

The Selective Disk Bispectrum and Its Inversion, with Application to Multi-Reference Alignment

Adele Myers Lantow
UC Santa Barbara
adele@ucsb.edu

Nina Miolane
UC Santa Barbara
ninamiolane@ucsb.edu

Abstract

In many computer vision and shape analysis tasks, practitioners are interested in learning from the shape of the object in an image, while disregarding the object’s orientation. To this end, it is valuable to define a rotation-invariant representation of images, retaining all information about that image, but disregarding the way an object is rotated in the frame. To be practical for learning tasks, this representation must be computationally efficient for large datasets and invertible, so the representation can be visualized in image space. To this end, we present the selective disk bispectrum: a fast, rotation-invariant representation for image shape analysis. While the translational bispectrum has long been used as a translational invariant representation for 1D and 2D signals, its extension to 2D (disk) rotational invariance on images has been hindered by the absence of an invertible formulation and its cubic complexity. In this work, we derive an explicit inverse for the disk bispectrum, which allows us to define a “selective” disk bispectrum, which only uses the minimal number of coefficients needed for faithful shape recovery. We show that this representation enables multi-reference alignment for rotated images—a task previously intractable for disk bispectrum methods. These results establish the disk bispectrum as a practical and theoretically grounded tool for learning on rotation-invariant shape data.

1. Introduction

The translation-invariant bispectrum is a higher-order spectral representation that originated in the mid-20th century for signal processing [4, 28, 37]. Unlike the power spectrum, which captures second-order (pairwise) correlations between Fourier coefficients, the bispectrum encodes third-order relationships, making it sensitive to asymmetries and nonlinear interactions in the signal. More importantly, unlike the power spectrum which drops the phase of the signal, the translation-invariant bispectrum conserves every-

thing about the signal except for its global phase, making it a *complete* translational invariant.

As a result, the translation-invariant bispectrum has found widespread applications across disciplines: oceanographers use it to analyze wave interactions [10], seismologists to identify complex waveforms [8], and biomedical engineers to study electroencephalograms [38] as well as heart rate variability [26]. It is similarly used in optical interferometry [14], radar communications [22] and other applications as a way to identify patterns in 1D signals (see [21] for a recent review). Such applications are only interested in translation invariance in 1D or 2D signals (images).

In contrast, other applications are interested in rotation invariance, and despite the widespread applicability and utilization of the 1D/2D translation-invariant bispectrum, the 2D rotation-invariant bispectrum has been under-developed. In many applications, key information lies in the *shape* of a 2D image, and its orientation is a nuisance variable. Without an easily computable rotation invariant representation, practitioners must register each image before analysis.

To this end, [40] proposed the disk bispectrum, which is invariant to 2D planar (disk) rotations. However, because they did not derive a bispectrum inversion, their analysis was restricted to classification problems in bispectrum space. For example, tasks like multi reference alignment (MRA) would be impossible. MRA estimates a ground truth signal by mapping shifted noisy versions of that signal into bispectrum space, averaging the signals in bispectrum space, and then inverting the average back to image space, which would be impossible without a bispectrum inversion. Additionally [40]’s disk bispectrum was computationally expensive. Authors reported a cubic complexity of $\mathcal{O}(m^3/N_m)$, where m is the number of “disk harmonic frequencies”, and N_m the maximum frequency. Lastly, they resorted to principal component analysis (PCA) in bispectrum space to reduce its dimensionality. Authors do not prove completeness of the disk bispectrum or their PCA reduction, meaning it was unclear whether the bispectrum actually described all information about the original image up to a rotation.

In this work, we propose the first computationally tractable disk bispectrum, the *selective disk bispectrum*, and its inversion, for use in image processing and shape analysis. We also provide a method for approximating the disk bispectrum, building on [25]. Then, following [3], we derive an MRA-specific disk bispectrum inversion, which corrects for noise propagation during disk bispectrum averaging, and demonstrate its efficacy on rotated noisy MNIST images.

Contributions

1. We introduce the *selective disk bispectrum*, an efficient formulation that proves and exploits the redundancy of disk bispectral coefficients. This reduces the space complexity from $\mathcal{O}(m^3/N_m)$ to $\mathcal{O}(m)$ and the time complexity from $\mathcal{O}(L^3 + m^3/N_m)$ to $\mathcal{O}(L^3 + m)$, where $L \times L$ is the image size, L^2 is the number of total pixels, m is the number of disk harmonic frequencies, and N_m the maximum frequency. The selective disk bispectrum offers an *efficient* rotation-invariant representation, making it scalable for large amounts of data.
2. We derive the first disk bispectrum inversion, mapping (selective) disk bispectrum coefficients back to image space. The value of this inversion is twofold. First, it enables practitioners to reconstruct an image from statistical results (such as computation of the mean) in disk bispectrum space, which is essential for interpretability. Second, the inversion proves that the selective disk bispectrum is a *complete rotational invariant* for images on the disk, i.e., it preserves all information about an image up to a disk rotation. By association, it also provides the first completeness proof for the full disk bispectrum.
3. We provide a numerical approximation of the selective disk bispectrum, that further reduces its time complexity from $\mathcal{O}(L^3 + m)$ to $\mathcal{O}(L^2 \log L + m)$. We prove precise accuracy guarantees on this approximation.
4. We provide an open-source, unit-tested implementation of the selective disk bispectrum and its inversion in Python, along with its application to multi-reference alignment. We use this implementation to illustrate the forward and inverse selective disk bispectrum map (between images and coefficients), and compare against the full disk bispectrum map. We empirically demonstrate that the selective disk bispectrum can be used for multi-reference alignment on rotated images.

In summary, the selective disk bispectrum offers a fast (scalable) and invertible (interpretable) tool for imaging applications involving rotation-invariant 2D shape analysis.

2. Related Works

Equivariant and Invariant Representations. As geometric learning has gained popularity [5], there has been increased interest in rotation-equivariant and rotation-invariant image representations. Put simply (and described

Method	Equiv.	Local Inv.	Global Inv.	Mag. Inv.	Invertible Inv. Rep.
SDB(ours)	✗	✓	✓	✓	✓
DH Tr. [24]	✓	✗	✗	✗	✗
Log-Polar Tr. [30]	✓	✗	✗	✓	✗
Fourier Mellin Tr. [30]	✓	✗	✗	✓	✗
SIFT / SURF / ORB / AKAZE					
[1, 2, 23, 31]	✗	✓	✗	✗	✗
Hu Moments [16]	✗	✓	✓	✗	✗
Zernike Moments [35]	✗	✗	✗	✓	✗
TFNs [36]	✓	✗	✗	✓	✗
Spherical convs. [7, 11]	✓	✗	✗	✗	✗
Capsule Networks [32]	✓	✗	✗	✗	✗

Table 1. Comparison of equivariance, invariance, and invertibility properties across image descriptors. Abbreviations: TFN Tensor Field Networks, Tr. Transform, Inv. Invariant, Equiv. Equivariant, SIFT Scale-Invariant Feature Transform, SURF Speeded-Up Robust Features, ORB Oriented FAST and Rotated BRIEF, AKAZE Accelerated KAZE.

later in more detail), if a representation is *equivariant* to rotation, the representation changes in a specific way when the original image is rotated. If a representation is rotation *invariant*, the representation does not change if the original image is rotated.

Table 1 summarizes rotation equivariant and rotation invariant representations, showing that only our selective disk bispectrum is rotation invariant and invertible, making it the only method capable of multi-reference alignment on rotated images. As an aside, we note that every rotation-equivariant neural network (and every rotation-equivariant layer) can be made rotation-invariant by adding an invariant aggregation function, like the sum or the max, after it. However, because the sum or max cannot be inverted to reconstruct the input, this invariant shape descriptor is not invertible.

G-bispectra and their inversions. Kakarala [17] proposed the *G*-Bispectrum, defined for a signal whose domain is a group. The *G*-Bispectrum gained interest as it is the lowest-degree spectral invariant that is complete. The translation-invariant bispectrum, invariant to the cyclic group $G = C_n$, is a classic and widely used signal processing tool, and is an example of a *G*-bispectrum. [12] proposed the first inversion scheme of the translation-invariant bispectrum, reconstructing a 1D signal analytically from translation bispectrum coefficients. [34] developed another analytical inversion algorithm that is more robust to noise, and [29] revisited this inversion problem, providing a gradient-based inversion procedure. [12] also extended his 1D inversion scheme to invert the 2D translation-invariant bispectrum, invariant to the product of cyclic groups $C_n \times$




Figure 1. We propose the selective disk bispectrum, which reduces the time complexity of the disk bispectrum from $\mathcal{O}(m^3/N_m)$ to $\mathcal{O}(m)$, where m is the number of disk harmonic frequencies, and N_m the maximum frequency. We provide the first inversion of the disk bispectrum from its coefficients back to image space.

C_n . The 2D translation-invariant bispectrum is another example of a G -Bispectrum. [27] generalized that proof for the product of any number of cyclic groups, *i.e.*, to any G -bispectrum, for signals defined over a finite and commutative group G . A reflection-invariant bispectrum was also proposed by Edidin and Katz [9], who formulated a gradient descent optimization for its inversion. Methods to invert the G -bispectrum of signals defined on a (continuous) Lie group G have been proposed in [18, 20]. Note, the disk bispectrum is not covered under any of the aforementioned proofs, as the signal is defined over a planar disk, which is not a group.

The G -bispectrum on homogeneous spaces, defined for a signal whose domain is homogeneous for a group G , was proposed by [19]. A space that is homogeneous for a group is one where for any x, y there is a group element $g \in G$ s.t. $g \circ x = y$. [19] also proposed a method to invert the G -bispectrum of signals defined on spaces homogeneous for a group G , specifically for the sphere S^2 , which is homogeneous for $SO(3)$. Existing rotation-invariant G -bispectra are defined for signals on the circle ($SO(2)$) or the sphere ($SO(3)$). Note, the disk bispectrum is not covered under any of these proofs, as the signal is defined over a planar disk, which is not homogeneous for the rotation group $SO(2)$.

Selective G -Bispectra. Across existing G -bispectra, several authors have shown that bispectrum coefficients are generally redundant. It is usually possible to recover the signal with significantly fewer bispectrum coefficients than those provided in the full bispectrum, thus enabling a *selective* version of that bispectrum. The work of [13, 17, 33] showed that for *some* finite, commutative groups, the G -Bispectrum can be computed with only $|G|$ space complexity, where $|G|$ is the number of elements in the group G . Maigne et al. [27] extended these results by showing this result for a wider range of finite groups, *i.e.*, all discrete commutative groups, and some non-commutative ones such as the dihedral groups of any order, the octahedral and full octahedral group. [27] further proposed inversion of selective bispectra for the following non-commutative groups:

any dihedral group, the octahedral group, and the full octahedral group. [27] estimates bispectra (for signals defined on a group) via a group convolution, and offers an algorithm that can be used to identify selective coefficients. In contrast, the disk bispectrum provides an analytic map from the image to the bispectrum coefficients and is defined for signals defined over the full disk, so that it can be applied to images. Additionally, our disk bispectrum inversion identifies the exact selective coefficients required for complete bispectrum inversion. In fact, this work is the first to introduce a selective bispectrum and inversion proof for a non-homogeneous bispectrum.

A Disk Bispectrum. Most relevant to our work is the work by Zhao and Singer [40], who proposed the first disk bispectrum. As described in introduction, this full disk bispectrum is impracticable due to its cubic complexity and lack of inversion, which is why we propose the selective disk bispectrum and its inversion.

Related to our work, Kakarala [18] proposes inverting a non-complete bispectrum-like function on the disk by introducing an auxiliary function that supplies the missing information needed for reconstruction and using the inversion method of the cyclic group. In contrast, our disk bispectrum inversion is complete and requires no external information to reconstruct the original image from bispectral coefficients.

Multi-reference alignment We will show the utility of our bispectrum and its inversion to perform multi-reference alignment (MRA). Traditional *translational* MRA utilizes the translation bispectrum, but special methods are required to adjust for a bias induced by noise in bispectrum space. Various techniques have been proposed for this task, such as noise bias estimation and correction [3], spectral initialization approaches to MRA optimization [6], and Gauss-Newton optimization schemes [15]. These corrections, however, have been derived only for the *1D translational bispectrum*. No such correction has been established for the *disk bispectrum*. In this work, we not only propose

the *selective disk bispectrum*, its inversion, and its approximation, but we also derive the corresponding MRA noise correction term in a similar manner to [3]. Notably, MRA with the disk bispectrum would not be possible without our inversion method, since reconstructing the mean bispectrum back to image space requires a well-defined inverse.

3. Background

Disk Harmonics. The 1D translation bispectrum (which is invariant to group action on the circle S^1) is composed of Fourier coefficients, which are constructed using Fourier modes, which are eigenfunctions of the Laplacian on S^1 . Thus, naturally, if we wish to compose a bispectrum invariant to disk rotations, we should do so with disk harmonic coefficients, which are constructed using disk harmonics, which are the eigenfunctions of the Laplacian on the unit disk $\mathbb{D} = \{x \in \mathbb{R}^2 : \|x\| < 1\}$. In polar coordinates $(r, \theta) \in [0, 1) \times [0, 2\pi)$, these disk harmonics are given by

$$\psi_{nk}(r, \theta) = c_{nk} J_n(\lambda_{nk} r) e^{in\theta}$$

where $(n, k) \in \mathbb{Z} \times \mathbb{Z}_{>0}$, J_n is the n -th order Bessel function of the first kind, λ_{nk} is its k -th root of the n -th Bessel function, and c_{nk} is a normalization constant. Thus, radially, $\psi_{nk}(r, \theta)$ looks like a Bessel function whose value is λ_{nk} on the edge of the disk, and angularly, $\psi_{nk}(r, \theta)$ looks like a complex sinusoid.

Disk Harmonic (DH) Transform. Let $f : \mathbb{D} \rightarrow \mathbb{R}$ be a real-valued square integrable function defined on the unit disk. Disk harmonics form an orthonormal basis for square-integrable functions (like Fourier modes), such that f can be expanded as $f(r, \theta) = \sum_{n \in \mathbb{Z}} \sum_{k \in \mathbb{Z}_{>0}} a_{nk} \psi_{nk}(r, \theta)$, where $a_{nk} \in \mathbb{C}$ are *disk harmonic (DH) coefficients*. The *disk harmonic (DH) transform* computes these coefficients, expanding f in the disk harmonic basis (akin to the Fourier transform expanding f in the Fourier mode basis). A single coefficient, given by

$$a_{nk} = \int_0^{2\pi} \int_0^1 f(r, \theta) \overline{\psi_{nk}(r, \theta)} r dr d\theta,$$

encodes the amplitude and phase of ψ_{nk} , necessary to reconstruct the signal from the expansion equation above. Importantly, DH coefficients are *equivariant* to rotation of f , meaning that a rotation in f will correspond to a rotation in the DH coefficients of f .

Proposition 3.1 (Equivariance). *More formally, denote $SO(2)$ the group of 2D rotations. Consider f a square-integrable function on the unit disk with DH coefficients $\{a_{n,k}^f\}_{n \in \mathbb{Z}, k \in \mathbb{Z}_{>0}}$. For every $\phi \in SO(2)$, if $f' = f \circ R_\phi$, then for all $n, k \in \mathbb{Z} \times \mathbb{Z}_{>0}$, we have: $a_{n,k}^{f'} = e^{in\phi} \cdot a_{n,k}^f$. Thus, the DH coefficients are equivariant to 2D rotations.*

In the general case, a function f is *equivariant* to a transformation group G if

$$f(T_g x) = T'_g f(x), \quad \forall g \in G,$$

where T_g acts on the input and T'_g is a corresponding action on the output. A function f is *invariant* to G if

$$f(T_g x) = f(x), \quad \forall g \in G.$$

Signal Discretization on images. While the disk harmonics are defined for functions $f : \mathbb{D} \rightarrow \mathbb{R}$, practical applications require discretization. We will work with discretized images, i.e., arrays of size $L \times L$ whose pixel values are interpreted as samples of f . Let $p = L^2$ be the total number of pixels such that locations and values of image intensities are denoted as $x_1, \dots, x_p$ and $f_1, \dots, f_p$, respectively. For images of shape $L \times L$ with spacing $\Delta x = 2/L$, corresponding to the domain $[-1, 1] \times [-1, 1]$, the sampling rate is $(\Delta x)^{-1} = L/2$, yielding a Nyquist frequency (i.e., bandlimit) of $L/4$. The *Nyquist frequency* is the highest frequency that can be well represented when the signal is sampled at a given rate.

Recall that λ_{nk} is the k -th root of the n -th order Bessel function, and that the disk harmonic ψ_{nk} does not extend radially past λ_{nk} . Therefore, λ_{nk} determines the number of Bessel function oscillations over the disk and is therefore a quasi-frequency, which is subject to Nyquist criterion. According to the Nyquist criterion, we should only keep the DH coefficients for which $\frac{\lambda_{nk}}{2\pi} \leq \frac{L}{4}$, as higher-frequency components correspond to features beyond the image resolution [39].

Truncation of DH coefficients. Following closely with [25], for a given bandlimit $\lambda > 0$, typically $\lambda = \frac{\pi L}{2}$, we keep the DH coefficients for indices (n, k) such that $\lambda_{nk} \leq \lambda$, and order them from lowest frequency to highest $\lambda_0 \leq \dots \leq \lambda_{m-1}$. The index $j = 0, \dots, m-1$ represents this order. As such, we now index λ_j or equivalently λ_{n_j, k_j} , where n_j, k_j are the n, k of the frequency at index j . Importantly, Bessel function roots do not repeat over order, so we can uniquely identify a disk harmonic via λ_{n_j, k_j} . As such, we also index ψ by j , denoting $\psi_j = \psi_{n_j, k_j}$.

Table 2 shows the indices $j = 0, \dots, m-1$ organized in a table labeled by n, k for an image of size 8×8 and $\lambda = \frac{\pi L}{2}$. Further, for a given bandwidth λ that determines a number $m = m_\lambda$, we define the maximum angular frequency (maximum order of the Bessel functions) as $N_m = \max\{n_j \in \mathbb{Z} : j \in \{0, \dots, m-1\}\}$, and the maximum Bessel root as $K_n = \max\{k \in \mathbb{Z}_{>0} : \lambda_{nk} \leq \lambda\}$ for $n \in \{-N_m, \dots, N_m\}$. For example, Table 2 has $m = 46$, $N_m = 10$, and $K_2 = 4$.

$n \rightarrow$ $k \downarrow$	-10	-9	-8	-7	-6	-5	-4	-3	-2	-1	0	1	2	3	4	5	6	7	8	9	10
1	44	38	30	25	19	15	10	6	3	1	0	2	4	7	11	16	20	26	31	39	45
2	-	-	-	48	40	32	23	17	12	8	5	9	13	18	24	33	41	49	-	-	-
3	-	-	-	-	-	-	42	34	27	21	14	22	28	35	43	-	-	-	-	-	-
4	-	-	-	-	-	-	-	-	46	36	29	37	47	-	-	-	-	-	-	-	-

Table 2. Indices $j = 0, \dots, m-1$, for $m = 46$ organized in a table labeled by n, k for an image of size $L \times L = 8 \times 8$ and $\lambda = \frac{\pi L}{2}$. The columns represent the Bessel orders, i.e., the angular frequencies. The rows represent the Bessel root indices, i.e., the radial frequencies.

Multi-Reference Alignment (MRA) In classical (translational) MRA, the goal is to recover a ground truth signal f from multiple noisy and translated observations of the signal $\tilde{f}_i = f + n_i$, for $i = 1, \dots, I$, and where the noise variables n_i are i.i.d. with $n \sim \mathcal{N}(0, \sigma^2)$. Bispectrum MRA does this by mapping all noisy translated observations to their bispectra $b^{f_i} = b_i \in \{b_1, \dots, b_I\}$ and computing the average bispectrum $\mathbb{E}[b_i] = b^f + \delta$, where b^f is the bispectrum of the true signal f and δ is a bias term introduced by noise propagation. Theoretically, the ground truth signal can be recovered by correcting for the bias term and mapping the corrected bispectrum back to image space.

[3] has derived the bias term for the 1D translation bispectrum so that MRA can be performed on translated 1D signals. In our theory section, we derive a bias term for the 2D disk bispectrum so that MRA can be performed on rotated 2D signals.

4. Theory: The Selective Disk Bispectrum and Its Inversion

We first introduce the existing disk bispectrum [40] and then we propose two modifications to improve its space and time complexity. *Proofs for all propositions and theorems are in the Appendix.*

4.1. The Disk Bispectrum

We recall the definition of the disk bispectrum originally introduced in [40].

Definition 4.1 (Disk Bispectrum [40]). Using notation and image setup established in the previous section, let $a_{n_j k_j}$ be the DH coefficients of a square integrable real-valued function $f : \mathbb{D} \rightarrow \mathbb{R}$ defined on the disk $\mathbb{D}$, for $j = 0, \dots, m-1$ such that $\lambda_0 \leq \dots \leq \lambda_{m-1}$ and n_j, k_j are the associated angular and radial indices. The disk bispectrum is defined as:

$$b_{j_1, j_2, k_3} = a_{n_{j_1}, k_{j_1}} \cdot a_{n_{j_2}, k_{j_2}} \cdot a_{n_{j_1} + n_{j_2}, k_3}^* \in \mathbb{C}, \quad (1)$$

where $*$ indicates complex conjugate, $j \in \{0, \dots, m-1\}$, $k \in \mathbb{Z}_{>0}$, $-N_m \leq n_{j_1} + n_{j_2} \leq N_m$, and $k_3 \leq K_{n_{j_1} + n_{j_2}}$.

The computational complexity of this disk bispectrum is prohibitive. First, it inherits the time complexity of the DHT, which is $O(L^3)$ where $L \times L$ is the size of the image.

Image size	$L = 8$	$L = 16$	$L = 28$	$L = 56$	$L = 112$
Num. DB coeffs.	244	2,334	13,715	122,910	1,057,052
Num. SDB coeffs.	27	105	315	1,247	4,957

Table 3. Number of coefficients in the disk bispectrum (DB) (first row) and in the selective disk bispectrum (ours, second row) as a function of the image size $L \times L$, where the selected bandlimit for the DH coefficients is $\lambda = \pi L/2$. Using the full disk bispectrum becomes quickly prohibitive.

While the space complexity can be reduced by realizing the symmetry $j_1 \leftrightarrow j_2$, it is still cubic in m , the number of DH coefficients. The next proposition summarizes these complexities.

Proposition 4.2 (Complexity of the disk bispectrum). *Denote $L \times L$ the size of the image, m the number of DH coefficients, N_m the maximum angular frequency. The space complexity of the disk bispectrum is $O(m^3/N_m)$. The time complexity of the disk bispectrum is $O(L^3 + m^3/N_m)$.*

The first row of Table 3 gives the exact number of coefficients in the disk bispectrum for various image sizes $L \times L$. Using the disk bispectrum becomes prohibitive even for traditional image sizes.

4.2. The Selective Disk Bispectrum (SDB) and Its Inversion

We propose the selective disk bispectrum (SDB), defined as a carefully chosen subset of the (full) disk bispectrum coefficients. The SDB improves the space complexity from $O(m^3/N_m)$ of the previous bispectrum (or $O(L^2)$ of the original image) to $O(m)$ and the time complexity from $O(L^3 + m^3/N_m)$ to $O(L^3 + m)$. We prove that the selected coefficients are sufficient for recovering the image f , up to rotation, meaning that most of the coefficients in the original disk bispectrum are, in fact, redundant.

Definition 4.3 (Selective Disk Bispectrum). Let $a_{n_j k_j}$ be the DH coefficients of a square integrable real-valued function $f : \mathbb{D} \rightarrow \mathbb{R}$ defined on the disk $\mathbb{D}$, for $j \in \{0, \dots, m-1\}$ such that $\lambda_1 \leq \dots \leq \lambda_m$ and n_j, k_j are the associated angular and radial indices. The selective disk bispectrum is

defined as

$$b^f = \begin{bmatrix} b_{0,0,1} \cdots b_{0,0,K_0} \\ b_{2,0,1} \cdots b_{2,0,K_0} \\ b_{2,N_m-1,1} \cdots b_{2,N_m-1,K_{N_m-1}} \end{bmatrix}, \quad (2)$$

where the entries correspond to selected coefficients of the disk bispectrum defined as:

$$\begin{cases} b_{0,0,k} = a_{0,1}^2 \cdot a_{0,k}^*, \\ \text{for } k \in \{1, \dots, K_0\}, \\ b_{2,n,k} = a_{1,1} \cdot a_{n,1} \cdot a_{n+1,k}^*, \\ \text{for } n \in \{0, N_m - 1\}, \text{ and } k \in \{1, \dots, K_{n+1}\}. \end{cases} \quad (3)$$

We note that the selective bispectrum coefficients are indexed by j_1, n_2, k_3 (limiting $j_1 \in \{0, 2\}$), as opposed to j_1, j_2, k_3 . Selecting only a subset of the disk bispectrum coefficients reduces its space and time complexities.

Proposition 4.4 (Complexity of the selective disk bispectrum). *Denote $L \times L$ the size of the image, m the number of DH coefficients, N_m the maximum angular frequency. The number of selective bispectrum coefficients is $N = \sum_{n=0}^{N_m} K_n$. The space complexity of the selective disk bispectrum is $O(m)$. Its time complexity is $O(L^3 + m)$.*

Despite discarding many coefficients from the full disk bispectrum, our main result shows that the remaining $N = \sum_{n=0}^{N_m} K_n$ coefficients are sufficient to reconstruct the signal, up to a global rotation. Fig. 2 illustrates the technique behind the proof, showing which set of bispectrum coefficients are necessary to reconstruct a specific DH coefficient.

Theorem 4.5 (Inversion). *Denote $L \times L$ the size of the image f , m the number of DH coefficients, N_m the maximum angular frequency, and a_{n_j, k_j} the DH coefficients for $j = 0, \dots, m-1$. Assume that $a_{n,1} \neq 0$ for $n \in \{0, \dots, N_m-1\}$. Then, its selective disk bispectrum b^f can be inverted, that is f can be reconstructed from b^f up to a disk rotation. The time complexity of the inversion is $O(m + L^3)$.*




Figure 2. Graphical illustration of the selective disk bispectrum inversion. The arrows indicate in which order the DH coefficients are reconstructed.

We note that our inversion of the selective disk bispectrum also yields an inversion method for the (full) disk bispectrum. We also note that nonzero coefficient assumptions (such as $a_{n,1} \neq 0$) are classic in bispectrum inversion literature to avoid division by zero. However, this assumption does not impact real-world application as long as there is some noise on the images.

Theorem 4.6 (Complete Invariance). *Let f and f' be a pair of real-valued, square-integrable functions on the disk. Let the bispectrum be defined as in Definition 4.3. Assume that the DH coefficients $a_{n,1}$ of f are non-zero for all $n \in \{0, \dots, N_m - 1\}$. Then $b^f = b^{f'}$ if and only if $f' = f \circ R_\phi$ for any 2D rotation $\phi \in SO(2)$. Thus, the selective bispectrum is a complete invariant.*

Part of the proof relies on the equivariance of the DH coefficients, which allow us to (re)prove the invariance of the (selective) disk bispectral coefficients used in [40]. The completeness of the SDB is proven by our inversion theorem, as the SDB contains all information about the original signal necessary for reconstruction.

4.3. Numerical Approximation of The Selective Disk Bispectrum and Its Inversion

To further reduce the time complexity of the selective bispectrum, we propose leveraging a recent numerical approximation of the DH coefficients [25] which computes the disk harmonic transform in $O(L^2 \log L)$. This yields the following result:

Proposition 4.7 (Complexity of the selective disk bispectrum). *Denote $L \times L$ the size of the image, m the number of DH coefficients. The time complexity of the approximate selective disk bispectrum is $O(L^2 \log L + m)$.*

Next, we prove that we still have precise guarantees on the accuracy of the bispectrum coefficient approximations, the invariance property, and the reconstruction of the original image.

Theorem 4.8. *Let $0 < \varepsilon \leq 1$ be a target accuracy, and assume that $\lambda \leq \sqrt{\pi p}$ and $|\log \varepsilon| \leq \sqrt{p}$, where λ is the bandlimit for the DH coefficients and $p = L^2$ the number of pixels in the image f of size $L \times L$. Let b^f and $\tilde{b}^f$ denote the exact and approximate selective disk bispectrum of f , where $\tilde{b}^f$ is computed using the approximate disk harmonic transform of Marshall et al. [25]. Then,*

$$\|\tilde{b}^f - b^f\|_\infty \leq C\varepsilon\|f\|_1,$$

for a constant C depending only on a bound of f on its compact disk support.

As a corollary, we find that the bispectrum coefficients are approximately invariant with controlled guarantees.




Figure 3. (a) The selective disk bispectrum (SDB) map is significantly faster than the full disk bispectrum map, especially for large images. (b) Bispectrum invariance error decreases with higher image resolution. (c-d) The SDB coefficient inversion successfully reconstructs images for resolution 28×28 (c) and 112×112 (d), and achieves a similar reconstruction error as the DH Transform [25] inversion.

Corollary 4.9. *Let f, f' be a pair of real-valued square integrable functions on the disk such that there exists a $\phi \in SO(2)$ such that $f' = f \circ R_\phi$. Under the notations and assumptions of Theorem 4.8, we have:*

$$\|\tilde{b}^f - \tilde{b}^{f'}\|_\infty \leq 2C\varepsilon\|f\|_1, \quad (4)$$

for a constant C depending only on a bound of f on its compact disk support.

Linking the accuracy to f is useful, as it allows us to control accuracy through number of pixels p . If $|f|$ and C are large, but we still want to reach a given accuracy, then we can choose p accordingly such that $\log(\varepsilon) < \sqrt{p}$ to match the desired accuracy.

4.4. MRA with the Selective Disk Bispectrum

We follow [3]’s assumption that $\mathbb{E}[b_i] = b^f + \delta$, where $b_i \in \{b_1, \dots, b_I\}$ is the bispectrum of one rotated noisy observation, b^f is the bispectrum of the true signal f , and δ is a bias term introduced by noise propagation in bispectrum space. In this setting, we derive δ in the setting where b^f is the disk bispectrum.

Proposition 4.10 (SDB MRA Correction). *When performing MRA using the selective disk bispectrum of noisy rotated images, the mean bispectrum must be corrected by subtracting δ before inversion to image space.*

$$\delta = \sigma^2[S_{0,k_3}^\dagger \delta_{n_{j_1},0} \delta_{n_{j_2},0} + S_{01} \delta_{1,k_3} \delta_{n_{j_2},0} + S_{01} \delta_{1,k_3} \delta_{n_{j_1},0}] \quad (5)$$

5. Tractability of Selective Disk Bispectrum

To illustrate the computational and representational properties of our proposed selective disk bispectrum, we conduct a series of experiments testing *speed*, *invertibility*, and *rotation invariance*. We compare against the full disk bispectrum proposed by [40] and described in Eq. 4.1. Both the full and selective bispectra are computed using the disk harmonic coefficient library implemented by [25] on a MacBook Pro with an Apple M2 Max chip.

5.1. Design of Tractability Experiments

We consider square grayscale palm tree images of length $L \in \{16, 28, 56, 112\}$, each perturbed with Gaussian noise of standard deviation $\sigma \in \{0, 0.01, 0.05, 0.1\}$. For each image, we compute both the full and selective disk bispectrum using the formulations in Eq. 4.1 (full) and Eq. 4.3 (selective). We repeat these experiments for all digits of MNIST, and show results in the Appendix.

Speed. We compare the runtime of the forward and inverse bispectrum transforms across image sizes L and noise levels σ . The full disk bispectrum is computationally expensive and impractical for many cases. In contrast, our *selective disk bispectrum* enables efficient computation, making both encoding and reconstruction feasible even for larger images.

Invertibility. In the previous section, we prove that the selective disk bispectrum is *complete*. Here, we visually confirm our theoretical results by constructing an image f from its bispectrum and measuring the relative reconstruction error $\varepsilon_{\text{rel}} = \frac{\|\hat{f} - f\|}{\|f\|}$, where $\hat{f}$ is the image reconstruction obtained by applying the inverse map to the disk bispectrum coefficients. We observe in Fig. 3 that the selective bispectrum retains the essential information needed to accurately reconstruct the original image, demonstrating its practical completeness.

Rotation Invariance. We verify the invariance of the disk bispectrum to disk rotations. Given an image f and its rotated version f' , we compute their respective bispectra b^f and $b^{f'}$. Since the bispectrum is designed to be invariant to rotations, we expect $b^f = b^{f'}$. To quantify invariance, we compute the relative error between the two bispectra $\varepsilon_{\text{rel}} = \frac{\|b^f - b^{f'}\|}{\|b^f\|}$. Low error values across test conditions confirm that both the full and selective bispectra exhibit strong empirical rotation invariance, as shown in Fig. 3.

5.2. Results of Tractability Experiments

We present the results of the tractability experiments. Full experimental results, including results on all MNIST im-




Figure 4. (a) Dataset samples and original image used in MRA experiments for MNIST 6. (b-e) Relative error compares the similarity between the MRA reconstructed image and the original image used to generate the noisy rotated dataset for (b) $\sigma^2 = 0$ (c) $\sigma^2 = 0.01$, (d) $\sigma^2 = 0.05$, (e) $\sigma^2 = 0.1$. Sample reconstructions are shown below each result plot.

ages, are available in the Appendix.

Speed. The selective disk bispectrum offers a $3\times$ speedup over the full bispectrum for 16×16 images, and a $17322\times$ speedup for 112×112 images. The inverse transform is equivalently fast for the full and selective bispectra, since the full bispectrum inversion utilizes only the subset of its coefficients included in the selective variant. Forward and inverse timing results are shown in Fig. 3 for $\sigma = 0$; run times are unaffected by noise level.

Invariance. Bispectrum invariance error decreases with higher image resolution, as shown in Fig. 3, and is as low as 3 percent for 112×112 images. Invariance error is largely unaffected by noise level.

Invertibility. As shown in Fig. 3, the selective bispectrum accurately reconstructs the original image. This reconstruction accuracy is identical to the full bispectrum’s. While reconstruction error increases slightly at lower tested resolutions, images are sufficiently reconstructed across all tested settings (see Appendix for reconstruction across bispectrum type, noise, and image resolution). Importantly, bispectrum inversions match the original image, up to a rotation. Since the bispectrum inversion reconstructs the original image up to a global rotation, reconstructions in Fig. 3 have been manually registered to match the orientation of the original image for clear comparison.

6. Application: Multi-Reference Alignment with the Selective Disk Bispectrum

We demonstrate the utility of our S. D.B. and its inversion for rotation-invariant multi-reference alignment (MRA), leveraging our new MRA correction term. We test rotation-invariant MRA on all MNIST digits, with results for the digit 6 shown in Fig. 4 and other digits in the Appendix. For each experiment, we begin with a noiseless MNIST image and generate $n_X \in \{10, 20, 30, 40, 50, 60, 100, 200, 300, 400, 500\}$

copies. Each copy is corrupted with i.i.d. Gaussian noise $n \sim \mathcal{N}(0, \sigma^2)$ for $\sigma^2 \in \{0, 0.01, 0.05, 0.1\}$, randomly rotated by $R \in [0, 2\pi]$, and masked outside the disk, yielding n_X images $\{f_i\}_{i=1}^{n_X}$. We map each image to bispectrum space to obtain $\{b_i\}_{i=1}^{n_X}$, average the bispectra, apply our correction term from Eq. 4.10, and map the corrected average back to image space using our disk bispectrum inversion method from Sec. 4. Each experiment is repeated ten times with different random seeds, and computes relative error between MRA estimate $\bar{f}$ and true signal f as $\varepsilon_{\text{rel}} = \frac{\|\bar{f} - f\|^2}{\|f\|^2}$. As shown in Fig. 4, reconstruction error decreases and stabilizes around 0.23 as n_X increases—consistent with the empirical baseline reconstruction and invariance errors reported for 28×28 images in Sec. 5. These results are the first demonstration of MRA using a rotation-invariant bispectrum.

Conclusion

We present the first disk bispectrum inversion, which reveals the minimal sufficient coefficients needed for accurate shape recovery. This enables the construction of our *selective* disk bispectrum, which is **17,322 $\times$ faster** to compute than the full bispectrum for (122×112) images. Crucially, our disk bispectrum inversion provides a mapping from disk bispectrum coefficients back to the image they represent, meaning that statistics can be done in bispectrum space and the result can be mapped back to image space. Without this, methods such as MRA would be impossible. Accordingly, we additionally derive the noise propagation that would occur for selective disk bispectrum MRA (in line with [3] and use it to showcase our bispectrum inverse in MRA experiments. The selective disk bispectrum’s speed and invertibility position it as a tractable and scalable shape representation for rotation-invariant learning.

References

- [1] Pablo F. Alcantarilla, Adrien Bartoli, and Andrew J. Davison. Kaze features. In *European Conference on Computer Vision (ECCV)*, pages 214–227. Springer, 2012. 2
- [2] Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European Conference on Computer Vision (ECCV)*, pages 404–417. Springer, 2006. 2
- [3] Tamir Bendory, Nicolas Boumal, William Leeb, and Amit Singer. Bispectrum inversion with application to multireference alignment. *IEEE Transactions on Signal Processing*, 66(4):1037–1050, 2018. 2, 3, 4, 5, 7, 8
- [4] David R. Brillinger. Some history of higher-order statistics and spectra. *Statistica Sinica*, 1(2):465–476, 1991. 1
- [5] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, 2021. 2
- [6] Hua Chen, Mona Zehni, and Zhizhen Zhao. A spectral method for stable bispectrum inversion with application to multireference alignment. *IEEE Signal Processing Letters*, 25(7):911–915, 2018. 3
- [7] Taco S. Cohen, Mario Geiger, Jonas Köhler, and Max Welling. Spherical cnns. *arXiv preprint arXiv:1801.10130*, 2018. Published as ICLR 2018. 2
- [8] Wen-Xuan Du, Clifford H. Thurber, and Donna Eberhart-Phillips. Earthquake relocation using cross-correlation time delay estimates verified with the bispectrum method. *Bulletin of the Seismological Society of America*, 94(3):856–866, 2004. 1
- [9] D. Edidin and J. Katz. The reflection-invariant bispectrum: Signal recovery in the dihedral model, 2025. 3
- [10] Steve Elgar and RT Guza. Observations of bispectra of shoaling surface gravity waves. *Journal of Fluid Mechanics*, 161:425–448, 1985. 1
- [11] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so(3) equivariant representations with spherical cnns. In *European Conference on Computer Vision (ECCV)*, pages 52–68. Springer, 2018. 2
- [12] D. Giannakis. Signal reconstruction from the bispectrum. *Journal of Sound and Vibration*, 133(2):199–211, 1989. 2
- [13] Georgios B. Giannakis. Signal reconstruction from multiple correlations: frequency- and time-domain approaches. *J. Opt. Soc. Am. A*, 6(5):682–697, 1989. 3
- [14] J. A. Gordon and D. F. Buscher. Detection noise bias and variance in the power spectrum and bispectrum in optical interferometry. *Astronomy & Astrophysics*, 541:A46, 2012. 1
- [15] James Herring, James Nagy, and Lars Ruthotto. Gauss-newton optimization for phase recovery from the bispectrum. *IEEE Transactions on Computational Imaging*, pages 1–1, 2019. 3
- [16] M.K. Hu. Visual pattern recognition by moment invariants. *IRE Transactions on Information Theory*, 8(2):179–187, 1962. 2
- [17] Ramakrishna Kakarala. *Triple correlation on groups*. PhD thesis, UC Irvine, 1992. 2, 3
- [18] R. Kakarala. The bispectrum as a source of phase-sensitive invariants for fourier descriptors: a group-theoretic approach. *Journal of Mathematical Imaging and Vision*, 44(3):341–353, 2012. 3
- [19] Ramakrishna Kakarala. The bispectrum as a source of phase-sensitive invariants for fourier descriptors: a group-theoretic approach, 2012. 3
- [20] R. Kakarala, B. M. Bennett, M. D’Zmura, and G. J. Iverson. Bispectral techniques for spherical functions. In *Proceedings of ICASSP ’93*, pages 237–240, Minneapolis, MN, 1993. IEEE. 3
- [21] Georgia Koukiou. Identifying system non-linearities by fusing signal bispectral signatures. *Electronics*, 13(7):1287, 2024. 1
- [22] Xuemin Liu, Yaoliang Song, Kuiyu Chen, Shihao Yan, Si Chen, and Baihua Shi. Modulation recognition of low-snr uav radar signals based on bispectral slices and ga-bp neural network. *Drones*, 7(7):472, 2023. 1
- [23] David G. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. 2
- [24] N. F. Marshall. Fast expansion into harmonics on the disk. *arXiv preprint arXiv:2207.13674*, 2022. Fast Disk Harmonics Transform (FDHT). 2
- [25] Nicholas F. Marshall, Oscar Mickelin, and Amit Singer. Fast expansion into harmonics on the disk: A steerable basis with fast radial convolutions. *SIAM Journal on Scientific Computing*, 45(5):A2431–A2457, 2023. 2, 4, 6, 7
- [26] Adrián Martín-Montero, Gonzalo C Gutiérrez-Tobal, David Gozal, Verónica Barroso-García, Daniel Álvarez, Félix Del Campo, Leila Kheirandish-Gozal, and Roberto Hornero. Bispectral analysis of heart rate variability to characterize and help diagnose pediatric sleep apnea. *Entropy*, 23(8):1016, 2021. 1
- [27] Simon Maigne, Johan Mathe, Sophia Sanborn, Christopher Hillar, and Nina Miolane. The selective g-bispectrum and its inversion: Applications to g-invariant networks. In *Advances in Neural Information Processing Systems*, pages 115682–115711. Curran Associates, Inc., 2024. 3
- [28] Chrysostomos L Nikias and Jerry M Mendel. Signal processing with higher-order spectra. *IEEE Signal processing magazine*, 10(3):10–37, 1993. 1
- [29] L. Pinilla, A. Mishra, and B. M. Sadler. Unique bispectrum inversion for signals with finite spectral/temporal support. In *Proceedings of ICASSP 2023*. IEEE, 2023. 2
- [30] B. S. Reddy and B. N. Chatterji. An fft-based technique for translation, rotation, and scale-invariant image registration. *IEEE Transactions on Image Processing*, 5(8):1266–1271, 1996. 2
- [31] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 2564–2571, 2011. 2
- [32] Sara Sabour, Nicholas Frosst, and Geoffrey E. Hinton. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3856–3866, 2017. 2

- [33] B.M. Sadler. Shift and rotation invariant object reconstruction using the bispectrum. In *Workshop on Higher-Order Spectral Analysis*, pages 106–111, 1989. [3](#)
- [34] G. Sundaramoorthy, M. R. Raghuveer, and S. A. Dianat. Bispectral reconstruction of signals in noise: amplitude reconstruction issues. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 38(7):1297–1306, 1990. [2](#)
- [35] M. R. Teague. Image analysis via the general theory of moments. *JOSA*, 70(8):920–930, 1980. [2](#)
- [36] Nathaniel Thomas, T. Smidt, S. Kearnes, L. Yang, L. Li, K. Kohlhoff, P. Riley, and K. Kohlhoff. Tensor field networks: Rotation- and translation-equivariant neural networks for 3d point clouds. In *NeurIPS*, 2018. [2](#)
- [37] J. Tukey. The spectral representation and transformation properties of the higher moments of stationary time series. *Reprinted in The Collected Works of John W. Tukey*, 1:165–184, 1953. [1](#)
- [38] Ji-Wu Zhang, Chong-Xun Zheng, and An Xie. Bispectrum analysis of focal ischemic cerebral eeg signal using third-order recursion method. *IEEE transactions on biomedical engineering*, 47(3):352–359, 2000. [1](#)
- [39] Zhizhen Zhao and Amit Singer. Fourier–bessel rotational invariant eigenimages. *Journal of the Optical Society of America A*, 30(5):871–877, 2013. [4](#)
- [40] Zhizhen Zhao and Amit Singer. Rotationally invariant image representation for viewing direction classification in cryo-em. *Journal of structural biology*, 186(1):153–166, 2014. [1](#), [3](#), [5](#), [6](#), [7](#)