

Steerable e PCA: Rotationally Invariant Exponential Family PCA

Zhizhen Zhao *Member, IEEE*, Lydia T. Liu, and Amit Singer

Abstract—In photon-limited imaging, the pixel intensities are affected by photon count noise. Many applications require an accurate estimation of the covariance of the underlying 2-D clean images. For example, in X-ray free electron laser (XFEL) single molecule imaging, the covariance matrix of 2-D diffraction images is used to reconstruct the 3-D molecular structure. Accurate estimation of the covariance from low-photon-count images must take into account that pixel intensities are Poisson distributed, hence the classical sample covariance estimator is highly biased. Moreover, in single molecule imaging, including in-plane rotated copies of all images could further improve the accuracy of covariance estimation. In this paper we introduce an efficient and accurate algorithm for covariance matrix estimation of count noise 2-D images, including their uniform planar rotations and possibly reflections. Our procedure, *steerable ePCA*, combines in a novel way two recently introduced innovations. The first is a methodology for principal component analysis (PCA) for Poisson distributions, and more generally, exponential family distributions, called e PCA. The second is steerable PCA, a fast and accurate procedure for including all planar rotations when performing PCA. The resulting principal components are invariant to the rotation and reflection of the input images. We demonstrate the efficiency and accuracy of steerable e PCA in numerical experiments involving simulated XFEL datasets and rotated face images from Yale Face Database B.

Index Terms—Poisson noise, X-ray free electron laser, steerable PCA, eigenvalue shrinkage, autocorrelation analysis, image denoising.

I. INTRODUCTION

X-ray free electron laser (XFEL) is an emerging imaging technique for elucidating the three-dimensional structure of molecules [1], [2]. Single molecule XFEL imaging collects two-dimensional diffraction patterns of single particles at random orientations. The images are very noisy due to the low photon count. The detector count-noise follows an approximately Poisson distribution. Since only one diffraction pattern is captured per particle and the particle orientations are unknown, it is challenging to reconstruct the 3-D structure at low signal-to-noise ratio (SNR). One approach is to use expectation-maximization (EM) [3], [4], but it has a high computational cost at low signal-to-noise ratio (SNR). Alternatively, assuming that the orientations of the particles are uniformly distributed over the special orthogonal group

$SO(3)$, Kam’s correlation analysis [5]–[12] bypasses orientation estimation and requires just one pass over the data, thus alleviating the computational cost. Kam’s method requires a robust estimation of the covariance matrix of the noiseless 2-D images. This serves as our main motivation for developing efficient and accurate covariance estimation and denoising methods for Poisson data. Nevertheless, the methods presented here are quite general and can be applied to other imaging modalities involving Poisson distributions.

Principal Component Analysis (PCA) is widely used for dimension reduction and denoising of large datasets [13], [14]. However, it is most naturally designed for Gaussian data, and there is no consensus on the extension to non-Gaussian settings such as exponential family distributions [13, Sec. 14.4]. For denoising with non-Gaussian noise, popular approaches reduce it to the Gaussian case by a wavelet transform such as the Haar transform [15]; by adaptive wavelet shrinkage [15], [16]; or by approximate variance stabilization such as the Anscombe transform [17]. The latter is known to work well for Poisson signals with large parameters, due to approximate normality. However, the normal approximation breaks down for Poisson distributions with a small parameter, in the case of photon-limited XFEL [18, Sec. 6.6]. Other methods are based on alternating minimization [19]–[21], singular value thresholding (SVT) [22], [23] and Bayesian techniques [24]. Many of aforementioned methods such as [19] are computationally intractable for large datasets and do not have statistical guarantees for covariance estimation. In particular, [21] applied the methodology of alternating minimization [19] to denoise a single image by performing PCA on clusters of patches extracted from the image. However, this approach is not suitable for our problem setting, where the goal is to simultaneously denoise a large number of images and estimate their covariance. Recently, [25] introduced exponential family PCA (e PCA), which extends PCA to a wider class of distributions. It involves the eigendecomposition of a new covariance matrix estimator, constructed in a deterministic and non-iterative way using moment calculations and shrinkage. e PCA was shown to be more accurate than PCA and its alternatives for exponential families. It has computational cost similar to that of PCA, substantial theoretical justification building on random matrix theory, and interpretable output. We refer readers to [25] for experiments that benchmark e PCA against previous methods.

In XFEL imaging, the orientations of the particles are uniformly distributed over $SO(3)$, so it is equally likely to observe any planar rotation of the given 2-D diffraction pattern. Therefore, it makes sense to include all possible in-plane rotations of the images when performing e PCA. To

ZZ is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61820 USA e-mail: zhizhenz@illinois.edu .

LTL is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA, 94720 e-mail: lydiatliu@berkeley.edu

AS is with the Department of Mathematics, and Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ, 08544 e-mail: amits@math.princeton.edu

this end, we incorporate steerability in e PCA, by adapting the steerable PCA algorithm, which avoids duplicating rotated images [26]. The concept of a steerable filter was introduced in [27] and various methods were proposed for computing data adaptive steerable filters [28]–[31]. We take into account the action of the group $O(2)$ on diffraction patterns by in-plane rotation (and reflection if needed). The resulting principal components are invariant to any $O(2)$ transformation of the input images.

The new algorithm, to which we refer as steerable e PCA, combines e PCA and steerable PCA in a natural way. Steerable e PCA is not an iterative optimization procedure. All steps involve only basic linear algebra operations. The various steps include expansion in a steerable basis, eigen-decomposition, eigenvalue shrinkage, and different normalization steps. The mathematical and statistical rationale for all steps is provided in Section II.

We illustrate the improvement in covariance matrix estimation by applying our method to image denoising in Section III. Specifically, we introduce a Wiener-type filtering using the principal components. Rotation invariance enhances the effectiveness of e PCA in covariance estimation and thus achieves better denoising. In addition, the denoised expansion coefficients are useful in building rotationally invariant image features (i.e. bispectrum-like features [32]). Numerical experiments are performed on simulated XFEL diffraction patterns and a natural image dataset—Yale Face Database B [33], [34]. As is the case for standard PCA, the computational complexity of the steerable e PCA is lower than that of e PCA.

An implementation of steerable e PCA in MATLAB is publicly available at github.com/zhizhenz/sepca/.

II. METHODS

The goal of steerable e PCA is to estimate the rotationally invariant covariance matrix from the images whose pixel intensities are affected by photon count noise. To develop this estimator, we combine our previous works on steerable PCA with e PCA in a novel way. The main challenge in combining e PCA with steerable PCA is that steerable PCA involves a (nearly orthogonal) change of basis transformation from a Cartesian grid to a steerable basis. While multidimensional Gaussian distributions are invariant to orthogonal transformations, the multidimensional Poisson distribution is not invariant to such transformations. Therefore, steerable PCA and e PCA cannot be naively combined in a straightforward manner. Algorithm 1 details the steps of the combined procedure that achieves the goal of estimating a rotationally invariant covariance matrix from Cartesian grid images with pixel intensities following a Poisson distribution. The following subsections explain the main concepts underlying e PCA and steerable PCA, and how they are weaved together in a manner that overcomes the aforementioned challenge. In the following subsections, we detail the associated concepts and steps for steerable e PCA.

A. The observation model and homogenization

We adopt the same observation model introduced in [25]. We observe n noisy images $Y_i \in \mathbb{R}^p$ (i.e., p is the number

of pixels), for $i = 1, \dots, n$. These are random vectors sampled from a hierarchical model defined as follows. First, a latent vector—or hyperparameter— $\omega \in \mathbb{R}^p$ is drawn from a probability distribution P with mean μ_ω and covariance matrix Σ_ω . Conditioning on ω , each coordinate of $Y = (Y(1), \dots, Y(p))^\top$ is drawn independently from a canonical one-parameter exponential family,

$$Y(j)|\omega(j) \sim p_{\omega(j)}(y), \quad Y = (Y(1), \dots, Y(p))^\top, \quad (1)$$

with density

$$p_{\omega(j)}(y) = \exp[\omega(j)y - A(\omega(j))] \quad (2)$$

with respect to a σ -finite measure ν on \mathbb{R} , where the j^{th} entry of the latent vector, $\omega(j) \in \mathbb{R}$, is the natural parameter of the family and $A(\omega(j)) = \log \int \exp(\omega(j)y) d\nu(y)$ is the corresponding log-partition function. The mean and variance of the random variable $Y(j)$ can be expressed as $A'(\omega(j))$ and $A''(\omega(j))$, where we denote $A'(\omega) = dA(\omega)/d\omega$. Therefore, the mean of Y conditioning on ω is

$$X := \mathbb{E}(Y|\omega) = (A'(\omega(1)), \dots, A'(\omega(p)))^\top = A'(\omega),$$

so the noisy data vector Y can be expressed as $Y = A'(\omega) + \tilde{\varepsilon}$, with $\mathbb{E}(\tilde{\varepsilon}|\omega) = 0$ and the marginal mean of Y is $\mathbb{E}Y = \mathbb{E}A'(\omega)$. Thus, one can think of Y as a noisy realization of the clean vector $X = A'(\omega)$. However, the latent vector ω is also random and varies from sample to sample. In the XFEL application, $X = A'(\omega)$ are the unobserved noiseless images, and their randomness stems from the random (and unobserved) orientation of the molecule. We may write¹ $Y = A'(\omega) + \text{diag}[A''(\omega)]^{1/2}\varepsilon$, where the coordinates of ε are conditionally independent and standardized given ω . The covariance of Y is given by the law of total covariance:

$$\begin{aligned} \text{Cov}[Y] &= \text{Cov}[\mathbb{E}(Y|\omega)] + \mathbb{E}[\text{Cov}[Y|\omega]] \\ &= \text{Cov}[A'(\omega)] + \mathbb{E} \text{diag}[A''(\omega)]. \end{aligned} \quad (3)$$

The off-diagonal entries of the covariance matrix of the noisy images are therefore the same as those of the clean images. However, the diagonal of the covariance matrix (i.e., the variance) of the noisy images is further inflated by the noise variance. Unlike white noise, the noise variance here often changes from one coordinate to another (i.e., there is heteroscedasticity). In e PCA, the homogenization step is a particular weighting method that improves the signal-to-noise ratio [25, Section 4.2.2]. Specifically, the homogenized vector is defined as

$$Z = \text{diag}[\mathbb{E}A''(\omega)]^{-1/2}Y = \text{diag}[\mathbb{E}A''(\omega)]^{-1/2}A'(\omega) + \varepsilon. \quad (4)$$

Then the corresponding homogenized covariance matrix is

$$\text{Cov}[Z] = \text{diag}[\mathbb{E}A''(\omega)]^{-1/2} \text{Cov}[A'(\omega)] \text{diag}[\mathbb{E}A''(\omega)]^{-1/2} + I. \quad (5)$$

This step is commonly called “prewhitening” in signal and image processing, while the term “homogenization” is more commonly used in statistics. The terms “prewhitened” and “homogenized” are synonyms in the context of this paper. In

¹ $\text{diag}[x]$ for $x \in \mathbb{R}^p$ denotes a $p \times p$ diagonal matrix whose diagonal entries are $x(j)$ for $j = 1, \dots, p$.

Section II-C, we discuss how to estimate the homogenized rotationally invariant covariance matrix.

For the special case of Poisson observations, where $Y \sim \text{Poisson}_p(X)$ and $X \in \mathbb{R}^p$ is random, we can write $Y = X + \text{diag}(X)^{1/2}\varepsilon$. The natural parameter is the vector ω with $\omega(j) = \log X(j)$ and $A'(\omega(j)) = A''(\omega(j)) = \exp(\omega(j)) = X(j)$. Therefore, we have $\mathbb{E}Y = \mathbb{E}X$, and

$$\text{Cov}[Y] = \text{Cov}[X] + \text{diag}[\mathbb{E}X]. \quad (6)$$

In other words, while the mean of the noisy images agrees with the mean of the clean images, their covariance matrices differ by a diagonal matrix that depends solely on the mean image. If we homogenize the observations by $Z = \text{diag}(\mathbb{E}[X])^{-1/2}Y$, then Eq. (6) becomes,

$$\begin{aligned} \text{Cov}[Z] &= \text{diag}(\mathbb{E}[X])^{-1/2} \text{Cov}[Y] \text{diag}(\mathbb{E}[X])^{-1/2} \\ &= \text{diag}(\mathbb{E}[X])^{-1/2} \text{Cov}[X] \text{diag}(\mathbb{E}[X])^{-1/2} + I. \end{aligned} \quad (7)$$

We can estimate the covariance of X from the homogenized covariance Z according to Eq. (7), i.e. $\text{Cov}[X] = \text{diag}(\mathbb{E}[X])^{1/2}(\text{Cov}[Z] - I) \text{diag}(\mathbb{E}[X])^{1/2}$. In Sec. II-E, we detail the corresponding recoloring step to estimate $\text{Cov}[X]$.

In Alg. 1, Steps 1 and 5, with $D_n = \text{diag}[\bar{Y}]$, correspond to the homogenization procedure in *ePCA*. We provide more details on how to estimate the homogenized rotationally invariant covariance matrix (Steps 1–5) in Sec. II-B and Sec. II-C.

B. Steerable basis expansion

Under the observation model in Sec. II-A, we develop a method that estimates the rotationally invariant $\text{Cov}[X]$ efficiently and accurately from the image dataset Y . We assume that a digital image I is composed of discrete samples from a continuous function f with band limit c . The Fourier transform of f , denoted $\mathcal{F}(f)$, can be expanded in any orthogonal basis for the class of square-integrable functions in a disk of radius c . For the purpose of steerable PCA, it is beneficial to choose a basis whose elements are products of radial functions with Fourier angular modes, such as the Fourier-Bessel functions, or 2-D prolate functions [35]. For concreteness, in the following we use the Fourier-Bessel functions given by

$$\psi_c^{k,q}(\xi, \theta) = \begin{cases} N_{k,q} J_k \left(R_{k,q} \frac{\xi}{c} \right) e^{ik\theta}, & \xi \leq c, \\ 0, & \xi > c, \end{cases} \quad (8)$$

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 . We also assume that the functions of interest are concentrated in a disk of radius R in real domain. In order to avoid aliasing, we only use Fourier-Bessel functions that satisfy the following criterion [26], [36]

$$R_{k,q+1} \leq 2\pi c R. \quad (9)$$

For each angular frequency k , we denote by p_k the number of components satisfying Eq. (9). The total number of components is $p = \sum_{k=-k_{\max}}^{k_{\max}} p_k$, where k_{\max} is the maximal possible

Algorithm 1: Steerable *ePCA* (*sePCA*) and denoising

Input: Image data Y that contains n images of size $L \times L$

Output: Rotationally invariant covariance estimator of noiseless images and denoised images

- 1 Compute the sample mean $\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i$ (**ePCA**)
 - 2 Estimate the support size R and band limit c for the mean image (**sPCA**)
 - 3 Compute the Fourier-Bessel expansion coefficients of $F(\bar{Y})$ and estimate the rotationally invariant sample mean \bar{f} as in Eq. (14) (**sePCA**)
 - 4 Compute the variance estimate $D_n = \text{diag}[\bar{f}]$ (**sePCA**)
 - 5 Prewhiten the image data $Z = D_n^{-1/2}Y$ (**ePCA**)
 - 6 Estimate the band limit c for whitened images (**sPCA**)
 - 7 Compute the truncated Fourier-Bessel expansion coefficients of $F(Z)$ and form the coefficients matrices $A^{(k)}$, for $k = 0, \dots, k_{\max}$ (**sPCA**)
 - 8 **for** $k = 0, 1, \dots, k_{\max}$ **do**
 - 9 Compute the prewhitened sample covariance matrix $S_h^{(k)}$ as in Eq. (16) and its eigendecomposition $S_h^{(k)} = \hat{U} \Lambda \hat{U}^*$ (**ePCA**)
 - 10 Shrink the eigenvalues $S_{h,\gamma_k}^{(k)} = \hat{U} \eta_{\gamma_k}(\Lambda) \hat{U}^*$ of top r_k eigenvalues according to Eq. (19) (**ePCA, sPCA**)
 - 11 Compute the recoloring matrix $B^{(k)}$ in eq. (23) and $D^{(k)}$ in eq. (26) (**sePCA**)
 - 12 Recolor the covariance matrix $S_{he}^{(k)} = (B^{(k)})^* \cdot S_{h,\gamma_k}^{(k)} \cdot B^{(k)}$ (**sePCA**)
 - 13 Compute the scaling coefficients $\hat{\alpha}$ in Eq. (27) and keep components with $\hat{\alpha} > 0$ (**sePCA**)
 - 14 Scale the covariance matrix $S_s^{(k)} = \sum \hat{\alpha}_i^{(k)} \hat{v}_i^{(k)} \left(\hat{v}_i^{(k)} \right)^*$, where the eigendecomposition of $S_{he}^{(k)}$ is $\sum \hat{v}_i^{(k)} \left(\hat{v}_i^{(k)} \right)^*$ (**ePCA**)
 - 15 Denoise $\{A^{(k)}\}_{k=0}^{k_{\max}}$ as in Eqs. (29) and (30) (**sePCA**)
 - 16 **end**
 - 17 The rotationally invariant covariance matrix estimator $\hat{S}((x, y), (x', y')) = G^{(0)}(x, y) S_s^{(0)} G^{(0)}(x', y')^* + 2 \sum_{k=1}^{k_{\max}} G^{(k)}(x, y) S_s^{(k)} G^{(k)}(x', y')^*$. (**sPCA**)
 - 18 Reconstruct the denoised image using Eq. (31) (**sPCA**)
-

value of k satisfying Eq. (9). We also denote $\gamma_k = \frac{p_k}{2n}$ for $k > 0$ and $\gamma_0 = \frac{p_0}{n}$.

The inverse Fourier transform (IFT) of $\psi_c^{k,q}$ is

$$\begin{aligned} \mathcal{F}^{-1}(\psi_c^{k,q})(r, \phi) &= \frac{2c\sqrt{\pi}(-1)^q R_{k,q} J_k(2\pi cr)}{i^k (2\pi cr)^2 - R_{k,q}^2} e^{ik\phi} \\ &\equiv g_c^{k,q}(r) e^{ik\phi}, \end{aligned} \quad (10)$$

where $g_c^{k,q}(r)$ is the radial part of the inverse Fourier transform of the Fourier-Bessel function. Therefore, we can approximate f using the truncated expansion

$$f(r, \phi) \approx \sum_{k=-k_{\max}}^{k_{\max}} \sum_{q=1}^{p_k} a_{k,q} g_c^{k,q}(r) e^{ik\phi}. \quad (11)$$

The approximation error in Eq. (11) is due to the finite truncation of the Fourier-Bessel expansion. For essentially band-limited functions, the approximation error is controlled using the asymptotic behavior of the Bessel functions, see [37, Section 2] for more details. We evaluate the Fourier-Bessel expansion coefficients numerically as in [26] using a quadrature rule that consists of equally spaced points in the angular direction $\theta_l = \frac{2\pi l}{n_\theta}$, with $l = 0, \dots, n_\theta - 1$ and a Gaussian quadrature rule in the radial direction ξ_j for $j = 1, \dots, n_\xi$ with the associated weights $w(\xi_j)$. Using the sampling criterion introduced in [26], the values of n_ξ and n_θ depend on the compact support radius R and the band limit c . Our previous work found that using $n_\xi = \lceil 4cR \rceil$ and $n_\theta = \lceil 16cR \rceil$ results in highly-accurate numerical evaluation of the integral to evaluate the expansion coefficients. To evaluate $a_{k,q}$, we need to sample the discrete Fourier transform of the image I , denoted $F(I)$, at the quadrature nodes,

$$F(I)(\xi_j, \theta_l) = \frac{1}{2R} \sum_{i_1=-R}^{R-1} \sum_{i_2=-R}^{R-1} I(i_1, i_2) \times \exp(-i2\pi(\xi_j \cos \theta_l i_1 + \xi_j \sin \theta_l i_2)), \quad (12)$$

which can be evaluated efficiently using the the nonuniform discrete Fourier transform [38], and we get

$$a_{k,q} \approx \sum_{j=1}^{n_\xi} N_{k,q} J_{k,q} \left(\frac{\xi_j}{c} \right) \widehat{F(I)}(\xi_j, k) \xi_j w(\xi_j), \quad (13)$$

where $\widehat{F(I)}(\xi_j, k)$ is the 1D FFT of $F(I)$ on each concentric circle of radius ξ_j . For real-valued images, it is sufficient to evaluate the coefficients with $k \geq 0$, since $a_{-k,q} = a_{k,q}^*$. In addition, the coefficients have the following properties: under counter-clockwise rotation by an angle α , $a_{k,q}$ changes to $a_{k,q} e^{-ik\alpha}$; and under reflection, $a_{k,q}$ changes to $a_{-k,q}$. The numerical integration error in Eq. (13) was analyzed in [26] and drops below 10^{-17} for the chosen values of n_ξ and n_θ .

The steerable basis expansion are applied in two parts of the steerable ePCA: (1) the rotationally invariant sample mean estimation in Step 3 of Alg. 1 and (2) the expansion of the whitened images in Step 7 of Alg. 1.

C. The sample rotationally-invariant homogenized covariance matrix

Suppose I_1, \dots, I_n are n discretized input images sampled from f_1, \dots, f_n . Here, the observation vectors Y_i for ePCA are simply $Y_i = I_i$. In ePCA, the first step is to prewhiten the data using the sample mean as suggested by Eqs. (4) and (6). However, the sample mean $\bar{Y} = \frac{1}{n} \sum_{i=1}^n Y_i$ is not necessarily rotationally invariant. With the estimated band limit and support size in Step 2, we compute the truncated Fourier-Bessel expansion coefficients of $F(\bar{Y})$, denoted by $\bar{a}_{k,q}$. The rotationally invariant sample mean can be evaluated from $\bar{a}_{k,q}$,

$$\bar{f}(r, \phi) = \frac{1}{2\pi} \int_0^{2\pi} \frac{1}{n} \sum_{i=1}^n f_i(r, \phi - \alpha) d\alpha \approx \sum_{q=1}^{p_0} \bar{a}_{0,q} g_c^{0,q}(r). \quad (14)$$

The approximation error in Eq. (14) follows directly from that of Eq. (11). The rotationally invariant sample mean is circularly symmetric. We denote by \bar{A} a vector that contains all the coefficients $\bar{a}_{0,q}$ ordered by the radial index q . Although the input images are non-negative, the finite truncation may result in small negative values in the estimated mean, so we threshold any negative entries to zero. As mentioned in Sec. II-A, the expected covariance matrix of the Poisson observations differs from the covariance matrix of clean data by a diagonal matrix, where the diagonal entries are equal to the mean image. Therefore, in Step 4 of Alg. 1, we have $D_n = \text{diag}[\bar{f}]$ and it is used in Step 5 to compute the homogenized vectors, similar to Eq. (4).

We prewhiten the images by the estimated mean image to create new images Z_1, \dots, Z_n as $Z_i(x, y) = \bar{f}(x, y)^{-1/2} Y_i(x, y)$, when $\bar{f}(x, y) > 0$, and $Z_i(x, y) = 0$ otherwise. The whitening step might change the band limit. Therefore, we estimate the band limit of the whitened images in Step 6. Combining Eqs. (9) and (13), we compute the truncated Fourier-Bessel expansion coefficients $a_{k,q}^i$ of $F(Z_i)$. Let us denote by $A^{(k)}$ the matrix of expansion coefficients with angular frequency k , obtained by putting $a_{k,q}^i$ into a matrix, where the columns are indexed by the image number i and the rows are ordered by the radial index q . The coefficient matrix $A^{(k)}$ is of size $p_k \times n$.

We use \bar{Z} to represent the rotationally invariant sample mean of the whitened images. Under the action of the group $O(2)$, i.e. counter-clock wise rotation by an angle $\alpha \in [0, 2\pi)$ and reflection $\beta \in \{+, -\}$, where ‘+’ indicates no reflection and ‘-’ indicates with reflection, the image Z_i is transformed to $Z_i^{\alpha, \beta}$. Since the truncated Fourier-Bessel transform is almost unitary [26], the rotationally invariant covariance kernel $\mathcal{S}((x, y), (x', y'))$ built from the whitened image data with all possible in-plane rotations and reflections, defined as,

$$\mathcal{S}((x, y), (x', y')) = \frac{1}{4\pi n} \sum_{i=1}^n \sum_{\beta \in \{+, -\}} \int_0^{2\pi} \left(Z_i^{\alpha, \beta}(x, y) - \bar{Z}(x, y) \right) \left(Z_i^{\alpha, \beta}(x', y') - \bar{Z}(x', y') \right) d\alpha, \quad (15)$$

can be computed in terms of the IFT 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 $a_{0,q}^i \leftarrow a_{0,q}^i - \frac{1}{n} \sum_{j=1}^n a_{0,q}^j$. In terms of the expansion coefficients, the rotationally invariant homogenized sample covariance matrix is $S_h = \bigoplus_{k=-k_{\max}}^{k_{\max}} S_h^{(k)}$, with

$$S_h^{(k)} = \frac{1}{n} \text{Re} \left\{ A^{(k)} \left(A^{(k)} \right)^* \right\}. \quad (16)$$

We further denote the eigenvalues and eigenvectors of $S_h^{(k)}$ by $\lambda_i^{(k)}$ and $\hat{u}_i^{(k)}$, that is,

$$S_h^{(k)} = \sum_{i=1}^{p_k} \lambda_i^{(k)} \hat{u}_i^{(k)} (\hat{u}_i^{(k)})^*. \quad (17)$$

The procedure of homogenization for steerable ePCA is detailed in Steps 1–5 in Alg. 1. Step 9 of Alg. 1 computes the rotationally invariant homogenized sample covariance matrix.

D. Eigenvalue shrinkage

For data corrupted by additive white noise (with noise variance 1 in each coordinate), previous work [39]–[42] showed that if the population eigenvalue ℓ is above the Baik, Ben Arous, Péché (BBP) phase transition, then the corresponding sample eigenvalue lies to the right of the Marčenko-Pastur distribution of the “noise” eigenvalues. The sample eigenvalue will converge to the value given by *the spike forward map* as $p_k, n \rightarrow \infty$, and $p_k/n = \gamma_k$:

$$\lambda(\ell; \gamma_k) = \begin{cases} (1 + \ell) \left(1 + \frac{\gamma_k}{\ell}\right) & \text{if } \ell > \sqrt{\gamma_k}, \\ (1 + \sqrt{\gamma_k})^2 & \text{otherwise.} \end{cases} \quad (18)$$

The underlying clean population covariance eigenvalues can be estimated by solving the quadratic equation in Eq. (18),

$$\hat{\ell} = \eta_{\gamma_k}(\lambda) = \begin{cases} \frac{(\lambda - 1 - \gamma_k + \sqrt{(\lambda - 1 - \gamma_k)^2 - 4\gamma_k})}{2}, & \lambda > (1 + \sqrt{\gamma_k})^2, \\ 0 & \lambda \leq (1 + \sqrt{\gamma_k})^2. \end{cases} \quad (19)$$

Shrinking the eigenvalues improves the estimation of the sample covariance matrix [43]. Since the homogenized sample covariance matrix S_h is decoupled into small sub-blocks $S_h^{(k)}$, the shrinkers are defined for each frequency k separately. The shrinkers $\eta_{\gamma_k}(\lambda)$ set all noise eigenvalues to zero for λ within the support of the Marčenko-Pastur distribution and reduce other eigenvalues according to Eq. (19). Then the denoised covariance matrices are

$$S_{h,\eta}^{(k)} = \sum_{i=1}^{r_k} \eta_{\gamma_k}(\lambda_i^{(k)}) \hat{u}_i^{(k)} (\hat{u}_i^{(k)})^*, \quad (20)$$

where r_k is the number of components with $\eta_{\gamma_k}(\lambda) > 0$. The empirical eigenvector $\hat{u}^{(k)}$ of $S_h^{(k)}$ is an inconsistent estimator of the true eigenvector. We can heuristically quantify the inconsistency based on results from the Gaussian standard spiked model, even though the noise is non-Gaussian. Under this model, the empirical and true eigenvectors have an asymptotically deterministic angle: $((u^{(k)})^* \hat{u}^{(k)})^2 \rightarrow c^2(\ell; \gamma_k)$ almost surely, where $c(\ell; \gamma_k)$ is the *cosine forward map* given by [41], [42]:

$$c^2(\ell; \gamma_k) = \begin{cases} \frac{1 - \gamma_k/\ell^2}{1 + \gamma_k/\ell} & \text{if } \ell > \sqrt{\gamma_k}, \\ 0 & \text{otherwise.} \end{cases} \quad (21)$$

Therefore, asymptotically for the population eigenvectors beyond the BBP phase transition, the sample eigenvectors have positive correlation with the population eigenvectors, but this correlation is less than 1 [44]–[46]. We denote by \hat{c} an estimate of c using the estimated clean covariance eigenvalues $\hat{\ell}$ in Eq. (19) and $\hat{s}^2 = 1 - \hat{c}^2$.

In short, Step 10 of Alg. 1 improves the estimation of the rotationally invariant homogenized covariance matrix through eigenvalue shrinkage.

E. Recoloring

Homogenization changes the direction of the clean eigenvectors. Therefore, after eigenvalue shrinkage, we recolor (heterogenize) the covariance matrix by conjugating the recoloring

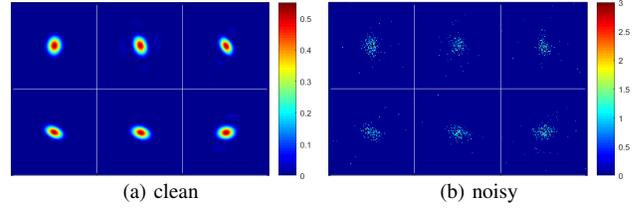


Fig. 1: Sample clean and noisy images of the XFEL dataset. Image size is 128×128 with mean per pixel photon count = 0.01.

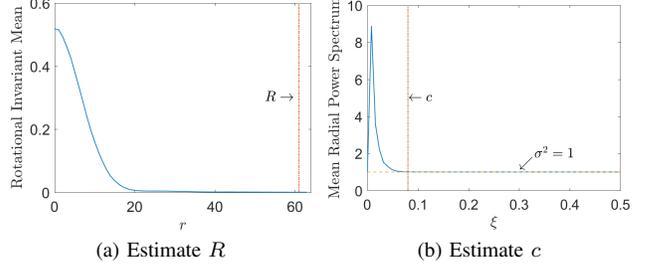


Fig. 2: Estimating R and c from $n = 7 \times 10^4$ simulated noisy XFEL diffraction intensity maps of lysozyme. Each image is of size 128×128 pixels. (a) The radial profile of the rotationally invariant sample mean image. The radius of compact support is chosen at $R = 61$. (b) Mean radial power spectrum of the whitened noisy images. The curve levels off at $\sigma^2 = 1$.

matrix B with $S_{h,\eta}$: $S_{he} = B^* \cdot S_{h,\eta} \cdot B$. The recoloring matrix is derived as,

$$\begin{aligned} B_{k_1, q_1; k_2, q_2} &= \int_0^R \int_0^{2\pi} \sqrt{\bar{f}(r)} \mathcal{F}^{-1}(\psi_c^{k_1, q_1})(r, \theta) \\ &\quad \times \left(\mathcal{F}^{-1}(\psi_c^{k_2, q_2})(r, \theta) \right) r dr d\theta \\ &= \int_0^R \sqrt{\bar{f}(r)} g_c^{k_1, q_1}(r) \overline{g_c^{k_2, q_2}(r)} r dr \\ &\quad \times \int_0^{2\pi} e^{i(-k_1 + k_2)\theta} d\theta \\ &= \delta_{k_1, k_2} \int_0^R \sqrt{\bar{f}(r)} g_c^{k_1, q_1}(r) \overline{g_c^{k_2, q_2}(r)} r dr, \end{aligned} \quad (22)$$

which has a block diagonal structure and is decoupled for each angular frequency, $B = \bigoplus_{k=-k_{\max}}^{k_{\max}} B^{(k)}$, with

$$B_{q_1, q_2}^{(k)} = \int_0^R \sqrt{\bar{f}(r)} g_c^{k, q_1}(r) \overline{g_c^{k, q_2}(r)} r dr. \quad (23)$$

The radial integral in Eq. (23) is numerically evaluated using the Gauss-Legendre quadrature rule [47, Chap. 4], which determines the locations of $n_r = \lceil 4cR \rceil$ points $\{r_j\}_{j=1}^{n_r}$ on the interval $[0, R]$ and the associated weights $w(r_j)$. The integral in Eq. (23) is thus approximated by

$$B_{q_1, q_2}^{(k)} \approx \sum_{j=1}^{n_r} \sqrt{\bar{f}(r_j)} g_c^{k, q_1}(r_j) \overline{g_c^{k, q_2}(r_j)} r_j w(r_j). \quad (24)$$

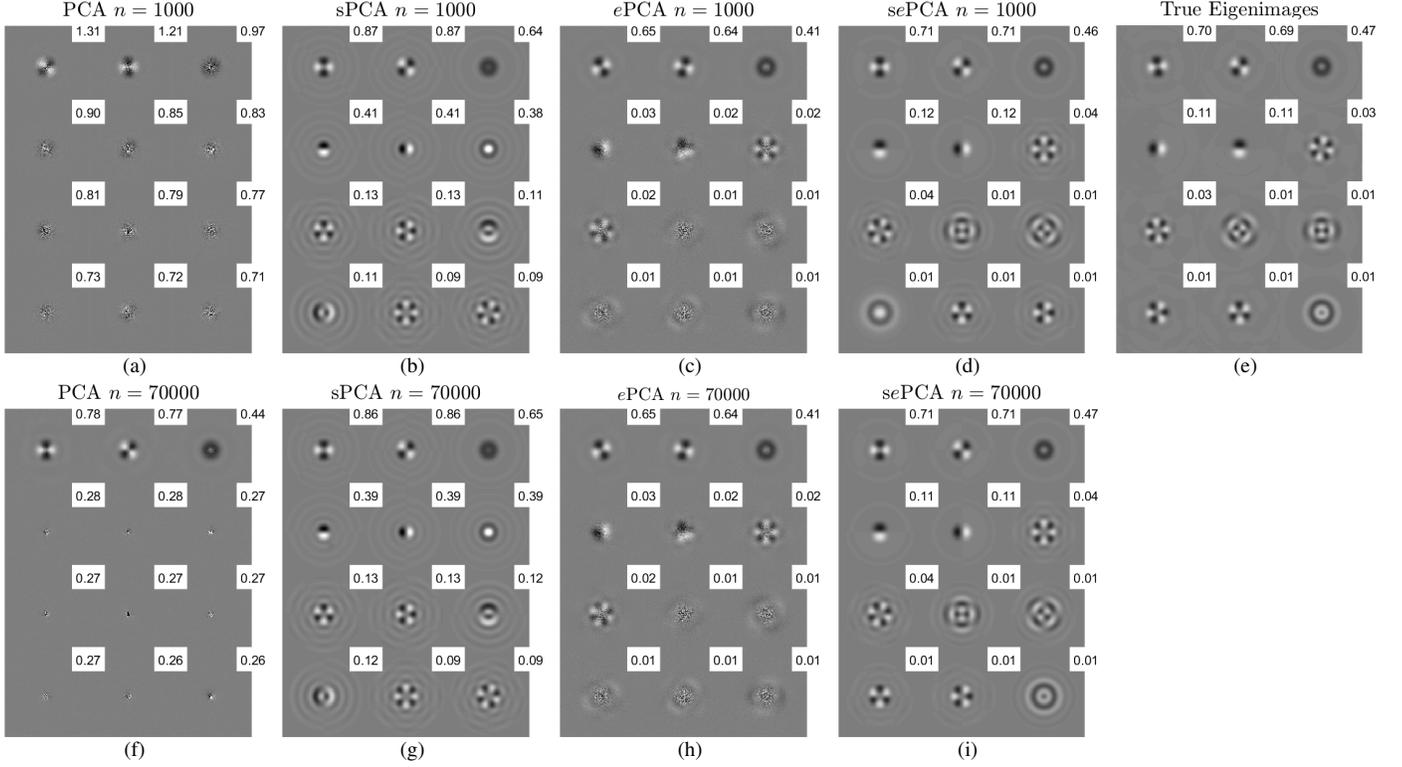


Fig. 3: Eigenimages estimated from noisy XFEL data using PCA, steerable PCA (sPCA), e PCA, and steerable e PCA (sePCA), ordered by eigenvalues. Input images are corrupted by Poisson noise with mean photon count 0.01 (shown in Figure 1b).

The recoloring step is also decoupled for each angular frequency sub-block. The heterogenized covariance estimators are

$$S_{he}^{(k)} = \left(B^{(k)} \right)^* \cdot S_{h,\eta_{\gamma_k}}^{(k)} \cdot B^{(k)}. \quad (25)$$

Similar to Eq. (23), we define $D^{(k)}$ which will be used to scale the heterogenized covariance matrix estimator (see Eq. (27)) and denoise the expansion coefficients (see Eqs. (29) and (30)),

$$D_{q_1, q_2}^{(k)} = \int_0^R \bar{f}(r) g_c^{k, q_1}(r) \overline{g_c^{k, q_2}(r)} r dr \\ \approx \sum_{j=1}^{n_r} \bar{f}(r_j) g_c^{k, q_1}(r_j) \overline{g_c^{k, q_2}(r_j)} r_j w(r_j). \quad (26)$$

In Alg. 1, Steps 11 and 12 summarize the recoloring procedure.

F. Scaling

The eigendecomposition of $S_{he}^{(k)}$ gives $S_{he}^{(k)} = \sum_{i=1}^{r_k} \hat{v}_i^{(k)} \left(\hat{v}_i^{(k)} \right)^*$. The empirical eigenvalues are $\hat{t} = \|\hat{v}^{(k)}\|^2$ which is a biased estimate of the true eigenvalues of the clean covariance matrix Σ_X . In [25, Sec. 4.2.3], a scaling rule was proposed to correct the bias. We extend it in Steps 13 and 14 in Alg. 1 to the steerable case and scale each eigenvalue of $S_{he}^{(k)}$ by a parameter $\hat{\alpha}^{(k)}$,

$$\hat{\alpha}^{(k)} = \begin{cases} \frac{1 - \hat{s}^2 \tau^{(k)}}{\hat{c}^2}, & \text{for } 1 - \hat{s}^2 \tau^{(k)} > 0 \text{ and } \hat{c}^2 > 0 \\ 0 & \text{otherwise} \end{cases} \quad (27)$$

where the parameter $\tau^{(k)} = \frac{\text{tr} D^{(k)}}{p_k} \cdot \frac{\hat{\ell}}{\|\hat{v}^{(k)}\|^2}$. The scaled covariance matrices are

$$S_s^{(k)} = \sum_{i=1}^{r_k} \hat{\alpha}_i^{(k)} \hat{v}_i^{(k)} \left(\hat{v}_i^{(k)} \right)^*. \quad (28)$$

The rotationally invariant covariance kernel $S((x, y), (x', y'))$ is well approximated by $\sum_{k=-k_{\max}}^{k_{\max}} G^{(k)}(x, y) S_s^{(k)} \left(G^{(k)}(x', y') \right)^*$, where $G^{(k)}$ contains IFT of all $\psi_c^{k, q}$ with angular frequency k (see Step 17) in Alg. 1). The computational complexity of steerable e PCA is $O(nL^3 + L^4)$, same as steerable PCA, and it is lower than the complexity of e PCA which is $O(\min(nL^4 + L^6, n^2L^2 + n^3))$.

In summary, Step 14 scales the heterogenized covariance matrix S_{he} . The covariance matrix in the original pixel domain is efficiently computed from $S_s^{(k)}$ in Step 17.

G. Denoising

As an application of steerable e PCA, we develop a method to denoise the photon-limited images. Since the Fourier-Bessel expansion coefficients are computed from the prewhitened images, we first recolor the coefficients by multiplying $B^{(k)}$ with $A^{(k)}$ and then we apply Wiener-type filtering to denoise the coefficients. For the steerable basis expansion coefficients $A^{(k)}$ with angular frequency $k \neq 0$,

$$\hat{A}^{(k)} = S_s^{(k)} \left(D^{(k)} + S_s^{(k)} \right)^{-1} B^{(k)} A^{(k)}. \quad (29)$$

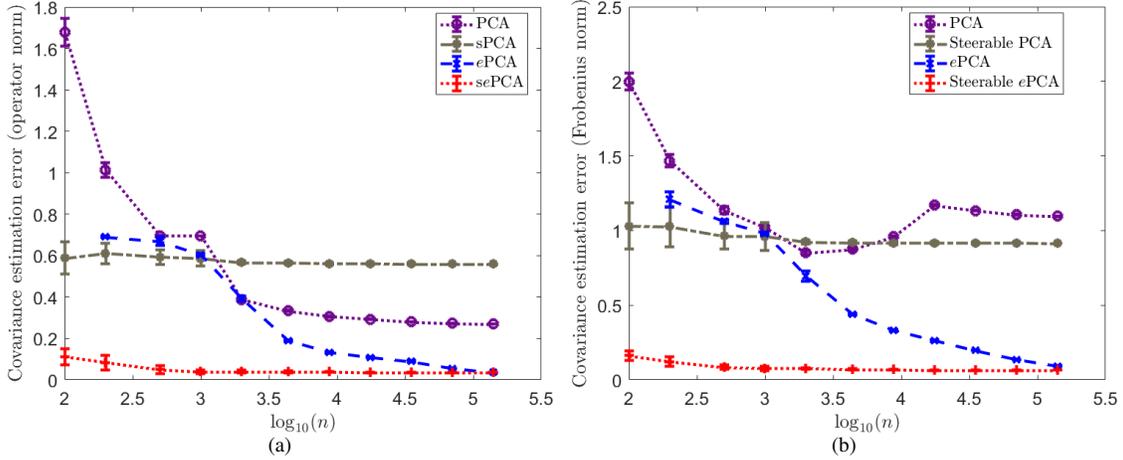


Fig. 4: Error of covariance matrix estimation, measured as the (a) operator norm and (b) Frobenius norm of the difference between each covariance estimate and the true covariance matrix. The sample size n ranges from 100 to 140,000.

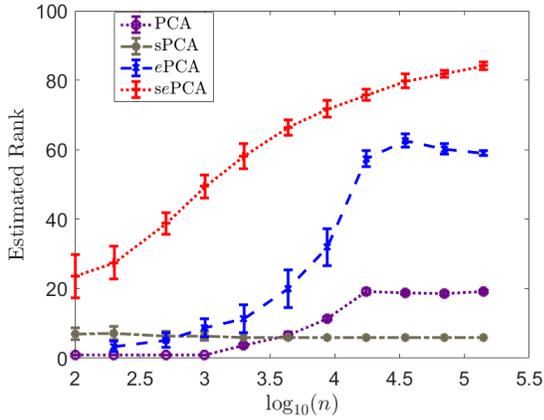


Fig. 5: The estimated number of signal principal components using PCA, sPCA, ePCA and sePCA. Images are corrupted by Poisson noise with mean per pixel photon count 0.01.

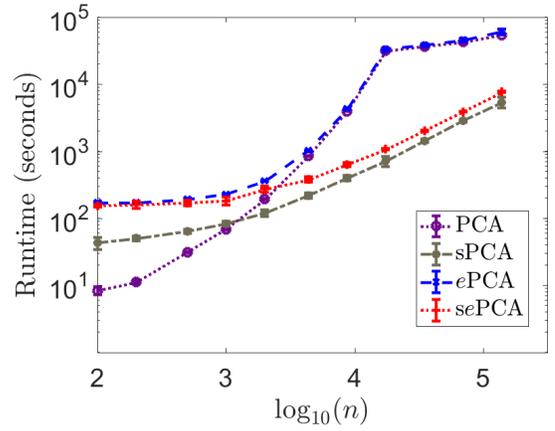


Fig. 6: The runtime for computing the principal components using PCA, sPCA, ePCA and sePCA. Images are corrupted by Poisson noise with mean per pixel photon count = 0.01.

For $k = 0$, we need to take into account the rotationally invariant mean expansion coefficients,

$$\hat{A}^{(0)} = S_s^{(0)} \left(D^{(0)} + S_s^{(0)} \right)^{-1} B^{(0)} A^{(0)} + D^{(0)} \left(D^{(0)} + S_s^{(0)} \right)^{-1} \bar{\mathbf{A}} \mathbf{1}_n^\top. \quad (30)$$

The denoised image sampled on the Cartesian grid (x, y) in real domain are computed from the filtered expansion coefficients $\hat{a}_{k,q}^i$,

$$\hat{X}_i(x, y) = \sum_{q=1}^{p_0} \hat{a}_{0,q}^i g_c^{0,q}(r_{x,y}) + 2\text{Re} \left[\sum_{k=1}^{k_{\max}} \sum_{q=1}^{p_k} \hat{a}_{k,q}^i g_c^{k,q}(r_{x,y}) e^{-ik\theta_{x,y}} \right], \quad (31)$$

where $r_{x,y} = \sqrt{x^2 + y^2}$ and $\theta_{x,y} = \tan^{-1} \left(\frac{y}{x} \right)$.

In essence, we first denoise the recolored steerable expansion coefficients in Step 15 of Alg. 1 according to Eqs. (29) and (30) and then reconstruct the image using Eq. (31) in Step 18 of Alg. 1.

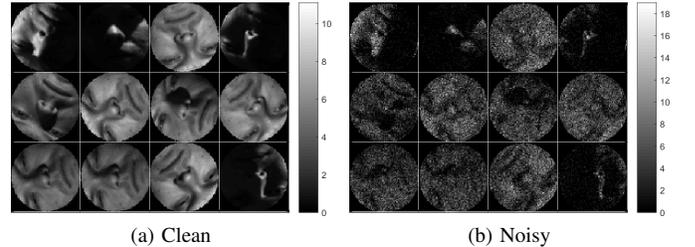


Fig. 7: Sample clean and noisy images of randomly rotated Yale Face Database B. Image size is 64×64 with mean intensity 2.3 photons per pixel.

III. NUMERICAL RESULTS

We apply PCA, ePCA, steerable PCA, and steerable ePCA to a simulated XFEL dataset and compare the results for covariance estimation and denoising. The algorithms are implemented in MATLAB on a machine with 60 cores, running at 2.3 GHz, with total RAM of 1.5TB. Only 8 cores were used

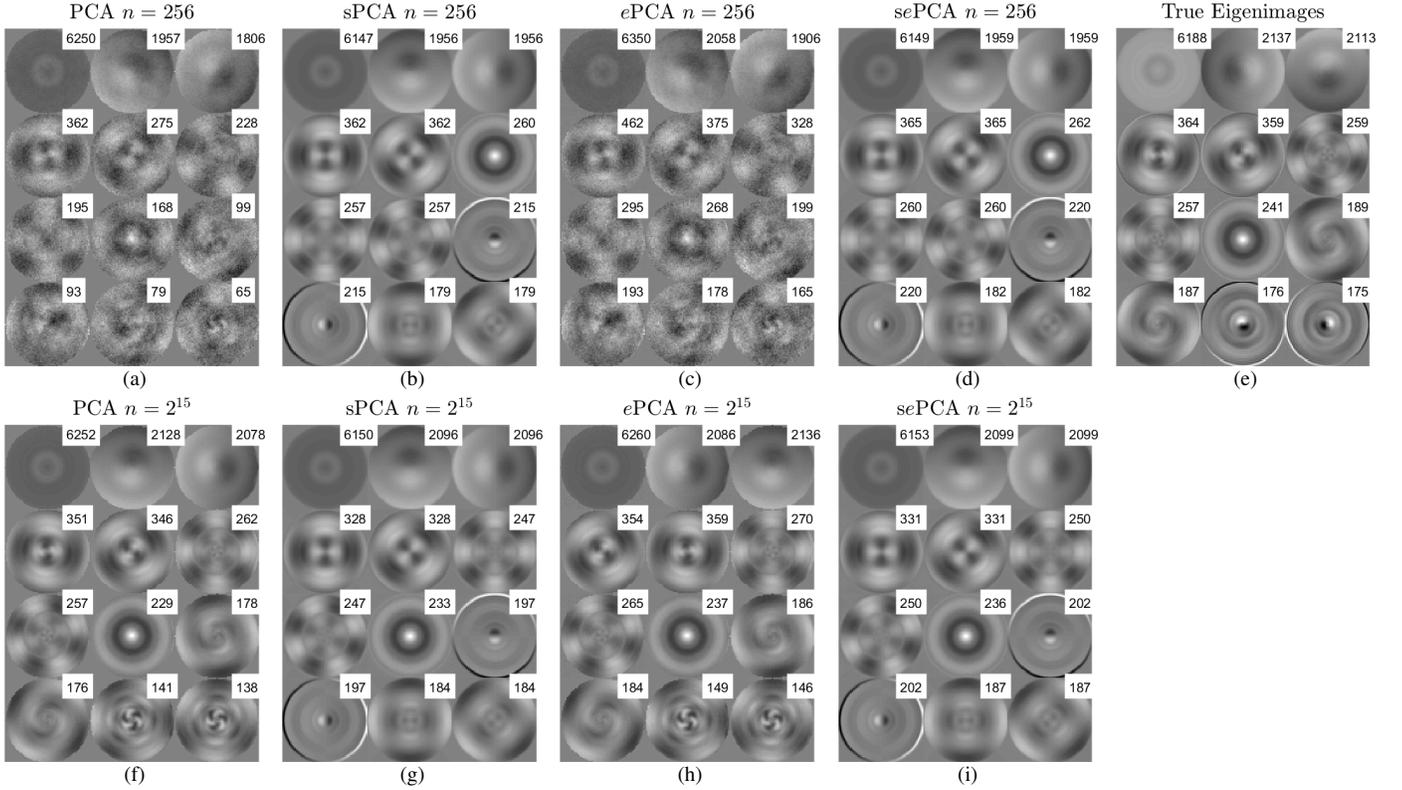


Fig. 8: Eigenimages estimated from noisy rotated Yale Face Database B using PCA, sPCA, ePCA, and sePCA, ordered by eigenvalues. Input images are corrupted by Poisson noise with mean intensity 2.3 photons per pixel (shown in Fig. 7b).

in our experiments.

We simulate $n = 140,000$ noiseless XFEL diffraction intensity maps of a lysozyme (Protein Data Bank 1AKI) with Condor [48]. The average pixel intensity is rescaled to be 0.01 for image size 128×128 pixels such that shot noise dominates [49]. To sample an arbitrary number n of noisy diffraction patterns, we sample an intensity map at random, and then sample the photon count of each detector pixel from a Poisson distribution whose mean is the pixel intensity. Figs. 1a and 1b illustrate the clean intensity maps and the resulting noisy diffraction patterns.

We estimate the radius of the concentration (R) of the diffraction intensities in real domain and the band limit (c) in Fourier domain from the noisy images in the following way. The data variance is proportional to the sample mean. The region with non-zero rotationally invariant sample mean contains the object of interest. We show the radial profile of the estimated sample mean in Fig. 2a, which is also the variance map of the dataset averaged in the angular direction. At large r , the radial part of the rotational invariant sample mean levels off at 0. We compute the cumulative variance by integrating the radial sample mean over r with a Jacobian weight rdr . For the XFEL data, the fraction of the cumulative mean exceeds $\eta = 99.9\%$ at $r = 61$, and therefore R was chosen to be 61 (see Fig. 2a). We compute the whitened projection images using the rotationally invariant sample mean. To estimate the parameter c for the whitened images, we compute the angular average of the mean 2D power spectrum. With a large number of images, a region in the Fourier domain contains signal information, if the corresponding radial mean power spectrum values are larger than the noise variance $\sigma^2 = 1$. The curve in Fig. 2b

levels off at the noise variance $\sigma^2 = 1$ when ξ is large. We use the same method as before to compute the cumulative radial power spectrum. The fraction reaches $\eta = 99.9\%$ at $\xi = 0.08$, therefore the band limit is chosen to be $c = 0.08$. The band limit c and support radius R are used in both steerable PCA and steerable ePCA. Both parameters can be chosen by the user through either setting the values of R and c directly, or controlling the threshold value η . The proposed parameter selection procedure leads to good empirical results for covariance estimation. The impact of different choices of c and R on denoising is illustrated in Fig. 13.

A. Covariance estimation and principal components

Fig. 3 shows the top 12 eigenimages for clean XFEL diffraction patterns (Fig. 3e), and noisy diffraction patterns with mean photon count per pixel 0.01 (Figs. 3c–3i) using PCA, steerable PCA, ePCA, and steerable ePCA. The true eigenimages in Fig. 3e are computed from 70000 clean diffraction patterns whose orientations are uniformly distributed over $SO(3)$. Figs. 3a–3d are computed from 1000 noisy images. Since the number of samples is much smaller than the size of the image and the noise type is non-Gaussian, PCA can only recover the first two or three components. ePCA improves the estimation and is able to extract the top 7 eigenimages. Moreover, steerable PCA and steerable ePCA achieve much better estimation of the underlying true eigenimages for a given sample size. Steerable ePCA achieves the best performance in estimating both the eigenvalues and eigenimages.

Furthermore, we compare the operator norms and Frobenius norms of the difference between the covariance estimates and

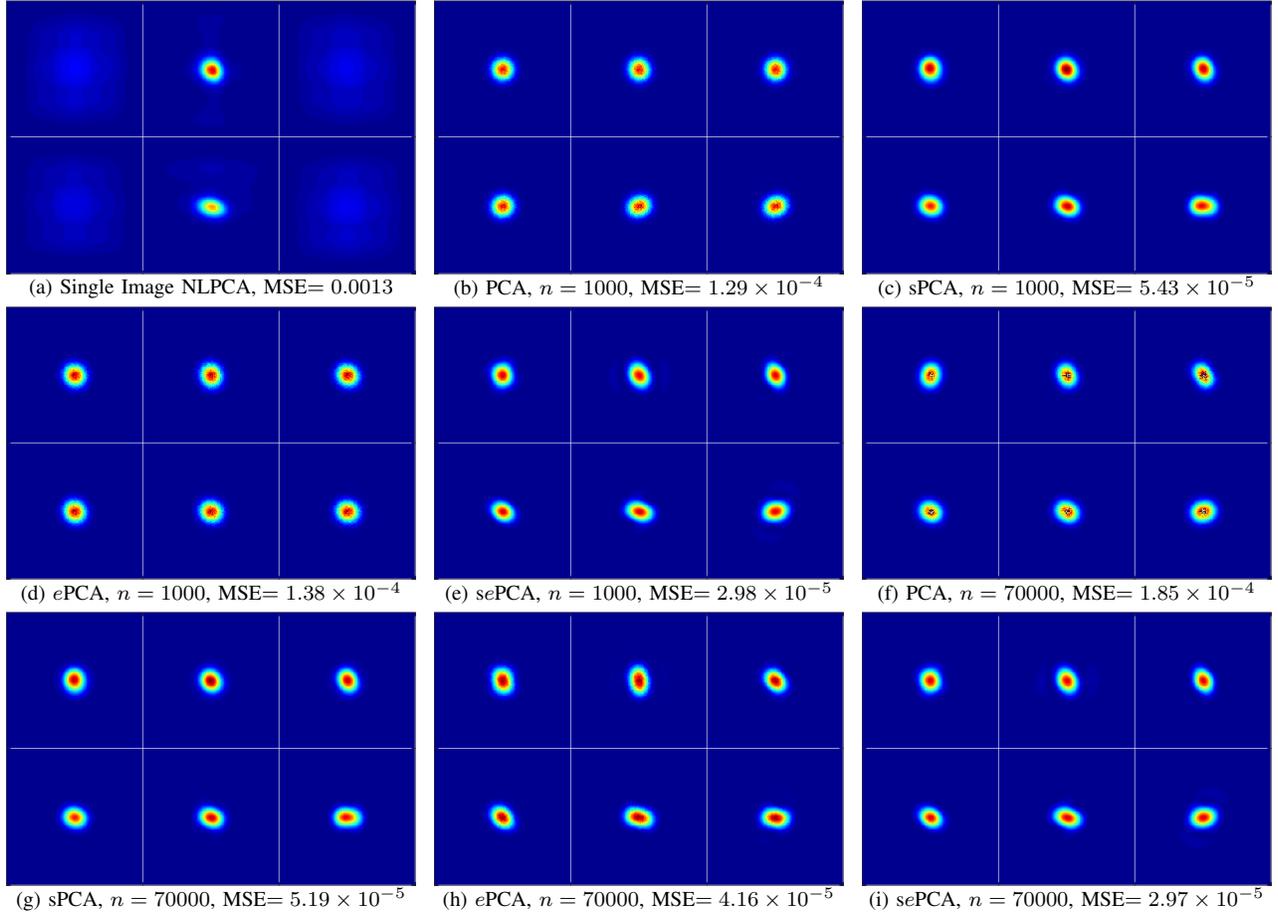


Fig. 9: Sample denoised images of the XFEL dataset illustrated in Fig. 1 using NLPCA, PCA, sPCA, ePCA, sePCA. Image size is 128×128 .

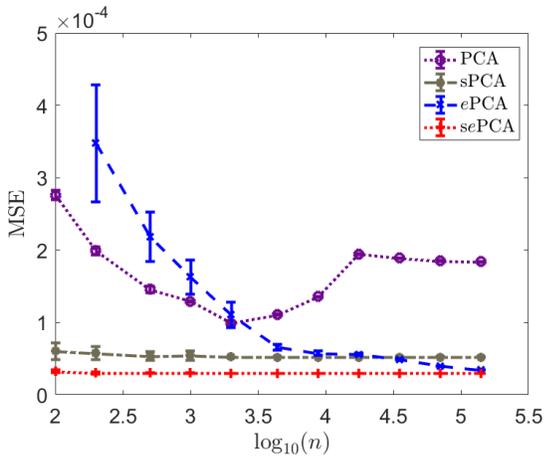


Fig. 10: Comparing the denoising quality of the XFEL dataset ($p = 16384$) with various number of images.

the true covariance matrix. Fig. 4 shows that steerable ePCA significantly improves the covariance estimation, especially when the sample size is small.

For experiments using ePCA, we use a permutation bootstrap-based method to estimate the rank of the covariance matrix, following e.g. [50]. By randomly permuting each

column of the mean-subtracted data matrix, we completely destroy structural information including linear structure, while the noise statistics remain the same (see [51], [52] for an analysis). Singular values of the randomly permuted matrices reveal what should be the largest covariance matrix eigenvalues that correspond to noise, up to a user-selected confidence level ρ . This can replace the other rank estimation methods that assume Gaussianity when the noise model is non-Gaussian, such as in our case. Empirically, we observe that $\rho = 0.1$ gives the best performance in covariance estimation. For steerable ePCA, we estimate the number of components using the right edge of the Marčenko-Pastur distribution for homogenized covariance matrices $S_h^{(k)}$ and include only the components whose scaling factor $\hat{\alpha}^{(k)}$ are above zero.

Steerable ePCA is able to recover more signal principal components from noisy images than PCA, steerable PCA, and ePCA (see Fig. 5). When the sample size $n = 1000$, the mean number of estimated components is 11 and 49 for ePCA and steerable ePCA respectively. For $n = 70000$, the estimated number of components is 59 and 81 for ePCA and steerable ePCA respectively. Fig. 6 shows that steerable ePCA is more efficient than ePCA and PCA. Because steerable ePCA contains extra steps such as prewhitening, recoloring, and scaling, its runtime is slightly longer than steerable PCA.

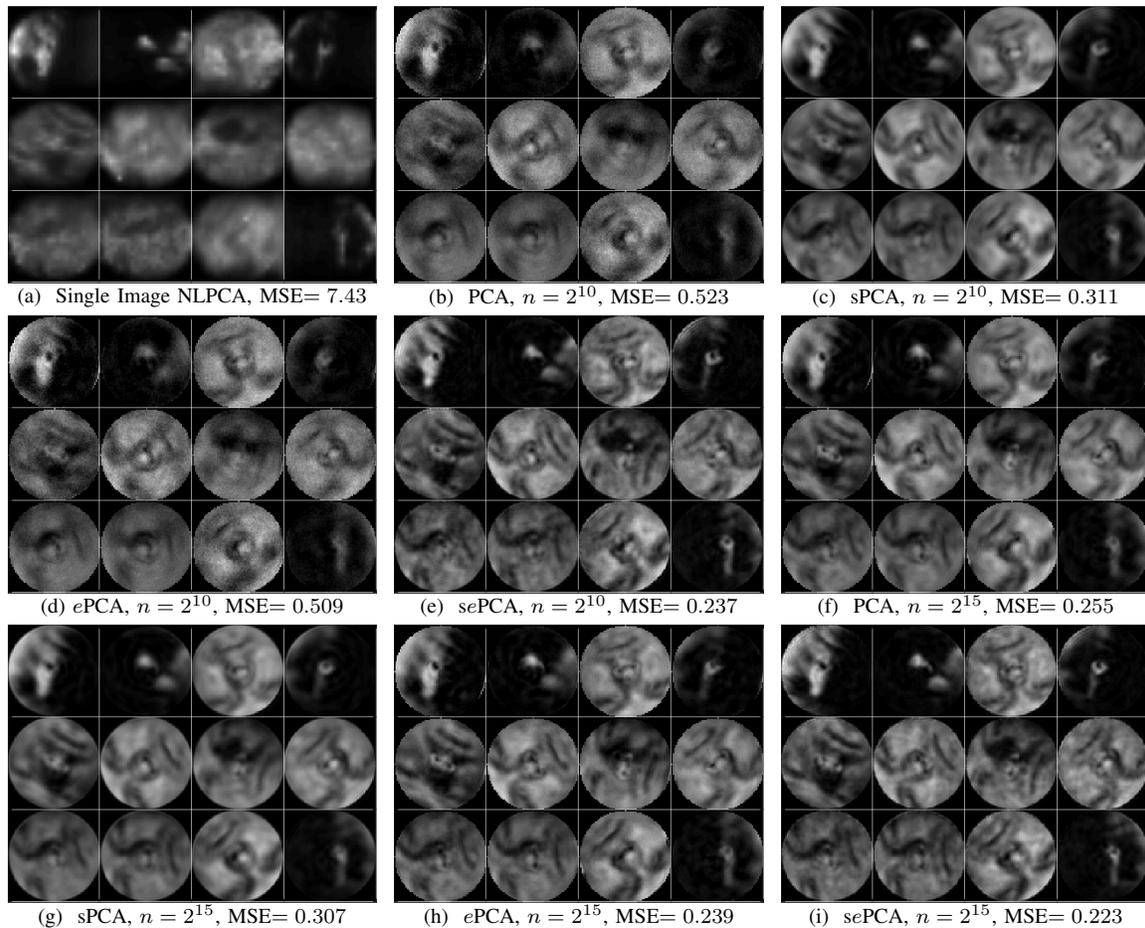


Fig. 11: Sample denoised images from Yale Face Database B illustrated in Fig. 7 using NLPCA, PCA, sPCA, ePCA, and sePCA. Image size is 64×64 . The intensity range is $[0, 2.6]$ for (a), and is $[0, 11.1]$ for (b)–(i).

When $n = 140000$, steerable ePCA is 8 times faster than ePCA.

In addition to the XFEL diffraction intensity data, we include a natural image dataset—Yale Face Database B [33], [34]—to illustrate the efficacy of the proposed method. The database contains 5760 single light source images of 10 subjects with 9 poses. For every subject in a particular pose, 64 images were captured under different ambient illuminations. The original images of a single face under different lighting conditions inhabit an approximately 9-dimensional linear space [53]. In the experiments, we take one subject at a particular pose with 64 different illumination conditions. We downsample the original images to the size of 64×64 pixels and scale the intensity such that the mean photon count for the whole dataset is 2.3 photons per pixel (i.e., “faces in the dark”). Since we need to rotate the images, we mask images by a disk of radius $R = 32$. We uniformly sample n clean images from the original images and apply a random in-plane rotation to each sample. Fig. 7 shows 12 samples of the clean data and the corresponding noisy observations. The corresponding principal components are illustrated in Fig. 8. The true eigenimages are evaluated using the clean images rotated at every 0.7 degrees.

B. Denoising

We compare the denoising effects of steerable ePCA with PCA, steerable PCA, ePCA, and patch-based single image non-local PCA (NLPCA) [21], by the mean squared error, $\text{MSE} := (pn)^{-1} \sum_{i=1}^n \|\hat{X}_i - X_i\|^2$. We perform “ePCA denoising” using the empirical best linear predictor (EBLP) [25], which had been shown to outperform “PCA denoising,” i.e., orthogonal projection onto sample or ePCA / heterogenized eigenimages, as well as the exponential family PCA method proposed by [19]. Note that in our implementation of ePCA Wiener-type filtering, to avoid inverting a singular matrix (when some coordinates have 0 sample mean and sample variance), we compute $\text{diag}[\bar{Y}] + S_s$ with regularization, $\text{diag}[\bar{Y}] + S_s \leftarrow S_s + (1 - \epsilon) \text{diag}[\bar{Y}] + \epsilon m I$ where $\epsilon = 0.1$, $m = \frac{1}{p} \bar{Y} \cdot \mathbf{1}$, and I is the $p \times p$ identity matrix. The number of components are estimated by the permutation rank estimation as described in the previous section.

Fig. 9 shows some examples of denoised XFEL images for sample size $n = 1000$ and 70000. For robustness, we repeated the numerical experiment for another dataset simulated from the small protein chignolin (Protein Data Bank entry 1UAO) and obtained qualitatively similar results. Steerable ePCA is able to recover the images with lower MSEs compared to PCA, steerable PCA, and ePCA (see Fig. 10), especially when the sample size n is small.

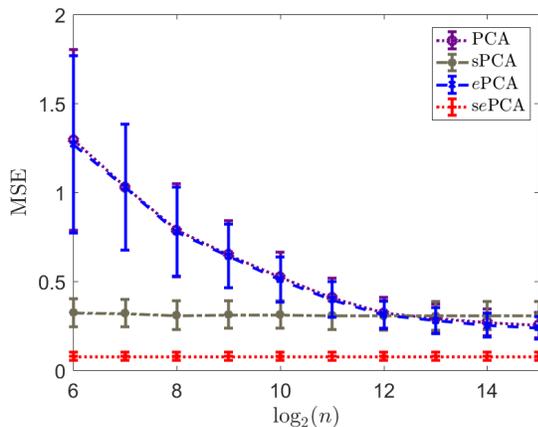


Fig. 12: Comparing the denoising quality of the Yale Face Database B ($p = 4096$) with various number of images.

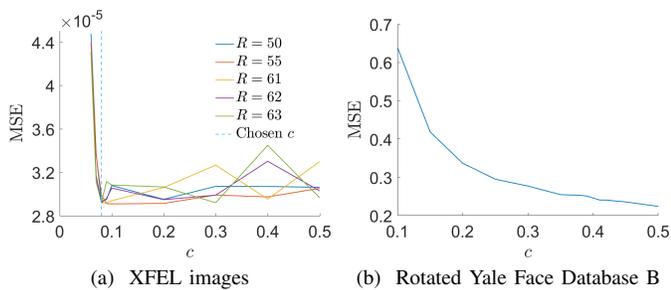


Fig. 13: Sensitivity of the denoising performance over the parameters c and R . For Yale Face Database B, we cropped the images by the disc of radius $R = 32$.

Fig. 11 shows examples of the denoised face images for sample size $n = 2^{10}$ and 2^{15} . We use the permutation rank estimation described in the previous section with confidence level $\rho = 0.1$ for PCA, sPCA, and ePCA. In this example, we see that when the number of images is small, PCA and ePCA can only recover rough contours of the faces, whereas sPCA and sePCA are able to recover finer details in the faces. The MSEs of the denoised images with varying number of samples are plotted in Fig. 12. In this experiment, the data only contains 64 unrotated clean images and the mean photon count is 2.3 per pixel. Compared to the diffraction patterns with mean photon count 0.01 per pixel. Therefore, sePCA is able to sufficiently capture the clean data subspace with only a small number of images, resulting in a much flatter MSE curve with respect to the number of images.

We also test how sensitive our method is to the choice of parameters, i.e. the band limit c and the support radius R . For the diffraction intensity data, we use $n = 10000$ diffraction patterns with mean photon count 0.01 per pixel. The underlying clean diffraction patterns are very smooth. We vary the band limit from $c = 0.06$ to 0.5 and the support radius R ranges from 50 to 63. With the chosen parameters indicated in Fig. 2, the MSE for the denoised image is small (see Fig. 13a). Although the MSEs vary when we increase the parameter c above the chosen band limit and change R

within the range of 50 to 63, the errors are still relatively small ($2.9 - 3.5 \times 10^{-5}$), compared to above 5×10^{-5} with ePCA and sPCA. For the face data, we compare the results with $n = 8192$ noisy images with 2.3 photons per pixel. Since we crop out the region within a disk of radius $R = L/2$, the parameter R is fixed to be 32 and we vary the band limit c from 0.1 to 0.5. We observe that choosing $c = 0.5$ provides the best denoising performance (see Fig. 13b), and this agrees with the fact that the face images contain more high-frequency information than the XFEL data.

IV. CONCLUSION

We presented steerable ePCA, a method for principal component analysis of a set of low-photon-count images and their uniform in-plane rotations. This work has been mostly motivated by its application to XFEL, but is relevant to other low-photon-count imaging applications. The computational complexity of the new algorithm is $O(nL^3 + L^4)$, whereas that of ePCA is $O(\min(nL^4 + L^6, n^2L^2 + n^3))$. Incorporating rotational invariance allows more robust estimation of the true eigenvalues and eigenvectors. Our numerical experiments showed that steerable ePCA more accurately estimates the covariance matrix and achieves better denoising results. Finally, we remark that the Fourier-Bessel basis can be replaced with other suitable bases, for example, the 2-D prolate spheroidal wave functions (PSWF) on a disk [54], [55]. Nevertheless, our method has certain limitations. For example, if the noise distribution is not specified in advance, it is hard to estimate the noise covariance $\mathbb{E} \text{diag}[A''(\omega)]$, which will affect the homogenization step of the proposed method. In addition, the optimal shrinkage function for the covariance estimation depends on how we measure the error, for example using Frobenius norm, operator norm, or other criteria. The corresponding theoretical guarantee in the non-Gaussian case is still an open problem. Thus, the statistical optimality is beyond the scope of this paper. For future work, we will study how sensitive the proposed method is to the misspecification of the noise distribution and the statistical optimality of our procedure in the non-Gaussian case.

ACKNOWLEDGMENTS

The authors would like to thank associate editor and the anonymous referees for their valuable comments. The authors would also like to thank Edgar Dobriban and Boris Landa for discussions. ZZ was partially supported by National Center for Supercomputing Applications Faculty Fellowship, University of Illinois at Urbana-Champaign College of Engineering Strategic Research Initiative, and NSF DMS-1854791. AS was partially supported by Award Number R01GM090200 from the NIGMS, FA9550-17-1-0291 from AFOSR, Simons Investigator Award, the Moore Foundation Data-Driven Discovery Investigator Award, and NSF BIGDATA Award IIS-1837992.

REFERENCES

- [1] V. Favre-Nicolin, J. Baruchel, H. Renevier, J. Eymery, and A. Borbély, "XTOP: high-resolution X-ray diffraction and imaging," *Journal of Applied Crystallography*, vol. 48, no. 3, pp. 620–620, 2015.

- [2] F. R. N. C. Maia and J. Hajdu, "The trickle before the torrent-diffraction data from X-ray lasers," *Scientific Data*, vol. 3, 2016.
- [3] S. H. W. Scheres, H. Gao, M. Valle, G. T. Herman, P. P. B. Eggermont, J. Frank, and J.-M. Carazo, "Disentangling conformational states of macromolecules in 3D-EM through likelihood optimization," *Nature Methods*, vol. 4, no. 1, pp. 27–29, 2007.
- [4] D. N.-T. Loh and V. Elser, "Reconstruction algorithm for single-particle diffraction imaging experiments," *Phys. Rev. E*, vol. 80, p. 026705, Aug 2009.
- [5] Z. Kam, "Determination of macromolecular structure in solution by spatial correlation of scattering fluctuations," *Macromolecules*, vol. 10, no. 5, pp. 927–934, 1977.
- [6] D. K. Saldin, V. L. Shneerson, R. Fung, and A. Ourmazd, "Structure of isolated biomolecules obtained from ultrashort X-ray pulses: exploiting the symmetry of random orientations," *Journal of Physics: Condensed Matter*, vol. 21, no. 13, 2009.
- [7] D. Starodub, A. Aquila, S. Bajt, M. Barthelmess, A. Barty, C. Bostedt, J. D. Bozek, N. Coppola, R. B. Doak, S. W. Epp *et al.*, "Single-particle structure determination by correlations of snapshot X-ray diffraction patterns," *Nature communications*, vol. 3, p. 1276, 2012.
- [8] K. Pande, P. Schwander, M. Schmidt, and D. K. Saldin, "Deducing fast electron density changes in randomly orientated uncrystallized biomolecules in a pump-probe experiment," *Phil. Trans. R. Soc. B*, vol. 369, no. 1647, p. 20130332, 2014.
- [9] R. P. Kurta, J. J. Donatelli, C. H. Yoon, P. Berntsen, J. Bielecki, B. J. Daurer, H. DeMirci, P. Fromme, M. F. Hantke *et al.*, "Correlations in scattered X-ray laser pulses reveal nanoscale structural features of viruses," *Physical Review Letters*, vol. 119, no. 15, p. 158102, 2017.
- [10] J. J. Donatelli, J. A. Sethian, and P. H. Zwart, "Reconstruction from limited single-particle diffraction data via simultaneous determination of state, orientation, intensity, and phase," *Proceedings of the National Academy of Sciences*, vol. 114, no. 28, pp. 7222–7227, 2017.
- [11] B. von Ardenne, M. Mechelke, and H. Grubmüller, "Structure determination from single molecule X-ray scattering with three photons per image," *Nat Commun.*, vol. 9, no. 1, p. 2375, 2018.
- [12] K. Pande, J. J. Donatelli, E. Malmerberg, L. Foucar, C. Bostedt, I. Schlichting, and P. H. Zwart, "Ab initio structure determination from experimental fluctuation x-ray scattering data," *Proceedings of the National Academy of Sciences*, vol. 115, no. 46, pp. 11772–11777, 2018. [Online]. Available: <http://www.pnas.org/content/115/46/11772>
- [13] I. Jolliffe, *Principal Component Analysis*. Wiley Online Library, 2002.
- [14] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis*. Wiley New York, 2003.
- [15] R. D. Nowak and R. G. Baraniuk, "Wavelet-domain filtering for photon imaging systems," *IEEE Transactions on Image Processing*, vol. 8, no. 5, pp. 666–678, 1999.
- [16] F. Luisier, T. Blu, and M. Unser, "A new sure approach to image denoising: Interscale orthonormal wavelet thresholding," *IEEE Transactions on image processing*, vol. 16, no. 3, pp. 593–606, 2007.
- [17] F. J. Anscombe, "The transformation of poisson, binomial and negative-binomial data," *Biometrika*, vol. 35, no. 3–4, p. 246, 1948.
- [18] J.-L. Starck, F. Murtagh, and J. M. Fadili, *Sparse image and signal processing: wavelets, curvelets, morphological diversity*. Cambridge university press, 2010.
- [19] M. Collins, S. Dasgupta, and R. E. Schapire, "A generalization of principal component analysis to the exponential family," *Advances in Neural Information Processing Systems (NIPS)*, 2001.
- [20] A. P. Singh and G. J. Gordon, "A unified view of matrix factorization models," in *Machine Learning and Knowledge Discovery in Databases*, W. Daelemans, B. Goethals, and K. Morik, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 358–373.
- [21] J. Salmon, Z. Harmany, C.-A. Deledalle, and R. Willett, "Poisson noise reduction with non-local PCA," *Journal of mathematical imaging and vision*, vol. 48, no. 2, pp. 279–294, 2014.
- [22] T. Furnival, R. K. Leary, and P. A. Midgley, "Denoising time-resolved microscopy image sequences with singular value thresholding," *Ultra-microscopy*, 2016.
- [23] Y. Cao and Y. Xie, "Low-rank matrix recovery in Poisson noise," in *Signal and Information Processing (GlobalSIP), 2014 IEEE Global Conference on*. IEEE, 2014, pp. 384–388.
- [24] M. Sonnleitner, J. Jeffers, and S. M. Barnett, "Local retrodiction models for photon-noise-limited images," in *SPIE Photonics Europe*. International Society for Optics and Photonics, 2016, pp. 98 960V–98 960V.
- [25] L. T. Liu, E. Dobriban, and A. Singer, "ePCA: High dimensional exponential family PCA," *Annals of Applied Statistics*, vol. 12, no. 4, pp. 2121–2150, 2018.
- [26] Z. Zhao, Y. Shkolnisky, and A. Singer, "Fast steerable principal component analysis," *IEEE Transactions on Computational Imaging*, vol. 2, no. 1, pp. 1–12, 2016.
- [27] W. T. Freeman and E. H. Adelson, "The design and use of steerable filters," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 9, pp. 891–906, 1991.
- [28] R. Hilai and J. Rubinstein, "Recognition of rotated images by invariant Karhunen-Loève expansion," *J. Opt. Soc. Am. A.*, vol. 11, no. 5, pp. 1610–1618, 1994.
- [29] P. Perona, "Deformable kernels for early vision," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 17, no. 5, pp. 488–499, 1995.
- [30] M. Uenohara and T. Kanade, "Optimal approximation of uniformly rotated images: Relationship between Karhunen-Loève expansion and discrete cosine transform," *IEEE Trans. Image Proces.*, vol. 7, no. 1, pp. 116–119, 1998.
- [31] C. Ponce and A. Singer, "Computing steerable principal components of a large set of images and their rotations," *IEEE Transactions on Image Processing*, vol. 20, no. 11, pp. 3051–3062, 2011.
- [32] Z. Zhao and A. Singer, "Rotationally invariant image representation for viewing direction classification in cryo-EM," *Journal of Structural Biology*, vol. 186, no. 1, pp. 153–166, 2014.
- [33] A. S. Georghiadis, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 6, pp. 643–660, 2001.
- [34] K.-C. Lee, J. Ho, and D. J. Kriegman, "Acquiring linear subspaces for face recognition under variable lighting," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 5, pp. 684–698, 2005.
- [35] B. Landa and Y. Shkolnisky, "Approximation scheme for essentially bandlimited and space-concentrated functions on a disk," *Applied and Computational Harmonic Analysis*, vol. 43, no. 3, pp. 381–403, 2017.
- [36] A. Klug and R. A. Crowther, "Three-dimensional image reconstruction from the viewpoint of information theory," *Nature*, vol. 238, pp. 435–440, 1972.
- [37] Q. Wang, O. Ronneberger, and H. Burkhardt, "Rotational invariance based on fourier analysis in polar and spherical coordinates," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 9, pp. 1715–1722, 2009.
- [38] L. Greengard and J. Lee, "Accelerating the Nonuniform Fast Fourier Transform," *SIAM Review*, vol. 46, no. 3, pp. 443–454, 2004.
- [39] J. Baik, G. Ben Arous, and S. Péché, "Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices," *Annals of Probability*, vol. 33, no. 5, pp. 1643–1697, 2005.
- [40] J. Baik and J. W. Silverstein, "Eigenvalues of large sample covariance matrices of spiked population models," *Journal of Multivariate Analysis*, vol. 97, no. 6, pp. 1382–1408, 2006.
- [41] D. Paul, "Asymptotics of sample eigenstructure for a large dimensional spiked covariance model," *Statistica Sinica*, vol. 17, no. 4, pp. 1617–1642, 2007.
- [42] F. Benaych-Georges and R. R. Nadakuditi, "The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices," *Advances in Mathematics*, vol. 227, no. 1, pp. 494–521, 2011.
- [43] D. L. Donoho, M. Gavish, and I. M. Johnstone, "Optimal shrinkage of eigenvalues in the Spiked Covariance Model," *Annals of Statistics*, vol. 46, no. 4, p. 1742, 2018.
- [44] I. M. Johnstone and A. Y. Lu, "On consistency and sparsity for principal components analysis in high dimensions," *Journal of the American Statistical Association*, vol. 104, no. 486, pp. 682–693, 2009.
- [45] S. Jung, J. S. Marron *et al.*, "Pca consistency in high dimension, low sample size context," *The Annals of Statistics*, vol. 37, no. 6B, pp. 4104–4130, 2009.
- [46] I. M. Johnstone and D. Paul, "Pca in high dimensions: An orientation," *Proceedings of the IEEE*, vol. 106, no. 8, pp. 1277–1292, 2018.
- [47] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical recipes in Fortran 90*. Cambridge university press Cambridge, 1996, vol. 2.
- [48] M. F. Hantke, T. Ekeberg, and F. R. N. C. Maia, "Condor: A simulation tool for flash X-ray imaging," *Journal of Applied Crystallography*, vol. 49, no. 4, pp. 1356–1362, 2016.
- [49] P. Schwander, D. Giannakis, C. H. Yoon, and A. Ourmazd, "The symmetries of image formation by scattering. II. Applications," *Opt. Express*, vol. 20, no. 12, pp. 12 827–12 849, Jun 2012.
- [50] J. Landgrebe, W. Wurst, and G. Welzl, "Permutation-validated principal components analysis of microarray data," *Genome Biology*, vol. 3, no. 4, 2002.
- [51] E. Dobriban, "Permutation methods for factor analysis and PCA," *arXiv preprint arXiv:1710.00479*, Oct. 2017.

- [52] E. Dobriban and A. B. Owen, “Deterministic parallel analysis: an improved method for selecting factors and principal components,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2018.
- [53] R. Basri and D. W. Jacobs, “Lambertian Reflectance and Linear Subspaces,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 2, pp. 218–233, 2003.
- [54] D. Slepian, “Prolate spheroidal wave functions, Fourier analysis, and uncertainty–IV: extensions to many dimensions, generalized prolate spheroidal wave functions,” *Bell System Technical Journal*, vol. 43, pp. 3009–3057, 1964.
- [55] B. Landa and Y. Shkolnisky, “Steerable principal components for space-frequency localized images,” *SIAM Journal on Imaging Sciences*, vol. 10, no. 2, pp. 508–534, 2017.