



HHS Public Access

Author manuscript

SIAM J Imaging Sci. Author manuscript; available in PMC 2017 October 27.

Published in final edited form as:

SIAM J Imaging Sci. 2017 ; 10(2): 508–534. doi:10.1137/16M1085334.

Steerable Principal Components for Space-Frequency Localized Images*

Boris Landa[†] and Yoel Shkolnisky[‡]

[‡]Department of Applied Mathematics, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel

Abstract

As modern scientific image datasets typically consist of a large number of images of high resolution, devising methods for their accurate and efficient processing is a central research task. In this paper, we consider the problem of obtaining the steerable principal components of a dataset, a procedure termed “steerable PCA” (steerable principal component analysis). The output of the procedure is the set of orthonormal basis functions which best approximate the images in the dataset and all of their planar rotations. To derive such basis functions, we first expand the images in an appropriate basis, for which the steerable PCA reduces to the eigen-decomposition of a block-diagonal matrix. If we assume that the images are well localized in space and frequency, then such an appropriate basis is the prolate spheroidal wave functions (PSWFs). We derive a fast method for computing the PSWFs expansion coefficients from the images' equally spaced samples, via a specialized quadrature integration scheme, and show that the number of required quadrature nodes is similar to the number of pixels in each image. We then establish that our PSWF-based steerable PCA is both faster and more accurate than existing methods, and more importantly, provides us with rigorous error bounds on the entire procedure.

Keywords

principal component analysis; prolate spheroidal wave functions; steerable filters; group invariance; band limited functions

1. Introduction

Principal component analysis (PCA), also known as Karhunen–Loeve transform, is a ubiquitous method for dimensionality reduction which is often utilized for compression, denoising, and feature extraction from datasets. Given a dataset, the basic idea behind classical PCA is to find the best linear approximation (in the least-squares sense) to the dataset using a set of orthonormal basis functions, thus allowing for processing methods adaptive to the dataset at hand. Due to increasing improvements in image acquisition and storage techniques, we often encounter the need to process very large datasets, which consist

[†]Corresponding author. Department of Applied Mathematics, School of Mathematical Sciences, Tel Aviv University, Tel Aviv 6997801, Israel (sboris20@gmail.com).

*Received by the editors July 20, 2016; accepted for publication (in revised form) December 7, 2016; published electronically April 13, 2017.

of many thousands of images of ever-growing resolutions. In addition, there exist image acquisition techniques that introduce known deformation types into the images, thus increasing the variability in the data.

In this work, we focus on the rotation deformation, and specifically, on the setting where each image was acquired through an unknown planar rotation. Therefore, it is only natural to include all planar rotations of all images when performing the PCA procedure. It is important to mention that, when handling large datasets, the naive approach of introducing a large number of rotated versions of all images into the dataset, and then performing standard PCA, is computationally prohibitive, and, moreover, is less accurate than considering the continuum of all rotations. In the literature, one can find numerous works concerning the study of deformations (such as rotations, translations, dilations, etc.) and their connections to invariant-feature extraction, through the theory of Lie groups and compact group theory [20, 4, 21, 22, 11]. In this context, we aim to incorporate the action of the group $SO(2)$ into the framework of PCA of image datasets. We will refer to such a procedure as steerable PCA. By the theory of Lie groups, and specifically the action of the group $SO(2)$, it is known that the resulting principal components (which best approximate all rotations of a set of images) are described by tensor products of radial functions and angular Fourier modes. Such functions are referred to as “steerable” [7], since they can be rotated (or “steered”) by a simple multiplication with a complex-valued constant, and hence the term “steerable” in steerable PCA. An early approach for computing the steerable PCA was proposed in [25], which used the SVD to obtain the principal components of a single template and its set of deformations. In [9], the authors presented an efficient steerable-PCA algorithm based on angular Fourier decomposition of the images, after they have been resampled to a polar grid. While this method considers the continuum of all planar rotations, it requires a nonobvious discretization of the images in the radial direction. Efficient methods for steerable PCA were also introduced in [27] and [10]. These methods considered a finite set of equiangular rotations of each image on a polar grid, which allows for efficient circulant matrix decompositions to be carried out when computing the principal components. Recently, [38] and [37] utilized Fourier–Bessel basis functions to expand the images, followed by applying PCA in the domain of the expansion coefficients, thus accounting for all (infinitely many) rotations. We also mention [35], where the authors present an accurate algorithm to obtain the steerable principal components of templates whose analytical form is known in advance.

As digital images are typically specified by their samples on a Cartesian grid, considering their rotations implicitly assumes that they were sampled from some underlying bivariate function. Rotation of a sampled image essentially interpolates this underlying function from the available Cartesian samples. We remark that while previous works provide algorithms for steerable PCA of discretized input images, they lack in describing rigorous connections between the results of the procedure and the images prior to discretization, i.e., the underlying bivariate functions. In this work, we assume that the digital images were sampled from bivariate functions that are essentially bandlimited and are sufficiently concentrated in an area of interest in space. These assumptions guarantee that if an image was sampled with a sufficiently high sampling rate, it can be reconstructed from its samples with high precision [26]. Such assumptions are very common in various areas of engineering and physics, and as all acquisition devices are essentially bandlimited and restricted in space/

time, are expected to hold for a wide range of image datasets. Since we are interested in processing images with arbitrary orientations, it is only natural to consider a circular support area, both in space and frequency, instead of the (classical) square. We note that this model was implicitly assumed to hold in [38] and [37] for single-particle cryo-electron microscopy (cryo-EM) images.

Under the model assumptions mentioned above, our goal is to develop a fast and accurate steerable-PCA procedure which considers the continuum of all planar rotations of all images in our dataset. Particularly appealing basis functions, for expanding bandlimited functions which are also concentrated in space, are the prolate spheroidal wave functions (PSWFs) [33, 16, 17, 32, 24], defined as the strictly bandlimited set of orthogonal functions, which maximize the ratio between their \mathcal{L}^2 norm inside some finite region of interest and their \mathcal{L}^2 norm over the entire Euclidean space. Recently, [13] described an approximation scheme for functions localized to disks in space and frequency, using a series of two-dimensional PSWFs. We therefore incorporate the methods introduced in [13] into the framework of steerable PCA, providing accurate, scalable, and efficient algorithms. Our approach is in the spirit of [37], resulting in a similar block-diagonal covariance matrix. However, replacing the Fourier–Bessel with PSWFs turns out to be advantageous in terms of accuracy, available error bounds, speed, and statistical properties.

The contributions of this paper are the following. By utilizing theoretical and computational tools related to PSWFs, we are able to provide accuracy guarantees for our steerable-PCA algorithm, under the assumptions of space-frequency localization. This accuracy is in part related to a rigorous truncation rule we provide for the PSWFs series expansion, in contrast to the series truncation rules used in [38] and [37]. In addition, using a quadrature integration scheme optimized for integrating bandlimited functions on a disk [31], we present an algorithm which is in theory (and in practice) faster than [37] by a factor between 2 and 4. Finally, we also show that under some conditions on the space-frequency concentration of the images at hand, the transformation to the PSWFs expansion coefficients is nearly orthogonal.

The organization of this paper is as follows. In section 2 we introduce the PSWFs and their usage in expanding space-frequency localized images specified by their equally spaced Cartesian samples. In particular, we review the results of [13] on expanding a function into a series of PSWFs, evaluating the expansion coefficients, and bounding the overall approximation error. In section 3 we present a fast algorithm for approximating the expansion coefficients from section 2 up to an arbitrary precision. In section 4 we formalize the procedure of steerable PCA for the continuous setting (similarly to [35]), and combine it with our PSWFs-based approximation scheme. Section 5 then summarizes all relevant algorithms, and analyses in detail the computational complexities involved. In section 6 we provide some numerical results on the spectrum of the transformation to the PSWFs expansion coefficients, as this spectrum is of particular interest for noisy datasets, and in section 7 we compare our algorithm to that of [37] in terms of running time and accuracy. Finally, in section 8 we provide some concluding remarks and some possible future research directions.

2. Image approximation based on PSWFs expansion

Let $f: \mathbb{R}^2 \rightarrow \mathbb{R}$ be a square integrable function on \mathbb{R}^2 . We define a function $f(x)$ as c -bandlimited if its two-dimensional Fourier transform, denoted by $F(\omega)$, vanishes outside a disk of radius c . Specifically, if we denote $\mathbf{D} \triangleq \{x \in \mathbb{R}^2, |x| \leq 1\}$, then f is bandlimited with bandlimit c if

$$f(x) = \left(\frac{1}{2\pi}\right)^2 \int_{c\mathbf{D}} F(\omega) e^{i\omega x} d\omega, \quad x \in \mathbb{R}^2. \quad (1)$$

Among all c -bandlimited functions, the prolate spheroidal wave functions (PSWFs) on \mathbf{D} (the unit disk) are the most energy concentrated in \mathbf{D} ; that is, they satisfy

$$\frac{\|\psi(x)\|_{\mathcal{L}^2(\mathbf{D})}}{\|\psi(x)\|_{\mathcal{L}^2(\mathbb{R}^2)}} \rightarrow \max_{\psi} \quad (2)$$

while constituting an orthonormal system over $\mathcal{L}^2(\mathbf{D})$. The two-dimensional PSWFs were derived and analyzed in [32], and were shown to be the solutions to the integral equation

$$\alpha \psi(x) = \int_{\mathbf{D}} \psi(\omega) e^{i c \omega x} d\omega, \quad x \in \mathbf{D}. \quad (3)$$

We denote the PSWFs with bandlimit c as $\psi_{N,n}^c(x)$, and their corresponding eigenvalues as $\alpha_{N,n}^c$, which together form the eigenfunctions and eigenvalues of (3), with $N \in \mathbb{Z}$ and $n \in \mathbb{N} \cup \{0\}$. In addition, it turns out that the functions $\psi_{N,n}^c(x)$ are orthogonal on both \mathbf{D} and \mathbb{R}^2 using the standard \mathcal{L}^2 inner products on \mathbf{D} and \mathbb{R}^2 , respectively, and are dense in both the class of $\mathcal{L}^2(\mathbf{D})$ functions and in the class of c -bandlimited functions on \mathbb{R}^2 . In polar coordinates, the functions $\psi_{N,n}^c(x)$ have a separation of variables and can be written in complex form as

$$\psi_{N,n}^c(r, \theta) = \frac{1}{\sqrt{2\pi}} R_{N,n}^c(r) e^{iN\theta}, \quad N \in \mathbb{Z}, n \in \mathbb{N} \cup \{0\}, \quad (4)$$

where $R_{N,n}(r)$ (defined explicitly in (66) in Appendix C) satisfies $R_{N,n}(r) = R_{|N|,n}(r)$, and the eigenfunctions $\psi_{N,n}^c(x)$ are normalized to have an $\mathcal{L}^2(\mathbf{D})$ norm of 1. The indices N and n are often referred to as the angular index and the radial index respectively. Equation (4) also tells us that the PSWFs are steerable [7]; that is, rotating $\psi_{N,n}^c(r, \theta)$ is equivalent to multiplying it by a complex constant. This property is important for handling datasets which include

rotations and in particular for the steerable-PCA procedure. A detailed numerical evaluation procedure for the two-dimensional PSWFs can be found in [31] and an illustration of the PSWFs for the several first index pairs (N, n) can be seen in Figure 1.

Since our images are assumed to be essentially bandlimited and sufficiently concentrated in a disk, the properties of the PSWFs mentioned above (and especially their optimal concentration property) make them suitable basis functions for expanding such images, as we consider next. Let us define a function $f(x)$ as (ν, μ) -concentrated if its \mathcal{L}^2 norm outside a disk of radius ν is upper bounded by μ , that is

$$\sqrt{\int_{x \notin \nu \mathbf{D}} |f(x)|^2 dx} \leq \mu. \quad (5)$$

Using this definition, the class of c -bandlimited functions is a subclass of (c, δ_c) -concentrated functions in the Fourier domain (for any $\delta_c > 0$). We consider the two-dimensional functions from which our images were sampled to be $(1, \varepsilon)$ -concentrated in space, with their Fourier transforms being (c, δ_c) -concentrated. This assumption always holds for some set of parameters c, δ_c, ε , and the general notion is that δ_c and ε are “small.” Since our images are given in their sampled form, we define the unit square $Q \triangleq [-1, 1] \times$

$[-1, 1]$, and assume that we are given the samples $\{f(x_k) : x_k = \frac{k}{L} \in Q\}$, where k is a two-dimensional integer index. These samples correspond to a Cartesian grid of $(2L + 1) \times (2L + 1)$ equally spaced samples, with sampling frequency of L in each dimension. We mention that in this setup, the Nyquist frequency corresponds to a bandlimit c of at most πL .

Given an image $I(x) \in \mathcal{L}^2(\mathbf{D})$, we can expand it as

$$I(x) = \sum_{N=-\infty}^{\infty} \sum_{n=0}^{\infty} a_{N,n} \psi_{N,n}^c(x), \quad x \in \mathbf{D}, \quad a_{N,n} = \int_{\mathbf{D}} I(x) (\psi_{N,n}^c(x))^* dx, \quad (6)$$

where $(\cdot)^*$ denotes complex conjugation. As numerical algorithms cannot use infinite expansions, the image $I(x)$ needs to be approximated by a finite sequence of PSWFs, while its expansion coefficients $\{a_{N,n}\}$ are approximated using only the available samples

$\{I(\frac{k}{L})\}_{\frac{k}{L} \in Q}$. To this end, we follow [13], and define the truncated PSWFs expansion that approximates $I(x)$ as

$$\hat{I}(x) \triangleq \sum_{N,n \in \Omega_T} \hat{a}_{N,n} \psi_{N,n}^c(x), \quad (7)$$

where $\hat{a}_{N,n}$ are the approximated coefficients (to be described shortly), and Ω_T is a finite set of indices that is determined by a truncation parameter T , and is defined by

$$\Omega_T \triangleq \left\{ (N, n): \sqrt{\frac{|\lambda_{N,n}^c|^2}{1 - |\lambda_{N,n}^c|^2}} > T \right\}, \quad T > 0, \quad \lambda_{N,n}^c \triangleq \frac{c}{2\pi} \alpha_{N,n}^c, \quad (8)$$

where $\alpha_{N,n}^c$ is the eigenvalue from (3) corresponding to the eigenfunction $\psi_{N,n}^c$. The properties of the index set Ω_T are determined by the “normalized” eigenvalues $\lambda_{N,n}^c$, whose behavior is exemplified in Figure 2. The coefficients $\hat{a}_{N,n}$ of (7) are defined by

$$\hat{a}_{N,n} \triangleq |\lambda_{N,n}^c|^2 \frac{\langle I, \psi_{N,n}^c \rangle}{L^2} = \frac{|\lambda_{N,n}^c|^2}{L^2} \sum_{\frac{k}{L} \in \mathbf{D}} I\left(\frac{k}{L}\right) \left(\psi_{N,n}^c\left(\frac{k}{L}\right)\right)^*, \quad (9)$$

where both I and $\psi_{N,n}^c$ stand for the appropriate vectors of samples of $I(x)$ and $\psi_{N,n}^c(x)$ inside the unit disk, respectively. We note that for real-valued images we have that $a_{N,n} = (a_{-N,n})^*$, and thus it is sufficient to compute the coefficients $\hat{a}_{N,n}$ only for $N \geq 0$.

Using our definitions in (7), (8), (9), and assuming that $I(x)$ is $(1, \varepsilon)$ -concentrated in space and (c, δ_c) -concentrated in Fourier domain (as defined earlier), we show in Appendix A that

$$\|I(x) - \sum_{N,n \in \Omega_T} \hat{a}_{N,n} \psi_{N,n}^c(x)\|_{\mathcal{L}^2(\mathbf{D})} \leq \left(\varepsilon + \frac{\delta_c}{2\pi}\right) (T+4) \triangleq \mathcal{E}(\varepsilon, \delta_c, T). \quad (10)$$

It is evident that the bound $\mathcal{E}(\varepsilon, \delta_c, T)$ in (10) depends only on the images' concentration properties and on the truncation parameter T , all of which can be determined a priori.

Lastly, it is also mentioned in [13] that the cardinality of the index set Ω_T of (8) is expected to be

$$|\Omega_T| = \frac{c^2}{4} - \frac{2}{\pi^2} c \log(c) \log(T) + o(c \log(c)). \quad (11)$$

The first term on the right hand-side of (11), which depends quadratically on c , is in fact (up to a constant of $1/\pi$) what is known as the Landau rate [14] for stable sampling and reconstruction, and it is the term which dominates asymptotically the required number of basis functions for the approximation (see also Figure 2). A table listing the values of $|\Omega_T|$ for various L and T can be found in [13]. We remark that for $c = \pi L$ (Nyquist sampling) and $T = 1$, which essentially means that the approximation error is of the order of the space-

frequency localization, we have that $|\Omega_T|$ is about $\pi^2 L^2/4$, which is smaller than the number of samples (pixels) in the unit disk.

3. Fast PSWFs coefficients approximation

Computing the expansion coefficients $\hat{a}_{N,n}$ of (9) directly results in a computational complexity of $\mathcal{O}(L^4)$ operations. This is due to the fact that each image contains $\mathcal{O}(L^2)$ samples, while there are about $\mathcal{O}(L^2)$ different basis functions in the expansion. Since we aim to process large datasets of high resolution images, in what follows we describe an asymptotically more efficient method for evaluating the expansion coefficients in $\mathcal{O}(L^3)$ operations.

Using (3), we can rewrite the PSWFs expansion coefficients (9) as

$$\hat{a}_{N,n} = \alpha_{N,n}^c \left(\frac{c}{2\pi L} \right)^2 \sum_{\frac{k}{L} \in \mathbf{D}} I \left(\frac{k}{L} \right) \left[\int_{\mathbf{D}} \psi_{N,n}^c \left(\frac{k}{L} \right) e^{ic\omega \frac{k}{L}} \right]^* \quad (12)$$

$$= \alpha_{N,n}^c \left(\frac{c}{2\pi L} \right)^2 \int_{\mathbf{D}} [\psi_{N,n}^c(\omega)]^* \phi^c(\omega) d\omega, \quad (13)$$

where

$$\phi^c(u) \triangleq \sum_{\frac{k}{L} \in \mathbf{D}} I \left(\frac{k}{L} \right) e^{-icu \frac{k}{L}}, \quad u \in \mathbb{R}^2. \quad (14)$$

Since both $\psi_{N,n}^c$ and ϕ^c are c -bandlimited functions, the product $\psi_{N,n}^c \cdot \phi^c$ defines a function which is $2c$ -bandlimited. Therefore, we can compute the integral in (13) using a quadrature formula for $2c$ -bandlimited functions on a disk. Such an integration scheme was proposed in [31] using specialized quadrature nodes and weights. Thus, it follows that, by using such a quadrature formula, we can approximate the expansion coefficients of (13) by

$$\tilde{a}_{N,n} \triangleq \alpha_{N,n}^c \left(\frac{c}{2\pi L} \right)^2 \sum_{\ell=1}^{\mathcal{N}_r} \sum_{j=1}^{\mathcal{N}_\theta^\ell} \mathcal{W}_{\ell,j}^{2c} [\psi_{N,n}^c(\omega_{\ell,j}^{2c})]^* \phi^c(\omega_{\ell,j}^{2c}), \quad (15)$$

where $\omega_{\ell,j}^{2c}$ and $\mathcal{W}_{\ell,j}^{2c}$ are the quadrature nodes and weights respectively, \mathcal{N}_r is the number of different radii of the quadrature nodes, and \mathcal{N}_θ^ℓ is the number of quadrature nodes per radius (which may vary for each radius). We use for $\omega_{\ell,j}^{2c}$ and $\mathcal{W}_{\ell,j}^{2c}$ the quadrature nodes and weights

of [31] corresponding to bandlimit $2c$. These quadrature nodes are equally spaced in the angular direction and nonequally spaced in the radial direction. In polar coordinates, the nodes and weights are given by

$$\omega_{\ell,j}^{2c} = (r_{\omega}^{\ell}, \theta_{\omega}^{\ell,j}), \quad \theta_{\omega}^{\ell,j} = \frac{2\pi j}{\mathcal{N}_{\theta}^{\ell}}, \quad \mathcal{W}_{\ell,j}^{2c} = \frac{2\pi}{\mathcal{N}_{\theta}^{\ell}} r_{\omega}^{\ell} \tilde{\mathcal{W}}_{\ell}^{2c}, \quad (16)$$

and the derivation of $\tilde{\mathcal{W}}_{\ell}^{2c}$ is detailed in [31]. A typical array of quadrature nodes is plotted in Figure 3.

By employing (15), (16), and the analytical expression of the PSWFs (4), we finally arrive at

$$\tilde{a}_{N,n} = \sqrt{2\pi} \alpha_{N,n}^c \left(\frac{c}{2\pi L} \right)^{2\mathcal{N}_r} \sum_{\ell=1}^{2\mathcal{N}_r} \tilde{\mathcal{W}}_{\ell}^{2c} \frac{R_{N,n}^c(r_{\omega}^{\ell}) r_{\omega}^{\ell}}{\mathcal{N}_{\theta}^{\ell}} \sum_{j=1}^{\mathcal{N}_{\theta}^{\ell}} \phi^c(\omega_{\ell,j}^{2c}) e^{-\frac{2\pi i j N}{\mathcal{N}_{\theta}^{\ell}}}, \quad (17)$$

which is particularly appealing for efficient numerical evaluation, as we explain later in this section.

Clearly, replacing the integral in (13) with a quadrature formula results in an approximation error. In this respect, it is desirable to choose \mathcal{N}_r and $\mathcal{N}_{\theta}^{\ell}$ such that this error is smaller than some prescribed accuracy ϑ_q (usually chosen as the machine precision). To this end, and according to [31], it is sufficient to satisfy the condition

$$\sum_{|j| > \mathcal{N}_{\theta}^{\ell}} J_j(2cr_{\omega}^{\ell}\rho) \leq C_1 \vartheta_q, \quad (18)$$

for all $1 \leq \ell \leq \mathcal{N}_r$ and $\rho \in [0, 1]$, where J_j is the Bessel function of the first kind and order j , as well as the condition

$$\sum_{k=2\mathcal{N}_r+1}^{\infty} \frac{|\lambda_{0,k}^{2c}|}{2c} \|R_{0,k}^{2c}(r)\|_{\infty} \|R_{0,k}^{2c}(r) \sqrt{r}\|_{\infty} \left(1 + \sum_{j=1}^{\mathcal{N}_r} |\tilde{W}_j^{2c}| \right) \leq C_2 \vartheta_q, \quad (19)$$

where $\lambda_{N,n}^c$ was defined in (8), $\|\cdot\|_{\infty}$ stands for the max-norm in $[0, 1]$, and both C_1 and C_2 are constants which depend on the bandlimit c . We mention that $|\lambda_{0,k}^{2c}|$ are ordered in a nonincreasing order with respect to the index k . In order to determine appropriate values for \mathcal{N}_r and $\mathcal{N}_{\theta}^{\ell}$, one can proceed to solve the inequalities in conditions (18) and (19) numerically by directly evaluating these expressions. However, in order to analyze the

computational complexity of the procedure described in this section, and to offer the reader simpler analytic expressions and some insight concerning conditions (18) and (19), we provide some further analysis of these conditions and the resulting number of quadrature nodes. In Appendix B, we prove that, in order to satisfy condition (18) for every $1 \leq \ell \leq \mathcal{N}_r$, it is sufficient to choose the integers \mathcal{N}_θ^ℓ as

$$\mathcal{N}_\theta^\ell = \lceil cr_\omega^\ell e + \log v_q^{-1} + \log \frac{2}{C_1} + 1 \rceil, \quad (20)$$

where r_ω^ℓ are the radial quadrature nodes from (16), e is the base of the natural logarithm, and $\lceil \cdot \rceil$ is the rounding up operation. Then, it is evident that the term $cr_\omega^\ell e$ dominates condition (20), i.e., $\mathcal{N}_\theta^\ell \sim cr_\omega^\ell e$, since all other factors (prescribed error v_q and the constant C_1) affect it only logarithmically. Therefore, we expect the overall number of quadrature nodes to be

$$\sum_{\ell=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \sim ce \sum_{\ell=1}^{\mathcal{N}_r} r_\omega^\ell \sim \frac{ce\mathcal{N}_r}{2}, \quad (21)$$

assuming that the quadrature points r_ω^ℓ are approximately symmetric about 0.5. We remark that numerical experiments reveal that there is a small difference between conditions (18) and (20), and specifically, choosing \mathcal{N}_θ^ℓ directly via the numerical evaluation of condition (18) results in about 20% fewer angular nodes compared to choosing \mathcal{N}_θ^ℓ according to (20).

With regard to condition (19), although it may seem somewhat daunting, the sum on the left-hand side is dominated by the values of $|\lambda_{0,k}^{2c}|$ and their decay properties. We argue in Appendix C that the number of non-negligible (relative to machine precision) values of $|\lambda_{0,k}^{2c}|$ is only about $2c/\pi$, that is

$$|\{k: |\lambda_{0,k}^{2c}| > \varepsilon\}| \sim \frac{2c}{\pi}, \quad (22)$$

for some small ε and a sufficiently large c . This observation stems from the fact that the values of $|\lambda_{0,k}^{2c}|$ become arbitrarily small (and decay at a superexponential rate) once k reaches $2c/\pi + O(\log c)$. We provide numerical evidence for this claim in Figure 4.

Therefore, it is evident that condition (19) can be satisfied if $|\lambda_{0,k}^{2c}|$ is sufficiently small for $k > 2\mathcal{N}_r$, which together with (22) implies that for $c = \pi L$ (Nyquist sampling)

$$\mathcal{N}_r \sim \frac{c}{\pi} = L. \quad (23)$$

In total, by (21) and (23), the number of quadrature nodes for $c = \pi L$ is

$$\sum_{\ell=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \sim \frac{e}{2\pi} c^2 = \frac{\pi e}{2} L^2 \approx 4.27L^2, \quad (24)$$

which is only slightly larger than the number of sampling points in the unit square.

Let us now turn our attention to the computational complexity of computing $\tilde{a}_{N,n}$ of (17).

The first step is to compute $\phi^c(\omega_{\ell,j}^{2c})$ from (14) for all quadrature nodes $\{\omega_{\ell,j}^{2c}\}$, which can be implemented efficiently using the nonequispaced fast fourier transform (NFFT) [3, 8, 28, 5]. The computational complexity of such algorithms is

$$O(L^2 \log L + \log \varepsilon^{-2} P), \quad (25)$$

where L is the sampling rate, ε is the required accuracy of the transform, and P is the number of evaluation points, which is equal to the number of quadrature nodes (24). Next, we notice that, for every value of ℓ the inner sum in (17)

$$C_\ell^N \triangleq \sum_{j=1}^{N_\theta^\ell} \phi^c(\omega_{\ell,j}^{2c}) e^{-\frac{2\pi i j N}{\mathcal{N}_\theta^\ell}} \quad (26)$$

can be computed for multiple values of N using the FFT in $O(\mathcal{N}_\theta^\ell \log \mathcal{N}_\theta^\ell)$ operations,

resulting in total of $O(\sum_{\ell=1}^{\mathcal{N}_r} \mathcal{N}_\theta^\ell \log \mathcal{N}_\theta^\ell)$ operations for all required values of ℓ . By (24), this results in a computational complexity of $O(L^2 \log L)$ for $c = \pi L$. Lastly, the approximated expansion coefficients (17) are given by the remaining outer sum

$$\tilde{a}_{N,n} = \sqrt{2\pi} \alpha_{N,n}^c \left(\frac{c}{2\pi L} \right)^{2\mathcal{N}_r} \sum_{\ell=1}^{\mathcal{N}_r} \tilde{\mathcal{W}}_\ell^{2c} \frac{R_{N,n}^c(r_\omega^\ell)}{\mathcal{N}_\theta^\ell} C_\ell^N, \quad (27)$$

which can be computed for all indices $\{N, n\} \in \Omega_T$ in $\mathcal{N}_r |\Omega_T|$ operations, where $|\Omega_T|$ denotes the cardinality of the index set Ω_T . From (11) and (23), the computational complexity of the last step is essentially $O(L^3)$, which thus governs the computational complexity of the entire procedure described in this section. We end this discussion with the

observation that the complexity of evaluating (27) can be further reduced by exploiting the rapid decay of the functions $R_{N,n}^c(r)$ with the radius r (see Figure 1b for an illustration), due to which a substantial number of the terms involving $R_{N,n}^\ell(r_\omega^\ell)$ can be discarded for certain sets of the radii $\{r_\omega^\ell\}$ and indices $\{(N, n)\}$. To exemplify this point, for $T = 1$, $L = 150$, and $c = \pi L$, we have that about 30% of the values $\{|R_{N,n}^c(r_\omega^\ell)|\}$ are below 10^{-12} , and therefore can be safely discarded from (27).

4. Steerable-PCA procedure

As discussed in the introduction, steerable PCA extends the classical PCA by artificially including in the analyzed dataset all (infinitely many) planar rotations of each image. In the previous sections, we have established a framework for expanding images using PSWFs, under the assumption that the images to be approximated are sufficiently concentrated in space and frequency. Since PSWFs constitute an excellent basis for expanding images which are localized in space and frequency (see (50), (10) and [13]), we use them in this section to construct optimized basis functions for a given set of images and their rotations.

Let us suppose that our dataset consists of M (sampled) images, where the m th image is

given by the samples $I_m\left(\frac{k}{L}\right)$ of some function $I_m(x) \in \mathcal{L}^2(\mathbf{D})$. We denote by $I_m^\varphi(x)$ the planar rotation of $I_m(x)$ by an angle $\varphi \in [0, 2\pi)$. Now, ideally, we would like to obtain basis functions that best approximate (in the sense of $\mathcal{L}^2(\mathbf{D})$) the images $I_m^\varphi(x)$ for all rotation angles $\varphi \in [0, 2\pi)$. However, we do not have access to the underlying images $I_m(x)$, nor their rotations $I_m^\varphi(x)$. Therefore, we first replace them with their PSWFs-based approximations, denoted by $\hat{I}_m(x)$ (see (7)) and $\hat{I}_m^\varphi(x)$. Then, we derive a steerable-PCA procedure for the set of approximated images, and finally we make the connection between the resulting steerable principal components of the approximated images and the original images.

The k th principal component $g_k(x) \in \mathcal{L}^2(\mathbf{D})$ for the set of approximated images $\{\hat{I}_m^\varphi\}$, $m = 0, \dots, M-1$, is defined as

$$g_k(x) = \arg \max_{g(x)} \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \left| \left\langle \hat{I}_m^\varphi(x) - \hat{\mu}(x), g(x) \right\rangle_{\mathcal{L}^2(\mathbf{D})} \right|^2 d\varphi, \quad (28)$$

$$\text{s.t. } \|g_k(x)\|_{\mathcal{L}^2(\mathbf{D})} = 1, \langle g_k(x), g_i(x) \rangle_{\mathcal{L}^2(\mathbf{D})} = 0, \quad 0 \leq i < k,$$

where $\langle \cdot, \cdot \rangle_{\mathcal{L}^2(\mathbf{D})}$ denotes the standard inner product on $\mathcal{L}^2(\mathbf{D})$, and $\hat{\mu}(x)$ is the average of all approximated images and their planar rotations, given by

$$\hat{\mu}(x) = \frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \hat{I}_m^\varphi(x) d\varphi. \quad (29)$$

In other words, the function $g_k(x)$ is expected to maximize the average projection norm (defined over all images and their rotations), such that $\{g_k(x)\}$ forms a set of orthonormal functions over $\mathcal{L}^2(\mathbf{D})$. The functions $\{g_k(x)\}$ are named the steerable principal components of our dataset. The formulation (28) differs from the classical formulation of PCA in the additional integration over the rotation angle φ , which has the interpretation of including all (infinitely many) rotations of all images in our data set.

After substituting (7) in (28) and exploiting the orthonormality and steerable structure of the PSWFs, the optimization problem (28) is reduced to a problem in the domain of the PSWFs expansion coefficients $\hat{a}_{N,n}$ for which the derivation in [38] reveals that the solution is given by the eigenvectors of the $|\Omega_{\mathcal{T}}| \times |\Omega_{\mathcal{T}}|$ Hermitian positive semidefinite matrix, whose entries are

$$C_{(N,n),(N',n')} = \begin{cases} \frac{1}{M} \sum_{m=0}^{M-1} \hat{c}_{N,n}^m \left(\hat{c}_{N',n'}^m \right)^*, & N=N', \\ 0, & N \neq N', \end{cases} \quad (30)$$

where

$$\hat{c}_{N,n}^m = \begin{cases} \hat{a}_{N,n}^m - \frac{1}{M} \sum_{m'=0}^{M-1} \hat{a}_{0,n}^{m'}, & N=0, \\ \hat{a}_{N,n}^m, & N \neq 0, \end{cases} \quad (31)$$

and $\hat{a}_{N,n}^m$, given by (9) (or (17)), are the approximated expansion coefficients for the m th image in our dataset, i.e., the coefficients of \hat{I}_m . We refer to C of (30) as our rotationally invariant covariance matrix (replacing the standard covariance matrix used in classical PCA). Since the matrix C enjoys a block-diagonal structure, we obtain its eigen-decomposition through the eigen-decomposition of its blocks. Let $\hat{\lambda}_1 \hat{\lambda}_2 \dots \hat{\lambda}_{|\Omega_{\mathcal{T}}|}$ be the eigenvalues of the matrix C and let \hat{g}_k be the eigenvector corresponding to eigenvalue $\hat{\lambda}_k$. For each pair $(\hat{\lambda}_k, \hat{g}_k)$, the entries of \hat{g}_k , denoted by $\hat{g}_{N,n}^k$, are nonzero only on some block of the matrix C corresponding to an angular index N_k and it can be shown that the functions $g_k(x)$ of (28) are recovered as

$$g_k(x) = \sum_{n=0}^{n_k} \hat{g}_{N_k,n}^k \psi_{N_k,n}(x), \quad (32)$$

where n_k stands for the largest index n such that $(N_k, n) \in \Omega_T$. Equation (32), together with (4), confirms that each $g_k(x)$ is a steerable function (as in [7]).

By the formulation of (28), it follows that the functions $g_k(x)$ form the optimal basis (in $\mathcal{L}^2(\mathbf{D})$) for expanding the approximated images $\{\hat{I}_m(x)\}$ and their rotations, such that if we define

$$\tilde{I}_{K,m}^\varphi(x) \triangleq \hat{\mu}(x) + \sum_{k=1}^K c_{k,m}^\varphi g_k(x), \quad c_{k,m}^\varphi \triangleq e^{-iN_k\varphi} \sum_{n=0}^{n_k} \hat{c}_{N_k,n}^m \left(\hat{g}_{N_k,n}^k \right)^*, \quad (33)$$

then we have that

$$\frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \|\hat{I}_m^\varphi(x) - \tilde{I}_{K,m}^\varphi(x)\|_{\mathcal{L}^2(\mathbf{D})}^2 d\varphi \leq \sum_{k=K+1}^{|\Omega|} \hat{\lambda}_k, \quad (34)$$

where $\hat{\lambda}_k$ is the k th eigenvalue of C .

Since the steerable principal components were computed for the set of approximated images $\{\hat{I}_m(x)\}$ and not for the set of underlying images $\{I_m(x)\}$, the bound in (34) holds only for the set of the approximated images. Nonetheless, the PSWFs approximation scheme from section 2 provides us with strict error bounds when expanding images localized in space and frequency. Therefore, when using (33) to approximate our images $I_m(x)$, the error norm (averaged over all images and rotations) can be bounded by joining (34) and (10). In particular, for a truncation parameter T , we have that

Algorithm 1. Evaluating PSWFs expansion coefficients (direct method).

- 1 **Required:** An image $\{I(\frac{k}{L})\}$ sampled on a Cartesian grid of size $(2L+1) \times (2L+1)$ with $\frac{k}{L} \in \mathbf{Q}$, and $\mathbf{Q} = [-1, 1] \times [-1, 1]$.
 - 2 **Precomputation:**
 1. Choose a bandlimit $c(\pi L)$ and a truncation parameter T .
 2. Evaluate the PSWFs $\psi_{N,n}^c(\frac{k}{L})$ and their eigenvalues $\alpha_{N,n}^c$ according to [31] for $(N, n) \in \Omega_T$ (see (8)), where $\frac{k}{L} \in \mathbf{D}$, and compute the normalized eigenvalues $\lambda_{N,n}^c = \frac{c}{2\pi} \alpha_{N,n}^c$.
 - 3 For all $(N, n) \in \Omega_T$, compute the expansion coefficients

$$\hat{a}_{N,n} = \frac{|\lambda_{N,n}^c|^2}{L^2} \sum_{\frac{k}{L} \in \mathbf{D}} I(\frac{k}{L}) (\psi_{N,n}^c(\frac{k}{L}))^* .$$
-

$$\frac{1}{M} \sum_{m=0}^{M-1} \frac{1}{2\pi} \int_0^{2\pi} \|I_m^\varphi(x) - \tilde{I}_{K,m}^\varphi(x)\|_{\mathcal{L}^2(\mathbb{D})}^2 d\varphi \leq \sum_{k=K+1}^{|\Omega_T|} \hat{\lambda}_{k+2} \mathcal{E}(\varepsilon, \delta_c, T) \sqrt{\sum_{k=K+1}^{|\Omega_T|} \hat{\lambda}_k + \mathcal{E}^2(\varepsilon, \delta_c, T)},$$

(35)

where $\tilde{I}_{K,m}^\varphi(x)$ is the expansion via the steerable principal components (33), and $\mathcal{E}(\varepsilon, \delta_c, T)$ is the approximation error term of (10) for the truncated series of PSWFs. In essence, (35) asserts that if the images are sufficiently localized in space and frequency, then for an appropriate truncation parameter T the error in expanding $I_m(x)$ using the steerable principal components computed from $\hat{I}_m(x)$ is close to the smallest possible error given by (34).

5. Algorithm summary and computational cost

We summarize the algorithms for evaluating the expansion coefficients $\hat{a}_{N,n}$ and $\tilde{a}_{N,n}$ (corresponding to the direct method from section 2 and the efficient method of section 3, respectively) in Algorithms 1 and 2 respectively. The steerable-PCA procedure described in section 4 is summarized in Algorithm 3.

We now turn our attention to the computational complexity of Algorithms 1, 2, and 3. We omit the precomputation steps from the complexity analysis as they can be performed only once per setup, and do not depend on the specific images.

Since we have $\mathcal{O}(L^2)$ PSWFs in our expansions (corresponding to the indices in the set Ω_T) and $\mathcal{O}(L^2)$ equally spaced Cartesian samples in the unit disk, computing the expansion coefficients of M images using the direct approach of Algorithm 1 requires $\mathcal{O}(ML^4)$ operations. On the other hand, Algorithm 2 allows us to obtain the expansion coefficients in $\mathcal{O}(ML^3)$ operations, since the NFFT in step 3 requires $\mathcal{O}(L^2 \log L)$ operations and step 4 can be implemented using $\mathcal{O}(L^2 \log L + L^3)$ operations, for each image.

Although Algorithm 2 and the method described in [37] (based on Fourier–Bessel basis functions) have the same order of computational complexity, our method enjoys a twofold asymptotic speedup in the coefficients' evaluation, and often it runs about three to four times faster. The twofold asymptotic speedup is because we need asymptotically only half the number of radial nodes for the numerical integration (see analysis in section 3 and Appendix C) compared to the Gaussian quadratures used in [37], which dictates the constant in the leading term of the computational complexity. The greater speedup observed in practice stems from the fact that practically, the most time-consuming operation in both methods is the NFFT, which depends heavily on the total number of quadrature nodes. In our integration scheme, the total number of quadrature nodes is about a quarter of the number of nodes in [37].

Algorithm 2. Evaluating PSWFs expansion coefficients (efficient method).

- 1 **Required:** An image $\{I(\frac{k}{L})\}$ sampled on a Cartesian grid of size $(2L+1) \times (2L+1)$ with $\frac{k}{L} \in \mathbf{Q}$, and $\mathbf{Q} = [-1, 1] \times [-1, 1]$.
- 2 **Precomputation:**
 1. Choose a bandlimit $c (\leq \pi L)$ and a truncation parameter T .
 2. Choose the number of radial nodes \mathcal{N}_r and angular \mathcal{N}_θ^ℓ nodes for $1 \leq \ell < \mathcal{N}_r$ (according to (18) and (19) or similar relaxed conditions).
 3. Compute the quadrature nodes $\omega_{\ell,j}^{2c} = (r_\omega^\ell, \theta_\omega^{\ell,j})$ and weights for $\mathcal{W}_{\ell,j}^{2c}$ for $\ell = 1, \dots, \mathcal{N}_r$ and $j = 1, \dots, \mathcal{N}_\theta^\ell$, as described in [31].
 4. Evaluate the radial part of the PSWFs $R_{N,n}^c(r)$ at the radial quadrature nodes for $r_\omega^\ell(N, n) \in \Omega_T$ (see (8)).
- 3 Compute $\phi^c(\omega_{\ell,j}^{2c})$ from (14) by NFFT.
- 4 For $(N, n) \in \Omega_T$, compute the expansion coefficients $\hat{a}_{N,n}$ via equations (26) and (27).

Algorithm 3. PSWFs-based steerable PCA.

- 1 **Required:** PSWFs expansion coefficients of M images $\{\hat{a}_{N,n}^m\}_{m=0}^{M-1}$ for $(N, n) \in \Omega_T$ and a bandlimit c .
 - 2 **Precomputation:** Evaluate the PSWFs $\psi_{N,n}^c(\frac{k}{L}, \frac{k}{L}) \in \mathbf{D}$ and their eigenvalues $\alpha_{N,n}^c$ for $(N, n) \in \Omega_T$ according to [31].
 - 3 Compute the expansion coefficients of the mean image for $\hat{\mu}_{0,n} = \frac{1}{M} \sum_{m=0}^{M-1} \hat{a}_{0,n}^m$ for $n = 0, \dots, n_0$ where n_0 is the largest index n such that $(0, n) \in \Omega_T$.
 - 4 Update $\hat{a}_{0,n}^m \leftarrow \hat{a}_{0,n}^m - \hat{\mu}_{0,n}$.
 - 5 Compute the eigenvalues $\hat{\lambda}_1, \dots, \hat{\lambda}_{|\Omega_T|}$ and eigenvectors $\hat{g}_1, \dots, \hat{g}_{|\Omega_T|}$ of the matrix C (from (30)) by diagonalizing each of its blocks separately.
 - 6 Compute the sampled basis functions $g_\ell(\frac{k}{L}) = \sum_{n=0}^{n_\ell} \hat{g}_{N_\ell,n}^\ell \psi_{N_\ell,n}^\ell(\frac{k}{L})$, where n_ℓ stands for the largest index n such that $(N_\ell, n) \in \Omega_T$, and $\hat{g}_{N_\ell,n}^\ell$ are the entries of the eigenvector \hat{g}_ℓ corresponding to the pair (N_ℓ, n) .
 - 7 Compute the coefficients of I_m in the steerable-PCA basis by $c_{\ell,m} = \sum_{n=0}^{n_\ell} \hat{a}_{N_\ell,n}^m (\hat{g}_{N_\ell,n}^\ell)^*$.
-

Next, as of the computational complexity of the steerable PCA (Algorithm 3), forming the blocks of the matrix C requires $\mathcal{O}(ML^3)$ operations, and then the eigen-decomposition of each block requires $\mathcal{O}(L^3)$ operations, where there are $\mathcal{O}(L)$ different blocks in the matrix. Therefore, obtaining the eigenvalues and eigenvectors in step 5 of Algorithm 3 requires $\mathcal{O}(ML^3 + L^4)$ operations. As pointed out in [38] and [37], the eigenvalues and eigenvectors can be also obtained from the SVD of the matrices of coefficients $\hat{c}_{N,n}^m$ from which the blocks of C are obtained.

Following (32), if we have $\mathcal{O}(L^2)$ basis functions $g_\ell(x)$, their evaluation on the Cartesian grid requires $\mathcal{O}(L^5)$ operations. Sometimes it may be more convenient to evaluate these basis

functions on a polar grid instead, in which case the computational complexity reduces to $\mathcal{O}(L^4)$.

Lastly, computing the expansion coefficients of the images in the steerable basis via (33) requires $\mathcal{O}(ML^3)$ operations, since $\mathcal{O}(L)$ operations are required to compute a single expansion coefficient for every image.

We point out though that often only a small fraction of the basis functions are chosen for subsequent processing (via their eigenvalues), so the contribution of steps 6 and 7 in Algorithm 3 to the overall running time of the steerable PCA procedure is usually negligible.

To summarize, the computational complexity of the entire procedure (computing PSWFs expansion coefficients + steerable PCA), when using Algorithm 2 to compute the expansion coefficients, is $\mathcal{O}(ML^3 + L^5)$ when sampling the steerable principal components on the Cartesian grid and $\mathcal{O}(ML^3 + L^4)$ when sampling them on a polar grid. It is important to note that, although Algorithm 1 for evaluating PSWFs expansion coefficients suffers from an inferior order of computational complexity (compared to Algorithm 2), it is simpler to implement and may still run faster (due to optimized implementations of the scalar product on CPUs and GPUs), particularly for small values of L .

6. Steerable PCA in the presence of noise

Up to this point, we have presented a method for computing the steerable PCA of a set of images localized in space and frequency, sampled on a Cartesian grid. In many practical settings however, the images are corrupted with noise. Therefore, it is beneficial to understand the impact this noise has on the PSWFs expansion coefficients, and in particular, it is generally convenient if the transformation to the expansion coefficients does not alter the spectrum of the noise. In this section, we demonstrate numerically that for sufficiently localized images in space-frequency, for which we can choose a truncation parameter $T \gg 1$ (see (8)), the transformation to the PSWFs expansion coefficients is essentially orthogonal. In particular, the higher the truncation parameter T , the closer is our steerable PCA to orthogonality.

Let us denote by I a column vector consisting of the clean (Cartesian) samples of an input image. Suppose that our image is corrupted by an additive noise such that

$$\tilde{I} = I + \xi, \quad (36)$$

where ξ is zero mean noise vector with covariance matrix R_ξ . From (9), our approximated expansion coefficients are given (up to a constant) by

$$(\hat{\psi}^c)^* \tilde{I} = (\hat{\psi}^c)^* I + (\hat{\psi}^c)^* \xi, \quad (37)$$

where the operator $(\cdot)^*$ stands for the conjugate transpose, and $\hat{\psi}^c$ denotes the matrix whose columns contain samples of $\psi_{N,n}^c(x)|\lambda_{N,n}^c|^2/L$ inside the unit disk, with different columns corresponding to different pairs of indices (N, n) . If we define the vector of additive noise in the expansion coefficients as

$$\hat{\xi} = (\hat{\psi}^c)^* \xi, \quad (38)$$

then its covariance matrix is provided by

$$R_{\hat{\xi}} = (\hat{\psi}^c)^* R_{\xi} \hat{\psi}^c. \quad (39)$$

Now, if the noise in (36) is white ($R_{\xi} = \sigma^2 \mathbf{I}$), and in order to preserve the covariance of the noise, we would require the matrix

$$H_c \triangleq (\hat{\psi}^c)^* \hat{\psi}^c \quad (40)$$

to be as close as possible to the identity matrix, or equivalently that the eigenvalues of the matrix H_c be as close as possible to 1. As the matrix H_c is Hermitian and positive semidefinite, it has non-negative real-valued eigenvalues $\nu_1, \nu_2, \dots, \nu_{QT}$. To determine how close H_c is to the identity matrix, we evaluated numerically the maximal distance (in absolute value) between the eigenvalues $\{\nu_k\}$ and 1, and the results are plotted in Figure 5 for various values of L, T , and the bandlimit c .

It is noteworthy that for $T = 10^6$ the eigenvalues of H_c differ from 1 by about 10^{-12} . We have also observed that the spectrum of the matrix H_c is related to the truncation parameter T by

$$\max_k |1 - \nu_k| \approx 2T^{-2}, \quad (41)$$

for $T \gg 10$. This observation is exemplified numerically in Figure 6. As the matrix H_c in (40) is essentially orthogonal for sufficiently large values of T , it is important to mention that, if the images are sufficiently localized in space and frequency, it is often possible to choose the value of T as to enjoy the orthogonality of the transform, while keeping the bound in (10) sufficiently small.

7. Numerical experiments

We begin by demonstrating the running times of our algorithm for large datasets. We generated several sets of 20,000 white noise images, where the sets consist of images with

different L for each set, and compared the running time of our algorithm with the FFBSPCA algorithm of [37], where the coefficients' evaluation for our method was carried out using both Algorithm 1 (direct method) and Algorithm 2 (efficient method). All of the algorithms were implemented in Matlab, and were executed on a dual Intel Xeon X5560 CPU (8 cores in total), with 96 GB of RAM running Linux. Whenever possible, all 8 cores were used simultaneously, either explicitly using Matlab's `PARFOR`, or implicitly, by employing Matlab's implementation of BLAS, which takes advantage of multicore computing. As for the NFFT implementation, we used the software package [12], with an oversampling of 2 and a truncation parameter $m = 6$ (which provides accuracy close to machine precision). The running times (in seconds) are shown in Table 1. As anticipated, Algorithm 1 runs faster than Algorithm 2 for small image sizes, but becomes significantly slower for larger values of L . As for the running time of our algorithm versus FFBSPCA [37], we have mentioned in section 5 that our algorithm is asymptotically two times faster, since the number of radial nodes in our quadrature scheme is asymptotically half that of FFBSPCA. In addition, the total number of quadrature nodes in our scheme is about a quarter of that of FFBSPCA, and since the NFFT procedure for evaluating the Fourier transform of the sampled images on a polar grid (see (14)) is the most time-consuming step of both algorithms, it is expected that our algorithm will be faster than FFBSPCA by a factor between 2 and 4, which indeed agrees with the results in Table 1. In all scenarios tested, most of the computation time was spent on evaluating the PSWFs expansion coefficients (Algorithm 2), and only a small fraction of the time on the eigen-decomposition of the rotationally invariant covariance matrix. Table 2 summarizes the time spent on the evaluation of the PSWFs expansion coefficients (using Algorithm 2) versus time spent on the eigen-decomposition of the rotationally invariant covariance matrix.

Next, we demonstrate our algorithms on simulated single-particle cryo-electron microscopy (cryo-EM) projection images. In single-particle cryo-EM, one is interested in reconstructing a three-dimensional model of a macromolecule (such as a protein) from its two-dimensional projection images taken by an electron microscope. The procedure begins by embedding many copies of the macromolecule in a thin layer of ice (hence the “cryo” in the name of the procedure), where due to the experimental setup the different copies are frozen at random unknown orientations. Then, an electron microscope acquires two-dimensional projection images of the electron densities of these macromolecules. This procedure of image acquisition can be modelled mathematically as the Radon transform of a volume function evaluated at random viewing directions. Due to the properties of the imaging procedure, each projection image generated by the electron microscope undergoes a convolution with a kernel, referred to as the “contrast transfer function” (CTF), which is known to have a Gaussian envelope [6]. Since the unknown volume is essentially limited in space, and since the behavior of the CTF dictates that all images are localized in Fourier domain, we conclude that projection images obtained by single-particle cryo-EM are essentially limited to circular domains in both space and frequency. We emphasize that although the goal in single particle cryo-EM is to reconstruct the three-dimensional structure of the macromolecule, the input to the reconstruction process is only the set of two-dimensional images. Now, since the in-plane rotation of each single-particle cryo-EM image in the detector plane is arbitrary and irrelevant for the reconstruction, the image processing

methods applied to the input image dataset, such as denoising and classification, should be invariant to these in-plane rotations. This observation explains why these images are suitable for exemplifying our steerable-PCA algorithm.

To demonstrate the accuracy of our method, we simulated 10,000 clean projection images from a noiseless three-dimensional density map (EMD-5578 from The Electron Microscopy Data Bank (EMDB) [19]) using the ASPIRE package [1], and obtained steerable principal components using both our fast algorithm and FFBSPCA [37]. Then, we used different numbers of principal components to reconstruct the images, and compared the theoretical error predicted by the residual eigenvalues (see (34)) with the empirical error obtained by comparing the reconstructions to the original images. Obviously, the difference between the two errors is due to the error incurred by the images' expansion scheme (PSWFs or Fourier–Bessel functions). Typical projection images from this dataset can be seen in Figure 7. To simplify the setting and the exposition, we used the projection images as they were obtained from the given volume, and we did not crop, filter, or process them beforehand. Therefore, we assumed a bandlimit $c = \pi L$ throughout the experiment. We note that in general it is possible to crop and filter the images (i.e., choose a smaller bandlimit c) without significantly degrading the quality of the images (up to some prescribed accuracy), thus reducing the computational burden of the algorithms. This can be accomplished either by power density estimation (as demonstrated in [37]), or by employing more sophisticated and dataset-specific estimation techniques for the B-factor (which governs the Gaussian envelope decay of the CTF); see for example [23]. Figure 8 depicts the first 12 eigenfunctions obtained by our steerable-PCA algorithm for this dataset. In Figure 9(a) we show the relative error norms for the FFBSPCA algorithm and our PSWFs-based algorithm, using several values of T and with different numbers of eigenfunctions in the reconstruction. As expected, the theoretical and empirical relative error norms coincide when a small number of principal components is used in the reconstruction, yet when a large number of principal components is used, the error incurred by the expansion of the images (either using PSWFs or Fourier–Bessel) comes into play and dominates the overall error. As guaranteed by (35) and (10) for space-frequency localized images, smaller values of T lead to smaller approximation errors, and we notice that our PSWFs-based method outperforms the FFBSPCA algorithm in terms of accuracy for $T = 10$ and $T = 10^{-3}$. It is also important to mention that the number of PSWFs taking part in the approximation of each image does not increase significantly when lowering T (see (11)). On the other hand, if one allows the approximation error to be of the order of 10^{-3} to 10^{-4} , then it is also possible to use a very large truncation parameter, such as $T = 10^5$, for which the PSWFs-based transform enjoys superior orthogonality properties (see Figure 5), as well as shorter expansions. Often, such scenarios arise when handling noisy datasets. In a way, the truncation parameter provides us with flexibility to adapt the approximation scheme to the specific setting. In Figure 9(b) we show the approximation errors obtained by expanding the images in the PSWFs basis and in the Fourier–Bessel basis (without the rotationally invariant orthogonalization), where we sorted the basis functions according to their contribution to the expansion. It is evident that PSWFs are more adapted to expanding our image dataset than Fourier–Bessel, which is reasonable due to the underlying model behind these images, whereas both bases provide

lower accuracy than the steerable principal components obtained from our PSWFs-based steerable PCA.

8. Summary and discussion

In this paper, we utilized PSWFs-based computational tools to construct fast and accurate algorithms for obtaining steerable principal components of large image datasets. The accuracy of our algorithms is guaranteed under the assumptions of localization of the images in space and frequency, which are natural assumptions for many datasets, particularly in the field of tomography. For M images, each sampled on a Cartesian grid of size $(2L + 1) \times (2L + 1)$, the computational complexity for obtaining the steerable principal components is $\mathcal{O}(ML^3)$ operations, and their accuracy is of the order of the localization of the images in space and frequency, i.e., the norm of the images outside the unit disk in space, and the norm of their Fourier transform outside a disk of radius c , where c is the chosen bandlimit. We have compared our method with the FFBSPCA algorithm [37], which is considered state of the art for performing steerable PCA on single-particle cryo-EM projection images, and have shown that our method is both faster and more accurate (for sufficiently small values of T). In addition, our method enjoys rigorous error bounds throughout its various steps, whereas in contrast the FFBSPCA algorithm provides no analytic guarantees on its accuracy. We mention that classical operations of windowing and filtering can be subsumed in our procedure to ensure that the images fulfil the requirements of space-frequency localization.

As image resolutions get higher, investigating more efficient methods for processing image datasets is an important ongoing research task. As the running times of our algorithms are mostly dominated by the task of computing the PSWFs expansion coefficients, reducing the computational complexity of this step from $\mathcal{O}(L^3)$ to $\mathcal{O}(L^2 \log L)$ for each image, resulting in the same asymptotic complexity as the two-dimensional FFT, will be a significant progress. Since the radial part of the PSWFs is evaluated to a prescribed accuracy using a finite series of special polynomials which admit a recurrence relation (see [31]), and since the expansion coefficients in terms of these polynomials are obtained from the eigenvectors of a tridiagonal matrix, it is possible to employ the methods described in [34] to derive an $\mathcal{O}(L^2 \log L)$ algorithm for computing the PSWFs expansion coefficients of a function. In addition, as mentioned earlier, the NFFT, which was employed to map the Cartesian grid samples to a polar grid, is a major time-consuming component. In this context, we mention that gridding methods (see for example [29]), when applied carefully, may accelerate such a mapping, and thus reduce the overall running time of our algorithms.

Acknowledgments

Funding: This research was supported by THE ISRAEL SCIENCE FOUNDATION grant No. 578/14, and by Award Number R01GM090200 from the NIGMS.

Appendix A. PSWFs expansion error bound

Let us define the restriction of an image $I(x)$ to the unit disk by

$$\bar{I}(x) \triangleq \begin{cases} I(x), & x \in \mathbf{D}, \\ 0, & x \notin \mathbf{D}, \end{cases} \quad (42)$$

and denote the Fourier transform of (x) by

$$\bar{I}^{\mathcal{F}}(\omega) \triangleq \mathcal{F}\{\bar{I}\}(\omega) = \int_{\mathbb{R}^2} \bar{I}(x) e^{i\omega x} dx. \quad (43)$$

Since the Fourier transform is a linear operator, we can write

$$\bar{I}^{\mathcal{F}}(\omega) = \mathcal{F}\{I - (I - \bar{I})\}(\omega) = I^{\mathcal{F}}(\omega) + \mathcal{F}\{I - \bar{I}\}(\omega), \quad (44)$$

where we defined the Fourier transform of $I(x)$ by $I^{\mathcal{F}}(\omega)$. Now, it is clear that

$$\|\bar{I}^{\mathcal{F}}(\omega)\|_{\mathcal{L}^2(\omega \notin c\mathbf{D})} \leq \|I^{\mathcal{F}}(\omega)\|_{\mathcal{L}^2(\omega \notin c\mathbf{D})} + \|\mathcal{F}\{I - \bar{I}\}(\omega)\|_{\mathcal{L}^2(\omega \notin c\mathbf{D})} \quad (45)$$

$$\leq \|I^{\mathcal{F}}(\omega)\|_{\mathcal{L}^2(\omega \notin c\mathbf{D})} + \|\mathcal{F}\{I - \bar{I}\}(\omega)\|_{\mathcal{L}^2(\mathbb{R}^2)}. \quad (46)$$

Next, by employing Parseval's identity we obtain

$$\|\mathcal{F}\{I - \bar{I}\}(\omega)\|_{\mathcal{L}^2(\mathbb{R}^2)} = 2\pi \|I(x) - \bar{I}(x)\|_{\mathcal{L}^2(\mathbb{R}^2)}, \quad (47)$$

and thus, by using our space-frequency concentration assumptions on the image $I(x)$, we have that

$$\|\bar{I}^{\mathcal{F}}(\omega)\|_{\mathcal{L}^2(\omega \notin c\mathbf{D})} \leq \delta_c + 2\pi\varepsilon. \quad (48)$$

Therefore, we conclude that (x) is $(1, \bar{\varepsilon})$ -concentrated in space, and its Fourier transform $I^{\mathcal{F}}(\omega)$ is $(c, \bar{\delta}_c)$ -concentrated, where

$$\bar{\varepsilon} = 0, \bar{\delta}_c = \delta_c + 2\pi\varepsilon. \quad (49)$$

Now, it follows from [13] that the approximation error of \bar{I} satisfies

$$\|\bar{I}(x) - \sum_{N,n \in \Omega_T} \hat{a}_{N,n} \psi_{N,n}^c(x)\|_{\mathcal{L}^2(\mathbf{D})} \leq \left(\bar{\varepsilon} + \frac{\bar{\delta}_c}{2\pi}\right) T + \eta \left(\frac{c}{2\pi L}\right)^2 \sqrt{\sum_{\frac{k}{L} \notin \mathbf{D}} \left|\bar{I}\left(\frac{k}{L}\right)\right|^2} + \frac{2\bar{\delta}_c}{\pi}, \quad (50)$$

where T is the truncation parameter defined in (8), and η is a constant no bigger than $2\pi^2 L/c$. Finally, since $I_m(x)$ and $\bar{I}_m(x)$ coincide on \mathbf{D} , we get using (50)

$$\|I_m(x) - \sum_{N,n \in \Omega_T} \hat{a}_{N,n}^m \psi_{N,n}^c(x)\|_{L^2(\mathbf{D})} \leq \left(\varepsilon + \frac{\delta_c}{2\pi}\right) (T+4). \quad (51)$$

Appendix B. A bound for $N\theta^\ell$

Recall that we choose the number of angular nodes per radius, denoted \mathcal{N}_θ^ℓ , such that it satisfies

$$\sum_{|j| > \mathcal{N}_\theta^\ell} |J_j(2cr_\omega^\ell \rho)| < C_1 \vartheta_q, \quad (52)$$

for every $\rho \in [0, 1]$. Using a bound for the Bessel functions [2] together with the fact that $|\rho| \leq 1$, we get

$$|J_j(2cr_\omega^\ell \rho)| \leq \left(\frac{2cr_\omega^\ell \rho}{2}\right)^j \frac{1}{\Gamma(j+1)} \leq (cr_\omega^\ell)^j \frac{1}{(j+1)!}, \quad (53)$$

and by using Stirling's approximation [2] (for $n > 1$)

$$n! \geq \sqrt{2\pi n} \left(\frac{n}{e}\right)^n > e \left(\frac{n}{e}\right)^n \quad (54)$$

we obtain (for $j > 0$)

$$|J_j(2cr_\omega^\ell \rho)| \leq \frac{1}{cr_\omega^\ell e} \left(\frac{cr_\omega^\ell e}{j+1}\right)^{j+1}. \quad (55)$$

Then, using the fact that the Bessel functions of the first kind satisfy

$$J_{-n}(x) = (-1)^n J_n(x), \quad (56)$$

we can write

$$\sum_{|j| > \mathcal{N}_\theta^\ell} \left| J_j(2cr_\omega^\ell \rho) \right| = 2 \sum_{j > \mathcal{N}_\theta^\ell} \left| J_j(2cr_\omega^\ell \rho) \right| \leq \frac{2}{cr_\omega^\ell e} \sum_{j > \mathcal{N}_\theta^\ell} \left(\frac{cr_\omega^\ell e}{j+1} \right)^{j+1}. \quad (57)$$

We define

$$\gamma_\ell \triangleq cr_\omega^\ell e \quad (58)$$

and choose the number of quadrature nodes per radius as

$$\mathcal{N}_\theta^\ell = \lceil \gamma_\ell \rceil + d, \quad (59)$$

where d is some positive integer (to be determined shortly) and $\lceil \cdot \rceil$ denotes the rounding up operation. It can be easily verified that

$$\frac{\gamma_\ell}{j+1} \leq \frac{\gamma_\ell}{\gamma_\ell + d} \quad (60)$$

whenever $j > \mathcal{N}_\theta^\ell$. Using (57) and (60) we get that

$$\sum_{|j| > \mathcal{N}_\theta^\ell} \left| J_j(2cr_\omega^\ell \rho) \right| \leq \frac{2}{\gamma_\ell} \sum_{j > \mathcal{N}_\theta^\ell} \left(\frac{\gamma_\ell}{\gamma_\ell + d} \right)^{j+1} \leq 2 \left(1 + \frac{d}{\gamma_\ell} \right)^{-\mathcal{N}_\theta^\ell} \leq 2 \left(1 + \frac{d}{\gamma_\ell} \right)^{-\gamma_\ell - d} \leq 2e^{-d} \quad (61)$$

where we have used the inequality

$$\left(1 + \frac{a}{b} \right)^{-b} < e^{-\frac{ab}{a+b}}, \quad (62)$$

from [2], with $a = (\gamma\ell + d)\frac{d}{\gamma\ell}$ and $b = \gamma\ell + d$. Finally, we can see that in order to satisfy (52) it is sufficient to choose

$$d = d(\vartheta_q) \geq \log \vartheta_q^{-1} + \log \frac{2}{C_1}, \quad (63)$$

and thus

$$\mathcal{N}_\theta^\ell \geq cr_\omega^\ell e + \log \vartheta_q^{-1} + \log \frac{2}{C_1} + 1. \quad (64)$$

Appendix C. Behavior and decay properties of $\lambda_{0,k}^{2c}$

We start by reviewing some known results on the behavior of the eigenvalues of the PSWFs in the one-dimensional setting, where they have been thoroughly investigated (see [24] and references therein). The most well-known characterization of these eigenvalues is that they can be divided into three distinct regions of behavior (as a function of their index n): a flat region, where the (normalized) eigenvalues are essentially 1, a transitional region, where they decay from values close to 1 to values close to 0, and a superexponential decay region, where they are very close to 0 and decay as $\sim e^{-m \log n}$. In addition, it is known that if we choose all eigenvalues that are greater than some small ε , then there are about $2c/\pi$

eigenvalues from the flat region, $O(\log(\frac{1-\varepsilon}{\varepsilon}) \log c)$ eigenvalues from the transitional region, and $O(\log c)$ eigenvalues from the decay region (see [18] for a precise formulation). Thus, the number of eigenvalues greater than ε is dominated by the number of eigenvalues in the flat region, which is $2c/\pi$. As for the eigenvalues of the two-dimensional PSWFs, results in [30] indicate that as in the one-dimensional setting, the eigenvalues can be similarly divided into three distinct regions: flat, transitional, and superexponential decay regions.

Correspondingly, the number of significant eigenvalues is dominated by the number of eigenvalues in the flat region. Since $\lambda_{0,k}^{2c} \propto \alpha_{0,k}^{2c}$ (see (8)), it is clear that in order to satisfy condition (19) we need to determine the number of terms in the flat region of $|\lambda_{0,k}^{2c}|$. To this end, we follow [15], which provides similar results for general non-Hermitian Toeplitz integral operators (see also [18] and [36]), and consider the sum

$$\sum_{k=0}^{\infty} |\lambda_{0,k}^{2c}|^2, \quad (65)$$

which is approximately equal to the number of values of $|\lambda_{0,k}^{2c}|$ which are close to 1 (denoted as the flat region). For simplicity of the presentation, we evaluate this sum for a bandlimit of

c , and eventually, replace c with $2c$. From [32], the radial functions $R_{N,n}^c(r)$ in (4) are real-valued, and are obtained as the solutions to the integral equation

$$\beta R(r) = \int_0^1 R(\rho) J_N(cr\rho) \rho d\rho, r \in [0, 1], \quad (66)$$

where $J_N(x)$ is the Bessel function of the first kind of order N . The eigenvalues $\alpha_{N,n}^c$ and $\beta_{N,n}^c$ of (3) and (66) are related by

$$\alpha_{N,n}^c = 2\pi \iota^N \beta_{N,n}^c. \quad (67)$$

By substituting $\varphi(r) = R(r)$ and $\gamma = \beta c$ into (66), we obtain the integral equation

$$\gamma \varphi(r) = \int_0^1 \varphi(\rho) J_N(cr\rho) \sqrt{cr\rho} d\rho, r \in [0, 1], \quad (68)$$

whose eigenvalues and eigenfunctions are $\gamma_{N,n}^c$ and $\varphi_{N,n}^c(r)$, respectively. Equation (68) was analyzed in [32], where it is established that the eigenfunctions $\varphi_{N,n}^c(r)$ constitute a complete orthonormal system in $\mathcal{L}^2[0, 1]$. Therefore, it follows that we have the identity

$$J_N(cr\rho) \sqrt{cr\rho} = \sum_{k=1}^{\infty} \gamma_{N,n}^c \varphi_{N,n}^c(r) \varphi_{N,n}^c(\rho), \quad (69)$$

for r and ρ both in $[0, 1]$. If we notice that $\lambda_{N,n}^c = \sqrt{c} \gamma_{N,n}^c$ and take $N=0$, we have

$$J_0(cr\rho) c \sqrt{r\rho} = \sum_{k=1}^{\infty} \lambda_{0,k}^c \varphi_{0,k}^c(r) \varphi_{0,k}^c(\rho). \quad (70)$$

Next, we take the squared absolute value of both sides of the equation above, followed by double integration (in r and ρ) to obtain

$$\int_0^1 \int_0^1 J_0^2(cr\rho) c^2 r \rho dr d\rho = \sum_{k=0}^{\infty} |\lambda_{0,k}^c|^2. \quad (71)$$

By evaluating the left-hand side of (71) using known integral identities of the Bessel functions [2], and after some manipulation, one can verify that

$$\sum_{k=0}^{\infty} |\lambda_{0,k}^c|^2 = \frac{c^2}{4} (J_0^2(c) - J_2(c)J_0(c) + 2J_1^2(c)). \quad (72)$$

Now, using the asymptotic approximation (see [2])

$$J_m(x) \sim \sqrt{\frac{2}{\pi x}} \cos\left(x - \frac{m\pi}{2} - \frac{\pi}{4}\right), \quad (73)$$

valid for $x \gg |m^2 - 1/4|$, we can write

$$\begin{aligned} \sum_{k=0}^{\infty} |\lambda_{0,k}^c|^2 &\sim \frac{c}{2\pi} \left(\cos^2\left(c - \frac{\pi}{4}\right) - \cos\left(c - \frac{\pi}{4}\right) \right. \\ &\quad \left. - \pi - \frac{\pi}{4} \cos\left(c - \frac{\pi}{4}\right) + 2\cos^2\left(c - \frac{\pi}{2} - \frac{\pi}{4}\right) \right) \\ &= \frac{c}{2\pi} \left(2\cos^2\left(c - \frac{\pi}{4}\right) - \frac{\pi}{4} + 2\sin^2\left(c - \frac{\pi}{4}\right) \right) = \frac{c}{\pi}. \end{aligned} \quad (74)$$

Therefore, the expression in (74) establishes that for large values of c , the number of terms in the set $\{|\lambda_{0,n}^{2c}|\}$ which are close to 1 is about $2c/\pi$ (see also Figure 4).

References

1. Algorithms for single particle reconstruction. <http://spr.math.princeton.edu/>
2. Abramowitz, M., Stegun, IA. Handbook of Mathematical Functions: With Formulas, Graphs, and Mathematical Tables. Courier Corporation; Chelmsford, MA: 1964.
3. Dutt A, Rokhlin V. Fast Fourier transforms for nonequispaced data. SIAM J Sci Comput. 1993; 14:1368–1393.
4. Ferraro M, Caelli TM. Relationship between integral transform invariances and Lie group theory. J Opt Soc Am A. 1988; 5:738–742.
5. Fessler JA, Sutton BP. Nonuniform fast Fourier transforms using min-max interpolation. IEEE Trans Signal Process. 2003; 51:560–574.
6. Frank, J. Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State. Oxford University Press; Oxford: 2006.
7. Freeman WT, Adelson EH. The design and use of steerable filters. IEEE Trans Pattern Anal Mach Intell. 1991; 13:891–906.

8. Greengard L, Lee JY. Accelerating the nonuniform fast Fourier transform. *SIAM Rev.* 2004; 46:443–454.
9. Hilai R, Rubinstein J. Recognition of rotated images by invariant Karhunen–Loève expansion. *J Opt Soc Am A.* 1994; 11:1610–1618.
10. Jogan M, Zagar E, Leonardis A. Karhunen–Loeve expansion of a set of rotated templates. *IEEE Trans Image Process.* 2002; 12:817–825.
11. Kanatani, K. *Springer Series in Information Sciences. Vol. 20.* Springer; New York: 2012. *Group-Theoretical Methods in Image Understanding.*
12. Keiner J, Kunis S, Potts D. Using NFFT 3 – a software library for various nonequispaced fast Fourier transforms. *ACM Trans Math Softw.* 2009; 36:19.
13. Landa B, Shkolnisky Y. Approximation scheme for essentially bandlimited and space-concentrated functions on a disk. *Appl Comput Harmon Anal.* 2016 accepted.
14. Landau HJ. Necessary density conditions for sampling and interpolation of certain entire functions. *Acta Math.* 1967; 117:37–52.
15. Landau HJ. On Szegő's eigenvalue distribution theorem and non-hermitian kernels. *Journal d'Analyse Mathématique.* 1975; 28:335–357.
16. Landau HJ, Pollak HO. Prolate spheroidal wave functions, Fourier analysis and uncertainty - II. *Bell Syst Tech J.* 1961; 40:65–84.
17. Landau HJ, Pollak HO. Prolate spheroidal wave functions, Fourier analysis and uncertainty - III: The dimension of the space of essentially time-and band-limited signals. *Bell Syst Tech J.* 1962; 41:1295–1336.
18. Landau HJ, Widom H. Eigenvalue distribution of time and frequency limiting. *J Math Anal Appl.* 1980; 77:469–481.
19. Lawson CL, Patwardhan A, Baker ML, Hryc C, Garcia ES, Hudson BP, Lagerstedt I, Ludtke SJ, Pintilie G, Sala R, Westbrook JD, Berman HM, Kleywegt GJ, Chiu W. EMDatabank unified data resource for 3DEM. *Nucleic Acids Res.* 2016; 44:D396–D403. [PubMed: 26578576]
20. Lenz R. Group-theoretical model of feature extraction. *J Opt Soc Am A.* 1989; 6:827–834.
21. Lenz R. Group invariant pattern recognition. *Pattern Recognition.* 1990; 23:199–217.
22. Lenz, R. *Group Theoretical Methods in Image Processing.* Springer; New York: 1990.
23. Mallick SP, Carragher B, Potter CS, Kriegman DJ. ACE: automated CTF estimation. *Ultramicroscopy.* 2005; 104:8–29. [PubMed: 15935913]
24. Osipov, A., Rokhlin, V., Xiao, H. *Applied Mathematical Sciences. Vol. 187.* Springer; New York: 2013. *Prolate Spheroidal Wave Functions of Order Zero.*
25. Perona P. Deformable kernels for early vision. *IEEE Trans Pattern Anal Mach Intell.* 1995; 17:488–499.
26. Petersen DP, Middleton D. Sampling and reconstruction of wave-number-limited functions in n -dimensional Euclidean spaces. *Inf Control.* 1962; 5:279–323.
27. Ponce C, Singer A. Computing steerable principal components of a large set of images and their rotations. *IEEE Trans Image Process.* 2011; 20:3051–3062. [PubMed: 21536533]
28. Potts, D., Steidl, G., Tasche, M. *Modern Sampling Theory.* Springer; New York: 2001. *Fast Fourier transforms for nonequispaced data: A tutorial;* p. 247-270.
29. Rosenfeld D. An optimal and efficient new gridding algorithm using singular value decomposition. *Magn Reson Med.* 1998; 40:14–23. [PubMed: 9660548]
30. Serkh, K. Technical Report TR-1519. Department of Mathematics, Yale University; New Haven, CT: 2015. *On generalized prolate spheroidal functions.*
31. Shkolnisky Y. Prolate spheroidal wave functions on a disc - integration and approximation of two-dimensional bandlimited functions. *Appl Comput Harmon Anal.* 2007; 22:235–256.
32. Slepian D. Prolate spheroidal wave functions, Fourier analysis and uncertainty - IV: Extensions to many dimensions; generalized prolate spheroidal functions. *Bell Syst Tech J.* 1964; 43:3009–3057.
33. Slepian D, Pollak HO. Prolate spheroidal wave functions, Fourier analysis and uncertainty -I. *Bell Syst Tech J.* 1961; 40:43–63.
34. Tygert M. Recurrence relations and fast algorithms. *Appl Comput Harmon Anal.* 2010; 28:121–128.

35. Vonesch C, Stauber F, Unser M. Steerable PCA for rotation-invariant image recognition. *SIAM J Imaging Sci.* 2015; 8:1857–1873.
36. Xiao H, Rokhlin V, Yarvin N. Prolate spheroidal wavefunctions, quadrature and interpolation. *Inverse Probl.* 2001; 17:805.
37. Zhao Z, Shkolnisky Y, Singer A. Fast steerable principal component analysis. *IEEE Trans Comput Imaging.* 2016; 2:1–12. [PubMed: 27570801]
38. Zhao Z, Singer A. Fourier–Bessel rotational invariant eigenimages. *J Opt Soc Am A.* 2013; 30:871–877.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

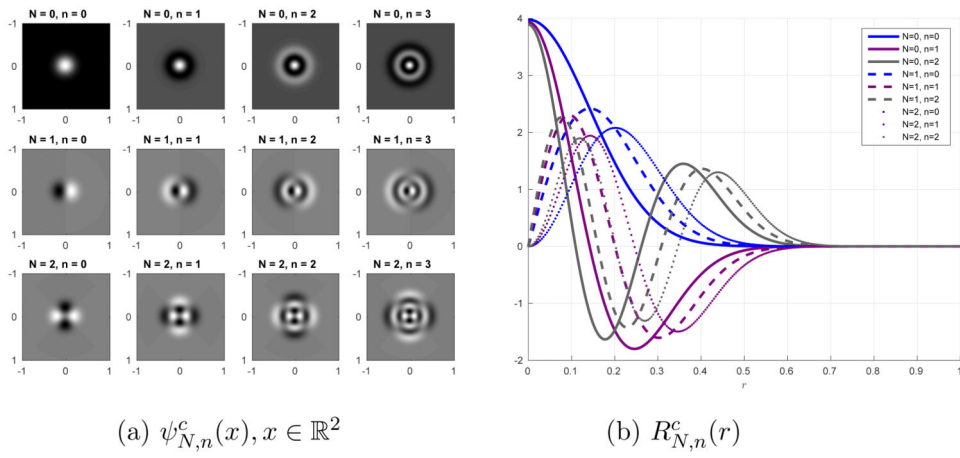


Figure 1. Illustration of the first few PSWFs (real part) with bandlimit $c = 16\pi$, ordered according to their eigenvalues (for every index N , the eigenvalues are ordered from largest to smallest as a function of the index n)

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

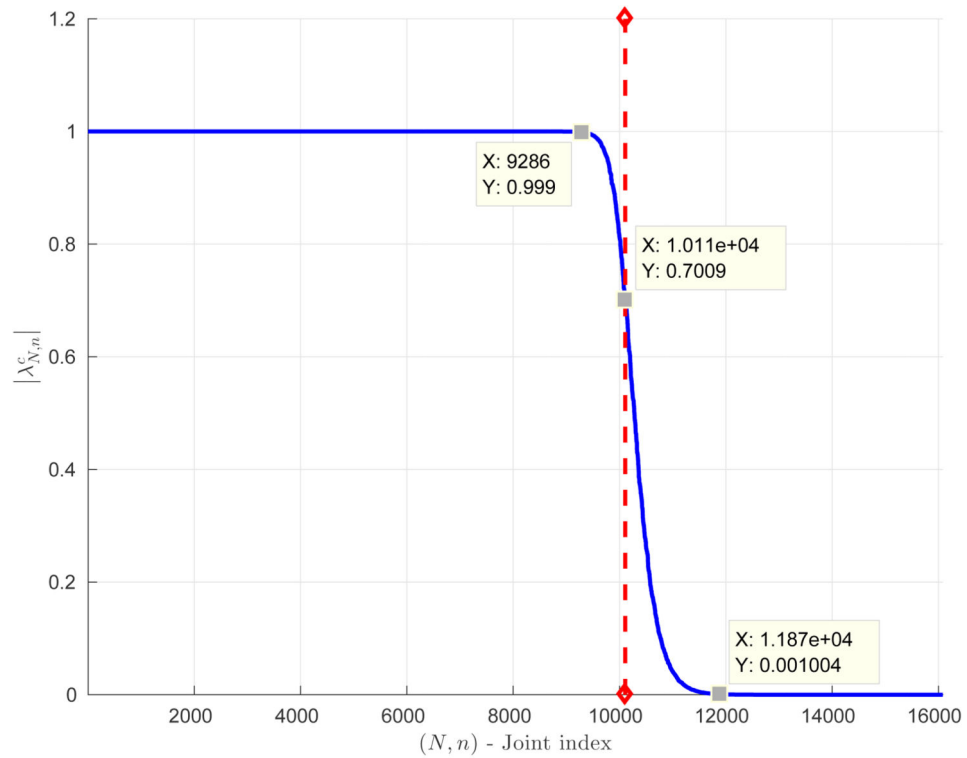


Figure 2.

Illustration of the normalized eigenvalues $|\lambda_{N,n}^c|$ for $L = 64$, $c = \pi L$, sorted in a nonincreasing order with a joint index k enumerating over (N, n) . It is evident that for $T \ll 1$ the set Ω_T consists of indices up to the right of the vertical dashed line, where the normalized eigenvalues rapidly decay to zero. On the other hand, for $T \gg 1$ the set Ω_T consists of indices up to the left of the vertical line, where the normalized eigenvalues are extremely close to 1. The dashed vertical line corresponds to $k = c^2/4$, which also approximately agrees with $|\lambda_{N,n}^c|^2 = 0.5$ (or equivalently with the rule $T = 1$).




Figure 3.

Quadrature nodes $\omega_{\ell,j}^{2c} = (r_{\omega}^{\ell}, \theta_{\omega}^{\ell,j})$ for $L = 16$, $c = \pi L$, and accuracy of the order of machine precision. The nodes are arranged to be equally spaced in the angular direction (the specific phase of the allocation is arbitrary) with different numbers of nodes for each radius (rings closer to the origin require fewer angular nodes). The number of radial nodes is about c/π , and the number of angular nodes per radius r_{ω}^{ℓ} is about $cr_{\omega}^{\ell}e$, where e is the base of the natural logarithm.

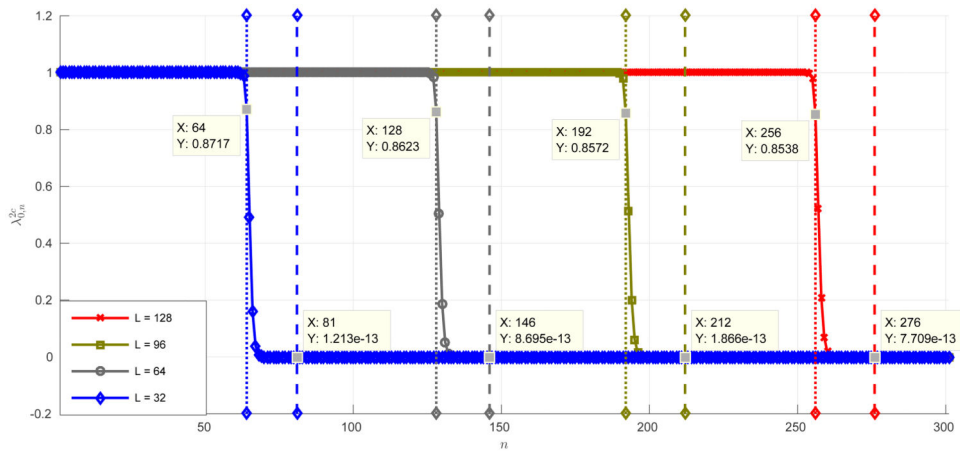


Figure 4.

Behavior of $|\lambda_{0,n}^{2c}|$ for $c = \pi L$ (Nyquist sampling) and several values of L . For each value of L , the dotted vertical line represents $n = 2c/\pi$, and the dashed vertical line represents the first index n for which $|\lambda_{0,n}^{2c}| \leq 10^{-12}$, which we denote by n_1^{2c} . First, we observe that $2c/\pi$ is the exact index for which $|\lambda_{0,n}^{2c}|$ begins its rapid decay, and, second, it is evident that the difference between n_1^{2c} and the estimate $2c/\pi$ grows extremely slowly with L , making the asymptotic estimate $n_1^{2c} \sim 2c/\pi$ sufficiently accurate even for moderate values of L .




Figure 5. Measured deviation of the eigenvalues of H_c from 1 for different values of the truncation parameter T , the bandlimit c , and the sampling resolution L . We notice that the deviation from orthogonality remains approximately constant for different values of L and c . Specifically, for $T = 10^6$ we notice that H_c is practically orthogonal.




Figure 6. Measured deviation of the eigenvalues of H_c from 1 as a function of T , versus the function $2T^{-2}$, for $L = 96$ and $c = \pi L$. We observe that the truncation parameter T gives us direct control over the spectrum of H_c with $\max_k |1 - \nu_k| \approx 2T^{-2}$ for $T \geq 10$.




Figure 7. Sample of 3 simulated noiseless projection images at a resolution of $L = 127$ pixels


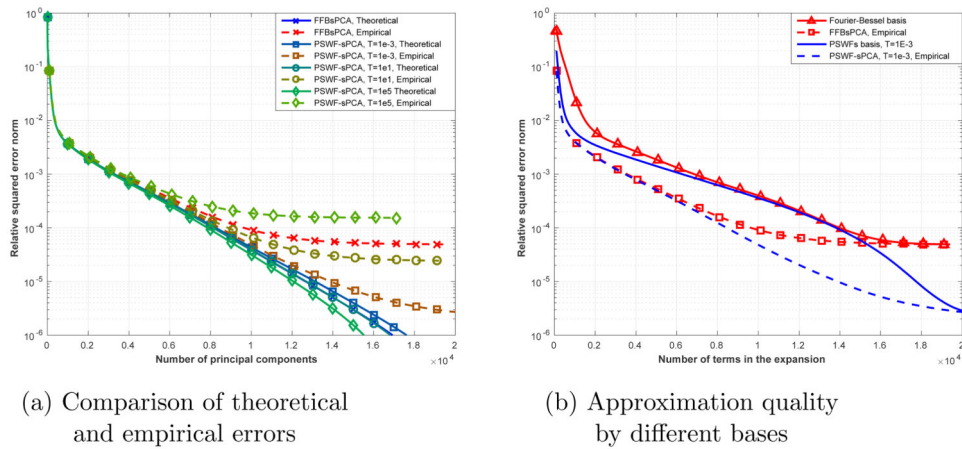


Figure 8. The first 12 eigenfunctions with largest eigenvalues, obtained for $T = 10^{-3}$



(a) Comparison of theoretical and empirical errors

(b) Approximation quality by different bases

Figure 9.

The ratio between the squared error norm and the squared norm of the images, when expanding the images with different numbers of basis functions. Panel (a) compares theoretical errors with empirical errors for our PSWFs-based algorithm (for several values of T) and the FFBsPCA algorithm of [37]. Curves which correspond to theoretical errors are obtained from the residual eigenvalues of the steerable-PCA procedure (see (34)), and empirical errors correspond to measured errors between the reconstructed and original projection images. Panel (b) illustrates the errors due to steerable-PCA expansions, PSWFs expansions, and Fourier–Bessel expansions. Note that the x axis, which counts the number of basis functions used in each expansion, counts only basis functions with non-negative angular indices (since these are sufficient for encoding real-valued images).

Table 1

Comparison of algorithms' running times, for 20,000 images consisting of white noise, for $T=10$ and several values of L . All timings are given in seconds.

L	PSWFs direct (Algs. 1+3)	PSWFs efficient (Algs. 2+3)	FFBsPCA [37]
32	6	25	47
64	95	59	151
96	491	117	363
128	1625	204	697

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Running times of coefficients' evaluation and eigen-decomposition for the efficient PSWFs-based method. All timings are given in seconds.

L	PSWFs coefficients evaluation (Alg. 2)	Eigen-decomposition (Alg. 3 steps 3-6)
32	24.5	0.5
64	56.5	2.5
96	111	6
128	193	11

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript