



Published in final edited form as:

IEEE Trans Comput Imaging. 2016 March ; 2(1): 1–12. doi:10.1109/TCI.2016.2514700.

Fast Steerable Principal Component Analysis

Zhizhen Zhao,

Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 USA

Yoel Shkolnisky, and

Department of Applied Mathematics, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel

Amit Singer

Department of Mathematics and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544 USA

Zhizhen Zhao: jzhao@cims.nyu.edu

Abstract

Cryo-electron microscopy nowadays often requires the analysis of hundreds of thousands of 2-D images as large as a few hundred pixels in each direction. Here, we introduce an algorithm that efficiently and accurately performs principal component analysis (PCA) for a large set of 2-D images, and, for each image, the set of its uniform rotations in the plane and their reflections. For a dataset consisting of n images of size $L \times L$ pixels, the computational complexity of our algorithm is $\mathcal{O}(nL^3 + L^4)$, while existing algorithms take $\mathcal{O}(nL^4)$. The new algorithm computes the expansion coefficients of the images in a Fourier–Bessel basis efficiently using the nonuniform fast Fourier transform. We compare the accuracy and efficiency of the new algorithm with traditional PCA and existing algorithms for steerable PCA.

Index Terms

Steerable PCA; group invariance; non-uniform FFT; denoising

I. Introduction

Principal component analysis (PCA) is widely used in image analysis and pattern recognition for dimensionality reduction and denoising. In particular, PCA is often one of the first steps [1] in the algorithmic pipeline of cryo-electron microscopy (cryo-EM) single particle reconstruction (SPR) [2] to compress and denoise the acquired 2D projection images in order to eventually determine the 3D structure of a macromolecule. The high level of noise in those images drastically deteriorates the performance of single-image based denoising methods, such as non-local means [3] and wavelet thresholding [4], and so the latter are outperformed by PCA. As any planar rotation of any given projection image is

equally likely to appear in the experiment, by either in-plane rotating the detector or the specimen, it makes sense to include all possible rotations of the projection images when performing PCA. The resulting decomposition, termed steerable PCA, consists of principal components which are tensor products of radial functions and angular Fourier modes [5], [6], [7], [8], [9]. Beyond cryo-EM, steerable PCA has many other applications in image analysis and computer vision [10].

The term “steerable PCA” comes from the fact that rotating the principal components is achieved by a simple phase shift of their angular part. The principal components are invariant to any in-plane rotation of the images, therefore finding steerable principal components is equivalent to finding in-plane rotationally invariant principal components.

In cryo-EM data processing, in addition to compression and denoising, steerable PCA is also useful in generating rotationally invariant image features (i.e. bispectrum-like features [11]). These are crucial for fast rotationally invariant nearest neighbors search used in efficient computation of class averages [11]. Rotational alignment between image pairs can also be computed more efficiently using the expansion coefficients in a steerable basis.

In this letter, we focus on the action of the group $O(2)$ on digital images by in-plane rotating and possibly reflecting them. The idea of using group actions for constructing group invariant features and filters has been previously proposed in [12], [13]. This group theoretical framework has been applied to $SO(3)$ and $SU(1, 1)$ in [14], [15]. The representation of finite groups, such as the dihedral groups, has been used for computing the Karhunen-Loève expansion of digital images in [16].

Various efficient algorithms for steerable PCA have been introduced [17], [8]. However, steerable PCA of modern cryo-EM datasets that contain hundreds of thousands of large images poses a computational challenge. Also, it is important to ensure that the steerable PCA algorithm is numerically accurate when the input images are noisy. In order to exploit the special separation of variables structure of the principal components in polar coordinates, most algorithms rely on resampling the images on a polar grid. However, the transformation from Cartesian to polar is non-unitary, and thus changes the statistics of the noise. In particular, resampling transforms uncorrelated white noise to colored noise that may lead to spurious principal components.

Recently, [9] addressed this issue by incorporating a sampling criterion into the steerable PCA framework and introduced an algorithm called Fourier-Bessel steerable PCA (FBsPCA). FBsPCA assumes that the underlying clean images (before being possibly contaminated with noise) are bandlimited and essentially compactly supported in a disk. This assumption holds, for example, for 2D projection images of a 3D molecule compactly supported in a ball. It also implies that the images can be expanded in an orthogonal basis for bandlimited functions, such as the Fourier-Bessel basis. In FBsPCA, the Fourier-Bessel expansion of each image is truncated into a finite series using a sampling criterion that was introduced by Klug and Crowther [18]. The sampling criterion ensures that the transformation from the Cartesian grid to the truncated Fourier-Bessel expansion is nearly unitary. Moreover, the covariance matrix built from the expansion coefficients of the images

and all their possible rotations has a block diagonal structure where the block size decreases as a function of the angular frequency. The computational complexity of FBsPCA is $\mathcal{O}(nL^4)$ operations for n images of size $L \times L$. Notice that, when $n > L^2$, the computational complexity of traditional PCA is $\mathcal{O}(nL^4 + L^6)$, where the first term corresponds to forming the $L^2 \times L^2$ covariance matrix and the second term corresponds to its eigen-decomposition. Although FBsPCA and PCA have a similar computational complexity, FBsPCA leads to better denoising as it takes into account all possible rotations and reflections. This makes FBsPCA more suitable than traditional PCA as a tool for 2D analysis of cryo-EM images [9]. With the enhancement of electron microscope detectors' resolution, a typical image size of a single particle can easily be over 300×300 pixels. Thus, FBsPCA is still not efficient enough to analyze a large number of images of large size (i.e. large n and large L). The bottleneck for this algorithm is the first step that computes the Fourier-Bessel expansion coefficients, whose computational complexity is $\mathcal{O}(nL^4)$.

In this letter we introduce a fast Fourier-Bessel steerable PCA (FFBsPCA) that reduces the computational complexity for FBsPCA from $\mathcal{O}(nL^4)$ to $\mathcal{O}(nL^3)$ by computing the Fourier-Bessel expansion coefficients more efficiently and accurately. This is achieved by first mapping the images from their Cartesian grid representation to a polar grid representation in the reciprocal (Fourier) domain using the non-uniform fast Fourier transform (NUFFT) [19], [20], [21], [22]. The polar grid representation enables to efficiently evaluate the Fourier-Bessel expansion coefficients of the images by 1D FFT on concentric circles followed by accurate evaluation of a radial integral with a Gaussian quadrature rule. The overall complexity of computing the Fourier-Bessel coefficients is reduced to $\mathcal{O}(nL^3)$ operations. The increased accuracy and efficiency in evaluating the Fourier-Bessel expansion coefficients are the main contributions of this letter.

We note that the Fourier-Bessel expansion coefficients can be computed in $\mathcal{O}(nL^2 \log L)$ operations using algorithms for rapid evaluation of special functions [23] or a fast analysis-based Fourier-Bessel expansion [24]. However, such "fast" algorithms may only lead to a marginal improvement for two reasons. First, the break even point for them compared to the direct approach is for relatively large L such as $L = 256$ or larger. Second, forming the covariance matrix from the expansion coefficients still requires $\mathcal{O}(nL^3)$ operations.

The letter is organized as follows: Section II contains the mathematical preliminaries of the Fourier-Bessel expansion, the sampling criterion, and the numerical evaluation of the expansion coefficients. The computation of the steerable principal components is described in Section III. We present the algorithm and give a detailed computational complexity analysis in Section IV. Various numerical examples concerning the computation time of FFBsPCA compared with FBsPCA and traditional PCA are presented in Section V. In the same section, we demonstrate the performance of FFBsPCA-based denoising using simulated cryo-EM projection images.

Reproducible research: The FFBsPCA is available in the SPR toolbox ASPIRE (<http://spr.math.princeton.edu/>). There are two main functions: *FBCoeff* computes the Fourier Bessel expansion coefficients and *sPCA* computes the steerable PCA basis and the associated expansion coefficients.

II. Fourier-Bessel Expansion of Bandlimited Images

We say that f has a band limit radius c if its Fourier transform

$$\mathcal{F}(f)(\xi_1, \xi_2) = \int_{\mathbb{R}^2} f(x, y) e^{-2\pi i(x\xi_1 + y\xi_2)} dx dy \quad (1)$$

satisfies $\mathcal{F}(f)(\xi_1, \xi_2) = 0$, for $\xi_1^2 + \xi_2^2 > c^2$. In our setup, a digital image I is obtained by sampling a squared-integrable bandlimited function f on a Cartesian grid of size $L \times L$, that

is, $I(i_1, i_2) = f(i_1, i_2)$, where $i_1, i_2 = -\lceil \frac{L-1}{2} \rceil, \dots, \lfloor \frac{L-1}{2} \rfloor$, and Δ is the pixel size.

For pixel size $\Delta = 1$, the Nyquist-Shannon sampling theorem implies that the Fourier transform of I is supported on the square $[-1/2, 1/2) \times [-1/2, 1/2)$. In many applications, the support size is effectively smaller due to other experimental considerations, for example, the exponentially decaying envelope of the contrast transfer function in electron microscopy.

Thus, we assume that the band limit radius of all images is $0 < c \leq \frac{1}{2}$. The scaled Fourier-Bessel functions are the eigenfunctions of the Laplacian in a disk of radius c with Dirichlet boundary condition and they are given by

$$\psi_c^{k,q}(\xi, \theta) = \begin{cases} N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) e^{ik\theta}, & \xi \leq c, \\ 0, & \xi > c, \end{cases} \quad (2)$$

where (ξ, θ) are polar coordinates in the Fourier domain (i.e., $\xi_1 = \xi \cos \theta$, $\xi_2 = \xi \sin \theta$, $\xi \geq 0$, and $\theta \in [0, 2\pi)$); $N_{k,q} = (c \sqrt{\pi} |J_{k+1}(R_{k,q})|)^{-1}$ is the normalization factor; J_k is the Bessel function of the first kind of integer order k ; and $R_{k,q}$ is the q th root of the Bessel function J_k . For a function f with band limit c that is also in $L^2(\mathbb{R}^2) \cap L^1(\mathbb{R}^2)$,

$$\mathcal{F}(f)(\xi, \theta) = \sum_{k=-\infty}^{\infty} \sum_{q=1}^{\infty} a_{k,q} \psi_c^{k,q}(\xi, \theta), \quad (3)$$

which converges pointwise. In Section II-A, we derive a finite truncation rule for the Fourier-Bessel expansion in Eq. (3).

A. Sampling Criterion

For digital implementations of Eq. (3), we must truncate it to a finite sum, namely to derive a sampling criterion for selecting k and q .

With the following convention for the 2D inverse polar Fourier transform of a function $g(\xi, \theta)$,

$$\mathcal{F}^{-1}(g)(r, \phi) = \int_0^{2\pi} \int_0^\infty g(\xi, \theta) e^{2\pi i r \xi \cos(\theta - \phi)} \xi d\xi d\theta, \quad (4)$$

the 2D inverse Fourier transform of the Fourier-Bessel functions, denoted $\mathcal{F}^{-1}(\psi_c^{k,q})$, is given in polar coordinates as

$$\mathcal{F}^{-1}(\psi_c^{k,q})(r, \phi) = \frac{2c \sqrt{\pi} (-1)^q R_{k,q} J_k(2\pi cr)}{i^k ((2\pi cr)^2 - R_{k,q}^2)} e^{i k \phi}. \quad (5)$$

The maximum of $|\mathcal{F}^{-1}(\psi_c^{k,q})(r, \phi)|$ in (5) is obtained near the circle $r = \frac{R_{k,q}}{2\pi c}$ and

$\mathcal{F}^{-1}(\psi_c^{k,q})(r, \phi)$ vanishes on concentric circles of radii $r = \frac{R_{k,q'}}{2\pi c}$ with $q' > q$. The smallest circle with vanishing $\mathcal{F}^{-1}(\psi_c^{k,q})$ that encircles the maximum of $|\mathcal{F}^{-1}(\psi_c^{k,q})|$ is of radius $r = \frac{R_{k,(q+1)}}{2\pi c}$.

We assume that the underlying clean images (before being possibly contaminated with noise) are essentially compactly supported in a disk of radius R . Therefore, we should rule out Fourier-Bessel functions for which the maximum of their inverse Fourier transform resides outside a disk of radius R . Otherwise, those functions introduce spurious information from noise. Notice that if the maximum is inside the disk, yet the zero after the maximum is outside the disk, then there is a significant spillover of energy outside the disk. We therefore require the more stringent criterion that the zero after the maximum is inside the disk, namely

$$\frac{R_{k,(q+1)}}{2\pi c} \leq R. \quad (6)$$

This sampling argument gives a finite truncation rule for the Fourier-Bessel expansion in Eq. (3), that is

$$R_{k,(q+1)} \leq 2\pi c R. \quad (7)$$

For each k , we denote by p_k the number of components satisfying Eq. (7). We also denote by $p = \sum_{k=-k_{\max}}^{k_{\max}} p_k$ the total number of components, where k_{\max} is the maximal possible value of k satisfying Eq. (7). The locations of Bessel zeros have been extensively studied, for example, in [25, p.517–521], [26, p.370], [27], [28], [29]. Several lower and upper bounds for Bessel zeros $R_{k,q}$ were proven by Breen in [29], such as

$$R_{k,q} > k + \frac{2}{3}|a_{q-1}|^{3/2}, \quad (8)$$

where a_q is the q th zero of the Airy function, shown to satisfy

$$\left[\frac{3}{8}\pi(4q - 1.4) \right]^{2/3} < |a_q| < \left[\frac{3}{8}\pi(4q - 0.965) \right]^{2/3}. \quad (9)$$

Using the lower bound for $|a_q|$ and the sampling criterion in Eq. (7), we have the following inequality for k and p_k ,

$$2\pi cR \approx R_{k,(p_k+1)} > k + \pi p_k - \frac{1.4\pi}{4}. \quad (10)$$

Breen also obtained

$$R_{k,q} < \left(\frac{k}{2} + q - \frac{0.965}{4} \right) \pi, \quad (11)$$

so we get another inequality for k and p_k ,

$$2\pi cR \approx R_{k,(p_k+1)} < \left(\frac{k}{2} + p_k + \frac{3.035}{4} \right) \pi. \quad (12)$$

Combining Eqs. (10) and (12), we have the following lower and upper bounds for p_k ,

$$2cR - \frac{k}{2} - \frac{3.035}{4} < p_k < 2cR - \frac{k}{\pi} + \frac{1.4}{4}. \quad (13)$$

The bound for the highest angular frequency k_{\max} is determined by setting $p_k = 1$ in Eq. (13), resulting in

$$4cR - 3.517 < k_{\max} < 2\pi cR - 2.042. \quad (14)$$

Equation (13) implies that as the angular frequency k increases, the number of components p_k decreases. Moreover, using the lower and upper bounds for p_k and k_{\max} in Eqs. (13) and (14), we derive that the total number of selected Fourier-Bessel basis functions is between

$8(cR)^2$ and $4\pi(cR)^2$. When c is the largest possible band limit, i.e. $c = \frac{1}{2}$, the number of basis functions is between $2R^2$ and πR^2 , where the latter is approximately the number of pixels

inside a disk of radius R . Also, whenever $c = \mathcal{O}(1)$ and $R = \mathcal{O}(L)$, we get that $p = \mathcal{O}(L^2)$ and $k_{\max} = \mathcal{O}(L)$.

Because the bandlimited function f is assumed to be essentially compactly supported, the infinite expansion in Eq. (3) is approximated by the finite expansion

$$P_{c,R}\mathcal{F}(f)(\xi, \theta) = \sum_{k=-k_{\max}}^{k_{\max}} \sum_{q=1}^{p_k} a_{k,q} \psi_c^{k,q}(\xi, \theta), \quad (15)$$

where $P_{c,R}$ is the orthogonal projection from $L^2(D_c)$ (the space of L^2 functions supported on a disk of radius c), to the space of functions spanned by a finite number of Fourier-Bessel functions that satisfy (7).

B. Numerical Evaluation of Fourier–Bessel Expansion Coefficients

Previously in [9], the evaluation of the expansion coefficients $a_{k,q}$ of Eq. (15) was done by least squares. Let Ψ be the matrix whose entries are evaluations of the Fourier-Bessel functions at the Cartesian grid points, with rows indexed by the grid points and columns indexed by angular and radial frequencies. Finding the coefficient vector a as the solution to $\min_a \|\Psi a - I\|_2^2$ requires the computation of $\Psi^* I$, which takes $\mathcal{O}(pL^2) = \mathcal{O}(L^4)$ operations, because $p = \mathcal{O}(L^2)$. In general $a = (\Psi^* \Psi)^{-1} \Psi^* I$, but here $\Psi^* \Psi$ is approximately the identity matrix, due to the orthogonality of the Fourier-Bessel functions.

We introduce here a method that computes the expansion coefficients $a_{k,q}$ in $\mathcal{O}(L^3)$ operations instead of $\mathcal{O}(L^4)$. The expansion coefficients in Eq. (15) are given analytically by

$$a_{k,q} = \int_0^{2\pi} \int_0^c \mathcal{F}(f)(\xi, \theta) \overline{\psi_c^{k,q}(\xi, \theta)} \xi d\xi d\theta = \int_0^c N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) \xi d\xi \int_0^{2\pi} \mathcal{F}(f)(\xi, \theta) e^{-ik\theta} d\theta. \quad (16)$$

We evaluate the last integral numerically using a quadrature rule that consists of equally spaced points in the angular direction and a Gaussian quadrature rule in the radial direction, that is, using the nodes, $\xi_1(j, l) = \xi_j \cos(2\pi l/n_\theta)$, $\xi_2(j, l) = \xi_j \sin(2\pi l/n_\theta)$, $j = 1, \dots, n_\xi$, $l = 0, \dots, n_\theta - 1$ (see Fig. 1). The values of n_ξ and n_θ depend on the compact support radius R and the band limit c and are derived later in the letter. To use this quadrature rule, we need to sample the Fourier transform of f at the quadrature nodes. This is approximated by the Fourier coefficients of the image I (consisting of samples of f on a Cartesian grid) at the given quadrature nodes, namely by the Fourier coefficients

$$F(I)(\xi_1, \xi_2) = \frac{1}{(2R)^2} \sum_{i_1=-R}^{R-1} \sum_{i_2=-R}^{R-1} I(i_1, i_2) e^{-i2\pi(\xi_1 i_1 + \xi_2 i_2)}, \quad (17)$$

which can be evaluated efficiently using the the nonuniform discrete Fourier transform. The angular integration in Eq. (16) is then sped up by 1D FFT on the concentric circles, followed by a numerical evaluation of the radial integral with a Gaussian quadrature rule.

As the samples on each concentric circle are equally-spaced, the natural quadrature weights

for the angular integral are $\frac{2\pi}{n_\theta}$, with the nodes taken at $\theta_l = \frac{2\pi l}{n_\theta}$ for $l = 0, \dots, n_\theta - 1$. The angular integration using one-dimensional FFT on each concentric circle thus yields

$$\widehat{F(I)}(\xi_j, k) = \frac{2\pi}{n_\theta} \sum_{l=0}^{n_\theta-1} F(I)(\xi_j, \theta_l) e^{-i \frac{2\pi k l}{n_\theta}}. \quad (18)$$

The radial integral is evaluated using the Gauss-Legendre quadrature rule [30, Chap. 4],

which determines the locations of n_ξ points $\{\xi_j\}_{j=1}^{n_\xi}$ on the interval $[0, c]$ and the associated weights $w(\xi_j)$. The integral in Eq. (16) is thus approximated by

$$a_{k,q} \approx \sum_{j=1}^{n_\xi} N_{k,q} J_k \left(R_{k,q} \frac{\xi_j}{c} \right) \widehat{F(I)}(\xi_j, k) \xi_j w(\xi_j). \quad (19)$$

Since I is real valued and $J_{-k}(x) = (-1)^k J_k(x)$, we get that $a_{-k,q} = a_{k,q}^*$ and thus we only need to evaluate coefficients with $k \geq 0$.

The procedure for numerical evaluation of the Fourier-Bessel expansion coefficients is illustrated in Fig. 1. In practice, we have observed that using $n_\xi = 4cR$ and $n_\theta = 16cR$ results in highly-accurate numerical evaluation of the integral in Eq. (16).

If our image can be expressed in terms of the truncated Fourier-Bessel expansion in Eq. (15), the approximation error in the radial integral comes from the numerical evaluation of the integrals

$$G(k, q_1, q_2) = \int_0^c J_k \left(R_{k,q_1} \frac{\xi}{c} \right) J_k \left(R_{k,q_2} \frac{\xi}{c} \right) \xi d\xi, \quad (20)$$

where the approximation error using n_ξ points is

$$E(k, q_1, q_2; n_\xi) = \left| \sum_{j=1}^{n_\xi} J_k \left(R_{k,q_1} \frac{\xi_j}{c} \right) J_k \left(R_{k,q_2} \frac{\xi_j}{c} \right) \xi_j w(\xi_j) - G(k, q_1, q_2) \right|. \quad (21)$$

Asymptotically, a Bessel function behaves like a decaying cosine function with frequency

$$\frac{R_{k,q}}{2\pi} \text{ for } R_{k,q} r \gg |k^2 - \frac{1}{4}| \text{ [26],}$$

$$J_k(R_{k,q}r) \sim \sqrt{\frac{2}{\pi R_{k,q}r}} \cos\left(R_{k,q}r - \frac{k\pi}{2} - \frac{\pi}{4}\right). \quad (22)$$

For a fixed n_ξ , the largest approximation error occurs when $k=0$ and $q_1 = q_2 = p_0$, since

$J_0\left(R_{0,p_0}\frac{\xi}{c}\right)$ is the most oscillatory function within the band limit. The Nyquist rate of $\xi J_0^2\left(R_{0,p_0}\frac{\xi}{c}\right)$ is $2\frac{2R_{0,p_0}}{2\pi} \approx 4cR$ and we need to sample at Nyquist rate, or higher. Therefore, we choose $n_\xi = \lceil 4cR \rceil$. Fig. 2a justifies this choice as the error decays dramatically to 10^{-17} before $n_\xi = \lceil 4cR \rceil$.

To choose n_θ , we computed the root mean squared error (RMSE) in evaluating the expansion coefficients for simulated images composed of white Gaussian noise with various R and n_θ , while $c = 1/2$. We oversampled on the radial lines by $n_\xi = \lceil 10cR \rceil$ and the ground truth for the angular integral in Eq. (16) was computed by Eq. (18) via oversampling in the angular direction by $n_\theta = 60cR$. We observe that when $n_\theta = 16cR$, the estimation error for the Fourier-Bessel expansion coefficients becomes negligible (see Fig. 2b). Notice that Eq. (14) implies that $k_{\max} < 2\pi cR$. The corresponding Nyquist rate is bounded by $4\pi cR$. We therefore sample at a slightly higher rate of $n_\theta = 16cR$ to ensure numerical accuracy.

Now that we are able to numerically evaluate $a_{k,q}$ with high accuracy, we can study the spectral behavior of the finite Fourier-Bessel expansion of the images. We define a as the vector that contains the expansion coefficients $a_{k,q}$ computed in Eq. (19) and denote by T^* the transformation that maps an image I to its finite Fourier-Bessel expansion coefficients through Eqs. (17), (18) and (19), that is,

$$a = T^*I. \quad (23)$$

Ideally we would like T^* to be a unitary transformation, that is $T^*T = I$, so that the transformation from the images to the coefficients preserves the noise statistics. Numerically, we observe that the majority of the eigenvalues of T^*T are 1 and the smallest eigenvalues are also close to 1 (see blue solid line in Fig. 3). The transformation T^* is close to unitary because it is a numerical approximation of an expansion in an orthogonal basis (Fourier-Bessel), and the sampling criterion prevents aliasing. In Fig. 3, the eigenvalues of $\Psi^*\Psi$ are also plotted for comparison. It can be observed that T^*T has fewer eigenvalues that deviate from 1. Although the Fourier-Bessel functions are orthogonal as continuous functions, their discrete sampled versions are not necessarily orthogonal, hence $\Psi^*\Psi$ deviates from the identity matrix. The fact that T^*T is closer to the identity than $\Psi^*\Psi$ implies that the numerical evaluation of the expansion coefficient vector a as T^*I is more accurate than estimating it as Ψ^*I . We compare the numerical accuracy explicitly with an example. We choose a signal f that satisfies $\mathcal{F}f(\xi, \theta) = \psi_c^{1,5}(\xi, \theta)$ for $c = 0.5$ and $R = 30$, $a_{k,q} = 1$, for $k = 1$ and $q = 5$, and otherwise, $a_{k,q} = 0$. The evaluation method from [9] is applied here in Fourier space. It first evaluates discrete samples of $\mathcal{F}(f)$ and the Fourier-Bessel basis

on a Cartesian grid of size $2R \times 2R$, and then projects the discrete samples onto the basis. The root mean squared error (RMSE) is 7.2×10^{-5} and the maximum absolute error is 4.0×10^{-4} . Using the numerical evaluation in Eq. (19), we get that $\text{RMSE} = 1.2 \times 10^{-16}$ and the maximum absolute error is 2.7×10^{-15} .

Computing the polar Fourier transform of an image of size $L \times L$ on a polar grid with $n_\xi \times n_\theta$ points in Eq. (17) is implemented efficiently using NUFFT [19], [20], [21], [22], whose computational complexity is $\mathcal{O}(L^2 \log L + n_\xi n_\theta)$. Since $n_\theta = 16cR = \mathcal{O}(L)$ and $n_\xi = 4cR = \mathcal{O}(L)$, $n_\xi \times n_\theta = \mathcal{O}(L^2)$ and the complexity of the discrete polar Fourier transform is $\mathcal{O}(L^2 \log L)$. The complexity of the 1D FFTs in Eq. (18) is $\mathcal{O}(n_\xi n_\theta \log n_\theta)$, because there are n_ξ concentric circles with n_θ samples on each circle. Both n_ξ and n_θ are of $\mathcal{O}(L)$, so the total complexity of the 1D FFTs is also $\mathcal{O}(L^2 \log L)$. Evaluating Eq. (19) (the quadrature rule for the radial integral in Eq. (16)) for all k and q requires a total of $\mathcal{O}(L^3)$ operations using a direct method, because there are $\mathcal{O}(L^2)$ basis functions to integrate, and each function is integrated using $\mathcal{O}(L)$ quadrature points. However, this complexity can be reduced to $\mathcal{O}(L^2 \log L)$ using a fast Bessel transform [23], [24]. In summary, the computational complexity of computing the Fourier-Bessel expansion coefficients of an image of size $L \times L$ is $\mathcal{O}(L^3)$ operations, or $\mathcal{O}(L^2 \log L)$ using a “fast” transform.

III. Steerable Principal Components

Given a dataset of n images $\{I_i\}_{i=1}^n$, we denote by f_i the underlying bandlimited function that corresponds to the i th image I_i . Under the action of the group $\mathcal{O}(2)$, the function f_i is transformed to $f_i^{\alpha, \beta}$, where $\alpha \in [0, 2\pi)$ is the counter-clockwise rotation angle and β denotes reflection and takes values in $\{+, -\}$. More specifically, $f_i^{\alpha, +}(r, \phi) = f_i(r, \phi - \alpha)$ and $f_i^{\alpha, -}(r, \phi) = f_i(r, \pi - (\phi - \alpha))$. The images $I_i^{\alpha, +}$ and $I_i^{\alpha, -}$ are obtained by sampling $f_i^{\alpha, +}$ and $f_i^{\alpha, -}$ respectively.

The Fourier transform of f_i commutes with the action of the group $\mathcal{O}(2)$, namely,

$$\mathcal{F}(f_i^{\alpha, +})(\xi, \theta) = (\mathcal{F}(f_i))^{\alpha, +}(\xi, \theta) = \mathcal{F}(f_i)(\xi, \theta - \alpha), \text{ and}$$

$\mathcal{F}(f_i^{\alpha, -})(\xi, \theta) = (\mathcal{F}(f_i))^{\alpha, -}(\xi, \theta) = \mathcal{F}(f_i)(\xi, \pi - (\theta - \alpha))$. The transformation of the images under rotation and reflection can be represented by the transformation of their Fourier-Bessel expansion coefficients in Eq. (3). Under counter-clockwise rotation by an angle α ,

$\mathcal{F}(f_i^{\alpha, +})$ is given by

$$\mathcal{F}(f_i^{\alpha, +})(\xi, \theta) = \sum_{k, q} a_{k, q}^i \psi_c^{k, q}(\xi, \theta - \alpha) = \sum_{k, q} a_{k, q}^i e^{-ik\alpha} \psi_c^{k, q}(\xi, \theta). \quad (24)$$

Therefore a planar rotation introduces a phase shift in the expansion coefficients. Under rotation and reflection,

$$\begin{aligned}
\mathcal{F}(f_i^{\alpha,-})(\xi, \theta) &= \sum_{k,q} a_{k,q}^i \psi_c^{k,q}(\xi, \pi - (\theta - \alpha)) \\
&= \sum_{k,q} a_{k,q}^i N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) e^{\iota k(\pi - \theta + \alpha)} \\
&= \sum_{k,q} a_{k,q}^i N_{k,q} (-1)^k J_k \left(R_{k,q} \frac{\xi}{c} \right) e^{\iota(-k)\theta} e^{\iota k \alpha} \\
&= \sum_{k,q} a_{k,q}^i e^{\iota k \alpha} \psi_c^{-k,q}(\xi, \theta) = \sum_{k,q} a_{-k,q}^i e^{-\iota k \alpha} \psi_c^{k,q}(\xi, \theta),
\end{aligned} \tag{25}$$

namely, the expansion coefficient $a_{k,q}^i$ changes to $a_{-k,q}^i e^{-\iota k \alpha}$.

If we augment the collection of bandlimited functions $\{f_i\}_{i=1}^n$ by all possible rotations and reflections, the Fourier transform of the sample mean of the augmented collection, denoted f_{mean} , becomes,

$$\mathcal{F}(f_{\text{mean}})(\xi, \theta) = \frac{1}{2n} \sum_{i=1}^n \sum_{\beta \in \{+, -\}} \frac{1}{2\pi} \int_0^{2\pi} \mathcal{F}(f_i^{\alpha, \beta})(\xi, \theta) d\alpha. \tag{26}$$

Using the properties in Eqs. (24) and (25), we have

$$\mathcal{F}(f_{\text{mean}})(\xi, \theta) = \frac{1}{2n} \sum_{i=1}^n \frac{1}{2\pi} \times \int_0^{2\pi} \sum_{k=-\infty}^{\infty} \sum_{q=1}^{\infty} [a_{k,q}^i + a_{-k,q}^i] e^{-\iota k \alpha} \psi_c^{k,q}(\xi, \theta) d\alpha = \sum_{q=1}^{\infty} \left(\frac{1}{n} \sum_{i=1}^n a_{0,q}^i \right) \psi_c^{0,q}(\xi, \theta). \tag{27}$$

As expected, the sample mean is radially symmetric, because $\psi_c^{0,q}$ is only a function of ξ but not of θ .

The rotationally invariant covariance kernel $\mathcal{C}((\xi, \theta), (\xi', \theta'))$ built from Fourier transformed functions with all their possible in-plane rotations and reflections is defined as

$$\begin{aligned}
\mathcal{C}((\xi, \theta), (\xi', \theta')) &= \frac{1}{4\pi n} \times \\
&\sum_{i=1}^n \sum_{\beta \in \{+, -\}} \int_0^{2\pi} \left(\mathcal{F}(f_i^{\alpha, \beta})(\xi, \theta) - \mathcal{F}(f_{\text{mean}})(\xi, \theta) \right) \\
&\times \overline{\left(\mathcal{F}(f_i^{\alpha, \beta})(\xi', \theta') - \mathcal{F}(f_{\text{mean}})(\xi', \theta') \right)} d\alpha.
\end{aligned} \tag{28}$$

From Eq. (27) it follows that if we express $\mathcal{F}(f_i)$ and $\mathcal{F}(f_{\text{mean}})$ in terms of the Fourier-Bessel basis and the associated expansion coefficients, subtracting the sample mean is equivalent to

subtracting $\frac{1}{n} \sum_{j=1}^n a_{0,q}^j$ from the coefficients $a_{0,q}^i$, while keeping other coefficients unchanged. Therefore, we first update the zero angular frequency coefficients by

$$a_{0,q}^i \leftarrow a_{0,q}^i - \frac{1}{n} \sum_{j=1}^n a_{0,q}^j \text{ and then}$$

$$\begin{aligned} \mathcal{C}((\xi, \theta), (\xi', \theta')) &= \frac{1}{4\pi n} \sum_{i=1}^n \int_0^{2\pi} \sum_{k=-\infty}^{\infty} \sum_{q=1}^{\infty} \sum_{k'=-\infty}^{\infty} \sum_{q'=1}^{\infty} \left(a_{k,q}^i \psi_c^{k,q}(\xi, \theta) \overline{a_{k',q'}^i \psi_c^{k',q'}(\xi', \theta')} \right. \\ &\quad \left. + a_{-k,q}^i \psi_c^{k,q}(\xi, \theta) \overline{a_{-k',q'}^i \psi_c^{k',q'}(\xi', \theta')} \right) e^{-\iota(k-k')\alpha} d\alpha \\ &= \sum_{k=-\infty}^{\infty} \sum_{q=1}^{\infty} \sum_{k'=-\infty}^{\infty} \sum_{q'=1}^{\infty} \psi_c^{k,q}(\xi, \theta) C_{(k,q),(k',q')} \overline{\psi_c^{k',q'}(\xi', \theta')}, \end{aligned} \quad (29)$$

where

$$C_{(k,q),(k',q')} = \frac{1}{4\pi n} \sum_{i=1}^n \int_0^{2\pi} \left(a_{k,q}^i \overline{a_{k',q'}^i} + a_{-k,q}^i \overline{a_{-k',q'}^i} \right) e^{-\iota(k-k')\alpha} d\alpha = \delta_{k,k'} \frac{1}{n} \sum_{i=1}^n \text{Re} \left\{ a_{k,q}^i \overline{a_{k',q'}^i} \right\}. \quad (30)$$

$\delta_{k,k'}$ comes from the integral over $\alpha \in [0, 2\pi)$. The covariance matrix in Eq. (30) is positive semi-definite and block diagonal because the non-zero entries of C correspond only to $k = k'$. Since the images are well approximated by the subspace spanned by a finite number of Fourier-Bessel basis functions (see Eq. (15)), $C_{(k,q),(k',q')}$ are close to zero when (k, q) or (k', q') do not satisfy the sampling criterion in Eq. (7). Therefore, we have a finite matrix representation C of \mathcal{C} . Moreover, it suffices to consider $k \geq 0$, because $C_{(k,q),(k',q')} = C_{(-k,q),(-k',q')}$. Thus, the covariance matrix in Eq. (30) can be written as the direct sum $C = \bigoplus_{k=0}^{k_{\max}} C^{(k)}$, where $C^{(k)}$ is by itself a sample covariance matrix of size $p_k \times p_k$, given by,

$$C_{q,q'}^{(k)} = \frac{1}{n} \sum_{i=1}^n \text{Re} \left\{ a_{k,q}^i \overline{a_{k,q'}^i} \right\}. \quad (31)$$

Let us denote by $A^{(k)}$ the matrix of expansion coefficients, obtained by putting the coefficients $a_{k,q}^i$ for all q and all i into a matrix, where the columns are indexed by the image number i and the rows are indexed by the radial index q . The coefficient matrix $A^{(k)}$ for $k \geq 0$ is of size $p_k \times n$ and the covariance matrix for $k \geq 0$ is,

$$C^{(k)} = \frac{1}{n} \text{Re} \left\{ A^{(k)} (A^{(k)})^* \right\}, \quad (32)$$

where A^* is the conjugate transpose ($A_{ij}^* = \bar{A}_{ji}$). The case $k=0$ is special because the expansion coefficients satisfy $a_{0,q} = \bar{a}_{0,q}$, and so $A^{(0)}$ is a matrix of size $p_0 \times n$ and

$$C^{(0)} = \frac{1}{n} A^{(0)} (A^{(0)})^*. \quad (33)$$

We compute the eigenvalues $\lambda_1^{(k)} \geq \lambda_2^{(k)} \dots \geq \lambda_{p_k}^{(k)}$ and eigenvectors $u_1^{(k)}, u_2^{(k)}, \dots, u_{p_k}^{(k)}$ of the covariance matrices $C^{(k)}$. Because \mathcal{C} and C are related through Eq. (29) and C is block diagonal as in Eq. (30), $\mathcal{C}((\xi, \theta), (\xi', \theta'))$ is well approximated by

$\sum_{k=-k_{\max}}^{k_{\max}} \Psi^{(k)}(\xi, \theta) C^{(k)} (\Psi^{(k)})^*(\xi', \theta')$, where $\Psi^{(k)}$ contains Fourier-Bessel functions with angular frequency k . Equation (29) reveals that the eigenfunctions of \mathcal{C} , which are the steerable principal components, can be expressed as linear combinations of the Fourier-Bessel functions with the coefficients given by the eigenvectors of the matrix C ,

$$g^{k,l}(\xi, \theta) = \sum_{q=1}^{p_k} \psi_c^{k,q}(\xi, \theta) u_l^{(k)}(q) = \sum_{q=1}^{p_k} N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) u_l^{(k)}(q) e^{i k \theta}. \quad (34)$$

Therefore the radial parts of the steerable principal components

$$f^{k,l}(\xi) = \sum_{q=1}^{p_k} N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) u_l^{(k)}(q) \quad (35)$$

are linear combinations of the Bessel functions within the same angular frequency. The associated expansion coefficients for I_i are

$$c_{k,l}^i = \sum_{q=1}^{p_k} a_{k,q}^i u_l^{(k)}(q), \quad \text{for } i=1, \dots, n. \quad (36)$$

The computational complexity for forming the matrix $C^{(k)}$ is $O(np_k^2)$. The complexity for eigendecomposition of $C^{(k)}$ is $O(p_k^3)$, since the size of the covariance matrix is $p_k \times p_k$. Using the upper and lower bounds for p_k in Eq. (13) and assuming $c = O(1)$ and $R = O(L)$, we get $\sum_k p_k^2 = O(L^3)$ and $\sum_k p_k^3 = O(L^4)$. Therefore, the complexity for forming the covariance matrix C is $O(n \sum_k p_k^2) = O(nL^3)$ and the complexity of its full eigendecomposition is $O(\sum_k p_k^3) = O(L^4)$. Equations (32) and (33) show that instead of

constructing the covariance matrices $C^{(k)}$ to compute the principal components, we can perform singular value decomposition (SVD) on the coefficient matrix $A^{(k)}$ directly and take the left singular vectors as the principal components. The computational complexity for full rank SVD on $A^{(k)}$ is $O(np_k^2)$ and the total complexity of SVD of all coefficient matrices is $O(n \sum_k p_k^2) = O(nL^3)$.

IV. Algorithm and Computational Complexity

The new algorithm introduced in this letter is termed fast Fourier-Bessel steerable PCA (FFBsPCA). The algorithm is composed of two steps. In the first step, Fourier-Bessel expansion coefficients are computed according to Algorithm 1. The input to the algorithm includes an image dataset, the band limit c , and the compact support radius R . The second step (Algorithm 2) takes the Fourier-Bessel expansion coefficients from Algorithm 1 as input and computes the steerable PCA radial functions and the expansion coefficients of the images in the new steerable basis. Algorithm 2 is the same as the corresponding part of the algorithm in [9].

Algorithm 1

Fast Fourier-Bessel Expansion

Require: n images I_1, \dots, I_n sampled on a Cartesian grid of size $L \times L$ with compact support radius R and band limit c .

- 1: (Precomputation) Select (k, q) 's that satisfy the sampling criterion of Eq. (7). Fix $n_\xi = \lceil 4cR \rceil$ and $n_\theta = \lceil 16cR \rceil$.
 - 2: (Precomputation) Find n_ξ Gaussian quadrature points and weights on the interval $[0, c]$ and evaluate $N_{k,q} J_k(R_{k,q} \frac{\xi_j}{c})$, $j = 1, \dots, n_\xi$, for all selected (k, q) 's.
 - 3: Compute $F(I_i)$ (Eq. (17)) on a polar grid of size $n_\xi \times n_\theta$ by NUFFT for each $i = 1, \dots, n$.
 - 4: For each $F(I_i)$, compute $a_{k,q}^i$ using Eqs. (18) and (19).
 - 5: **return** $a_{k,q}^i$ for all selected (k, q) 's.
-

The analysis of the computational complexity of FFBsPCA is as follows. The precomputation that generates all radial basis functions requires $O(L^3)$ operations because there are $O(L^2)$ basis functions, each of which is sampled over $O(L)$ points. Computing the Fourier-Bessel expansion coefficients $a_{k,q}^i$ in Eq. (19) for all images takes $O(nL^3)$ operations (or $O(nL^2 \log L)$ with a fast Bessel transform) as discussed in Section II-B.

The complexity of constructing the covariance matrix C and computing its full eigendecomposition is $O(nL^3 + L^4)$ as described in Section III. Another method for computing the principal components is by SVD of the coefficient matrices.

Algorithm 2

Steerable PCA

Require: Fourier-Bessel expansion coefficients $a_{k,q}^i$ for n images and the maximum angular frequency k_{\max} .

- 1: Compute the coefficient vector of the mean image

$$a_{0,q}^{\text{mean}} = \frac{1}{n} \sum_j a_{0,q}^j$$
 Then, set $a_{0,q}^i \leftarrow a_{0,q}^i - a_{0,q}^{\text{mean}}$.
- 2: **for** $k = 0, 1, \dots, k_{\max}$ **do**
- 3: Construct the coefficient matrix $A^{(k)}$.
- 4: Compute the covariance matrix $C^{(k)}$, its eigenvalues $\lambda_1^{(k)} \geq \lambda_2^{(k)} \dots \geq \lambda_{p_k}^{(k)}$, and eigenvectors, $u_1^{(k)}, \dots, u_{p_k}^{(k)}$; or perform SVD of $A^{(k)}$ and take the left singular vectors $u_1^{(k)}, \dots, u_{p_k}^{(k)}$.
- 5: Compute the radial eigenvectors $f^{k,l}(\xi_j)$ for $j = 1, \dots, n_{\xi}$ using Eq. (35).
- 6: Compute the expansion coefficients of the images in the new steerable basis $c_{k,l}^i$ using Eq. (36).
- 7: **end for**
- 8: **return** For all (k, l) , $u_l^{(k)}$, $\lambda_l^{(k)}$, $f^{k,l}$, and $c_{k,l}^i$
 $i = 1, \dots, n$.

Full rank SVD on all coefficient matrices requires $\mathcal{O}(nL^3)$ floating point operations (see Section III).

To generate the new steerable basis, we take linear combinations of the Bessel functions as in line 5 of Algorithm 2, which takes $\mathcal{O}(L^4)$ operations. Computing the steerable PCA expansion coefficients $c_{k,l}^i$ for $i = 1 \dots, n$ (line 6 in Algorithm 2) requires $\mathcal{O}(nL^3)$ operations by taking linear combinations of the Fourier-Bessel expansion coefficients as in Eq. (36). Therefore the total computational complexity of FFBSPCA is $\mathcal{O}(nL^3 + L^4)$.

The complexity of FBSPCA introduced in [9] is $\mathcal{O}(nL^4)$. Thus, FFBSPCA is faster than FBSPCA. For PCA, when the number of images is smaller than the number of pixels in the compact support disk, we form $X^T X$ and compute its eigendecomposition and the complexity is $\mathcal{O}(n^2 L^2 + n^3)$. However, as the number of images grows, the complexity of PCA switches to $\mathcal{O}(nL^4 + L^6)$ since it becomes more efficient to compute the eigendecomposition of XX^T . Therefore the computational complexity of traditional PCA, without taking into account all rotations and reflections is $\mathcal{O}(\min\{n^2 L^2 + n^3, nL^4 + L^6\})$. When $n > \mathcal{O}(L)$, FFBSPCA is more efficient than the traditional PCA.

FFBSPCA is easily adapted for parallel computation. The computation of Fourier-Bessel expansion coefficients in Algorithm 1 can run on multiple workers in parallel, where each worker is allocated with a subset of the images and Fourier-Bessel radial basis functions. In

addition, in Algorithm 2, the radial eigenfunctions and the steerable PCA expansion coefficients can also be efficiently computed in parallel for each angular index k .

V. Numerical Experiments

We compare the running times of FFBSPCA, FBSPCA and traditional PCA, where the latter is computed without the images' in-plane rotations and reflections. The algorithms are implemented in MATLAB on a machine with 60 cores, running at 2.3 GHz, with total RAM of 1.5TB.

We first simulated $n = 24,000$ images with different radii of compact support R , while the band limit is fixed at $c = 1/2$. For small R , since FFBSPCA performs polar Fourier transformation, it appears slightly slower than FBSPCA. However when R increases, FFBSPCA is computationally more efficient (see Tab. I). We next fixed the size of the images while using $R = 150$ and $c = 1/2$, and varied the number of images n . Table II shows that the running time of FBSPCA and FFBSPCA grows linearly with n .

To show that our new algorithm can handle large datasets efficiently, we simulated a large dataset with 10^5 images of size 300×300 pixels. The images consist entirely of Gaussian noise with mean 0 and variance 1. We assume that the compact support in the image domain is $R = 150$ and the band limit in Fourier domain is $c = 1/2$. In Table III, the total running time is divided into three parts: precomputation, Fourier-Bessel expansion (Algorithm 1), and steerable PCA (Algorithm 2). Fourier Bessel expansion took about 24 minutes, during which 91% of the time was spent on mapping images to polar Fourier grid, where we used the software package [22] downloaded from <https://www-user.tu-chemnitz.de/potts/nfft/potts/nfft/>. Numerical evaluation of the angular integration by 1D FFT and the radial integration by a direct method took 6.4% and 2.6% of the time respectively. Steerable PCA took 42 seconds.

In our third experiment, we simulated $n = 10^5$ clean projection images from a reconstructed volume of a human mitochondrial large ribosomal subunit, downloaded from the electron microscopy data bank [32] (EMDB-2762). The original volume in the data bank is of size $320 \times 320 \times 320$ voxels. We preprocessed the volume such that its center of mass is at the origin and cropped out a volume of size $240 \times 240 \times 240$ voxels that contains the particle. Each projection image is of size 240×240 pixels. We simulated both the vanishing behavior of the CTF at low frequencies and the blurring effect due to the Gaussian envelope of the CTF. This was done by convolving the images with the inverse Fourier transform of

$$\min(\pi\lambda z f^2 + a, 1) \exp(-Bf^2), \quad (37)$$

where f is the frequency, λ is the wavelength of the electron beam, z is the defocus, and a is the phase of the CTF introduced by microscope. This stems from the analytic form of the CTF given by $\sin(\pi\lambda z f^2 + a) \exp(-Bf^2)$. For the simulations we chose $\lambda = 0.0197\text{\AA}$, $z = 2.5\ \mu\text{m}$, $a = 0.1\text{rad}$, and $B = 100\text{\AA}^2$. Our clean images (see Fig. 4a) are the projection images

filtered by the filter in Eq. (37) and they were then corrupted by additive white Gaussian noise at SNR= 1/30, corresponding to noise variance of $\sigma^2 = 9$ (see Fig. 4b).

We estimated the radius of compact support of the particle in real domain and the band limit in Fourier domain from the noisy images in the following way. We first subtracted the mean image of the dataset from each image. Then we computed the 2D variance map of the dataset averaged in the angular direction, to get the mean radial variance (see Fig. 5a). At large r , the mean radial variance levels off at 9, which corresponds to the noise variance. We subtracted the noise variance from the estimated mean radial variance and computed the cumulative variance by integrating the mean radial variance over r with a Jacobian weight rdr . The fraction of the cumulative variance reaches 99.9% at $r = 98$, and therefore R was chosen to be 98. In the Fourier domain, we computed the angular average of the mean 2D power spectrum. The curve in Fig. 5b also levels off at the noise variance when ξ is large. We used the same method as before to compute the cumulative radial power spectrum. The fraction reaches 99.9% at $\xi = 0.196$, therefore the band limit is chosen to be $c = 0.196$.

The radial functions of the top nine principal components are shown in Fig. 6. Each radial function is indexed by k and l , where k determines the angular Fourier mode and l is the order of the radial function within the same k . Taking the tensor product of the radial functions and their corresponding angular Fourier modes gives the two dimensional principal components in Fourier domain. It took about 9 minutes in total to get the steerable PCA radial components and the associated expansion coefficients. In particular, Fourier-Bessel expansion coefficients were computed in 9 minutes and the steerable PCA took 12 seconds.

We computed the traditional PCA and FBsPCA on the same dataset in real image domain (see Fig. 8 for PCA components), which took 60 minutes and 16 minutes respectively. In order to compare the principal components computed by FFBsPCA with those computed by traditional PCA, we take the inverse Fourier transform of the FFBsPCA components. We do not compute the inverse polar Fourier transform directly, since such a transform is ill-conditioned. Instead, since the FFBsPCA components are linear combinations of the Fourier-Bessel functions as in Eq. (34), we evaluate the steerable principal components on a Cartesian grid in real space using the linear combinations of $\mathcal{F}^{-1}(\psi_c^{k,q})$, given by Eq. (5). Those principal components are shown in Fig. 7. Some of the top sixteen principal components computed from traditional PCA and FFBsPCA look very similar, for example, the first three and the last four principal components (see Fig. 7 and Fig. 8). Because the gap between the eigenvalues of the traditional PCA is very small for the components in the middle two rows of Fig. 8, those components become degenerate and therefore look different from the corresponding components in Fig. 7.

In our simulation, each noisy projection image I is obtained by contaminating the clean image I_c with additive white Gaussian noise of variance $\sigma^2 = 9$. Given the noise level, we would like to automatically select the appropriate principal components to compress and denoise the noisy images. Since the transformation T^* is nearly unitary, the coefficient matrices can be modeled approximately as $A^{(k)} = A_c^{(k)} + \varepsilon^{(k)}$, where $\varepsilon^{(k)}$ is white Gaussian

noise with variance σ^2 and $A_c^{(k)}$ is the coefficient matrix for the clean images. In the case when there is no signal, that is $A_c^{(k)}=0$, all eigenvalues of the covariance matrix $C^{(k)}$ from Eqs. (32) and (33) converge to σ^2 as n goes to infinity, while p_k is fixed. When $A_c^{(k)} \neq 0$, components with eigenvalues larger than σ^2 correspond to the underlying clean signal. In the non-asymptotic regime of a finite number of images, the eigenvalues of the sample covariance matrix from white Gaussian noise spread around σ^2 . The empirical density of the eigenvalues can be approximated by the Mar enko-Pastur distribution with parameter γ_k ,

where $\gamma_0 = \frac{p_0}{n}$ and $\gamma_k = \frac{p_k}{2n}$ for $k > 0$ and the eigenvalues of $C^{(k)}$ are supported on $[\lambda_-, \lambda_+]$, with $\lambda_{\pm} = \sigma^2(1 \pm \sqrt{\gamma_k})^2$. The principal components corresponding to eigenvalues larger than λ_+ correspond to signal information beyond noise level. Therefore, with the estimated noise variance $\hat{\sigma}^2$, we denote by $\lambda_1^{(k)} \geq \lambda_2^{(k)} \geq \dots \geq \lambda_{p_k}^{(k)}$ the eigenvalues of the covariance matrix $C^{(k)}$, and select the components with eigenvalues

$$\lambda_l^{(k)} > \hat{\sigma}^2(1 + \sqrt{\gamma_k})^2, l=1, \dots, p_k. \quad (38)$$

Various ways of selecting principal components from noisy data have been proposed. We refer to [33] for an automatic procedure for estimating the noise variance and the number of components beyond the noise level. For the simulated ribosomal subunit projections images, there are 966 steerable principal radial components above the threshold in Eq. (38), whereas considerably fewer principal components (391) with the traditional PCA were selected.

Moreover, we filter the expansion coefficients to get better denoising. To first order approximation, when $n \gg p_k$, the noise simply shifts all eigenvalues upward by σ^2 and this calls for soft thresholding of the sample covariance eigenvalues: $(\lambda - \sigma^2)_+$. To correct for the finite sample effect, we can apply more sophisticated shrinkage to the eigenvalues, such as the methods proposed in [34], [35]. Specifically, we applied the shrinkage method in [34] to the coefficients computed by FFBsPCA, FBsPCA, and PCA. Because we were able to use more principal components with FFBsPCA, we recovered finer details of the clean projection images, comparing Fig. 10c and Fig. 10f.

In addition to using data-adaptive bases, we also used a non-isotropic directional multiscale transform, i.e., Curvelet transform [36] with complex block thresholding and cycle spinning, to denoise the images. An example of a denoised image using PCA, Curvelet, FBsPCA, and FFBsPCA is shown in Fig. 9. The steerable PCA basis captures the variance of the clean dataset with fewer components than non-adaptive bases, such as Fourier-Bessel basis or Curvelets (see Fig. 11).

We computed the mean squared error (MSE) and Peak SNR (PSNR) to quantify the denoising effects in Tab. IV and Tab. V. Comparing with the traditional PCA, FFBsPCA reduced the MSE by more than 25% and increased the PSNR by over 1.3 dB. When the images are of low SNR, Curvelets are unable to outperform data adaptive bases, such as

PCA, FBsPCA and FFBSPCA (see Tab. IV and Tab. V). This experiment shows that FFBSPCA is an efficient and effective procedure for denoising large image datasets.

Finally, we show that steerable PCA denoising is robust to small shifts. We simulated clean data with random shifts in the $\pm x$ and $\pm y$ directions with maximum shifts equal to 0 (centered images), 5, 10, 15, and 20 pixels. The clean images are corrupted with additive white Gaussian noise of variance 9. As shown in Tab. VI, the denoising performance using FFBSPCA (measured in PSNR) is almost unaffected. The denoising results for centered images in Tab. V and Tab. VI are slightly different because we used different support sizes to evaluate PSNRs.

VI. Conclusion

In this letter we presented a fast Fourier-Bessel steerable PCA method that reduces the computational complexity with respect to the size of the images so that it can handle larger images. The complexity of the new algorithm is $\mathcal{O}(nL^3 + L^4)$ compared with $\mathcal{O}(nL^4)$ of the steerable PCA introduced in [9]. The key improvement is through mapping the images to a polar Fourier grid using NUFFT and evaluating the Fourier-Bessel expansion coefficients by angular 1D FFT and accurate radial integration.

This work has been mostly motivated by its application to cryo-EM single particle reconstruction. Besides compression and denoising of the experimental images required for 2D class averaging [11] and common-lines based 3D ab-initio modeling, FFBSPCA can also be applied in conjunction with Kam's approach [37] that requires the covariance matrix of the 2D images [38]. The method developed here can also be extended to perform fast principal component analysis of a set of 3D volumes and their rotations. For this purpose, the Fourier-Bessel basis is replaced with the spherical-Bessel basis, and the expansion coefficients can be evaluated by performing the angular integration using a fast spherical harmonics transform [39] followed by radial integration.

Our numerical experiments show that an adaptive basis is necessary for denoising images with very low SNR. Steerable PCA is able to recover more signal components than PCA and achieves better denoising results. It is definitely possible to improve the denoising obtained by just using steerable PCA. For example, we can have more sophisticated dictionary denoising schemes, in which part of the dictionary is made of the steerable principal components and another part of the dictionary is made of wavelets. As these methods require the computation of steerable PCA, computing steerable PCA fast would be useful also for more advanced denoising schemes.

Finally, we remark that the Fourier-Bessel basis can be replaced in our framework with other suitable bases, for example, the 2D prolate spheroidal wave functions (PSWF) on a disk [40]. The 2D prolates also have a separation of variables form which makes them convenient for steerable PCA. A possible advantage of using 2D prolates is that they are optimal in terms of the size of their support.

Acknowledgments

The authors would like to thank Leslie Greengard, Michael O'Neil, and Alex Townsend for valuable discussions.

This work was supported by the NIGMS under award number R01GM090200, in part by the Simons Foundation under award number LTR DTD 06-05-2012, in part by the Moore Foundation DDD Investigator Award, and in part by the Israel Science Foundation under Grant 578/14. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Alessandro Foi.

REFERENCES

1. van Heel M, Frank J. Use of multivariate statistics in analysing the images of biological macromolecules. *Ultramicroscopy*. 1981; 6(2):187–194. [PubMed: 7268930]
2. Frank, J. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. London, U.K: Oxford Univ. Press; 2006.
3. Buades A, Coll B, Morel J-M. A non-local algorithm for image denoising. *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recog. (CVPR)*. 2005; 2:60–65.
4. Donoho DL. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory*. 1995 May; 41(3):613–627.
5. Hilai R, Rubinstein J. Recognition of rotated images by invariant Karhunen-Loève expansion. *J. Opt. Soc. Amer. A*. 1994; 11(5):1610–1618.
6. Perona P. Deformable kernels for early vision. *IEEE Trans. Pattern Anal. Mach. Intell*. 1995 May; 17(5):488–499.
7. Uenohara M, Kanade T. Optimal approximation of uniformly rotated images: Relationship between Karhunen-Loève expansion and discrete cosine transform. *IEEE Trans. Image Process*. 1998 Jan. 7(1):116–119. [PubMed: 18267385]
8. Ponce C, Singer A. Computing steerable principal components of a large set of images and their rotations. *IEEE Trans. Image Process*. 2011 Nov.20(11):3051–3062. [PubMed: 21536533]
9. Zhao Z, Singer A. Fourier-Bessel rotational invariant eigenimages. *J. Opt. Soc. Amer. A*. 2013; 30(5):871–877.
10. Vonesch C, Stauber F, Unser M. Design of steerable filters for the detection of micro-particles. *Proc. 10th IEEE Int. Symp. Biomed. Imag. (ISBI'13)/Nano Macro*. 2013:934–937.
11. Zhao Z, Singer A. Rotationally invariant image representation for viewing direction classification in cryo-em. *J. Struct. Biol*. 2014; 186(1):153–166. [PubMed: 24631969]
12. Lenz R. Group-theoretical model of feature extraction. *J. Opt. Soc. Amer. A*. 1989; 6:827–834.
13. Lenz R. Group invariant pattern recognition. *Pattern Recognit*. 1990; 23:199–217.
14. Lenz, R. Steerable filters and invariant recognition in spacetime; *Proc. Int. Conf. Acoust. Speech Signal Process*; 1998. p. 2737-2740.
15. Lenz R, Bui TH, Hernández-Andrés J. Group theoretical structure of spectral spaces. *J. Math. Imag. Vis*. 2005; 23:297–313.
16. Lenz R. Investigation of receptive fields using representations of the dihedral groups. *J. Vis. Commun. Image Represent*. 1995; 6(3):209–277.
17. Jogan M, Zagar E, Leonardis A. Karhunen-Loève expansion of a set of rotated templates. *IEEE Trans. Image Process*. 2003 Jul.12(7):817–825. [PubMed: 18237956]
18. Klug A, Crowther RA. Three-dimensional image reconstruction from the viewpoint of information theory. *Nature*. 1972; 238:435–440.
19. Dutt A, Rokhlin V. Fast Fourier transforms for nonequispaced data. *SIAM J. Sci. Comput*. 1993; 14(6):1368–1393.
20. Fessler JA, Sutton BP. Nonuniform fast Fourier transforms using min-max interpolation. *IEEE Trans. Signal Process*. 2003 Feb.51(2):560–574.
21. Greengard L, Lee J. Accelerating the nonuniform fast Fourier transform. *SIAM Rev*. 2004; 46(3): 443–454.
22. Fenn M, Kunis S, Potts D. On the computation of the polar FFT. *Appl. Comput. Harmon. Anal*. 2007; 22:257–263.

23. O'Neil M, Woolfe F, Rokhlin V. An algorithm for the rapid evaluation of special function transforms. *Appl. Comput. Harmon. Anal.* 2010; 28:203–226.
24. Townsend A. A fast analysis-based discrete Hankel transform using asymptotic expansions. 2015 arxiv.org/abs/1501.01652.
25. Watson, GN. *Theory of Bessel Functions*. Cambridge, U.K: Cambridge Univ. Press; 1944.
26. Abramowitz, M.; Stegun, IA. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York, NY, USA: Dover; 1964.
27. Olver FWJ. A further method for the evaluation of zeros of Bessel functions and some new asymptotic expansions for zeros of functions of large order. *Math. Proc. Cambridge Philosoph. Soc.* 1951; 47:699–712.
28. Elbert Á. Some recent results on the zeros of Bessel functions and orthogonal polynomials. *J. Comput. Appl. Math.* 2001; 133(1–2):65–83.
29. Breen S. Uniform upper and lower bounds on the zeros of Bessel functions of the first kind. *J. Math. Anal. Appl.* 1995; 196:1–17.
30. Press, WH.; Flannery, BP.; Teukolsky, SA.; Vetterling, WT. *Numerical Recipes in FORTRAN 77: The Art of Scientific Computing*. plus Cambridge, U.K: Cambridge Univ. Press; 1986.
31. Rattey PA, Lindgren AG. Sampling the 2-D radon transform. *IEEE Trans. Acoust. Speech Signal Process.* 1981 Oct.(5):994–1002. vol. ASSP-29.
32. Brown A, et al. Structure of the large ribosomal subunit from human mitochondria. *Science*. 2014; 346(6210):718–722. [PubMed: 25278503]
33. Kritchman S, Nadler B. Determining the number of components in a factor model from limited noisy data. *Chemometr. Intell. Lab.* 2008; 94:19–32.
34. Singer A, Wu H-T. Two-dimensional tomography from noisy projections taken at unknown random directions. *SIAM J. Imag. Sci.* 2013; 6(1):136–175.
35. Donoho DL, Gavish M, Johnstone IM. Optimal shrinkage of eigenvalues in the spiked covariance model. 2014 arxiv.org/abs/1311.0851.
36. Candes EJ, Demanet L, Donoho DL, Ying L. Fast discrete curvelet transforms. *SIAM Multiscale Model. Simul.* 2005; 5(3):861–899.
37. Kam Z. The reconstruction of structure from electron micrographs of randomly oriented particles. *J. Theor. Biol.* 1980; 82:15–39. [PubMed: 7401655]
38. Bhamre T, Zhang T, Singer A. Orthogonal matrix retrieval in cryo-electron microscopy. *Proc. 12th IEEE Int. Symp. Biomed. Imag. (ISBI'15)*. 2015:1048–1052.
39. Rokhlin V, Tygert M. Fast algorithms for spherical harmonic expansions. *SIAM J. Sci. Comput.* 2006; 27(6):1903–1928.
40. Slepian D. Prolate spheroidal wave functions, Fourier analysis, and uncertainty—IV: Extensions to many dimensions, generalized prolate spheroidal wave functions. *Bell Syst. Tech. J.* 1964; 43:3009–3057.

Biographies



Zhizhen Zhao received the B.A. and M.Sc. degrees in physics from Trinity College, Cambridge University, Cambridge, U.K., the Ph.D. degree in physics from Princeton University, Princeton, NJ, USA, in 2008 and 2013, respectively. Since September 2013, she has been with the Courant Institute of Mathematical Sciences, New York University, New York, NY, USA. She is currently a Courant Instructor with NYU. Her research interests

include applied and computational harmonic analysis, signal processing, data analysis and the applications in structural biology and atmospheric, and oceanic sciences.



Yoel Shkolnisky received the B.Sc. degree in mathematics and computer science, and the M.Sc. and Ph.D. degrees in computer science from Tel-Aviv University, Tel-Aviv, Israel, in 1996, 2001, and 2005, respectively. From July 2005 to July 2008, he was a Gibbs Assistant Professor in applied mathematics with the Department of Mathematics, Yale University, New Haven, CT, USA. Since October 2009, he has been with the Department of Applied Mathematics, School of Mathematical Sciences, Tel-Aviv University. His research interests include computational harmonic analysis, scientific computing, and data analysis.



Amit Singer received the B.Sc. degree in physics and mathematics and the Ph.D. degree in applied mathematics from Tel-Aviv University, Tel-Aviv, Israel, in 1997 and 2005, respectively. He is a Professor of Mathematics and a Member of the Executive Committee of the Program in Applied and Computational Mathematics (PACM), Princeton University, Princeton, NJ, USA. He joined Princeton University as an Assistant Professor in 2008. From 2005 to 2008, he was a Gibbs Assistant Professor in Applied Mathematics with the Department of Mathematics, Yale University, New Haven, CT, USA. His research interests include applied mathematics focuses on theoretical and computational aspects of data science, and on developing computational methods for structural biology. He served in the Israeli Defense Forces from 1997 to 2003. He was the recipient of the Moore Investigator in Data-Driven Discovery (2014), the Simons Investigator Award (2012), the Presidential Early Career Award for Scientists and Engineers (2010), the Alfred P. Sloan Research Fellowship (2010), and the Haim Nessyahu Prize for Best PhD in Mathematics in Israel (2007).

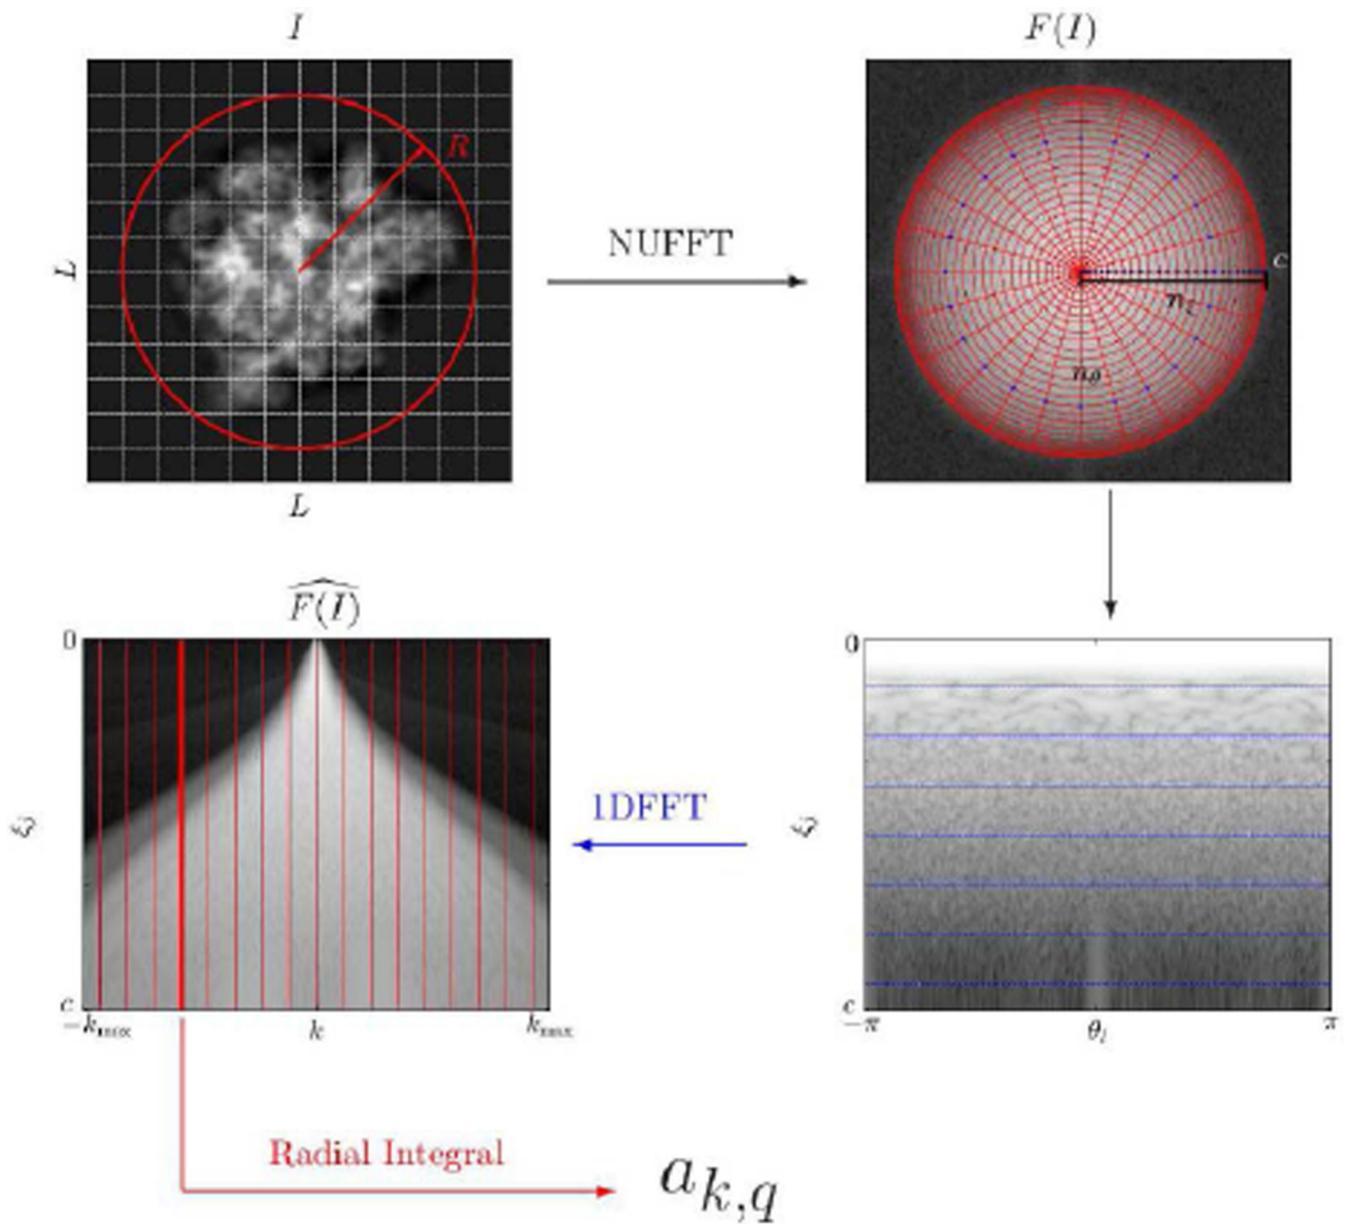
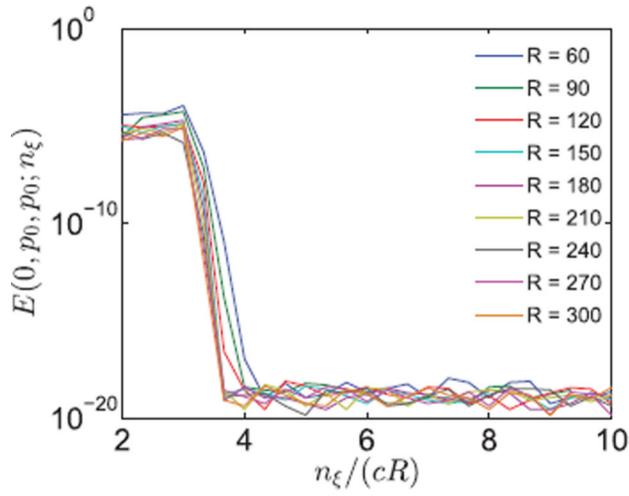
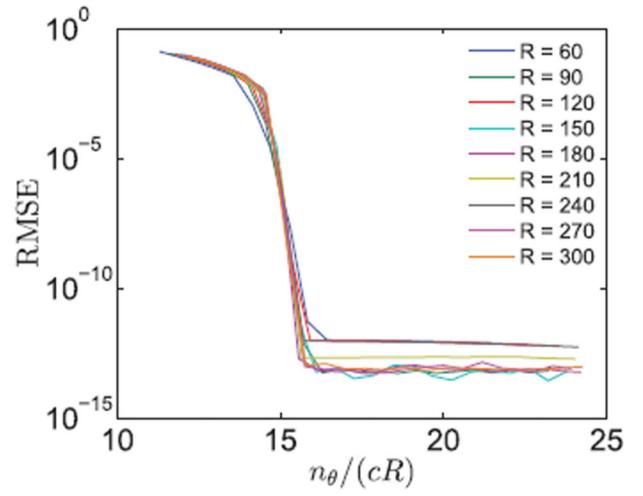


Fig. 1. Pictorial summary of the procedure for computing the Fourier-Bessel expansion coefficients. The original image (top left) is resampled on a polar Fourier grid (Eq. (17)) using NUFFT (top right and bottom right) followed by 1D FFT (Eq. (18)) on each concentric circle. The evaluation of the radial integral (Eq. (19)) gives the expansion coefficients $a_{k,q}$. The bow-tie phenomenon illustrated in bottom-left was discussed in [31].



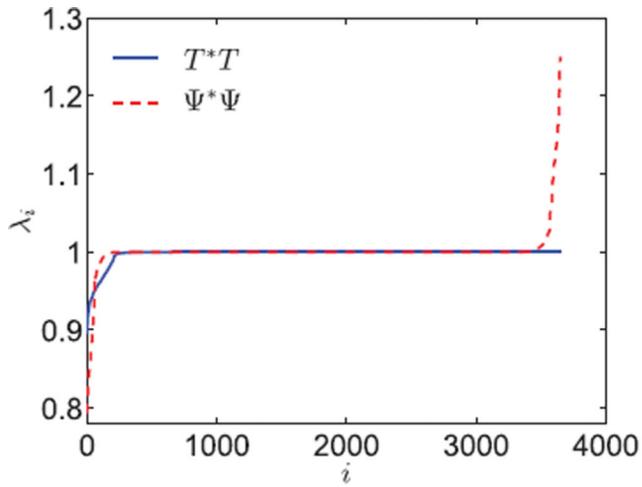
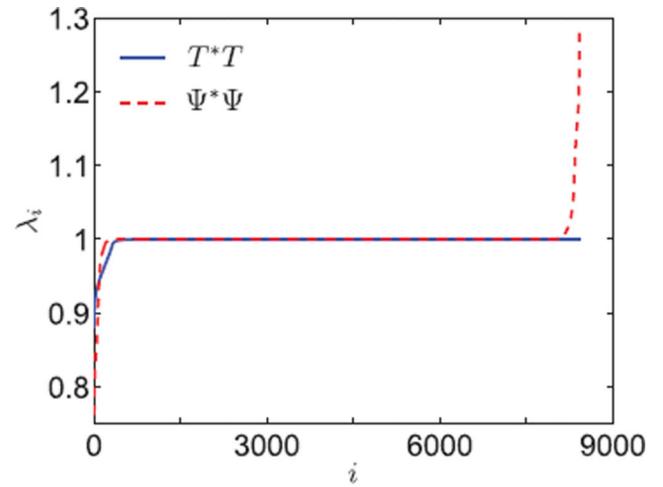
(a)



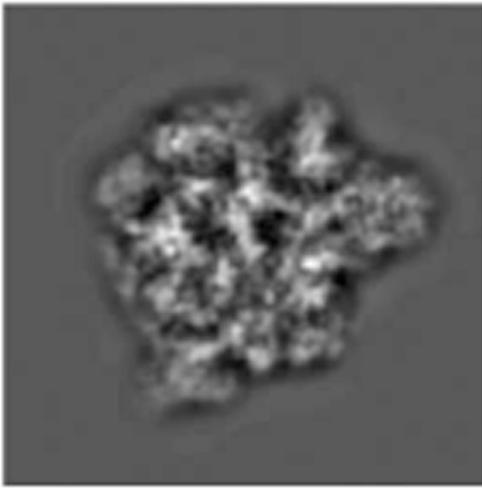
(b)

Fig. 2.

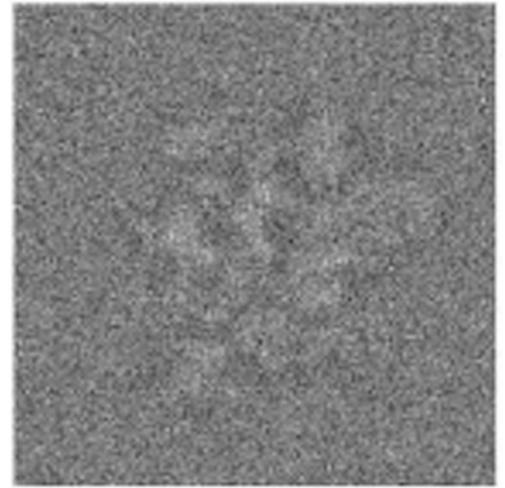
(a) Error, as a function of n_ξ , in the numerical evaluation of the integral $G(0, p_0, p_0)$ in Eq. (20). (b) Error, as a function of n_θ , in the evaluation of the integral in Eq. (16).

(a) $R = 60, c = 1/3, L = 121$ (b) $R = 90, c = 1/3, L = 181$ **Fig. 3.**

Eigenvalues of T^*T and $\Psi^*\Psi$, where T^* and Ψ^* are the truncated Fourier-Bessel transforms using numerical integration and least squares respectively. These are also the spectra of the population covariance matrices of transformed white noise images. Most eigenvalues are close to 1, indicating that the truncated Fourier-Bessel transform is almost unitary. Thus white noise remains approximately white.

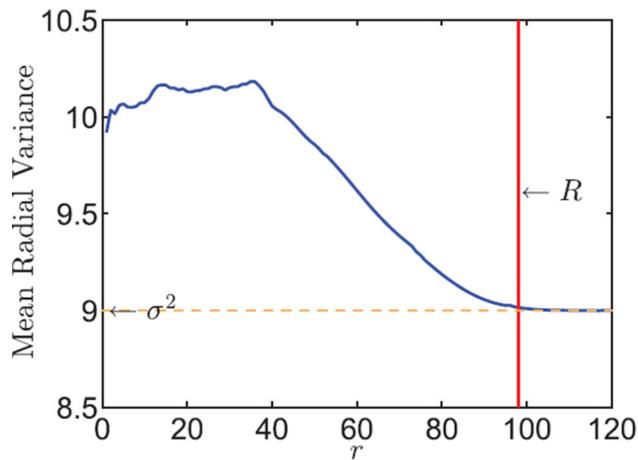
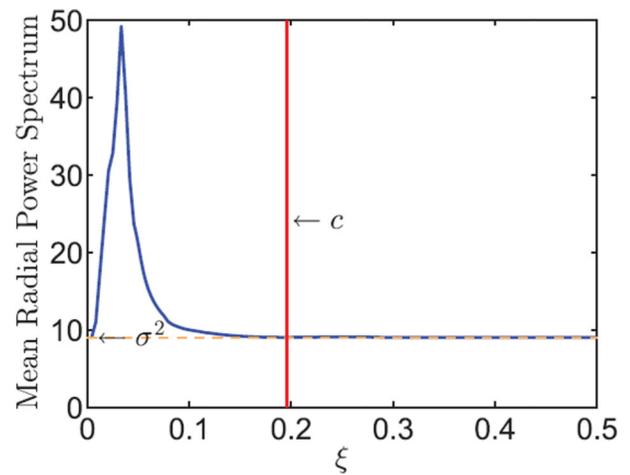


(a) Clean

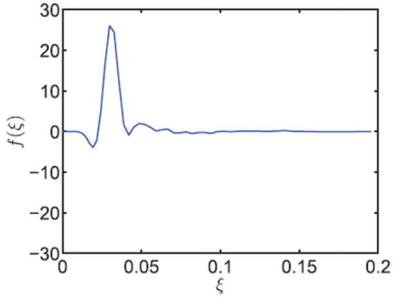


(b) SNR = 1/30

Fig. 4. Simulated projection images of the human mitochondrial large ribosomal subunit. Image size is 240×240 pixels.

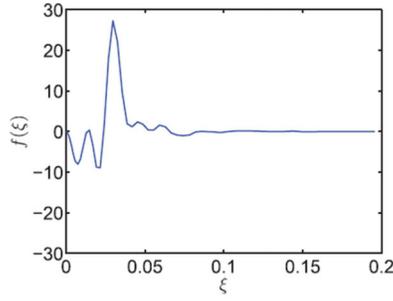
(a) Estimating R (b) Estimating c **Fig. 5.**

Estimating R and c from $n = 10^5$ simulated noisy projection images of a human mitochondrial large ribosomal subunit. Each image is of size 240×240 pixels. (a) Mean radial variance of the images. The curve levels off at about $\sigma^2 = 9$ when $r = 98$. The radius of compact support is chosen as $R = 98$. (b) Mean radial power spectrum. The curve levels off at $\sigma^2 = 9$ when $\xi = 0.196$. The band limit is chosen as $c = 0.196$.



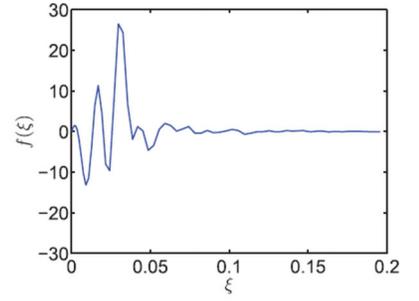
$$k = 8, l = 1$$

$$\lambda = 163.8$$



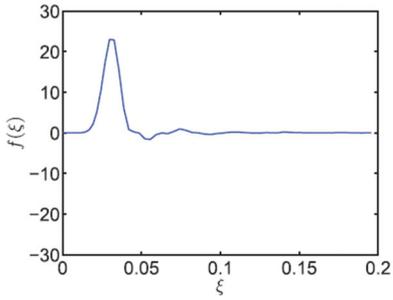
$$k = 2, l = 1$$

$$\lambda = 160.0$$



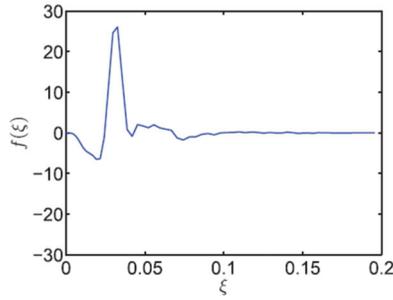
$$k = 1, l = 1$$

$$\lambda = 158.1$$



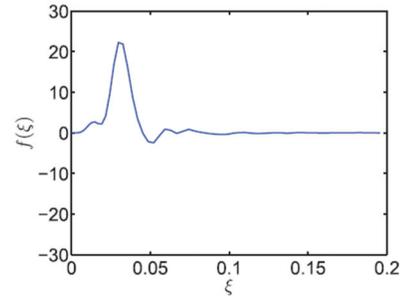
$$k = 9, l = 1$$

$$\lambda = 153.9$$



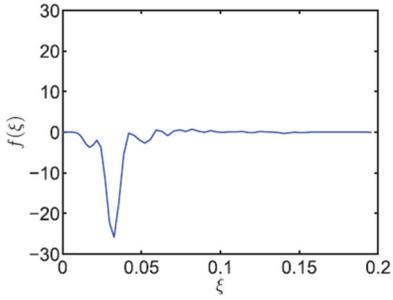
$$k = 4, l = 1$$

$$\lambda = 153.4$$



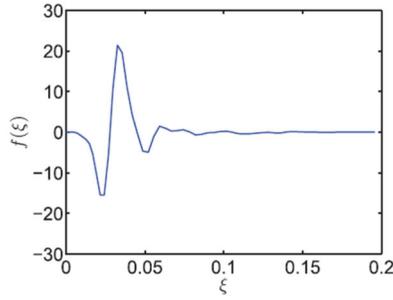
$$k = 5, l = 1$$

$$\lambda = 153.2$$



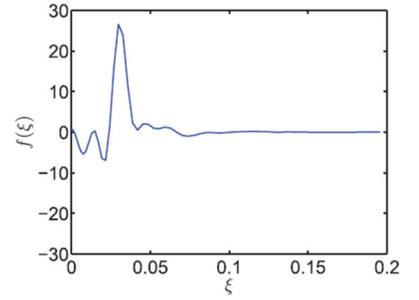
$$k = 6, l = 1$$

$$\lambda = 150.3$$



$$k = 3, l = 1$$

$$\lambda = 144.7$$



$$k = 0, l = 1$$

$$\lambda = 143.0$$

Fig. 6.

FFBsPCA principal radial functions in Fourier domain. The dataset contains $n = 10^5$ simulated human mitochondrial large ribosomal subunit projection images corrupted by additive white Gaussian noise with SNR= 1/30. Image size is 240×240 pixels, $R = 98$, $c = 0.196$. Each radial function is labeled with angular index k , radial order l , and eigenvalue λ .

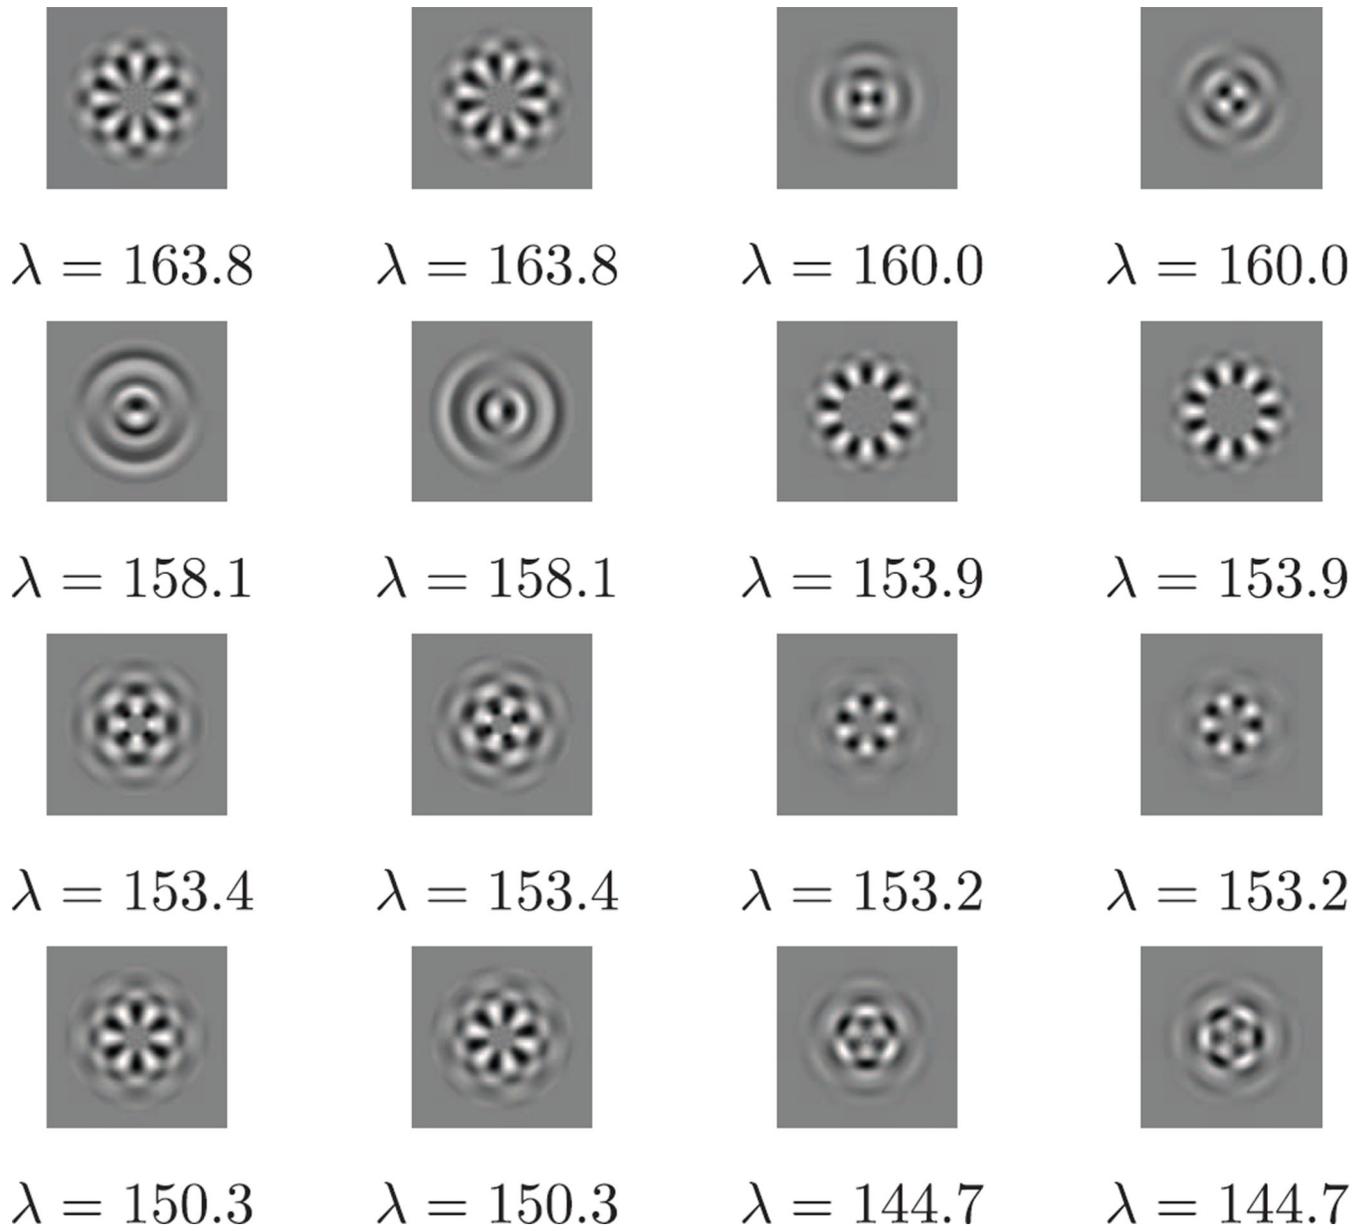


Fig. 7. FFBSPCA principal components (eigenimages in real domain) corresponding to Figure 6.

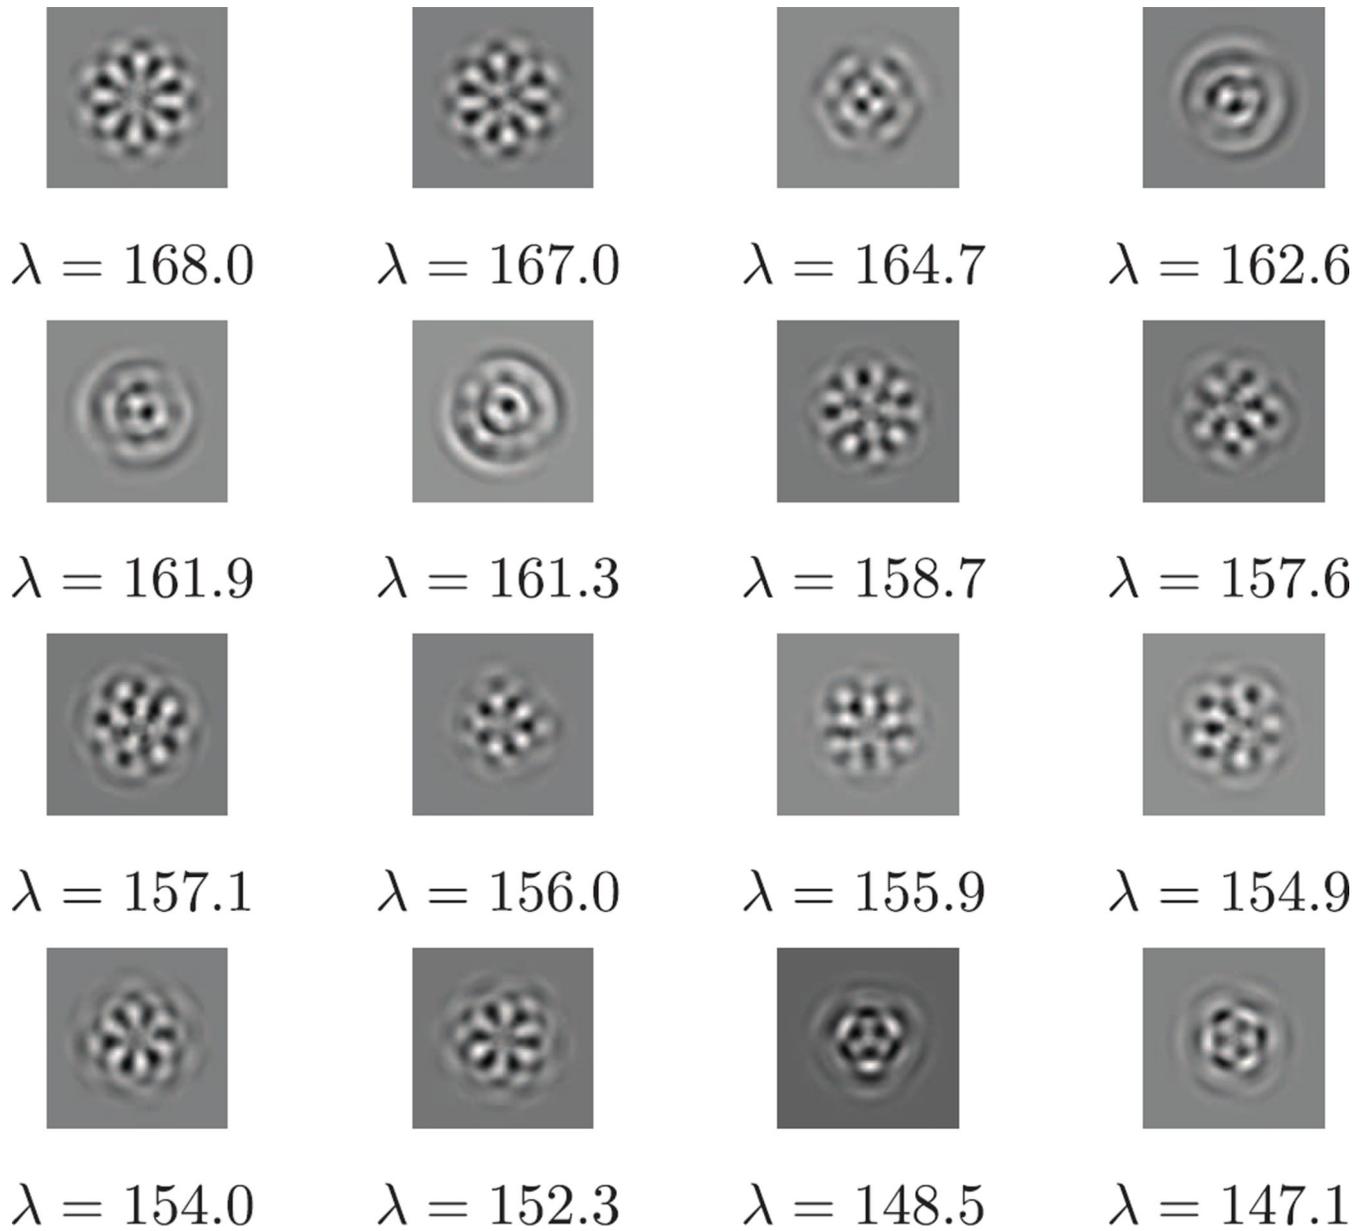


Fig. 8. Traditional PCA principal components in real image domain for the same dataset used in Figures 6 and 7.

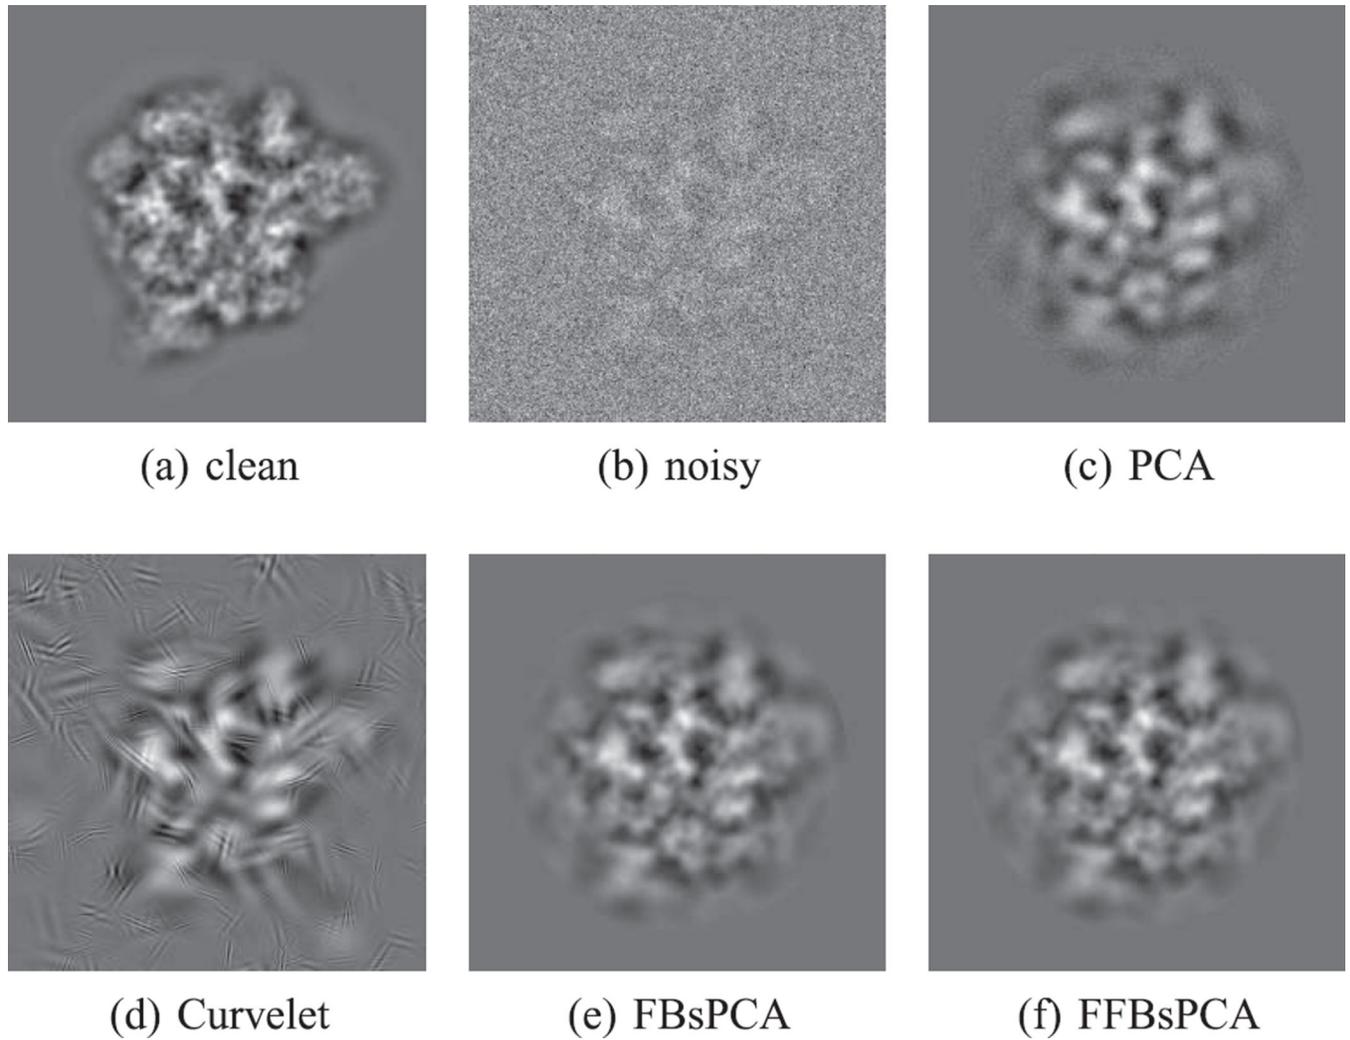


Fig. 9. Denoising simulated projection images. (a) clean projection image, (b) noisy projection image with SNR= 1/30, (c) denoised projection image using traditional PCA, (d) denoised projection image using Curvelet transform, complex block thresholding and cycle spinning, (e) denoised image using FBsPCA, and (f) denoised image using FFBSPCA.

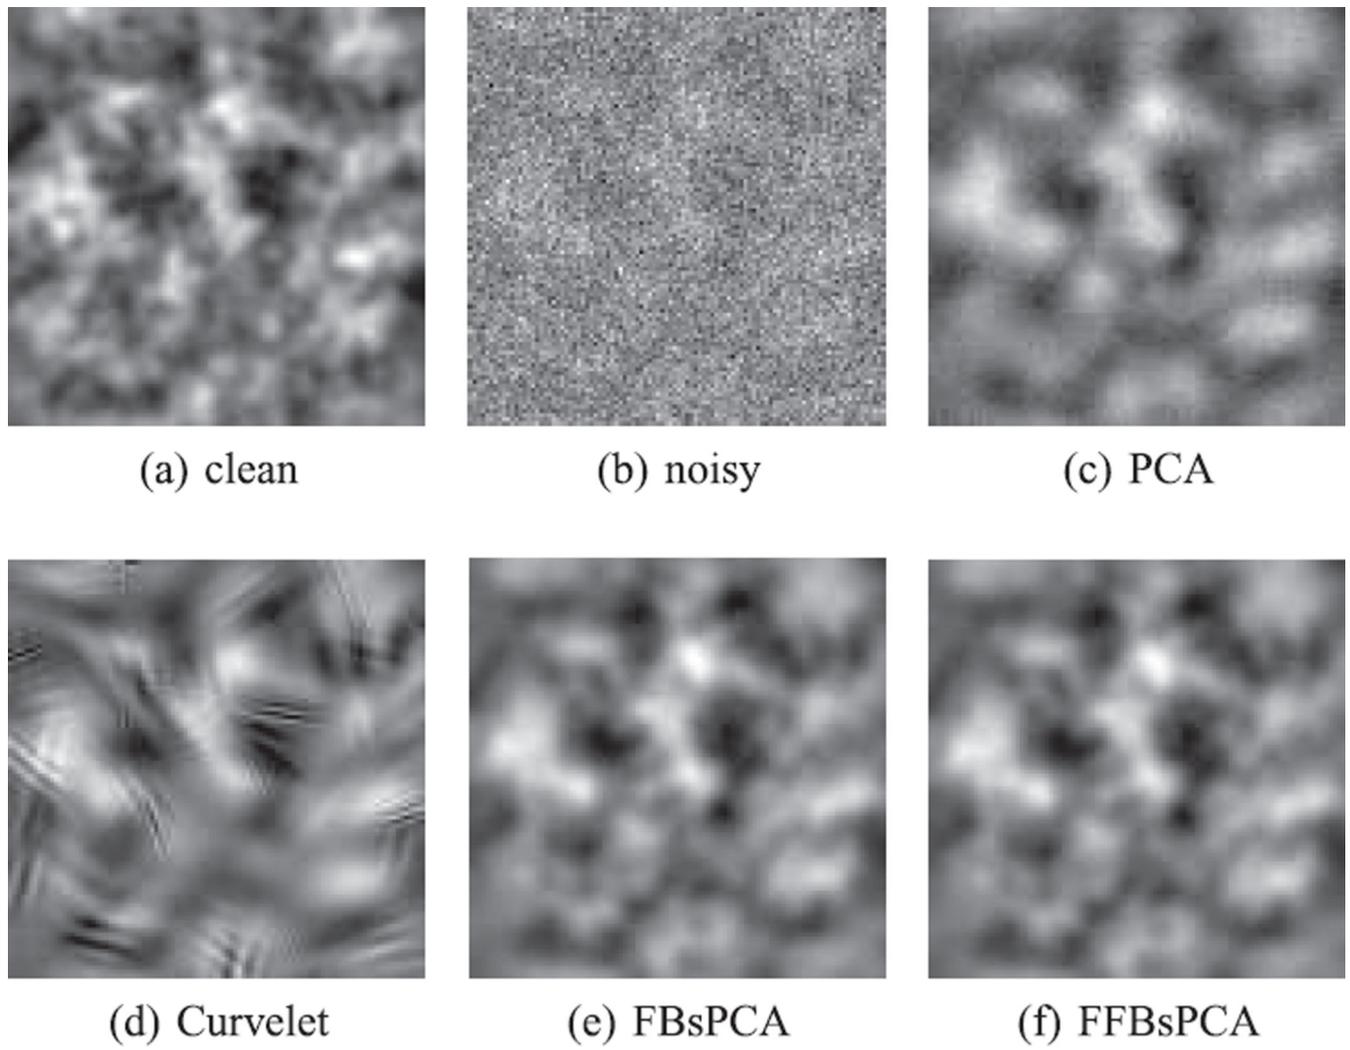


Fig. 10.
Enlarged view of 100×100 pixels box at the center of the images in Figure 9.

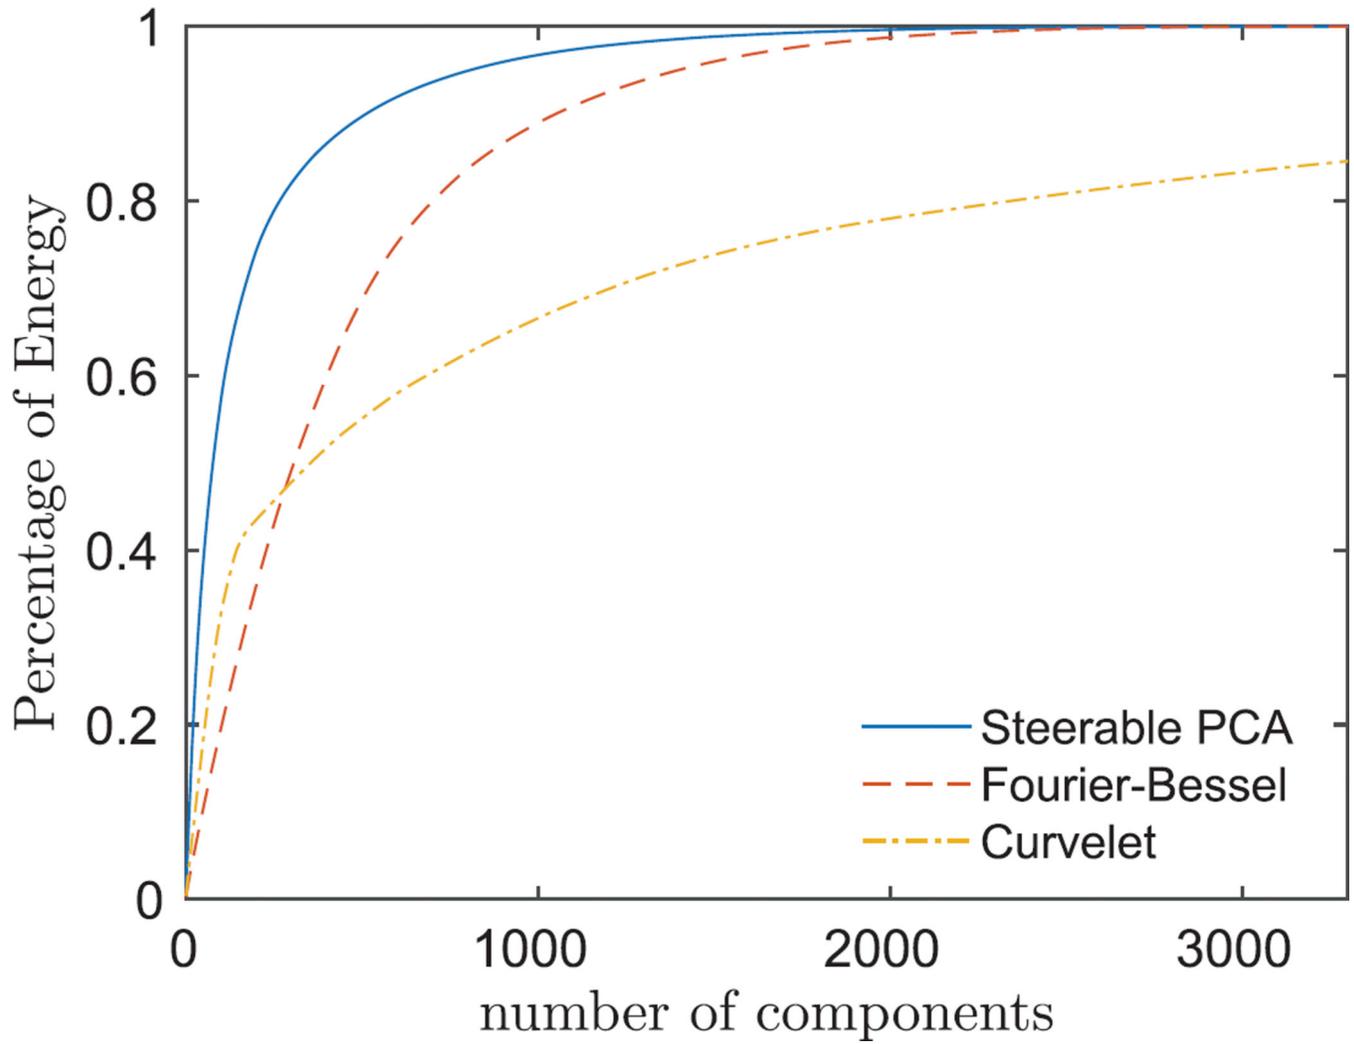


Fig. 11. Cumulative variance of FFBSPCA, Fourier-Bessel and Curvelet expansion coefficients of simulated clean ribosome projection images as in Fig. 4a.

TABLE IRunning Times (in seconds) as a Function of R for $n = 2.4 \times 10^4$, $c = 1/2$, and $L = 2R$

R	PCA	FBsPCA	FFBsPCA
30	8	7	51
60	214	50	87
90	1,636	168	148
120	1,640	413	234
150	1808	757	371
180	1,988	1,437	657
210	2,106	2,274	695
240	2,188	3,827	892

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE II

Running Times (in Minutes) as a Function of n for Image Size 300×300 pixels ($L = 300$), with $R = 150$ and $c = 1/2$

$n (\times 10^3)$	PCA	FBsPCA	FFBsPCA
1	0.05	1.2	1.1
2	0.1	2.1	1.3
4	0.3	3.5	1.8
8	1.3	4.3	2.4
16	9.8	8.7	4.6
32	59.1	17.9	8.0
64	424.7	35.7	14.4
128	653.7	74.2	30.6

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE III

Timing for FFBsPCA on a Large Dataset With $n = 10^5$ Images. Each Image is of Size 300×300 Pixels, $R = 150$ and $c = 1/2$. We Computed the Full Eigendecomposition in Algorithm 2

Steps	Time (sec)
Precomputation	7
NUFFT and Fourier-Bessel Expansion	1,438
Steerable PCA	42
Total	1487

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE IV

MSE of Denoised Images Using PCA, Curvelets, FBsPCA and FFBSPCA, all Computed Using Pixels Within $R = 98$

	MSE (10^{-5})			
	Curvelet	PCA	FBsPCA	FFBSPCA
Image 1	1.38	1.10	0.77	0.77
Image 2	1.63	1.29	0.95	0.96
Image 3	1.58	1.17	0.85	0.85

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE V

PSNR of Denoised Images Using PCA, Curvelets, FBSPCA and FFBSPCA, All Computed Using Pixels Within $R = 98$

	PSNR (dB)			
	Curvelet	PCA	FBsPCA	FFBsPCA
Image 1	18.10	19.06	20.62	20.63
Image 2	17.90	18.93	20.26	20.23
Image 3	18.68	19.99	21.35	21.35

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE VI

FFBsPCA Denoising of Images With Maximum Shifts of 0, 5, 10, 15, and 20 pixels. PSNRs Are Computed With Pixels Within $R = 110$. The Estimated Compact Support R Increases With Maximum Shift

max shifts (pixels)	R (pixels)	PSNR (dB)		
		Image 1	Image 2	Image 3
0	98	21.61	21.16	22.34
5	99	21.53	21.26	22.32
10	102	21.42	21.13	22.11
15	107	21.59	21.31	22.18
20	110	21.31	21.30	22.07

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript