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}\left((\mathbf{b}, R)^{-1}(\mathbf{b}_1, I)\right) \\ &= \kappa_{\tilde{\lambda}} f\left(\Psi_{\tilde{\lambda}}\left((\mathbf{b}, R)^{-1}(\mathbf{b}_1, I)\right) \mathbf{n}\right). \end{aligned}$$

Similarly, the second term satisfies

$$(2.19) \quad \begin{aligned} \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f(\mathbf{x}) &= \kappa_{\tilde{\lambda}} f((\exp(\mathbf{b})R)^\top \exp(\mathbf{b}_1)\mathbf{n}) \\ &= \kappa_{\tilde{\lambda}} f(\Psi_{\tilde{\lambda}}((\mathbf{b}, R))^{-1} \exp(\mathbf{b}_1)\mathbf{n}) \\ &= \kappa_{\tilde{\lambda}} f\left(\Psi_{\tilde{\lambda}}\left((\mathbf{b}, R)^{-1}\right) \Psi_{\tilde{\lambda}}((\mathbf{b}_1, I)\mathbf{n})\right), \end{aligned}$$

where the last transition follows from the first half of [Theorem 2.1](#), which does not require any restrictions on λ . Now, note that $(\mathbf{b}, R) = (0, R)(\mathbf{b}, I)$ and that $\Psi_{\tilde{\lambda}}(\mathbf{b}, R) = \Psi_{\tilde{\lambda}}(0, R)\Psi_{\tilde{\lambda}}(\mathbf{b}, I)$. Therefore, substituting $\mathbf{b} = \mathbf{0}$ into (2.18) and (2.19) shows that the left-hand side of both equations are equal if $(\mathbf{b}, R)$ is a purely rotational element of $\text{SE}(2)$. Thus, the integrand in (2.17) is zero on S^2 , which proves (2.10).

In order to prove (2.11), let $\alpha = \Psi_{\tilde{\lambda}}((\mathbf{b}, R)^{-1}(\mathbf{b}_1, I))$ and $\beta = \Psi_{\tilde{\lambda}}((\mathbf{b}, R)^{-1})\Psi_{\tilde{\lambda}}(\mathbf{b}_1, I)$. Since $\kappa_{\tilde{\lambda}} f$ is a smooth function with a compact domain S^2 , it is Lipschitz in the sense that there is $L \geq 0$ satisfying

$$(2.20) \quad |\kappa_{\tilde{\lambda}} f(\mathbf{x}) - \kappa_{\tilde{\lambda}} f(\mathbf{y})| \leq L d_{S^2}(\mathbf{x}, \mathbf{y}) \text{ for all } \mathbf{x}, \mathbf{y} \in S^2,$$

where $d_{S^2}(\mathbf{x}, \mathbf{y}) = \arccos(\mathbf{y}^\top \mathbf{x})$ is the standard metric on the sphere (great-circle distance).

Therefore,

$$\begin{aligned}
 (2.21) \quad & |\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f)(\mathbf{x}) - \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f(\mathbf{x})| \\
 &= |\kappa_{\tilde{\lambda}} f(\alpha \mathbf{n}) - \kappa_{\tilde{\lambda}} f(\beta \mathbf{n})| \quad \text{By (2.18) and (2.19)} \\
 &\leq L d_{S^2}(\alpha \mathbf{n}, \beta \mathbf{n}) \quad \text{By (2.20)} \\
 &\leq L \sqrt{\pi} \|\alpha \mathbf{n} - \beta \mathbf{n}\| \quad \text{By Lemma 2.14} \\
 &\leq L \sqrt{\pi} \|\alpha - \beta\| \quad \text{Since } \|\mathbf{n}\| = 1 \\
 &\leq \frac{2L\sqrt{\pi}C(\|\mathbf{b}\|, \|\mathbf{b}_1\|)}{\tilde{\lambda}^2} \quad \text{By Corollary 2.13,}
 \end{aligned}$$

where C is defined in (2.15).

From Lemma 2.8, it is easy to show that if we treat $\eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}$ as a map from $\mathbb{R}^2$ to the sphere, it can be represented in polar coordinates as

$$\eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}(r, \phi) = \left(\cos \phi \sin \left(\frac{r}{\tilde{\lambda}} \right), \sin \phi \sin \left(\frac{r}{\tilde{\lambda}} \right), \cos \left(\frac{r}{\tilde{\lambda}} \right) \right)^\top.$$

Furthermore, it is obvious that it is a smooth parameterization of $S^2 \setminus \{-\mathbf{n}\}$ with a Gram determinant

$$\det G = \frac{1}{\tilde{\lambda}^2} \sin^2 \left(\frac{r}{\tilde{\lambda}} \right).$$

Finally, since by hypothesis $(\mathbf{b}, R) \bullet f$ is supported within $B_{\lambda\pi}$, both $\Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f$ and $\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f)$ are supported within $\eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}(B_{\lambda\pi})$. Overall, we can write (2.17) as

$$K^2 = \frac{1}{\tilde{\lambda}} \int_0^{\lambda\pi} \int_0^{2\pi} |\kappa_{\tilde{\lambda}}((\mathbf{b}, R) \bullet f)(\mathbf{y}) - \Psi_{\tilde{\lambda}}((\mathbf{b}, R)) \bullet \kappa_{\tilde{\lambda}} f(\mathbf{y})|^2 \sin \left(\frac{r}{\tilde{\lambda}} \right) d\phi dr,$$

where $\mathbf{y} = \eta_{\text{SO}(3)} \circ \Psi_{\tilde{\lambda}}(\phi, r)$. From (2.21) we obtain

$$\begin{aligned}
 K^2 &\leq \left(\frac{2L\sqrt{\pi}}{\tilde{\lambda}^2} \right)^2 \frac{1}{\tilde{\lambda}} \int_0^{\lambda\pi} \int_0^{2\pi} C(\|\mathbf{b}\|, r)^2 \sin \left(\frac{r}{\tilde{\lambda}} \right) d\phi dr \\
 &= \left(\frac{2L\sqrt{\pi}}{\tilde{\lambda}^2} \right)^2 \frac{2\pi}{\tilde{\lambda}} \int_0^{\lambda\pi} C(\|\mathbf{b}\|, r)^2 \sin \left(\frac{r}{\tilde{\lambda}} \right) dr.
 \end{aligned}$$

It is clear that C in (2.15) satisfies

$$0 \leq C(\|\mathbf{b}\|, r) \leq 2e^{\|\mathbf{b}\|+r} \leq 2e^{\|\mathbf{b}\|+\lambda\pi}$$

for $0 \leq r \leq \lambda\pi$. It follows that

$$\begin{aligned} K^2 &\leq \left(\frac{2L\sqrt{\pi}}{\tilde{\lambda}^2} \right)^2 \frac{2\pi}{\tilde{\lambda}} \cdot 2e^{2(\|\mathbf{b}\|+\lambda\pi)} \cdot \int_0^{\lambda\pi} \sin\left(\frac{r}{\tilde{\lambda}}\right) dr \\ &= \left(\frac{2L\sqrt{\pi}}{\tilde{\lambda}^2} \right)^2 \cdot 2\pi \cdot 2e^{2(\|\mathbf{b}\|+\lambda\pi)} \cdot \left(-\cos\left(\frac{r}{\tilde{\lambda}}\right) \right) \Big|_0^{\lambda\pi} \\ &= \left(\frac{2L\sqrt{\pi}}{\tilde{\lambda}^2} \right)^2 \cdot 2\pi \cdot 2e^{2(\|\mathbf{b}\|+\lambda\pi)} \cdot \left(1 - \cos\left(\frac{\lambda\pi}{\tilde{\lambda}}\right) \right). \end{aligned}$$

Taking the square root of both sides, we obtain

$$K \leq \frac{4L\pi e^{\|\mathbf{b}\|+\lambda\pi}}{\tilde{\lambda}^2} \left(1 - \cos\left(\frac{\lambda\pi}{\tilde{\lambda}}\right) \right)^{1/2}. \quad \blacksquare$$

2.5. The spherical bispectrum as an $\text{SO}(3)$ -invariant representation of functions on the sphere. [Theorem 2.9](#) suggests an approximately $\text{SE}(2)$ -invariant representation for functions on the plane: project the function onto the sphere and then compute an orbit-characterizing $\text{SO}(3)$ -invariant representation of the projected function. To that end, we present the spherical bispectrum, an orbit-characterizing $\text{SO}(3)$ -invariant representation of functions on the sphere. The theory of the spherical bispectrum was largely developed by [12, 13, 14]. These works heavily relied on tools from harmonic analysis over groups. Here we merely draw on their results and do not revisit their theoretical underpinnings. Additionally, when needed, we drew material on spherical harmonics from [2, Chap. 2].

Let $L^2(S^2)$ be the Hilbert space of all square-integrable complex-valued spherical functions equipped with the inner product

$$(2.22) \quad \langle f, g \rangle_{S^2} = \int_{S^2} f(\mathbf{x})g(\mathbf{x})^* dS^2(\mathbf{x}).$$

It induces a norm $\|f\|_{S^2} = \sqrt{\langle f, f \rangle}$. Here dS^2 is the standard volume element on S^2 . $L^2(S^2)$ is the direct sum of the spaces of spherical harmonics of arbitrary degree. To make this statement more precise, let $\mathcal{H}_\ell$ be the space of spherical harmonics of degree $\ell \in \mathbb{Z}$; that is,

$$(2.23) \quad \mathcal{H}_\ell = \text{Span} \{ Y_{\ell,m} : S^2 \rightarrow \mathbb{C} \mid m = -\ell, -\ell+1, \dots, \ell \},$$

where the spherical harmonic function of degree ℓ and order m are given by

$$(2.24) \quad Y_{\ell,m}(\theta, \phi) = \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_{\ell,m}(\cos \theta) e^{im\phi}.$$

Here, $P_{\ell,m}(x)$ is the associated Legendre function of degree ℓ and order m . The collection of all spherical harmonic functions is an orthonormal basis of $L^2(S^2)$. Also, $R \bullet \mathcal{H}_\ell \subseteq \mathcal{H}_\ell$ for all $R \in \text{SO}(3)$ and ℓ .

Fix $f \in L^2(S^2)$. Since spherical harmonics (2.24) are orthonormal with respect to inner product (2.22), we can expand f in terms of (2.24) by taking as coefficients $f_{\ell,m} = \langle f, Y_{\ell,m} \rangle_{S^2}$.

We say f has bandlimit L if $f_{\ell,m} = 0$ for all $\ell > L$. Combining the results of [12] and [14], the spherical bispectrum of f is the set of all numbers

$$(2.25) \quad b_f[\ell_1, \ell_2, \ell] = \sum_{m=-\ell}^{\ell} \sum_{m_1=-\ell_1}^{\ell_1} C_{\ell_1, m_1, \ell_2, m-m_1}^{\ell, m} f_{\ell, m} f_{\ell_1, m_1}^* f_{\ell_2, m-m_1}^*,$$

where ℓ_1 and ℓ_2 are nonnegative integers and $|\ell_1 - \ell_2| \leq \ell \leq \ell_1 + \ell_2$. The numbers $C_{\ell_1, m_1, \ell_2, m_2}^{\ell, m}$ are the Clebsch–Gordan coefficients. These are real numbers related to the underlying group structure of $\text{SO}(3)$. We detail a numerical algorithm for their calculation in subsection 3.2.

The spherical bispectrum (2.25) is intimately related to the triple correlation function of f , which is defined by

$$T_f(R_1, R_2) = \int_{\text{SO}(2)} f(\mathbf{x})^* f(R_1 \mathbf{x}) f(R_2 \mathbf{x}) d\mu(\mathbf{x}),$$

where $d\mu$ is the Haar measure on $\text{SO}(2)$. Kakarala and Mao [14, Thm. 4.1(a)] proved that calculating (2.25) is equivalent to calculating the Fourier transform on the group $\text{SO}(3) \times \text{SO}(3)$ of the triple correlation of f . Furthermore, the Fourier transform of a triple correlation of a bandlimited real-valued spherical function was shown in [13] to uniquely determine this function up to a rotation, provided it is nonzero, that is, if for every $0 \leq \ell \leq L$ there is $-\ell \leq m \leq \ell$ such that $f_{\ell, m} \neq 0$, $b_g = b_f$ if and only if $g = R \bullet f$ for some $R \in \text{SO}(3)$. Thus, the spherical bispectrum (2.25) determines the orbit of a real-valued function of bandlimit L with nonzero projection onto $\mathcal{H}_\ell$ for all $0 \leq \ell \leq L$.

The specific form of the bispectrum we use in (2.25) is owed to Kondor [15], but is easily derivable from [14, eq. (24)]. Furthermore, Kakarala and Mao [14, Thm. 4.1(b)] showed that for every (ℓ_1, ℓ_2, ℓ) we have $b_f[\ell_1, \ell_2, \ell] = \zeta b_f[\ell_2, \ell_1, \ell]$ for some $\zeta \in \mathbb{C}$ independent of f . Thus, it is sufficient to compute (2.25) for $\ell_2 \leq \ell_1$. Also, it is obvious that if $\max\{\ell_1, \ell_2, \ell\} > L$, then (2.25) is zero for functions of bandlimit L .

Overall, given a real-valued function $f \in L^2(S^2)$ of bandlimit L which has a nonzero projection over $\mathcal{H}_\ell$ for all $0 \leq \ell \leq L$, we can represent it up to rotation by its spherical bispectrum (2.25) over all triplets (ℓ_1, ℓ_2, ℓ) satisfying

$$(2.26) \quad 0 \leq \ell_1 \leq L, \quad 0 \leq \ell_2 \leq \ell_1, \quad \ell_1 - \ell_2 \leq \ell \leq \min\{L, \ell_1 + \ell_2\}.$$

This representation uniquely determines f up to rotation; that is, it characterizes the orbit $\text{SO}(3) \bullet f$.

3. Compactification of functions on the plane: Computational aspects. We developed the theory in section 2 for functions on the plane. Here we detail several computational aspects that arise when working with discrete samples of such functions. In subsection 3.1, we explain how we use discrete samples of a function on the plane to estimate the spherical harmonics coefficients of the projection of this function onto the sphere. In subsection 3.2, we explain how the spherical bispectrum (2.25) is evaluated for bandlimited functions on the sphere. Most importantly, we present a method, introduced by [21], for the evaluation of the Clebsch–Gordan coefficients. We demonstrate the accuracy of this method. We show

numerically that the spherical bispectrum also echoes the results of [Theorem 2.9](#). In particular, for sufficiently small translation sizes, the spherical bispectrum of an image and its translation are approximately the same.

3.1. Projecting an image onto the sphere. Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a function compactly supported within $[-\zeta, \zeta]^2$, where $\zeta = \cos \frac{\pi}{4} = \sin \frac{\pi}{4}$. This implies it is compactly supported within B_π . Let $G = \{(x_i, y_j)\}_{i=0, j=0}^{n-1}$ ($n > 1$) be a regular equidistant grid on $[-\zeta, \zeta]^2$:

$$(3.1) \quad x_i = -\zeta + \frac{2\zeta}{n-1}i, \quad i = 0, 1, \dots, n-1, \quad \text{and} \quad y_j = -\zeta + \frac{2\zeta}{n-1}j, \quad j = 0, 1, \dots, n-1.$$

Let I be a real $n \times n$ matrix with $I_{i,j} = f(x_i, y_j)$. Given a fixed $\lambda \geq 1$, our objective is to use the discrete data available to us to approximate the spherical harmonics coefficients of $\kappa_\lambda f$ under the assumption that it has a bandlimit L .

Since the spherical harmonics [\(2.24\)](#) are orthonormal, our task amounts to approximating the integral

$$(3.2) \quad \langle \kappa_\lambda f, Y_{\ell,m} \rangle = \int_{S^2} \kappa_\lambda f(\mathbf{x}) Y_{\ell,m}(\mathbf{x})^* dS^2(\mathbf{x}).$$

Integrals of this form can be approximated well using spherical designs. A spherical t -design is set of spherical points $\mathcal{S}_t = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{N_t}\}$ satisfying

$$\int_{S^2} Y_{\ell,m}(\mathbf{x}) dS^2(\mathbf{x}) = \frac{1}{N_t} \sum_{n=1}^{N_t} Y_{\ell,m}(\mathbf{x}_n), \quad 0 \leq \ell \leq t.$$

Since we assume $\kappa_\lambda f$ has a bandlimit L and since $Y_{\ell_1,m_1} \cdot Y_{\ell_2,m_2}^* \in \mathcal{H}_{\ell_1+\ell_2}$ [\[2\]](#), we approximate [\(3.2\)](#) using a spherical $2L$ -design, thus

$$(3.3) \quad f_{\ell,m} \approx \frac{1}{N} \sum_{n=1}^N \kappa_\lambda f(\theta_n, \phi_n) Y_{\ell,m}(\theta_n, \phi_n)^*,$$

where (θ_n, ϕ_n) is $\mathbf{x}_n$ in spherical coordinates. This can be written in matrix form. Let $\mathbf{Y} = \mathbf{Y}(L, \mathcal{S}_{2L})$ be the complex matrix with rows index by $1, \dots, N_t$ and columns indexed by (ℓ, m) defined by

$$(3.4) \quad \mathbf{Y}_{n,(\ell,m)} = Y_{\ell,m}(\theta_n, \phi_n).$$

Our estimate of the spherical harmonics coefficients of f is then $\frac{1}{N_t} \mathbf{Y}^\dagger \mathbf{v}$, where $\mathbf{v}$ are the values of $\kappa_\lambda f$ on the spherical design $\mathcal{S}_{2L}$.

Remark 3.1. Approximating [\(3.2\)](#) by evaluating [\(3.3\)](#) on a spherical $2L$ -design might be numerically unstable, especially for large bandlimits L . In order to deal with that, using a K -design with $K > 2L$ is advisable, at least when L is large.

In order to calculate [\(3.3\)](#), one needs a spherical design and the value of $\kappa_\lambda f$ on it. Tables of spherical designs were obtained from [\[26, 25\]](#). The values of $\kappa_\lambda f$ on these points

are estimated by employing (2.8) and interpolation. At the heart of (2.8) is a mapping $(\theta, \phi) \mapsto (\lambda\theta \cos \phi, \lambda\theta \sin \phi)$ from the sphere to the plane. This mapping can easily be shown to be invertible as mapping from $S^2 \setminus \{(0, 0, -1)^\top\}$ onto $B_{\lambda\pi}$. We map the spherical design points onto the plane and isolate those that fall inside $[-\zeta, \zeta]^2$. We then interpolate the known values of f on the grid G onto these spherical design points. By (2.8), the values of f on any other spherical design point is necessarily zero.

Three points are important to note here. First, relying on interpolation complicates the handling of noisy data, since it introduces correlations in the noise. We deal with that in detail in subsection 4.1. Second, evaluating spherical harmonics is computationally costly. Thus, when calculating $\mathbf{Y}$ in (3.4), it is advisable to exploit the fact the values of $\kappa_\lambda f$ are assumed to be zero on spherical design points that are projected to outside of $[-\zeta, \zeta]^2$. We therefore remove said points from the spherical design before calculating (3.4). For brevity, we will still refer to this subset of the original spherical design as $\mathcal{S}_{2L}$. Third, if the image includes additive noise, one would want to consider how the projection affects it. In Appendix C, we discuss that for the noise model we use throughout the paper. In particular, we empirically demonstrate that additive white noise in the image approximately results in an additive white noise in the estimated spherical harmonics coefficients.

3.2. Calculating the spherical bispectrum. The procedure described in subsection 3.1 uses discrete samples of a function on the plane to approximate the spherical harmonics coefficients up to bandlimit L of its projection onto the sphere. Our objective now is to use the spherical harmonics coefficients to calculate its spherical bispectrum, its $\text{SO}(3)$ -invariant representation that we presented in subsection 2.5.

Let f be a bandlimited function on the sphere with spherical harmonics coefficients $\{f_{\ell,m}\}_{\ell=0}^L$ up to bandlimit L . The bispectrum of f is calculated by evaluating (2.25). Since the Clebsch–Gordan coefficients $C_{\ell_1,m_1,\ell_2,m_2}^{\ell,m}$ are nonzero only when $m_1 + m_2 = m$ [23, Chap. 8, eq. (2)], it follows that $C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m}$ is nonzero only when $\underline{m_1} \leq m_1 \leq \overline{m_1}$ for $\underline{m_1} = \max\{-\ell_1, m - \ell_2\}$ and $\overline{m_1} = \min\{\ell_1, m + \ell_2\}$. Thus, (2.25) becomes

$$(3.5) \quad b_f[\ell_1, \ell_2, \ell] = \sum_{m=-\ell}^{\ell} f_{\ell,m} \sum_{m_1=\underline{m_1}}^{\overline{m_1}} C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m} f_{\ell_1,m_1}^* f_{\ell_2,m-m_1}^*.$$

This expression is evaluated for all triplets (ℓ_1, ℓ_2, ℓ) in (2.26) and the resulting values are saved in a linear array in lexicographical ordering of the triplets. We refer to this array as the bispectrum vector and denote it by $\mathbf{b}_f$.

In order to evaluate (3.5), the Clebsch–Gordan coefficients need to be computed. We follow the method elaborated in [21] and compute all relevant Clebsch–Gordan coefficients in advance. For a fixed (ℓ_1, ℓ_2, ℓ) triplet and a fixed $-\ell \leq m \leq \ell$, [21] showed that $\mathbf{c} = (C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m})_{m_1=\underline{m_1}}^{\overline{m_1}}$ is a solution of the linear equation $\mathbf{C}\mathbf{c} = \mathbf{0}$. Here, $\mathbf{C}$ is the tridiagonal $(n-1) \times n$ real matrix, where $n = \overline{m_1} - \underline{m_1} + 1$ and its diagonals are

$$\begin{aligned} \mathbf{C}_{i,i} &= \ell_1(\ell_1 + 1) + \ell_2(\ell_2 + 1) + 2m_{1,i}m_{2,i} - \ell(\ell + 1), \\ \mathbf{C}_{i,i+1} &= \mathbf{C}_{i+1,i} = \sqrt{\ell_1(\ell_1 + 1) - m_{1,i}m_{1,i+1}} \sqrt{\ell_2(\ell_2 + 1) - m_{2,i}m_{2,i+1}}, \end{aligned}$$

where

$$m_{1,i} = \underline{m_1} + i - 1 \quad \text{and} \quad m_{2,i} = m - \underline{m_1} - i + 1.$$

This equation does not have a unique solution. However, the Clebsch–Gordan coefficients must have norm of 1 and by convention satisfy $C_{\ell_1, \underline{m_1}, \ell_2, m - \underline{m_1}}^{\ell, m} > 0$ (see [21, 3]). We therefore find a solution for this linear equation by choosing $C_{\ell_1, \underline{m_1}, \ell_2, m - \underline{m_1}}^{\ell, m} = \frac{1}{\underline{m_1} - \underline{m_1} + 1}$ and determining the remaining coordinates in the solution vector by backward substitution. The resulting vector is then normalized to be a unit vector.

We tested how well the results of Theorem 2.9 are echoed in the invariance of the bispectrum. To that effect, we generated a random image using the following procedure. First, we generated a random set of spherical harmonics coefficients of a real-valued function on the sphere f of bandlimit 14. The coefficients were sampled uniformly from the set of coefficients satisfying

$$(3.6) \quad \sum_{m=-\ell}^{\ell} |f_{\ell, m}|^2 = 1$$

for all $0 \leq \ell \leq 14$. Given a square-integrable function $f : S^2 \rightarrow \mathbb{R}$, its spherical harmonics coefficients are


$$f_{\ell, m} = \langle f, Y_{\ell, m} \rangle = \int_{S^2} f(\mathbf{x}) Y_{\ell, m}(\mathbf{x})^* dS^2(\mathbf{x}).$$

Since $Y_{\ell, m}(\mathbf{x})^* = (-1)^m Y_{\ell, -m}(\mathbf{x})$ [23] and f is real-valued, it follows that


$$(3.7) \quad f_{\ell, m}^* = (-1)^m f_{\ell, -m}.$$

Thus, the coefficients were sampled uniformly from those satisfying both (3.6) and (3.7). The resulting function was back-projected onto a 101×101 regular equidistant grid in $[-\zeta, \zeta]^2$ by evaluating (2.8) on all grid points for scaling parameter $\lambda = 1$. In the resulting images, the first and last 20 rows and columns of pixels were set to zero and then smoothed using a Gaussian filter. This was done in order to create an image that has approximately zero in its margin, with enough room to apply translations to. The resulting image was projected onto the sphere by estimating its spherical harmonics coefficients up to bandlimit 16 with $\lambda = 1$. The final image is produced by back-projecting this spherical function onto the same grid with $\lambda = 1$. Hereafter, we refer to these images as random images. Figure 3(a) shows eight examples of random images generated using this procedure.

For an image generated as described above, we performed three separate experiments. First, we tested whether its bispectrum is invariant to rotations. The original image was rotated, the rotated image was projected onto the sphere, and its coefficients up to bandlimit 16 were estimated. Its bispectrum was calculated and compared with the bispectrum of the projection of the original image onto the sphere. As Figure 4(a) (left) shows, the relative bispectrum error is negligible and is consistent with the numerical perturbations expected the



(a) Random images from subsection 3.2




(b) Simulated cryo-EM images


Figure 3. (a) Examples of random images produced using the procedure described in subsection 3.2. (b) Examples of projections in random directions produced using ASPIRE [1].

various approximations we use throughout, e.g., our use of interpolation (see subsection 3.1). This is in keeping with Theorem 2.9(i).

We next tested whether the bispectrum is invariant to translations. In polar coordinates, every translation is represented by its magnitude and direction. The image was translated by translations of increasing size. For each translation size, the image was translated 10^3 times, its bispectrum computed and compared with the bispectrum of the original image. The translation direction was independently, uniformly sampled. Figure 4(a) (middle) shows the relative bispectrum error averaged over the translation directions, as well as the symmetric region around the mean where 95 percent of the samples were. In Figure 4(a) (right), we followed a similar procedure, except the image was also rotated. For every translation size,



(a) Random image from subsection 3.2



(b) Simulated cryo-EM image

Figure 4. *The bispectrum is approximately invariant to rotations and translations of an image.* The relative bispectrum error is below 1 percent when the image is only rotated (left) and scales with the translation size when the image is translated only (middle) or translated and rotated (right). The blue curve in the middle and right subfigures is the relative bispectrum error averaged over translation directions and translation directions and rotation angles, respectively. The gray region is a symmetric empirical 95 percent confidence interval around the mean.

the image was rotated and translated $2 \cdot 10^3$ times. The translation direction and the rotation angle were independently, uniformly sampled. The results indicate the bispectrum is indeed approximately invariant to translations. The error clearly scales with the translation size, as expected from Theorem 2.9(ii). Of particular interest to our future numerical experiments, the relative error did not exceed 5 percent for translation sizes of 10 pixels (about 10 percent of the image size) and remained in single digits for even larger translation sizes.

We repeated this experiment with a simulated cryo-EM image. The image was of size 101×101 pixels and was simulated as a tomographic projection of a volume in a uniformly random direction using the built-in simulation module of the MATLAB package ASPIRE [1]. Examples of such images appear in Figure 3(b). Changing the images required increasing the bandlimit in order to accurately represent the projected images on the sphere (see subsection 3.3 for details). We used a bandlimit of 70. For a such a large bandlimit, calculating the bispectrum is substantially more time-consuming than for bandlimit 16 we used in the experiment of Figure 4(a). Therefore, in order to keep execution time reasonable, in this experiment we sampled

much less translation directions, rotation angles in the translation only, and translation and rotation experiments (dozens instead of thousands). The results, which appear in Figure 4(b), are substantially the same as the results obtained for a random image, in that the bispectrum is approximately invariant to both rotations and translations, provided the translation size is small enough.

3.3. Choice of scaling parameters and bandlimit. According to Theorem 2.9 two key parameters, λ and $\tilde{\lambda}$, determine how well the orbits of $\text{SE}(2)$ are approximated by the orbits of $\text{SO}(3)$. The first, λ , affects the embedding of the image in $\mathbb{R}^2$, in that it is assumed throughout that the image is compactly supported within $B_{\lambda\pi}$. The second, $\tilde{\lambda}$, is the parameter of the contraction map. As Figure 2 illustrates, for a fixed λ , increasing $\tilde{\lambda}$ maps $B_{\lambda\pi}$ to a progressively smaller spherical cap around the north pole of the sphere. This implies that a larger $\tilde{\lambda}$ makes the fine details of the projected image finer still. In order to compensate for this, it is necessary to use a larger bandlimit when computing the projection, because the bandlimit determines the resolution of a spherical harmonics expansion.

We demonstrate this phenomenon empirically. We generated a dataset of 50 images and projected them onto the sphere as described in subsection 3.1, with various choices of scaling parameter and bandlimit. Recall that it is assumed there that every image is a discrete sample on a regular grid of a function, and note that this regular grid is within B_1 . In particular, it is assumed the function is supported within B_π . In terms of Theorem 2.9, we fixed $\lambda = 1$, as we have in all the numerical experiments throughout the paper, and varied $\tilde{\lambda}$. For every scaling parameter, the projected images were then projected back onto the grid and the error relative to the ground truth was calculated. For every pair of bandlimit and scaling parameter, the results were averaged over all 50 images in the dataset. Figure 5 shows the results for




Figure 5. Trade-off between the scaling parameter and bandlimit. Error of back-projected image relative to the original image averaged over two sets of 50 images, one with the random images described in subsection 3.2 and one from a simulation of cryo-EM. As the scaling parameter increases, a larger bandlimit is required to maintain the same relative error.

a dataset of random images of the kind we described in [subsection 3.2](#) and for a dataset of simulated cryo-EM images. As expected, as the scaling parameter increases, a larger bandlimit is required to maintain a similar relative error.

One has to balance between the need to accurately represent projected images on the sphere and the time-complexity of the calculation of the bispectrum. As shown in [Figure 5](#), for bandlimit 50 and 70 with scaling parameter of 1, the mean relative error of recovery for simulated cryo-EM images is about 13 and 6 percent, respectively. However, despite the larger than 10 percent error, for most of our classification experiments on simulated cryo-EM images we used a bandlimit of 50 and achieved decent results (see [subsection 5.3](#)). This difference is important computationally, since the bispectrum vector has about 68000 elements for bandlimit 50 and about 182000 for bandlimit 70.

In this paper our concern is demonstrating in principle the viability of compactification in image processing, rather than optimizing it for real-world use. In keeping with this motivation, we chose the bandlimit and the scaling parameter based on empirical criteria. For every type of simulated dataset, we sought a bandlimit and scaling parameter pair that fulfilled two requirements. First, it had a decent relative error recovery rate in experiments like the ones showed in [Figure 5](#). For MRA, we generally preferred those with less than 10 percent relative error, but as noted above this could be relaxed for classification. Second, it resulted in an approximately rotationally and translationally invariant bispectrum in experiments such as the ones shown in [Figure 4](#).

To a large extent, the need to compromise for the sake of computational efficiency is the result of our use of spherical harmonics. On the one hand, using spherical harmonics is convenient. Because of their role in harmonic analysis on the sphere and representation theory of $SO(3)$, they yield an expression for the bispectrum that can be calculated in practice, as discussed in [subsection 3.2](#). This is a property other bases of functions on the sphere generally do not share. On the other hand, while our function is localized around the north pole of the sphere, spherical harmonics are generally not. Thus, in order to represent the fine details of an image, we have to take a relatively large bandlimit, resulting in a greater amount of spherical harmonics coefficients needed to represent an image and a greater inefficiency in calculating an invariant representation of them. We are currently seeking to address this issue by utilizing a locally supported basis on the sphere (for example, see [\[24\]](#)).

Finally, we note that in its current form our approach is probably inadvisable for image processing problems involving images with discontinuities or images with nonzero pixels at their boundary. Using spherical harmonics to represent the projection of such images onto the sphere is expected to make the trade-off mentioned above worse, as a larger bandlimit will be required to account for the discontinuities. A locally supported basis on the sphere might be useful to address this, if it can be used to represent a local feature of an image with higher resolution. Thus, it may allow a more accurate local representation of the discontinuities, without substantially increasing the number of basis elements required to represent the image as a whole. In [\[24\]](#), a basis capable of such local representations is elaborated upon, and perhaps their approach can be adapted for this purpose. However, we do not intend to pursue this direction here or in the future. Our motivation is cryo-EM, where all images can be assumed to be smooth and our main interest here is in demonstration of the potential of our approach, rather than its optimization for real-world use.

4. Application I: Compactification in multireference alignment over SE(2). We apply compactification to an MRA problem over SE(2). Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a smooth function compactly supported within $[-\zeta, \zeta]^2$, where $\zeta = \cos \frac{\pi}{4} = \sin \frac{\pi}{4}$. We refer to f as the ground truth function. We are given N discrete samples of f of the form

$$(4.1) \quad \mathbf{I}_j = \mathbf{D}(g_j \bullet f) + \boldsymbol{\varepsilon}_j, \quad j = 1, 2, \dots, N.$$

Here, $g_j = (\mathbf{b}_j, R_j) \in \text{SE}(2)$ is such that $g_j \bullet f$ is also compactly supported within $[-\zeta, \zeta]^2$. Furthermore, $\mathbf{b}_j = r_j (\cos \alpha_j, \sin \alpha_j)^\top$ where $\{r_j\}$ are independently sampled from a uniform distribution on $[0, T_{\max}]$ and α_j is likewise sampled from $[0, 2\pi)$. Also, R_j is a rotation of the plane by $\theta_j \in [0, 2\pi)$ radians counterclockwise and $\{\theta_j\}$ are independently sampled from a uniform distribution on $[0, 2\pi)$. $\mathbf{D}$ is a discretization operator sampling a function on a regular equidistant $n \times n$ grid centered at the origin of $\mathbb{R}^2$. $\mathbf{I}_j$ and $\boldsymbol{\varepsilon}_j$ are $n \times n$ real matrices representing the samples of the rotated and translated f on said grid and noise, respectively. The elements of the noise matrices are Gaussian independent and identically distributed (i.i.d.), $\varepsilon_{j,m,k} \sim \mathcal{N}(0, \sigma^2)$. Stated roughly, our objective is to estimate f up to an SE(2) action, from sampled images (4.1).

Our approach to the problem is a variant of the invariants approach to MRA described in the introduction. We compactify the samples. By that we mean we project the images onto the sphere to produce bandlimited functions on the sphere, as detailed in subsection 3.1. The spherical bispectrum of each projection is calculated as described in subsection 3.2. We use these to estimate the spherical bispectrum of the projection of f onto the sphere and then use it to estimate both the projection of f onto the sphere and $\mathbf{D}f$.

In the following sections we flesh out this approach. In subsection 4.1, we explain how we estimate the spherical bispectrum of the projection of f onto the sphere from samples (4.1). Then, in subsections 4.2 and 4.3, we explain how this estimate is used to estimate both the projection of f onto the sphere and $\mathbf{D}f$.

4.1. Spherical bispectrum estimation. We begin by estimating the spherical harmonics coefficients of the projection of the sampled images onto the sphere and calculating their corresponding bispectrum vectors, using the methods detailed in section 3. We denote the estimate of the former by $\mathbf{s}_j$ and the corresponding bispectrum vector by $\mathbf{b}_{\mathbf{s}_j}$. An unbiased estimator of the spherical bispectrum of the projection of f onto the sphere is

$$(4.2) \quad \hat{\mathbf{b}}[\ell_1, \ell_2, \ell] = \frac{1}{N} \sum_{j=1}^N \left(\mathbf{b}_{\mathbf{s}_j}[\ell_1, \ell_2, \ell] - \sum_{k=1}^3 K_{\mathbf{s}_j, k}[\ell_1, \ell_2, \ell] \right),$$

where K_1 , K_2 , and K_3 are sparse linear transformations acting on $\mathbf{s}_j$, considered as a real vector. The derivation of (4.2) and the full definition of $K_{\mathbf{s}_j, 1}$, $K_{\mathbf{s}_j, 2}$, and $K_{\mathbf{s}_j, 3}$ can be found in Appendix B. Calculating this estimator requires only a single pass on the data. It can also be continuously updated when new data becomes available and therefore is memory-efficient. Its calculation can also be easily parallelized.

4.2. Inversion of the spherical bispectrum. Let $\mathbf{b}$ be a bispectrum vector of a function on the sphere of bandlimit L . Our objective is to find a spherical function of bandlimit L

which has bispectrum $\mathbf{b}$. We do that by solving the unconstrained nonlinear least-squares problem

$$(4.3) \quad \min_{g \in L_L^2(S^2)} \|\mathbf{b}_g - \mathbf{b}\|^2,$$

where $L_L^2(S^2)$ is the space of spherical functions of bandlimit L represented by their spherical harmonics coefficients, and $\mathbf{b}_g$ for $g \in L_L^2(S^2)$ is its bispectrum vector. In general, (4.3) is a nonconvex optimization problem. This is a recurring feature of bispectrum inversion by optimization in various contexts, e.g., [6, 17]. Despite that, in the past this approach has worked well and yielded very good estimates. It was also found to be stable in the sense that inverting perturbed bispectrum vectors did not significantly magnify the errors.

We solve (4.3) for our estimator (4.2) of the bispectrum of F ; that is, $\mathbf{b} = \hat{\mathbf{b}}$. We denote the solution by $\hat{\mathbf{s}}$ and refer to it as the estimator of $\kappa_\lambda f$ on the sphere, up to 3D rotation. Though we defined the bispectrum vector as a vector of complex numbers, when solving this optimization problem we treat it as a vector of real numbers using the embedding of $\mathbb{C}^n$ into $\mathbb{R}^{2n}$. Similarly, we solve for the real numbers $\{\Re g_{\ell,m}\} \cup \{\Im g_{\ell,m}\}$, the real and imaginary parts of the spherical harmonics coefficients of g .

In our implementation, we solve (4.3) using the MATLAB implementation of a trust-regions algorithm in its built-in `lsqnonlin` function. This is an iterative algorithm which requires initialization. In [17], an approach similar to ours was used to estimate an image up to 2D rotations only. We project this estimate of the underlying image onto the sphere and use the resulting spherical harmonics coefficients to initialize our the trust-region iteration. In addition, iterative optimization algorithms often perform better when an explicit derivative of the objective function is provided. We therefore provide an explicit expression for the derivative of $g \mapsto \mathbf{b}_g - \hat{\mathbf{b}}$:

$$\frac{\partial b_g[\ell_1, \ell_2, \ell]}{\partial \Re g_{\omega, s}} = E_1 + E_2 + E_3 \quad \text{and} \quad \frac{\partial b_g[\ell_1, \ell_2, \ell]}{\partial \Im g_{\omega, s}} = i(E_1 + E_2 + E_3),$$

where

$$\begin{aligned} E_1 &= \delta_{\omega, \ell} \delta_{s, m} \sum_{m_1=\max\{-\ell_1, s-\ell_2\}}^{\min\{\ell_1, s+\ell_2\}} C_{\ell_1, m_1, \ell_2, s-m_1}^{\ell, s} g_{\ell_1, m_1}^* g_{\ell_2, s-m_1}^*, \\ E_2 &= \delta_{\omega, \ell_1} \delta_{s, m_1} \sum_{m=\max\{-\ell, s-\ell_2\}}^{\min\{\ell, s+\ell_2\}} C_{\ell_1, s, \ell_2, m-s}^{\ell, m} g_{\ell, m} g_{\ell_2, s-m_1}^*, \\ E_3 &= \delta_{\omega, \ell_2} \delta_{s, m_2} \sum_{m=\max\{-\ell, s-\ell_1\}}^{\min\{\ell, s+\ell_1\}} C_{\ell_1, m-s, \ell_2, s}^{\ell, m} g_{\ell, m} g_{\ell_1, m-s}^*, \end{aligned}$$

where $\delta_{k,m}$ is the Kronecker delta function defined as $\delta_{k,m} = 1$ if $k = m$ and otherwise 0. These expressions show the derivative is sparse, which reduces both memory and computational requirements.

4.3. Alignment of two bandlimited functions on the sphere. Recall the spherical bispectrum defines a function on the sphere up to 3D rotation. We wish to use our explicit projection formula (2.8) to project $\hat{\mathbf{s}}$ back onto the plane, onto the grid on which the original dataset was defined. However, $\hat{\mathbf{s}}$ might not be concentrated around the north pole $\mathbf{n} = (0, 0, 1)^\top$, where said grid is projected to by (2.8).

To solve this problem, we allow ourselves to exploit the fact that in our simulations we know the ground truth and align $\hat{\mathbf{s}}$ with the known $\kappa_\lambda F$. This is in keeping with our desire to demonstrate in principle the utility and viability of compactification in MRA problems, rather than optimize it for real-world use. We align the two spherical functions by brute force. We use a deterministic sequence of rotations with ideal uniform distribution in $\text{SO}(3)$, that was developed in [27]. This sequence has the property that $\{R_i\}_{i=1}^M \subseteq \text{SO}(3)$ is uniformly "denser" in $\text{SO}(3)$ as M increases.

For each rotation R in this sequence, we rotate $\hat{\mathbf{s}}$ by R . This is done by rotating a spherical design. We evaluate the spherical function represented by $\hat{\mathbf{s}}$ on a spherical t -design with sufficiently large t . The spherical design points are then rotated by R . The resulting set of spherical points is also a spherical t -design [26]. We use the values of $\hat{\mathbf{s}}$ as the values of the rotated function on the rotated design and estimate its spherical harmonics coefficients. We then measure the correlation between the resulting spherical function and the known ground truth spherical function, using the measure

$$\text{Corr}(f, g) = \frac{\Re \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} f_{\ell,m} g_{\ell,m}}{\sqrt{\sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} |f_{\ell,m}|^2} \sqrt{\sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} |g_{\ell,m}|^2}}.$$

From among the first M rotations in the sequence, we choose the rotation R that maximizes the correlation and use it to rotate $\mathbf{s}$ so as to align it to $\kappa_\lambda f$. In our experiments, we found that using the $M = 72 \cdot 2^8$ yields good alignment. Once they are aligned, we evaluate its values on the spherical points corresponding by (2.8) to the grid on the plane. We denote the resulting image by $\hat{\mathbf{I}}$. It is our estimate of the image $\mathbf{D}f$. Since we align $\hat{\mathbf{s}}$ to the spherical harmonics coefficients of the ground truth image, $\hat{\mathbf{I}}$ and $\mathbf{D}f$ are already aligned. We measure the error relative to $\mathbf{D}f$ in the Frobenius norm.

4.4. Numerical experiments. We tested our approach on a synthetic dataset of images. We first generated a ground truth 101×101 image using the procedure described in subsection 3.2. We generated a dataset of images by sampling rotations, translations, and noise according to (4.1) with $T_{\max} = 5$ and applying them to this image. The noise variance σ^2 was chosen to yield SNR of 0.5 and is assumed to be known. We projected these images onto the sphere with scaling parameter $\lambda = 1$ and bandlimit 16, estimated the bispectrum of the image, inverted it, and aligned the resulting function with the projection of the original image onto the sphere. We then projected the aligned function back onto the 101×101 equidistant grid on the plane and compared it with the original image. This was done for varying sample sizes, repeated 15 times for each with different noise realization. The estimation error was measured using the Frobenius norm for all our estimators.

The results are shown in Figures 6(a). We first draw the reader's attention to the bispectrum estimator. It was closer to the bispectrum of the original image as the sample size




Figure 6. Relative error of estimating the bispectrum and the image using a method of [17] which takes only rotations into account (initial guess) and using our approach (image). The back-projection bound is an empirical lower bound on the relative error of our image estimator, which results from a loss of information when projecting an image onto the sphere. All results are averaged over 15 repeats with different noise realization. (a) $\text{SNR} = 0.5$ and $\lambda = 1$. The bispectrum estimator comports decently with the expected behavior of an unbiased estimator. The initial guess was estimated by taking only rotations into account and shows no scaling with the sample size. The image estimator approaches the back-projection bound. (b) Sample size 10^4 , $\lambda = 1$. All estimators perform worse when the SNR is reduced, except the initial guess. It is less sensitive to decreasing SNR but also does not show much improvement as it is increased and is outperformed by our approach for large SNRs.

increased. Its slope in a logarithmic scale was -0.43 , which comports decently with the expected -0.5 slope dictated by the law of large numbers. We also draw the reader's attention to the behavior of our estimator of the image. It improves with an increase in the sample size. In contrast, the initial guess in the bispectrum inversion, which is based on an approach that takes only rotations into account, barely improves as the sample size increases. This is consistent with the fact it only takes into account rotations.

Finally, we note the behavior of our image estimator near what we refer to as the back-projection bound. Projecting an image onto the sphere is not lossless. An empirical estimate of the lost information can be obtained by projecting the ground truth image $\mathbf{DF}$ from (4.1) onto the sphere and then back-projecting it back onto the plane. The back-projected image error relative to the original image is a lower bound on the achievable relative error in our estimate of the image. This lower bound, which we call the back-projection bound, came out in our experiment at about 0.02 . This implies that our method is not expected to be able to estimate the image with less than about 2 percent relative error. Indeed, while our estimator of the image improves with an increase in the sample size, the improvement is approximately linear (in a log-log plot) away from the back-projection bound, but stagnates as the error of the estimator approaches the back-projection bound.

We further tested how our approach responds to varying SNRs. We generated a dataset of images as described above with a fixed sample size $N = 10^4$ and noise realization matching different SNRs. The noise variance σ^2 is always assumed to be known. For every SNR, we

performed 15 trials. The results appear in Figure 6(b). All estimators, except the initial guess estimator, performed markedly worse in lower SNRs. While the initial guess is less sensitive to decreasing SNR, it also does not show much improvement when it is increased and is outperformed by our approach when the SNR increases. Combined with the results in Figure 6(a), it appears our approach may yield a better estimate of the image than the initial guess estimator for lower SNRs, when provided with a sufficiently large dataset. However, such a dataset may be prohibitively large. In order to reduce the required sample size, one might need to refine our approach, perhaps by adding denoising steps prior to projecting the image onto the sphere.

5. Application II: Compactification in classification invariant to SE(2). We apply compactification to the following classification problem invariant to SE(2). Each sample belongs to one of m classes, generated from functions $f_1, \dots, f_m : \mathbb{R}^2 \rightarrow \mathbb{R}$ compactly supported within B_π . Our samples are assumed to be of the form

$$(5.1) \quad \mathbf{I}_j = \mathbf{D}(g_j \bullet f_{k_j}) + \varepsilon_j, \quad j = 1, 2, \dots, N, \quad k_j \sim \mathbf{p}.$$

Here, $\mathbf{D}$, g_j , and ε_j are as described in (4.1). Also, k_j are drawn i.i.d. from a distribution $\mathbf{p} = (p_1, \dots, p_m)^\top$. Finally, given a dataset of the form (5.1), let $L(j) = f_{k_j}$ be the label of the j th image. Our objective is to identify for every image its K -closest images up to rotation and translation.

In Zhao and Singer [30], a similar problem was dealt with in the context of cryo-EM. There a method that takes only rotations into account was used. In our approach and theirs, every image in a dataset is transformed into a rotation-invariant representation. We describe a representation similar to theirs in subsection 5.1 using the approximately SE(2)-invariant representation we developed in sections 2 and 3. A K -nearest neighbors graph is then constructed by measuring the distance between the invariant representation of the images, rather than the images themselves, as detailed in subsection 5.2.

In subsection 5.3, we numerically test how our approach, which takes into account both rotations and translations, performs compared to an approach that takes only rotations into account. We show that the similarity graphs match the classification function of the dataset better when translations are taken into account, even for relatively large maximal translation size $T_{\max}$ in (5.1).

5.1. Steerable basis and a rotation-invariant representation of a function. We use the MATLAB package ASPIRE [1] to estimate the expansion coefficients of $g_j \bullet F_{k_j}$ in a truncated Fourier–Bessel basis [28, 29]. Every member of this basis has the form

$$u_{k,q}(r, \phi) = f_{k,q}(r)e^{ik\phi},$$

where (r, ϕ) are elements of $\mathbb{R}^2$ in polar coordinates, $f_{k,q}$ is called the radial part, and $\phi \mapsto e^{ik\phi}$ is called the angular part. We assume all functions we expand in this basis are bandlimited in the sense they have a finite expansion. In particular, $|k| \leq k_{\max}$ and $0 \leq q \leq Q_k$.

A basis with this form is said to be steerable, meaning that it diagonalizes the action of $\text{SO}(2)$. In particular, if $R = R(\varphi)$ is the rotation of $\mathbb{R}^2$ by φ counterclockwise and

$$(5.2) \quad F = \sum_{k,q} F_{k,q} u_{k,q},$$

then

$$R \bullet F(r, \theta) = F(r, \theta + \varphi) = \sum_{k,q} F_{k,q} f_{k,q}(r) e^{ik(\theta+\varphi)} = \sum_{k,q} e^{ik\varphi} F_{k,q} f_{k,q}(r) e^{ik\theta}.$$

Thus, the expansion of $R \bullet F$ in the Fourier–Bessel basis satisfies $(R \bullet F)_{k,q} = e^{ik\varphi} F_{k,q}$.

In [30], this property was used to define a rotation-invariant representation of a function F of the form (5.2),

$$(5.3) \quad b_F[k_1, k_2, q_1, q_2, q_3] = F_{k_1, q_1} F_{k_2, q_2} F_{k_1+k_2, q_3}^*,$$

for all tuples of radial and angular indices $(k_1, k_2, q_1, q_2, q_3)$ satisfying

$$(5.4) \quad |k_1 + k_2| \leq k_{\max}, \quad 0 \leq q_1 \leq Q_{k_1}, \quad 0 \leq q_2 \leq Q_{k_2}, \quad \text{and} \quad 0 \leq q_3 \leq Q_{k_1+k_2}.$$

In keeping with the terminology of [30], we refer to (5.3) as the rotational bispectrum. When F is a Fourier–Bessel expansion estimated from an image, we refer to (5.3) as the rotational bispectrum of that image.

We wish to represent every image in our dataset using its rotational bispectrum. Unfortunately, even when $k_{\max} = 5$, the number of tuples satisfying (5.4) is extremely large, on the order of tens of thousands at least, and grows to millions even for modest $k_{\max}$. Therefore, even for modestly sized datasets of images, we cannot store in memory the rotational bispectrum of all images. On the other hand, the size of the rotational bispectrum suggests it contains a lot of redundancy. Thus, like in [30], we apply dimensionality reduction to the collection of rotational bispectra of all images in a dataset.

In particular, let $\mathbf{b}_j$ denote the rotational bispectrum of the j th image. Let $\mathbf{B}$ be a matrix with N columns, the j th column of which is the rotational bispectrum of $\mathbf{I}_j$ in lexicographical order of the indices (5.4). We apply to $\mathbf{B}$ a randomized algorithm, designed to approximate the best rank m approximation of a large matrix without forming it in memory [9]. We use it to approximate $\mathbf{B}$ by an $m \times N$ dimensional matrix. We refer to the j th column of the resulting matrix either as the reduced rotational bispectrum of $\mathbf{I}_j$ or as its reduced rotation-invariant representation.

5.2. Constructing a similarity graph from invariant representations. Given a dataset of images of the form (5.1), assume we represent it either by the reduced rotation-invariant representation introduced in subsection 5.1 or by the $\text{SE}(2)$ -invariant representation introduced in section 3, possibly after going through denoising by a linear filter. Denote by $\tilde{\mathbf{I}}_j$ the invariant representation of $\mathbf{I}_j$. We define two invariant distance metric on our dataset by

$$(5.5) \quad d(\mathbf{I}_{j_1}, \mathbf{I}_{j_2}) = \left\| \tilde{\mathbf{I}}_{j_1} - \tilde{\mathbf{I}}_{j_2} \right\|_F.$$

When $\tilde{\mathbf{I}}_{j_1}$ and $\tilde{\mathbf{I}}_{j_2}$ are the reduced rotation-invariant representations of the corresponding images, we refer to this metric as the rotation-invariant metric. Similarly, when these are our approximately SE(2)-invariant representations, we refer to it as the SE(2)-invariant metric.

We use an invariant distance metric to find the K -nearest images of every image in our dataset. Denote by G_j the set of indices of the K images closest to $\mathbf{I}_j$. For every node, we then measure the proportion of its K neighbors belonging to its class. That is, we calculate

$$(5.6) \quad s_j = \frac{|\{j_1 \in G_j \mid j_1 \neq j, L(j_1) = L(j)\}|}{K}.$$


We refer to (5.6) as the node score of the K -nearest neighbors graph we construct using the metric (5.5). We consider its distribution as a measure of the quality of the classification. The classification is considered better the more concentrated the distribution of node score is close to 1.

5.3. Numerical experiments. We tested our approach numerically as follows. We first generated seven class representatives using the procedure described in subsection 3.2. We generated a dataset of $5 \cdot 10^3$ 101×101 images following (5.1) with $m = 7$ and $p_1 = p_2 = \dots = p_m = \frac{1}{7}$. We then generated two 50-nearest neighbors graphs, one using the rotation-invariant metric described in subsection 5.1 and another using our SE(2)-invariant metric. For each graph we measured the node score (5.6) of every node. We performed this procedure for maximal translation sizes $T_{\max} = 0, 2.5, 5, 7.5, 10$. The labels, the rotation angle, the translation directions, and the translation sizes of the images in our dataset were sampled in advanced; the latter was sampled uniformly from $[0, 1]$ and then rescaled by multiplying it by the appropriate $T_{\max}$. For every $T_{\max}$, we repeated this experiment 10 times, each with different noise realization, all with SNR = 1. The noise variance σ^2 is always assumed to be known. For every $T_{\max}$, we averaged the node score over the 10 noise realizations and plotted the histogram of the average.


Figures 7(a) shows the histogram of the average node score for varying maximum translation sizes $T_{\max}$, when using a rotation-invariant metric and our SE(2)-invariant metric. The more the histogram is concentrated close to 1, the better the distribution of node score is. Both metrics yield similar results in the absence of translations. However, our metric is far less sensitive to the translation size and even yields excellent node score distribution for $T_{\max} = 10$.

We repeated this experiment with simulated cryo-EM projections of the type shown in Figures 3(b). Each class representative was generated using the built-in simulation module of the MATLAB package ASPIRE [1]. In accordance with the results of subsection 3.3, a bandlimit of 50 was chosen. Because of the computational load of using such a large bandlimit, the experiment was done with a single noise realization. For comparison, we present the node score obtained by the algorithm of [30], which takes only rotations into account. Specifically, we used an optimized implementation of it which is included in ASPIRE [1], which also includes a denoising stage via steerable PCA [28]. As the results in Figures 7(b) indicate, classification using our approximate invariant was more accurate for simulated cryo-EM images.

Finally, we repeated this experiment once on a dataset that resembles more the kind of datasets encountered in cryo-EM applications. We used a dataset of 10^4 images of size 101×101 . The maximal translation size was set to 10 pixels, there were 100 class representatives,



(a) Random images from subsection 3.2



(b) Simulated cryo-EM images

Figure 7. Node score for rotationally invariant and approximately rotationally and translationally invariant metrics. A histogram of the node score for maximal translation sizes $T_{\max} = 0, 2.5, 5, 7.5, 10$ on a dataset generated from seven class representatives. For (a), the node score is averaged over 10 different noise realizations. The data in (b) is based on a single noise realization. Both metrics yield similar results in the absence of translations. Taking into account translations significantly reduces the sensitivity of the distribution of node score to translations.




Figure 8. Node score for rotationally invariant and approximately rotationally and translationally invariant metrics on a dataset of simulated cryo-EM images containing a hundred class representatives. A histogram of the node score for maximal translation size 10 on a dataset generated from a hundred class representatives. Taking into account translations yielded a median node score of 1, whereas using a rotations-only invariant yielded a median node score of 0.22.

rather than 7, and the bandlimit used in our algorithm was 70. We again compared the distribution of node score achieved by our approach with that obtained by the same optimized implementation of the algorithm of [30] in ASPIRE [1]. The results are shown in Figure 8. Our algorithm clearly had a better distribution of node score. Consistent with that, our approach yielded a median node score of 1, compared with 0.22 using the rotations-only invariant.

Appendix A. Proofs for section 2. The various lemmas were not stated in the main text in the order in which they're proved. In particular, the proofs make it clear that Lemmas 2.4 and 2.5 depend on Lemma 2.8. We altered the order in the main text to improve the exposition.

Proof of Lemma 2.8. Recall the identity (2.4). Since $\mathfrak{so}(3)$ is a Lie subalgebra of $\mathfrak{gl}(3)$, one can write for every $\mathbf{x} \in \mathbb{R}^2$,

$$(A.1) \quad \exp_{\text{SO}(3)}(\mathbf{x}) = \sum_{n=0}^{\infty} \frac{1}{n!} \mathbf{X}^n, \quad \text{where } \mathbf{X} = \begin{bmatrix} \mathbf{0}_{2 \times 2} & \mathbf{x} \\ -\mathbf{x}^\top & 0 \end{bmatrix}.$$

From this point, the proof relies primarily on the Taylor series of sine and cosine. Note that

$$\mathbf{X}\mathbf{n} = \begin{pmatrix} \mathbf{x} \\ 0 \end{pmatrix} \quad \text{and} \quad \mathbf{X}^2\mathbf{n} = \mathbf{X} \begin{pmatrix} \mathbf{x} \\ 0 \end{pmatrix} = -\|\mathbf{x}\|^2 \mathbf{n}.$$

Therefore,

$$\mathbf{X}^{2n}\mathbf{n} = (-1)^n \|\mathbf{x}\|^{2n} \mathbf{n} \quad \text{and} \quad \mathbf{X}^{2n+1}\mathbf{n} = \left(-\|\mathbf{x}\|^2\right)^n \mathbf{X}\mathbf{n} = (-1)^n \|\mathbf{x}\|^{2n} \begin{pmatrix} \mathbf{x} \\ 0 \end{pmatrix}.$$

This implies the z -coordinate of (A.1) is

$$(A.2) \quad \sum_{n=0}^{\infty} \frac{(-1)^n \|\mathbf{x}\|^{2n}}{(2n)!} = \cos \|\mathbf{x}\|,$$

and the x - and y -coordinates are of the form

$$(A.3) \quad x \sum_{n=0}^{\infty} \frac{(-1)^n \|\mathbf{x}\|^{2n}}{(2n+1)!} = \frac{x}{\|\mathbf{x}\|} \sum_{n=0}^{\infty} \frac{(-1)^n \|\mathbf{x}\|^{2n+1}}{(2n+1)!} = \frac{x}{\|\mathbf{x}\|} \sin \|\mathbf{x}\|,$$

where x is the x -coordinate or the y -coordinate of $\mathbf{x}$, respectively. To conclude the proof, substitute $\mathbf{x}$ with $\mathbf{x}/\lambda$ in (A.2) and (A.3). ■

Proof of Lemmas 2.4 and 2.5. $\eta_{\text{SE}(2)}$ is a surjective map and it is an orbit map of $\mathbf{0}_2$ of the group action of $\text{SE}(2)$ on $\mathbb{R}^2$; that is, $\eta_{\text{SE}(2)}(\mathbf{x}, R) = (\mathbf{x}, R) \bullet \mathbf{0}_2$. It is therefore a smooth submersion (cf. [16, Prop. 7.26, p. 83] and [16, Thm. 4.14, p. 166]). In order to prove $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ is a smooth submersion, define $P(\zeta, \theta, \phi) = (\zeta \cos \theta, \zeta \sin \theta, R(\phi))^\top$ as a mapping $P : \mathbb{R}^3 \rightarrow \text{SE}(2)$, where $R(\phi)$ is a rotation of the plane by ϕ radians counterclockwise. Also define the standard spherical coordinates transformation $S(\varphi, \delta) = (\cos \varphi \sin \delta, \sin \varphi \sin \delta, \cos \delta)^\top$ as a mapping $S : \mathbb{R}^2 \rightarrow S^2$. P and S are surjective local diffeomorphisms.

By Lemma 2.8,

$$\eta_{\text{SO}(3)} \circ \Psi_\lambda \circ P(\zeta, \theta, \phi) = \left(\cos \theta \sin \left(\frac{\zeta}{\lambda} \right), \sin \theta \sin \left(\frac{\zeta}{\lambda} \right), \cos \left(\frac{\zeta}{\lambda} \right) \right)^\top.$$

Since S is a local diffeomorphism about S^2 , it is locally invertible around every point of S^2 . This implies

$$(A.4) \quad S^{-1} \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda \circ P(\zeta, \theta, \phi) = \left(\theta, \frac{\zeta}{\lambda} \right)^\top.$$

Thus, $S^{-1} \circ \eta_{\text{SO}(3)} \circ \Psi_\lambda \circ P(\zeta, \theta, \phi)$ is locally defined and locally linear, implying $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ is a smooth submersion. Finally, (A.4), the properties of the spherical coordinates and polar coordinates easily yield the characterization of the fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ in (2.5), (2.6), and (2.7). ■

Proof of Lemma 2.6. Since f is a smooth function on a homogeneous space of $\text{SE}(2)$, we can lift it to a function on $\text{SE}(2)$ that is constant on left cosets of $\text{SO}(2)$. Thus, without loss of generality, assume $f : \text{SE}(2) \rightarrow \mathbb{R}$. Now note that the fibers of $\eta_{\text{SO}(3)} \circ \Psi_\lambda$ of forms (2.5) and (2.7) intersect $B_{\lambda\pi} \times \text{SO}(2)$ only once and that the fiber of the form (2.6) only intersects $\partial B_{\lambda\pi} \times \text{SO}(2)$. Define $\tilde{f}$ on fibers of the form (2.5) and (2.7) to be the value of f on their intersection with $B_{\lambda\pi} \times \text{SO}(2)$ and zero on fibers (2.6).

The uniqueness of $\tilde{f}$ is immediate from our construction. It remains to prove $\tilde{f}$ is smooth. Let $B^{(n)} := (B_{\lambda\pi n} \setminus \{0\}) \setminus \overline{B}_{\lambda\pi(n-1)}$. Define a map $p : B_{\lambda\pi} \setminus \{0\} \rightarrow B^{(n)}$ as $p(\zeta, \theta) = (\zeta + \lambda\pi(n-1), \theta)$. This is obviously a (local) diffeomorphism. Note that for $(\mathbf{x}, R) \in B^{(n)} \times \text{SO}(2) \subset \text{SE}(2)$ we have $\tilde{f}(\mathbf{x}, R) = f \circ p^{-1}(\mathbf{x})$. Since both f and p are smooth, it follows that $\tilde{f}$ is smooth in a neighborhood of $\mathbf{x}$. It remains to prove $\tilde{f}$ is smooth for all $(\mathbf{x}, R) \in \partial B_{\lambda\pi n} \times \text{SO}(2)$. Since f is compactly supported in $B_{\lambda\pi}$, it follows that it is identically zero in a neighborhood V of $\mathbf{x}$. Therefore, making V smaller if necessary, we conclude that $\tilde{f}$ is identically zero on $V \times \text{SO}(2)$ and so it is smooth there as well. ■

Proof of Lemma 2.14. Recall the identity $2\sin^2\left(\frac{x}{2}\right) = 1 - \cos x$ and the inequality $\sin x \geq \frac{2}{\pi}x$ for all $x \in [0, \pi/2]$. Combining them it follows that

$$(A.5) \quad 2(1 - \cos x) = 4\sin^2\left(\frac{x}{2}\right) \geq 4 \cdot \left(\frac{2}{\pi}x\right)^2 = \frac{16}{\pi^2}x^2 \geq \frac{x^2}{\pi}.$$

Now, let $\mathbf{x}, \mathbf{y} \in S^2$ and denote $\theta = d_{S^2}(\mathbf{x}, \mathbf{y})$. Thus

$$\begin{aligned} \|\mathbf{x} - \mathbf{y}\|^2 &= \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\mathbf{y}^\top \mathbf{x} \\ &= 2(1 - \cos \theta) && \text{Because } x, y \in S^2 \\ &\geq \frac{\theta^2}{\pi} && \text{From (A.5).} \end{aligned}$$

The result easily follows. ■

Proof of Lemma 2.12. Let A and B be real $n \times n$ matrices. From the definition of the exponential map on matrix groups we obtain

$$\begin{aligned} T := \exp(A + B) - \exp(A)\exp(B) &= \sum_{n=0}^{\infty} \frac{(A + B)^n}{n!} - \left(\sum_{n=0}^{\infty} \frac{A^n}{n!}\right) \left(\sum_{n=0}^{\infty} \frac{B^n}{n!}\right) \\ &= \sum_{n=0}^{\infty} \frac{(A + B)^n}{n!} - \sum_{k=0}^{\infty} \sum_{m=0}^{\infty} \frac{A^k B^m}{k!m!} \\ &= \sum_{n=2}^{\infty} \frac{(A + B)^n}{n!} - \sum_{k=1}^{\infty} \sum_{m=1}^{\infty} \frac{A^k B^m}{k!m!} \\ &= \sum_{n=2}^{\infty} \frac{(A + B)^n}{n!} - \left(\sum_{k=1}^{\infty} \frac{A^k}{k!}\right) \left(\sum_{m=1}^{\infty} \frac{B^m}{m!}\right), \end{aligned}$$

and so

$$\|T\| \leq \sum_{n=2}^{\infty} \frac{(\|A\| + \|B\|)^n}{n!} + \left(\sum_{k=1}^{\infty} \frac{\|A\|^k}{k!}\right) \left(\sum_{m=1}^{\infty} \frac{\|B\|^m}{m!}\right).$$

Substituting $A = \frac{X}{\lambda}$ and $B = \frac{Y}{\lambda}$ and using the fact $\lambda \geq 1$,

$$\begin{aligned} &\left\| \exp\left(\frac{X}{\lambda} + \frac{Y}{\lambda}\right) - \exp\left(\frac{X}{\lambda}\right) \exp\left(\frac{Y}{\lambda}\right) \right\| \\ &\leq \frac{1}{\lambda^2} \sum_{n=2}^{\infty} \frac{(\|X\| + \|Y\|)^n}{n!} + \frac{1}{\lambda^2} \left(\sum_{k=1}^{\infty} \frac{\|X\|^k}{k!}\right) \left(\sum_{m=1}^{\infty} \frac{\|Y\|^m}{m!}\right) \\ &= \frac{1}{\lambda^2} \left(e^{\|X\| + \|Y\|} - \|X\| - \|Y\| - 1 + (e^{\|X\|} - 1)(e^{\|Y\|} - 1)\right) \\ &= \frac{1}{\lambda^2} \left(2e^{\|X\| + \|Y\|} - e^{\|X\|} - e^{\|Y\|} - \|X\| - \|Y\|\right). \end{aligned} \quad \blacksquare$$

Appendix B. Debiasing the spherical bispectrum estimator. Given a sample of N images as in (4.1) and their projections $\{\mathbf{s}_j\}_{j=1}^N$ onto the sphere, our objective is to derive an unbiased estimator $\hat{\mathbf{b}}$ of the bispectrum of f . We begin by establishing some notation, making the derivation easier. Denote by $\mathbf{P} = \mathbf{P}(L, \mathcal{S}_t, \lambda)$ the interpolation operator used when projecting an image onto the space of functions on the sphere of bandlimit L . This is a linear operator. Also, recall the definition of $\mathbf{Y} = \mathbf{Y}(L, \mathcal{S}_t)$ in (3.4).

Taking into account (4.1), the spherical harmonics coefficients we estimate for the j th image are of the form

$$(B.1) \quad \mathbf{s}_j = \mathbf{U}\mathbf{J}_j + \mathbf{U}\boldsymbol{\varepsilon}_j, \text{ where } \mathbf{U} = \mathbf{Y}^\dagger \mathbf{P} \text{ is a complex matrix, } \mathbf{J}_j = \mathbf{D}(g_j \bullet f)$$

and $\mathbf{s}_j$ is the concatenation of all spherical harmonics coefficients of the projection onto the sphere of $\mathbf{J}_j$ arranged in lexicographical order. Note that due to the indexing convention we used for $\mathbf{Y}$, $\mathbf{U}$ has rows indexed by (ℓ, m) . Since $\mathbf{P}$ operates on images of $n \times n$ pixels, the columns are indexed by $1, \dots, n^2$.

In order to obtain an unbiased estimator, we need to calculate the expected value of $b_{\mathbf{s}_j}[\ell_1, \ell_2, \ell]$. As we noted in subsection 3.1, our use of interpolation when projecting an image onto the sphere introduces correlations in the noise. In order to deal with them, we first observe that the spherical bispectrum (2.25) is multilinear as a function of $\mathbf{f}_\ell$, $\mathbf{f}_{\ell_1}^*$, and $\mathbf{f}_{\ell_2}^*$, where $\mathbf{f}_k = (f_{k,-k}, f_{k,-k+1}, \dots, f_{k,k})^\top$. In particular, for a fixed (ℓ_1, ℓ_2, ℓ) triplet satisfying (2.26), we can write $b_{\mathbf{s}_j}[\ell_1, \ell_2, \ell] = B[\mathbf{s}_j^*, \mathbf{s}_j^*, \mathbf{s}_j]$ for some multilinear transformation B . This enables us to write

$$(B.2) \quad \begin{aligned} \mathbb{E}[b_{\mathbf{s}_j}[\ell_1, \ell_2, \ell]] &= \mathbb{E}[B[\mathbf{s}_j^*, \mathbf{s}_j^*, \mathbf{s}_j]] \\ &= \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\mathbf{J}_j]] + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\mathbf{J}_j]] \\ &\quad + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]] + \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]] \\ &\quad + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\mathbf{J}_j]] + \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\mathbf{J}_j]] \\ &\quad + \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]] + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]]. \end{aligned}$$

Consider the form of B (see (2.25)) and elements in the sum on the right-hand side of (B.2) that depend on either exactly one or exactly three $\mathbf{U}\boldsymbol{\varepsilon}_j$ (the last four terms in the sum). Since every element of $\boldsymbol{\varepsilon}_j$ is Gaussian i.i.d. with zero mean, we have

$$\mathbb{E}[\boldsymbol{\varepsilon}_{j,i_1,i_2}] = \mathbb{E}[\boldsymbol{\varepsilon}_{j,i_1,i_2}\boldsymbol{\varepsilon}_{j,i_3,i_4}\boldsymbol{\varepsilon}_{j,i_5,i_6}] = 0.$$

Therefore, said elements in the sum in (B.2) are zero and we obtain the following:

$$(B.3) \quad \begin{aligned} \mathbb{E}[b_{\mathbf{s}_j}[\ell_1, \ell_2, \ell]] &= \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\mathbf{J}_j]] + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\mathbf{J}_j]] \\ &\quad + \mathbb{E}[B[(\mathbf{U}\boldsymbol{\varepsilon}_j)^*, (\mathbf{U}\mathbf{J}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]] + \mathbb{E}[B[(\mathbf{U}\mathbf{J}_j)^*, (\mathbf{U}\boldsymbol{\varepsilon}_j)^*, \mathbf{U}\boldsymbol{\varepsilon}_j]]. \end{aligned}$$

Now, denote the elements on the right-hand side of (B.3) by $K_0 = K_{\mathbf{s}_j,0}[\ell_1, \ell_2, \ell]$, $K_1 = K_{\mathbf{s}_j,1}[\ell_1, \ell_2, \ell]$, $K_2 = K_{\mathbf{s}_j,2}[\ell_1, \ell_2, \ell]$, and $K_3 = K_{\mathbf{s}_j,3}[\ell_1, \ell_2, \ell]$ by order of their appearance. K_0 does not depend on the noise. It depends only on $\mathbf{U}\mathbf{J}_j$ the projection onto the sphere of $\mathbf{J}_j$,

a discretization of the rotated and translated f , the ground truth function. We assume the translation is small enough that $\mathbf{U}\mathbf{J}_j$ is approximately $\kappa_\lambda f$ rotated on the sphere. Thus,

$$(B.4) \quad K_0 \approx b_{\kappa_\lambda f} [\ell_1, \ell_2, \ell].$$

In order to handle K_1 , note that by (2.25) it has the form

$$(B.5) \quad K_{\mathbf{s}_j,1} = \sum_{m=-\ell}^{\ell} \sum_{m_1=-\ell_1}^{\ell_1} C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m} f_{\ell,m} \mathbb{E} \left[(\mathbf{U}\boldsymbol{\varepsilon}_j)_{\ell_1,m_1}^* (\mathbf{U}\boldsymbol{\varepsilon}_j)_{\ell_2,m-m_1}^* \right].$$

Using the fact $\boldsymbol{\varepsilon}_{j,k,m} \sim \mathcal{N}(0, \sigma^2)$ i.i.d., we get

$$\begin{aligned} \mathbb{E} \left[(\mathbf{U}\boldsymbol{\varepsilon}_j)_{\ell_1,m_1}^* (\mathbf{U}\boldsymbol{\varepsilon}_j)_{\ell_2,m-m_1}^* \right] &= \mathbb{E} \left[\left(\sum_{\mathbf{k}_1} \mathbf{U}_{(\ell_1,m_1),\mathbf{k}_1}^* \boldsymbol{\varepsilon}_{j,\mathbf{k}_1} \right) \left(\sum_{\mathbf{k}_2} \mathbf{U}_{(\ell_2,m-m_1),\mathbf{k}_2}^* \boldsymbol{\varepsilon}_{j,\mathbf{k}_2} \right) \right] \\ &= \sum_{\mathbf{k}_1} \sum_{\mathbf{k}_2} \mathbf{U}_{(\ell_1,m_1),\mathbf{k}_1}^* \mathbf{U}_{(\ell_2,m-m_1),\mathbf{k}_2}^* \mathbb{E} [\boldsymbol{\varepsilon}_{j,\mathbf{k}_1} \boldsymbol{\varepsilon}_{j,\mathbf{k}_2}] \\ &= \sum_{\mathbf{k}_1} \sum_{\mathbf{k}_2} \mathbf{U}_{(\ell_1,m_1),\mathbf{k}_1}^* \mathbf{U}_{(\ell_2,m-m_1),\mathbf{k}_2}^* \cdot \sigma^2 \delta_{\mathbf{k}_1,\mathbf{k}_2} \\ &= \sigma^2 \sum_{\mathbf{k}_2} \mathbf{U}_{(\ell_1,m_1),\mathbf{k}}^* \mathbf{U}_{(\ell_2,m-m_1),\mathbf{k}}^* \\ &= \sigma^2 \left(\mathbf{U}\mathbf{U}^\top \right)_{(\ell_1,m_1),(\ell_2,m-m_1)}^*. \end{aligned}$$

Substituting that into (B.5), it follows that

$$(B.6) \quad K_1 = \sigma^2 \sum_{m=-\ell}^{\ell} \mathbf{s}_{j,(\ell,m)} \sum_{m_1=\max\{-\ell_1,m-\ell\}}^{\min\{\ell_1,m+\ell_2\}} C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m} \left(\mathbf{U}\mathbf{U}^\top \right)_{(\ell_1,m_1),(\ell_2,m-m_1)}^*.$$

Using a similar approach and remembering that the Clebsch–Gordan coefficients are nonzero only if $m_1 + m_2 = m$, one can show that

$$(B.7) \quad K_2 = \sigma^2 \sum_{m_2=-\ell_2}^{\ell_2} \mathbf{s}_{j,(\ell_2,m_2)}^* \sum_{m=\max\{-\ell,m_1-\ell_1\}}^{\min\{\ell,m+\ell_2\}} C_{\ell_1,m-m_2,\ell_2,m_2}^{\ell,m} \left(\mathbf{U}\mathbf{U}^\dagger \right)_{(\ell,m),(\ell_1,m_1)},$$

$$(B.8) \quad K_3 = \sigma^2 \sum_{m_1=-\ell_1}^{\ell_1} \mathbf{s}_{j,(\ell_1,m_1)}^* \sum_{m=\max\{-\ell,m_1-\ell_1\}}^{\min\{\ell,m+\ell_2\}} C_{\ell_1,m_1,\ell_2,m-m_1}^{\ell,m} \left(\mathbf{U}\mathbf{U}^\dagger \right)_{(\ell,m),(\ell_2,m_1)}.$$

Combining (B.3) and (B.4) and using the explicit expressions of K_1 , K_2 , and K_3 in (B.6), (B.7), and (B.8), we obtain an unbiased estimator of $\hat{\mathbf{b}}[\ell_1, \ell_2, \ell]$:

$$\hat{\mathbf{b}}[\ell_1, \ell_2, \ell] = \frac{1}{N} \sum_{j=1}^N \left(\mathbf{b}_{\mathbf{s}_j}[\ell_1, \ell_2, \ell] - \sum_{k=1}^3 K_{\mathbf{s}_j,k}[\ell_1, \ell_2, \ell] \right).$$

When implementing this estimator, it is important to note that $K_{\mathbf{s}_j,k}$ is a linear transformation, when applied to the realized version of $\mathbf{s}_j$. It is also sparse. Implementing it as such enables one to easily implement this estimator as a simple matrix multiplication and can considerably speed up its calculation.

Appendix C. Noise statistics of projected images. We wished to study the effects of the projection onto the sphere on noise. Specifically, given an image with additive white noise, $\mathbf{Y} = \mathbf{I} + \boldsymbol{\varepsilon}$, we expect the spherical harmonics coefficients to have additive noise, because the projection onto the sphere is linear. We sought to test empirically whether the noise on the sphere can still be regarded as white noise.

To that end, we generated 10^4 white noise images of size 101×101 . Specifically, each pixel was Gaussian i.i.d. with zero mean and variance 1. The power spectrum of every image was calculated and averaged over the corresponding pixels. We then projected every image onto the sphere as described in [subsection 3.1](#), estimating the spherical harmonics coefficients of the projection up to bandlimit 50. Then, we calculated the normalized power spectrum of every projected image and averaged them. Note that the normalized power spectrum P of a function f on the sphere for the ℓ th frequency is given by

$$(C.1) \quad P[f]_\ell = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |f_{\ell,m}|^2.$$

The results of this experiment are shown in [Figure 9](#). The average power spectrum of the original images and the normalized power spectrum on the sphere behaved as expected from the power spectra of white noise. In both cases, the power spectrum was approximately constant over all frequencies.




Figure 9. Left: the power spectrum of a white noise image, averaged over 10^4 noise realizations. The frequencies are ordered lexicographically. Right: the power spectrum (C.1) of the corresponding projection of the image up to bandlimit 50, averaged over the same 10^4 noise realizations projected onto the sphere.

Acknowledgment. We thank Nicolas Boumal for his major contribution to this project's conception and first steps; this project is also his brainchild.

REFERENCES

- [1] *ASPIRE—Algorithms for Single Particle Reconstruction*, <http://spr.math.princeton.edu/>.
- [2] K. ATKINSON AND W. HAN, *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, Lecture Notes in Math. 2044, Springer, New York, 2012.
- [3] A. S. BANDEIRA, B. BLUM-SMITH, J. KILEEL, A. PERRY, J. WEED, AND A. S. WEIN, *Estimation Under Group Actions: Recovering Orbits from Invariants*, preprint, <https://arxiv.org/abs/1712.10163>, 2017.
- [4] A. S. BANDEIRA, J. NILES-WEED, AND P. RIGOLLET, *Optimal rates of estimation for multi-reference alignment*, Math. Statistics Learning, 2 (2020), pp. 25–75.
- [5] T. BENDORY, A. BARTESAGHI, AND A. SINGER, *Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities*, IEEE Signal Processing Magazine, 37 (2020), pp. 58–76.
- [6] T. BENDORY, N. BOUMAL, C. MA, Z. ZHAO, AND A. SINGER, *Bispectrum inversion with application to multireference alignment*, IEEE Trans. Signal Processing, 66 (2018), pp. 1037–1050, <https://doi.org/10.1109/TSP.2017.2775591>.
- [7] L. CARLONE, R. TRON, K. DANIILIDIS, AND F. DELLAERT, *Initialization techniques for 3D slam: A survey on rotation estimation and its use in pose graph optimization*, in Proceedings of the International Conference on Robotics and Automation, IEEE, 2015, pp. 4597–4604.
- [8] A. H. DOOLEY, *Contractions of Lie groups and applications to analysis*, in Topics in Modern Harmonic Analysis, Istitutio Nazionale di alta Matematica, 1984, pp. 483–515.
- [9] N. HALKO, P.-G. MARTINSSON, Y. SHKOLNISKY, AND M. TYGERT, *An algorithm for the principal component analysis of large data sets*, SIAM J. Sci. Comput., 33 (2011), pp. 2580–2594, <https://doi.org/10.1137/100804139>.
- [10] S. HELGASON, *Differential Geometry, Lie Groups, and Symmetric Spaces*, Academic Press, New York, 1978.
- [11] E. INONU AND E. P. WIGNER, *On the contraction of groups and their representations*, Proc. Natl. Acad. Sci. USA, 39 (1953), pp. 510–524, <https://doi.org/10.1073/pnas.39.6.510>.
- [12] R. KAKARALA, *Triple Correlation on Groups*, Ph.D. thesis, The University of California, Irvine, 1992.
- [13] R. KAKARALA, B. M. BENNETT, G. J. IVERSON, AND M. D'ZMURA, *Bispectral techniques for spherical functions*, in Proceedings of the International Conference on Acoustics, Speech, and Signal Processing, Vol. 4, IEEE, 1993, pp. 216–219, <https://doi.org/10.1109/ICASSP.1993.319633>.
- [14] R. KAKARALA AND D. MAO, *A theory of phase-sensitive rotation invariance with spherical harmonic and moment-based representations*, in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2010, pp. 105–112, <https://doi.org/10.1109/CVPR.2010.5540222>.
- [15] R. KONDOR, *A Novel Set of Rotationally and Translationally Invariant Features for Images Based on the Non-commutative Bispectrum*, 2007, <https://ui.adsabs.harvard.edu/abs/2007cs.....1127K>.
- [16] J. M. LEE, *Introduction to Smooth Manifolds*, 2nd ed., Springer, New York, 2013.
- [17] C. MA, T. BENDORY, N. BOUMAL, F. SIGWORTH, AND A. SINGER, *Heterogeneous multireference alignment for images with application to 2D classification in single particle reconstruction*, IEEE Trans. Image Processing, 29 (2020), pp. 1699–1710, <https://doi.org/10.1109/TIP.2019.2945686>.
- [18] O. ÖZYİSİL, N. SHARON, AND A. SINGER, *Synchronization over cartan motion groups via contraction*, SIAM J. Appl. Algebra Geometry, 2 (2018), pp. 207–241, <https://doi.org/10.1137/16M1106055>.
- [19] A. PERRY, J. WEED, A. S. BANDEIRA, P. RIGOLLET, AND A. SINGER, *The sample complexity of multireference alignment*, SIAM J. Math. Data Sci., 1 (2019), pp. 497–517.
- [20] D. M. ROSEN, L. CARLONE, A. S. BANDEIRA, AND J. J. LEONARD, *SE-Sync: A certifiably correct algorithm for synchronization over the special euclidean group*, Internat. J. Robotics Research, 38 (2019), pp. 95–125.
- [21] W. O. STRAUB, *Efficient Computation of Clebsch-Gordan Coefficients*, <http://vixra.org/abs/1403.0263>, 2014.

- [22] R. TRON, X. ZHOU, AND K. DANIILIDIS, *A survey on rotation optimization in structure from motion*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, 2016, pp. 77–85.
- [23] D. A. VARSHALOVICH, A. N. MOSKALEV, AND V. K. KHERSONSKII, *Quantum Theory of Angular Momentum*, World Scientific, River Edge, NJ, 1988.
- [24] Y. G. WANG, Q. T. LE GIA, I. H. SLOAN, AND R. S. WOMERSLEY, *Fully discrete needlet approximation on the sphere*, Appl. Comput. Harmon. Anal., 43 (2017), pp. 292–316, <https://doi.org/10.1016/j.acha.2016.01.003>.
- [25] R. S. WOMERSLEY, *Efficient Spherical Designs on the Sphere s^2 with Good Geometric Properties*, <https://web.maths.unsw.edu.au/~rsw/Sphere/EffSphDes/sf.html>, 2015.
- [26] R. S. WOMERSLEY, *Efficient Spherical Designs with Good Geometric Properties*, <https://arxiv.org/abs/1709.01624>, 2017.
- [27] A. YERSHOVA, S. JAIN, S. M. LAVALLE, AND J. C. MITCHELL, *Generating uniform incremental grids on $SO(3)$ using the Hopf fibration*, Internat. J. Robotics Research, 29 (2009), pp. 801–812, <https://doi.org/10.1177/0278364909352700>.
- [28] Z. ZHAO, Y. SHKOLNISKY, AND A. SINGER, *Fast steerable principal component analysis*, IEEE Trans. Comput. Imaging, 2 (2016), pp. 1–12, <https://doi.org/10.1109/TCI.2016.2514700>.
- [29] Z. ZHAO AND A. SINGER, *Fourier–bessel rotational invariant eigenimages*, JOSA A, 30 (2013), pp. 871–877.
- [30] Z. ZHAO AND A. SINGER, *Rotationally invariant image representation for viewing direction classification in cryo-em*, J. Structural Biology, 186 (2014), pp. 153–166, <https://doi.org/10.1016/j.jsb.2014.03.003>.