Published in final edited form as:

IEEE Trans Image Process. 2016 February ; 25(2): 540–552. doi:10.1109/TIP.2015.2504901.

Removal of Vesicle Structures From Transmission Electron Microscope Images

Katrine Hommelhoff Jensen

Department of Computer Science, University of Copenhagen, Copenhagen 2100, Denmark,
katrine@di.ku.dk

Fred J. Sigworth

Department of Cellular and Molecular Physiology, Yale School of Medicine, New Haven, CT
06520 USA, fred.sigworth@yale.edu

Sami Sebastian Brandt

Department of Computer Science, University of Copenhagen, Copenhagen 2100, Denmark,
sbrandt@di.ku.dk

Abstract

In this paper, we address the problem of imaging membrane proteins for single-particle cryo-electron microscopy reconstruction of the isolated protein structure. More precisely, we propose a method for learning and removing the interfering vesicle signals from the micrograph, prior to reconstruction. In our approach, we estimate the subspace of the vesicle structures and project the micrographs onto the orthogonal complement of this subspace. We construct a 2d statistical model of the vesicle structure, based on higher order singular value decomposition (HOSVD), by considering the structural symmetries of the vesicles in the polar coordinate plane. We then propose to lift the HOSVD model to a novel hierarchical model by summarizing the multidimensional HOSVD coefficients by their principal components. Along with the model, a solid vesicle normalization scheme and model selection criterion are proposed to make a compact and general model. The results show that the vesicle structures are accurately separated from the background by the HOSVD model that is also able to adapt to the asymmetries of the vesicles. This is a promising result and suggests even wider applicability of the proposed approach in learning and removal of statistical structures.

Keywords

Statistical learning; image processing; biomedical image processing; molecular imaging; transmission electron microscopy; statistical analysis

I. Introduction

In structural biology, a long term target has been to image membrane proteins in their natural environment and to understand the functional organisation of the protein complex by

reconstructing its 3d structure. In certain cases, the imaging of membrane proteins has been achieved by electron crystallography [33], however, many proteins are difficult to crystallize. More recently, the method of single-particle reconstruction (SPR) [12] has grown powerful enough to even reconstruct atomic structures. In SPR, cryo-electron microscopic projection images of similar copies of the protein are used for tomographic 3d reconstruction of the protein structure. A drawback of this method is the fact that viewing angles of the projection images are unknown that, together with a very low signal-to-noise ratio (SNR), makes the reconstruction problem particularly challenging.

In order to reconstruct membrane proteins, one additionally needs to separate the signal of the protein from the overlapping membrane vesicle. In [41] and [42], a new method, the *random spherically constrained* single-particle reconstruction, was introduced. There, the membrane proteins were reconstituted into lipid vesicles, frozen in vitreous ice and imaged by the microscope. This setup is illustrated in Fig. 1 and a sample micrograph with vesicles illustrated in Fig. 2. Membrane proteins were identified on the vesicle surfaces and the associated vesicle structures were removed using a scalable, geometric model, prior to the reconstruction. The protein reconstruction is however affected by the residual of the membrane vesicle structures near the proteins. This degrades the reconstruction quality, since the vesicle signal is much stronger than that of the protein, that even weak artifacts have a strong influence on the protein structure estimate. The reconstruction is additionally affected by other types of structural noise such as fragments of ice and random background noise.

To extract the membrane proteins from the micrographs, we propose a unique approach where we consider the problem of separating three overlapping signals: 1) the vesicle, 2) the protein and 3) the background noise. The vesicle signal is dominant over those of the protein though there are strong intensity variations over all the micrographs that complicates the vesicle detection. The protein signal is weak and close to the noise level, and thus challenging to detect. Additionally, one has to automatically determine how complex the structures are. To solve the problems above, we must thus consider the three general problems addressed in the literature: signal separation, noise removal, and model selection.

In general, *signal separation*, aiming at recovering an unobserved signal from observed mixtures, has been widely studied in the literature. In *blind* signal separation, one assumes that one only observes mixtures of signals while the underlying components are unknown apart from the assumption that they are, for instance, minimally correlated, or maximally independent. Blind signal separation can be achieved for instance by principal component analysis (PCA) [19] its variants [6], [44], singular value decomposition (SVD) [39], independent component analysis (ICA) [16], independent subspace analysis (ISA) [17], multilinear PCA (MPCA) or higher order SVD (HOSVD) [24], [26]. Another approach is to impose structural constraints on the source signals, such as non-negativity imposing sparsity in the non-negative matrix factorisation (NMF) [7] and variance-difference maximising separation in common spatial pattern (CSP) [25]. If one of the signals to be separated is noise, the separation can be used as *noise reduction technique*, where the observed signal is decomposed to signal and noise subspaces and projected onto the signal subspace to remove the noise.

There is also a vast amount of other noise reduction techniques in the literature. The most primitive approaches are based on image filtering by simple linear or non-linear filters, such as local averaging or median filtering [31], [32]. These kinds of approaches, however, also obscure fine, low contrast details in the underlying image. The classic statistical approach for noise reduction is the Wiener filter [5], which is the optimal linear filter under the assumption of Gaussian noise. Edge preserving classic choices for noise reduction include total variation regularisation [36] and anisotropic diffusion [43]. More recent approaches are based on sparsity regularisation like those of Wavelet shrinkage [2], [8], [29], sparse code shrinkage [10], [15], and techniques relying on compressed sensing [9] and dictionary learning [11]. There are also information theoretical approaches based on the minimum description length principle (MDL) [34].

Model selection is the problem of picking up the best model from a set of candidate models of different complexity. A good model selection technique will find the balance between the goodness of fit and model complexity by implementing the Occam's razor principle that states that the simplest model is the most likely one of otherwise equal candidates. There are both frequentist and Bayesian methods for model selection such as F tests for nested models [28], Mallows's C_p [30], Akaike information criterion (AIC) [1], exhaustive search and stepwise/forward/backwards selection procedures [20], crossvalidation [23], Bayes factors of various types [22], Bayesian information criterion BIC [37], the generalisation of AIC and BIC for hierarchical models, the deviance information criterion (DIC) [3], Bayes model averaging [14], and the focused information criterion FIC [13]. An information theoretic approach is the minimum description length (MDL) principle, of which the most recent version is based on the normalised maximum likelihood (NML) criterion [35].

Our approach for separating the protein signal from the micrographs is to estimate the subspace of the vesicle structures in two dimensions and project the micrographs onto the orthogonal complement of this subspace. We chose to model the vesicles structures in two dimensions since this makes the removal simpler from both the statistical modelling and computational point of view, when compared to reconstructing an explicit three dimensional generative model of the vesicle structures including deformations. More precisely, the 2d model, in addition to the vesicle structure, contains the effect of the contrast transfer function (CTF) on the vesicle appearance in the micrographs. The separation of the noise component from the protein signal will be later addressed as part of the reconstruction problem, by the reconstruction algorithm that is beyond the scope of this work.

As the vesicles show a high degree of rotational symmetry, we chose to model the vesicle subspace by the higher order singular value decomposition on the polar coordinate plane, see Fig. 3. Moreover, in contrast to PCA-based approaches such as [6], [19], and [44], HOSVD provides a compact, orthogonal two-dimensional basis characterising the vesicle structure variations since the radial and angular parts are analysed separately. Due to the separation of the radial and angular parts, and the fact that they demonstrate different kinds of variation, we obtain a huge reduction in dimensionality of the model and thus in overfitting, when compared to a standard PCA. To find the best summary of the two-dimensional HOSVD basis patches, we propose another layer onto the HOSVD model by summarising the coefficients by principal component analysis. To select the complexity of the vesicle model,

i.e., the number of radial, angular, and summarised components, we use Morozov's discrepancy principle.

In order to make an accurate modelling of the vesicles, we must properly normalise the vesicle data in position, size and intensity range. The geometric normalisation is carried out by centering the vesicles around the centre point estimates and scaling the sizes using the radius estimates [41].¹ By intensity normalisation we adjust the vesicle signal amplitude and background level to compensate for global illumination and contrast changes, caused by different microscope settings during micrograph acquisition, as illustrated in Fig. 4.

The summary of our vesicle removal method is displayed in Fig. 5. The detailed organisation of this paper is as follows. The vesicle modelling theory is described in Section II. The HOSVD model is derived in Section III. Then, the implementation of the model training and vesicle removal is described in Section IV, whereas the model selection method is reported in Section V. The experiments are in Section VI and conclusions in Section VII.

II. Vesicle Modelling

A. Definition of the Objective

In mathematical terms, our goal is to remove the vesicle structures from micrographs, or, to project the micrographs onto the orthogonal complement of the affine subspace describing vesicle structures. In other words, defining the original micrograph as $f_{\text{micr}} \in \mathcal{C}^0(\mathbb{R}^2)$, where $\mathcal{C}^0(\mathbb{R}^2)$ is the class of compact, continuous functionals in $\mathbb{R}^2$, the projection $\mathcal{P}_\perp : \mathcal{C}^0(\mathbb{R}^2) \rightarrow \mathcal{C}^0(\mathbb{R}^2)$ onto the orthogonal complement is

$$f_{\text{proj}} = \mathcal{P}_\perp f_{\text{micr}}. \quad (1)$$

In practice, our objective is to construct the projection as a sequence of projections over the N vesicles, or,

$$\mathcal{P}_\perp = \mathcal{P}_\perp^{(N)} \circ \mathcal{P}_\perp^{(N-1)} \circ \dots \circ \mathcal{P}_\perp^{(1)}, \quad (2)$$

where $\mathcal{P}_\perp^{(n)}$ is the projection onto the orthogonal complement of the vesicle n .² We need a separate projection operator for each vesicle as they vary in position, size, shape and intensity. However, the projection operators will coincide in the *normalised* coordinate frame.

Let $\mathcal{T}^{(n)} \equiv T_{\text{geom}}^{(n)} \circ T_{\text{win}}^{(n)} \circ \mathcal{T}_{\text{int}}^{(n)}$ be an affine normalisation operator composed of linear operators $T_{\text{geom}}^{(n)}$ for the normalisation of the vesicle position and size and $T_{\text{win}}^{(n)}$ for windowing the vesicle pixels, and the affine intensity normalisation operator $\mathcal{T}_{\text{int}}^{(n)}$. Assume

¹We obtain a better, more compact basis if we normalise over the vesicle size. This is despite the fact that the vesicle wall shows somewhat constant absolute thickness in nature and the size normalisation makes the thickness variant in the normalised coordinate frame.

²We further assume that the affine subspaces of the vesicles on the micrograph are disjoint. When this is the case the projection (2) commutes, i.e., the order of the vesicles does not matter in the removal.

further that $T_{\text{geom}}^{(n)}$ and $\mathcal{T}_{\text{int}}^{(n)}$ are invertible. Let us define the normalised vesicle image as $f^{(n)} = \mathcal{T}^{(n)} f_{\text{micr}}$. The following theorem characterises the nature of the projection operator $\mathcal{P}_{\perp}^{(n)}$.

Theorem 1: Let $P_{\perp}$ be the projection onto the linear subspace orthogonal to the affine subspace of windowed and normalised vesicle structures. The projection of the micrograph f_{micr} onto the orthogonal complement of the affine subspace representing the vesicle structures is

$$\mathcal{P}_{\perp}^{(n)} f_{\text{micr}} \equiv \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ \left(T_{\text{geom}}^{(n)} \right)^{-1} P_{\perp} \left(f^{(n)} - g_0 \right) + \left(\left(I - T_{\text{win}}^{(n)} \right) \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right\}, \quad (3)$$

where g_0 is the offset of the affine subspace, and $I - T_{\text{win}}^{(n)}$ is the windowing operator corresponding to the background window $1 - w^{(n)}$.

Proof: Let $T_{\text{win}}^{(n)} f \equiv w^{(n)} f$, $f \in \mathcal{C}^0(\mathbb{R}^2)$, where $0 \leq w^{(n)}(x, y) \leq 1, \forall x, y \in \mathbb{R}$. The window functional $w^{(n)}$ divides the micrograph into the foreground and background of the vesicle n such that

$$f_{\text{micr}} \equiv \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ w^{(n)} \left(\mathcal{T}_{\text{int}}^{(n)} f_{\text{micr}} \right) + \left(1 - w^{(n)} \right) \left(\mathcal{T}_{\text{int}}^{(n)} f_{\text{micr}} \right) \right\} \quad (4)$$

where $1 - w^{(n)}$ is the background window. This is equivalent

$$\begin{aligned} & \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ \left(T_{\text{win}}^{(n)} \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} + \left(\left(I - T_{\text{win}}^{(n)} \right) \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right\} \\ &= \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ \left(T_{\text{geom}}^{(n)} \right)^{-1} \left(T_{\text{geom}}^{(n)} \circ T_{\text{win}}^{(n)} \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right. \\ & \quad \left. + \left(\left(I - T_{\text{win}}^{(n)} \right) \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right\} \\ &= \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ \left(T_{\text{geom}}^{(n)} \right)^{-1} f^{(n)} + \left(\left(I - T_{\text{win}}^{(n)} \right) \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right\}. \end{aligned} \quad (5)$$

Now, let us make a subspace approximation for the normalised, windowed image $f^{(n)}$, such that

$$f^{(n)} = \hat{f}^{(n)} + r^{(n)}, \quad (6)$$

where $\hat{f}^{(n)} = g_0 + P \left(f^{(n)} - g_0 \right)$, and P is the projection onto the linear subspace of the mean corrected vesicle structures. The residual after the removal of $\hat{f}^{(n)}$ is equivalent to

$$\begin{aligned} r^{(n)} &= f^{(n)} - \hat{f}^{(n)} = (f^{(n)} - g_0) - P(f^{(n)} - g_0) \\ &= P_{\perp}(f^{(n)} - g_0) \end{aligned} \quad (7)$$

where $P_{\perp} = I - P$. The removal of the vesicle is thus obtained by replacing

$f^{(n)}$ by $r^{(n)} = f^{(n)} - \hat{f}^{(n)}$ in (5) that yields the projection operator

$$\begin{aligned} \mathcal{P}_{\perp}^{(n)} f_{\text{micr}} &\equiv \left(\mathcal{T}_{\text{int}}^{(n)} \right)^{-1} \left\{ \left(T_{\text{geom}}^{(n)} \right)^{-1} P_{\perp} (f^{(n)} - g_0) \right. \\ &\quad \left. + \left((I - T_{\text{win}}^{(n)}) \circ \mathcal{T}_{\text{int}}^{(n)} \right) f_{\text{micr}} \right\} \end{aligned} \quad (8)$$

and the claim follows.

B. Vesicle Normalisation

In this section, we first describe the normalisation of the vesicle position and size and the windowing. Thereafter, we describe the normalisation of the intensity range to remove illumination differences.

1) Shape Normalisation: Let $f_{\text{norm}}^{(n)}$ be the intensity normalised micrograph $\mathcal{T}_{\text{int}}^{(n)} f_{\text{micr}}$. The composite operator $T_{\text{geom}}^{(n)} \circ T_{\text{norm}}^{(n)} : f_{\text{norm}}^{(n)} \mapsto f^{(n)}$ representing the geometric normalisation and windowing as Fig. 6 illustrates, is defined as

$$\begin{aligned} f^{(n)}(x, y) &= f_{\text{norm}}^{(n)}(x_{\text{micr}}, y_{\text{micr}}) \\ &\quad \times w^{(n)}(x_{\text{micr}} - x_{\text{micr}}^{(n)}, y_{\text{micr}} - y_{\text{micr}}^{(n)}), \end{aligned} \quad (9)$$

where $(x_{\text{micr}}^{(n)}, y_{\text{micr}}^{(n)})$ is the vesicle center in the micrograph, $w^{(n)}$ the window functional, and

$$\begin{pmatrix} x_{\text{micr}} \\ y_{\text{micr}} \end{pmatrix} = \frac{R^{(n)}}{\max_{n'} R^{(n')}} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} x_{\text{micr}}^{(n)} \\ y_{\text{micr}}^{(n)} \end{pmatrix}. \quad (10)$$

The ratio $\frac{R^{(n)}}{\max_{n'} R^{(n')}} R^{(n')}$ is the relative vesicle radius in the original image, normalised by the maximum radius over the n' training vesicles.

Representing the normalised vesicle in polar coordinates

$f^{(n)}(\rho, \varphi) \equiv f^{(n)}(\rho(x, y), \varphi(x, y))$, we define the vesicle window as

$$\begin{aligned} w^{(n)}(x_{\text{micr}} - x_{\text{micr}}^{(n)}, y_{\text{micr}} - y_{\text{micr}}^{(n)}) \\ \equiv w_{\text{Tuk}}^{(n)}\left(r(x_{\text{micr}} - x_{\text{micr}}^{(n)}, y_{\text{micr}} - y_{\text{micr}}^{(n)})\right) \end{aligned} \quad (11)$$

where

$$w_{\text{Tuk}}^{(n)}(r) = \begin{cases} 1, & \text{if } r \leq R^{(n)} \\ \frac{1}{2} + \frac{\cos \pi \alpha (r - R^{(n)})}{2}, & \text{if } R^{(n)} < r < R^{(n)} + \frac{1}{\alpha} \\ 0, & \text{otherwise;} \end{cases} \quad (12)$$

is the *Tukey window*. The parameter α controls the steepness of transition region of the window (see Fig. 6).

2) Intensity Normalisation: In the intensity normalisation, we aim at recovering the differences in contrast (Fig. 4) and background level (Fig. 7). We define the pixel intensity normalisation as the operator $\mathcal{T}_{\text{int}}^{(n)}: f_{\text{micr}} \mapsto f_{\text{norm}}^{(n)}$ as

$$f_{\text{norm}}^{(n)}(x, y) = \frac{f_{\text{micr}}(x, y) - b(x, y)}{s(x, y)} \quad (13)$$

where $b(x, y)$ is the background reference level and $s(x, y)$ the vesicle signal level.

Due to global changes of the background level (Fig. 7), we use a local median estimate for each vesicle, serving as a robust reference level estimator for the vesicle signal. For the vesicle n , the background level is thus estimated as

$$\hat{b}^{(n)} = \text{median}_{(x', y') \in \mathcal{U}^{(n)}} f_{\text{micr}}(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)}) \quad (14)$$

where the local background area is defined as

$$\mathcal{U}^{(n)} = \left\{ (x', y') \mid w_{\text{loc}}(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)}) = 1 \right. \\ \left. \wedge w_{\text{bg}}(x, y) = 1 \right\}, \quad (15)$$

whereas the local support

$$w_{\text{loc}}(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)}) = \begin{cases} 1, & \text{if } d(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)}) < R_{\text{loc}} \\ 0, & \text{otherwise,} \end{cases} \quad (16)$$

where $R_{\text{loc}} > R^{(n)} + \frac{1}{\alpha}$ is the radius of the local support; and the background is windowed by

$$w_{\text{bg}}(x, y) = \begin{cases} 1 & \text{if } \forall n \quad d(x - x_{\text{micr}}^{(n)}, y - y_{\text{micr}}^{(n)}) \geq R^{(n)} + \frac{1}{\alpha} \\ 0 & \text{otherwise} \end{cases} \quad (17)$$

The local background is illustrated in Fig. 8, where the radius of the local support is set to half the width of the image.

Finally, the vesicle signal level is robustly estimated by the median absolute deviation (MAD) estimator or

$$\hat{s}^{(n)} = \text{median}_{(x', y') \in \Omega^{(n)}} \left| f_{\text{micr}} \left(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)} \right) \times w^{(n)} \left(x' - x_{\text{micr}}^{(n)}, y' - y_{\text{micr}}^{(n)} \right) - \hat{b}^{(n)} \right| \quad (18)$$

where $\Omega^{(n)}$ is the support of the vesicle n .

C. Polar Coordinate Representation

In the construction of the projection operator (1), (3), our next goal is to construct the vesicle model $\hat{f}^{(n)}$ for vesicles $n = 1, 2, \dots, N$ by statistical modelling. We are searching for a reasonable fit in the sense of the energy functional

$$E = \sum_{n=1}^N \int_{\Omega^{(n)}} \left(f^{(n)}(\mathbf{r}) - \hat{f}^{(n)}(\mathbf{r}) \right)^2 d\mathbf{r}, \quad (19)$$

where $\mathbf{r} = (x, y)$ is a position in the normalised image and the support $\Omega^{(n)}$ of each vesicle $f^{(n)}$ is known. We seek for the models that provide both a good fit, but do not overfit, and are compact, i.e., they provide a reasonable energy level and small number of parameters and hence separate vesicle signal from the background noise.

We assume that the vesicles lie in an orthogonal subspace to the noise subspace on the image plane. So our aim is to find an *orthonormal basis* $\{g_k(\mathbf{r})\}_{k=1}^K$ spanning the vesicle subspace, or,

$$\langle g_k, g_k \rangle = 1, \quad k=1, \dots, K \quad (20)$$

and

$$\langle g_k, g_l \rangle = 0 \quad \text{for all } k \neq l. \quad (21)$$

Vesicles have a high level of symmetry in the polar coordinates that should be taken into account to make the representation compact. The least squares analysis in Cartesian (x, y) and polar coordinates (ρ, φ) are equivalent in the following sense.

Lemma 2: The L_2 minimisation on the image domain is equivalent to L_2 minimisation on a weighted polar coordinate frame, where the weighting function is $\sqrt{\rho(x, y)}$.

Proof: An L_2 approximation $\hat{f}(x, y)$ of the image $f(x, y) = f(x(\rho, \varphi), y(\rho, \varphi)) = f(\rho, \varphi)$ minimises

$$\begin{aligned} E &= \iint (f - \hat{f})^2 dx dy \\ &= \iint \rho (f(\rho, \varphi) - \hat{f}(\rho, \varphi))^2 d\rho d\varphi \\ &= \iint (\sqrt{\rho} f(\rho, \varphi) - \sqrt{\rho} \hat{f}(\rho, \varphi))^2 d\rho d\varphi \end{aligned} \quad (22)$$

which is the L_2 -minimisation between $\sqrt{\rho}f(\rho, \varphi)$ and $\sqrt{\rho}\hat{f}(\rho, \varphi)$.

The lemma gives us the freedom to make the analysis in the polar coordinates as long as we take the weighting $\sqrt{\rho}$ into consideration.

To make the representation compact, we assume that the basis is *separable* in the polar coordinates or

$$g_k(\mathbf{r}(\rho, \varphi)) = \tilde{u}_k^{(1)}(\rho) u_k^{(2)}(\varphi) \equiv \frac{u_k^{(1)}(\rho)}{\sqrt{\rho}} u_k^{(2)}(\varphi). \quad (23)$$

The separability implies the following useful lemma, since we want to generate an orthogonal basis in the vesicle subspace.

Lemma 3: Let $\{g_k\}_{k=1}^K$ be a separable basis in the polar coordinates. The basis is orthogonal if and only if either the radial or the angular part of the basis is orthogonal.

Proof: For $k \neq l$, from the separability it follows

$$\begin{aligned} \langle g_k, g_l \rangle &= \left(\int \rho \tilde{u}_k^{(1)}(\rho) \tilde{u}_l^{(1)}(\rho) d\rho \right) \\ &\quad \times \left(\int u_k^{(2)}(\varphi) u_l^{(2)}(\varphi) d\varphi \right) \\ &= \left\langle \sqrt{\rho} \tilde{u}_k^{(1)}, \sqrt{\rho} \tilde{u}_l^{(1)} \right\rangle \left\langle u_k^{(2)}, u_l^{(2)} \right\rangle \\ &= \left\langle u_k^{(1)}, u_l^{(1)} \right\rangle \left\langle u_k^{(2)}, u_l^{(2)} \right\rangle. \end{aligned} \quad (24)$$

where $u_k^{(1)}(\rho)$ is a weighted radial basis function and $u_k^{(2)}(\varphi)$ an angular basis function. The right hand side vanishes if and only if either the radial or angular part is orthogonal.

The message here is thus the follows.

Corollary 4: Aiming at the least squares fit and an orthogonal basis for the vesicle subspace on the image plane, it is equivalent to inspect the ρ weighted vesicle images in the polar coordinate plane and construct an orthogonal bases for the radial and angular part separately.³

Assuming $f^{(n)}(x, y)$ is a normalised, windowed vesicle image (9) we define the operator

$T_{\rho\varphi}^{(n)}: f^{(n)} \mapsto f_{\rho\varphi}^{(n)}$ transform $f^{(n)}(x, y)$ into the weighted, polar coordinates as

$$f_{\rho\varphi}^{(n)}(\rho, \varphi) = \sqrt{\rho} f^{(n)}(\rho, \varphi) = \sqrt{\rho(x, y)} f^{(n)}(x, y). \quad (25)$$

In the following section, we describe how we construct orthogonal bases using the higher-order singular value decomposition (HOSVD).

³Here the considerations are derived under the assumption of continuous images. The discretisation should be optimally performed in a way that the mapping between the discretised representations were unitary [44] so that the noise statistics would be preserved. Here the discretisation noise is however regarded to be negligible in practice.

III. HOSVD

A. Basic Model

Let us have the discretised vesicles stored in the three-way array $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ in the weighted polar coordinate frame; let us also introduce the reference array $\mathcal{A}_0 \in \mathbb{R}^{I_1 \times I_2 \times I_3}$, where the dimensions $d = 1, 2, 3$ correspond to the weighted radial, angular, and vesicle index, respectively, and $I_3 \equiv N$ represents the number of vesicle training images. The structure of $\mathcal{A}$ is illustrated in Fig. 9. We look for the approximation $\hat{\mathcal{A}}$ for $\mathcal{A} - \mathcal{A}_0$. We are going to use the following theorem.

Theorem 5: There is a decomposition

$$\mathcal{A} - \mathcal{A}_0 = \mathcal{S} \times_1 \mathbf{U}^{(1)} \times_2 \mathbf{U}^{(2)} \times_3 \mathbf{U}^{(3)}, \quad (26)$$

where $\mathbf{U}^{(d)}$ are orthonormal $I_d \times I_d$ matrices, $d = 1, 2, 3$, $\times_d$ is the d -way product, and $\mathcal{S} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ is the core array whose subarrays are (1) all-orthogonal, i.e., for all possible values of $d, \alpha, \beta, \alpha \neq \beta$, $\langle \mathcal{S}_{i_d=\alpha}, \mathcal{S}_{i_d=\beta} \rangle = 0$; (2) sorted so that the d^{th} order singular values $\sigma_{\alpha}^{(d)} = \|\mathcal{S}_{i_d=\alpha}\|_{\text{fro}}$ are sorted as $\sigma_1^{(d)} \geq \sigma_2^{(d)} \geq \dots \geq \sigma_{I_d}^{(d)}, d=1, 2, 3$.

The proof is in [26]. This decomposition is known as higher-order singular value decomposition (HOSVD). It is useful due to the following approximation property, which is the higherorder counter part to the approximation property of the SVD.

Theorem 6: Let $\hat{\mathcal{A}}$ be the array obtained by truncating the basis formed by the column vectors of $\sigma_{\tilde{I}_{d+1}}^{(d)}, \sigma_{\tilde{I}_{d+2}}^{(d)}, \dots, \sigma_{I_d}^{(d)}, d=1, 2, 3$ Then,

$$\|(\mathcal{A} - \mathcal{A}_0) - \hat{\mathcal{A}}\|_{\text{fro}} \leq \sum_{d=1}^3 \sum_{i_d=\tilde{I}_{d+1}}^{I_d} \sigma_{i_d}^{(d)}. \quad (27)$$

The proof can likewise found in [26]. Though the HOSVD approximation is good enough for practical purposes, there may be another basis that provides tighter approximation in the least square sense with the same number of basis vectors $\tilde{I}_d, d=1, 2, 3$, retained.

The practical computation of the HOSVD is easy with the help of the ordinary SVD. The HOSVD can be implemented by making a sequence of matrix unfoldings for $\mathcal{A} - \mathcal{A}_0$ and computing the singular value decomposition for each. In detail, let the d^{th} unfolding be denoted as $\mathbf{A}_{(d)} - \mathbf{A}_{0,(d)}, d = 1, 2, 3$. The $d = 1$ and $d = 2$ unfoldings are illustrated in Fig. 9. $\mathbf{A}_{(1)}$ allows for the analysis of the weighted radial components and $\mathbf{A}_{(2)}$ for the angular components, one component a column, independently of any other variation in the remaining dimensions of the data. The d^{th} mode singular matrices $\mathbf{U}^{(d)}$ are found as the left singular matrices of the corresponding d^{th} mode matrix unfolding, or $\mathbf{A}_{(d)} - \mathbf{A}_{0,(d)} = \mathbf{U}^{(d)} \Sigma^{(d)} \mathbf{V}^{(d)\text{T}}$. We can alternatively state as follows.

Corollary 7: The d -fold singular vectors are the eigenvectors of the d -fold correlation matrix

$$\mathbf{R}^{(d)} = \frac{1}{I_d} \left(\mathbf{A}_{(d)} - \mathbf{A}_{0,(d)} \right) \left(\mathbf{A}_{(d)} - \mathbf{A}_{0,(d)} \right)^T \quad (28)$$

Corollary 7 reveals the relation of the HOSVD to the principal component analysis (PCA), and the reason why we introduced the reference array $\mathcal{A}_0$. In principal component analysis one computes the eigenvectors of the sample covariance matrix where the observations have been centred around the mean observation. In the multidimensional case, it is not possible since there is generally no such a reference array that would produce the means of each unfoldings by just taking the corresponding unfolding of the reference array. However, we may select the reference array so that its $d=1$ and $d=2$ unfoldings are consistent, by setting the reference array to be a stack of the mean vesicle copies computed over the vesicle population $d=3$ that is a good compromise.

The vesicle bases are thus generated using the $d=1$ and $d=2$ unfoldings of the array $\mathcal{A} - \mathcal{A}_0$ to compute the weighted radial and angular basis vectors as the columns of $\mathbf{U}^{(1)}$ and $\mathbf{U}^{(2)}$. We then truncate the corresponding d -mode singular values by retaining $\tilde{I}_1$ and $\tilde{I}_2$ singular vectors in $\tilde{\mathbf{U}}^{(1)}$ and $\tilde{\mathbf{U}}^{(2)}$, respectively. We generate the orthogonal basis $\mathbf{G} = (\mathbf{g}_k)$, corresponding to (20), where $k = k(i_1, i_2)$ and $i_1 = 1, 2, \dots, \tilde{I}_1$, $i_2 = 1, 2, \dots, \tilde{I}_2$ are the column indices of $\tilde{\mathbf{U}}^{(1)}$ and $\tilde{\mathbf{U}}^{(2)}$, respectively. In other words, all weighted radial basis vectors are paired with all angular basis vectors (see Fig. 10). A vesicle is hence modelled as

$$\hat{\mathbf{f}} = \mathbf{g}_0 + \mathbf{G}\xi, \quad (29)$$

where $\mathbf{g}_0$ is the mean of all normalised, discretised images $\mathbf{f}^{(n)}$, $\xi = \mathbf{G}^T (\mathbf{f}^{(n)} - \mathbf{g}_0)$ are the coordinates at the vesicle in the basis formed by the column vectors of $\mathbf{G}$.

B. Hierarchical Model

The basic HOSVD model does not directly reveal whether any of the generated 2d bases images is irrelevant and could be dropped. We solve this problem by adding a PCA layer to the HOSVD model as follows.

Let $\xi^{(i_3)}$ be the coefficient vector of the normalised vesicle image $\mathbf{f}^{(i_3)}$. For all coefficient vectors, we compute the mean $\bar{\xi}$ and the covariance matrix

$$\mathbf{C} = \sum_{i_3} \left(\xi^{(i_3)} - \bar{\xi} \right) \left(\xi^{(i_3)} - \bar{\xi} \right)^T. \text{ We may now write}$$

$$\xi = \bar{\xi} + \sum_{k=1}^K \lambda_k \langle \mathbf{v}_k, \xi \rangle \mathbf{v}_k = \bar{\xi} + \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \xi \quad (30)$$

where λ_k , $\mathbf{v}_k$, $k = 1, 2, \dots, K$ are the eigenvalues and the eigenvectors of $\mathbf{C}$, all ordered in the descending order of the eigenvalues. Substituting (30) into (29) yields the ordered model,

$$\begin{aligned}
\hat{\mathbf{f}} &= \mathbf{g}_0 + \mathbf{G} \left(\bar{\xi} + \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \xi \right) \\
&= \mathbf{g}_0 + \mathbf{G} \bar{\xi} + \mathbf{G} \mathbf{V} \mathbf{\Lambda} \mathbf{V}^T \xi \\
&= \mathbf{g}'_0 + \mathbf{G}' \xi'
\end{aligned} \tag{31}$$

for the combined model offset

$$\mathbf{g}'_0 = \mathbf{g}_0 + \mathbf{G} \bar{\xi} \tag{32}$$

where $\mathbf{G} \bar{\xi}$ is the vesicle model estimate of the mean coefficient, $\xi' = \mathbf{\Lambda} \mathbf{V}^T \xi$ and $\mathbf{G}' = \mathbf{G} \mathbf{V}$. In other words, $\mathbf{g}'_0$ represents the offset and $\mathbf{G}'$ a new set of ordered basis vectors. The $\tilde{K}$ smallest eigenvalues and the associated new basis vectors $\tilde{\mathbf{G}}'$ can then be truncated to achieve a compressed model. The construction of the PCA layer is illustrated in Fig. 11.

IV. Implementation

In Section IV-A, we describe the complete model training procedure, after which, in Section IV-B, we report how vesicles are removed from micrographs.

A. Training

Given a set of micrographs containing particle-free vesicles, with known vesicle center positions and radii, the model is trained as follows. The vesicle images are first intensity normalised (13), windowed, shape normalised (9), and transformed onto the polar coordinate plane (25). Then, the polar coordinate patches are stacked into the three-way array $\mathcal{A}$. The array $\mathcal{A}$ is decomposed by the higher-order singular value decomposition (26) of which the smallest singular vectors corresponding to the radial and angular dimensions are truncated. As the last step, the hierarchical model (31) is constructed and the smallest singular vectors truncated to form the final basis $\tilde{\mathbf{G}}'$. The whole training algorithm is summarised in Algorithm 1.

Algorithm 1 Vesicle Model Training

Input: Micrographs $\{\mathbf{f}_{\text{micr}}\}$, vesicle centres $(x_{\text{micr}}^{(n)}, y_{\text{micr}}^{(n)})$, vesicle radii $R^{(n)}$, $n = 1, 2, \dots, N$ and truncation parameters $(\tilde{I}_1, \tilde{I}_2, \tilde{K})$.

Output: The vesicle basis $(\mathbf{g}'_0, \tilde{\mathbf{G}}')$.

for each vesicle $n = 1 \dots N$ **do**

$\mathbf{f}_{\text{norm}}^{(n)} = \mathbf{T}_{\text{int}}^{(n)}(\mathbf{f}_{\text{micr}}^{(n)})$ // Intensity norm. (13)

$\mathbf{f}_{\text{win}}^{(n)} = \mathbf{T}_{\text{win}}^{(n)}(\mathbf{f}_{\text{norm}}^{(n)})$ // Windowing (9)

$\mathbf{f}^{(n)} = \mathbf{T}_{\text{geom}}^{(n)}(\mathbf{f}_{\text{win}}^{(n)})$ // Shape norm. (9)

$\mathbf{f}_{\rho\varphi}^{(n)} = \mathbf{T}_{\rho\varphi}^{(n)}(\mathbf{f}^{(n)})$ // To polar coord. (25)

$\mathcal{A} \leftarrow \mathbf{f}_{\rho\varphi}^{(n)}$ // Stack into the three-way array

end for

$\{\mathbf{U}^{(1)}, \mathbf{U}^{(2)}, \mathbf{U}^{(3)}, \mathbf{S}\} \leftarrow \mathcal{A} - \mathcal{A}_0$ // HOSVD

$\tilde{\mathbf{U}}^{(1)} \leftarrow \mathbf{U}^{(1)}, \tilde{\mathbf{U}}^{(2)} \leftarrow \mathbf{U}^{(2)}$ // Truncate

$\mathbf{G} \leftarrow \{\tilde{\mathbf{U}}^{(1)}, \tilde{\mathbf{U}}^{(2)}\}$ // Generate orthogonal basis

$\xi^{(n)} = \mathbf{G}^T(\mathbf{f}^{(n)} - \mathbf{g}_0)$, $n = 1, 2, \dots, N$ // Proj.

$\{\tilde{\xi}, \mathbf{C}\} \leftarrow \{\xi^{(n)}\}$ // Second layer PCA

$\mathbf{g}'_0 = \mathbf{g}_0 + \mathbf{G}\tilde{\xi}$ // Compute the offset

$\tilde{\mathbf{G}}' = \mathbf{G}\tilde{\mathbf{V}}$ // Compute the final truncated basis

B. Vesicle Removal

Given the trained affine vesicle basis $(\mathbf{g}'_0, \tilde{\mathbf{G}}')$ and a novel micrograph with vesicle center positions and radii, the vesicle structures are removed as follows. Each vesicle is intensity normalised, windowed and shape normalised, after which they are projected onto the residual subspace by the orthogonal projection matrix (cf. (7)),

Algorithm 2 Vesicle Removal

Input: A novel micrograph $\mathbf{f}_{\text{micr}}$, vesicle centres $(x_{\text{micr}}^{(n)}, y_{\text{micr}}^{(n)})$, vesicle radii $R^{(n)}$, $n = 1, 2, \dots, N$ and vesicle basis $(\mathbf{g}'_0, \tilde{\mathbf{G}}')$.

Output: Micrograph $\tilde{\mathbf{f}}_{\text{micr}}$ with the vesicles removed.

for each vesicle $n = 1 \dots N$ **do**

$\mathbf{f}_{\text{norm}}^{(n)} = \mathbf{T}_{\text{int}}^{(n)}(\mathbf{f}_{\text{micr}}^{(n)})$ // Intensity norm. (13)

$\mathbf{f}_{\text{win}}^{(n)} = \mathbf{T}_{\text{win}}^{(n)}(\mathbf{f}_{\text{norm}}^{(n)})$ // Windowing (9)

$\mathbf{f}^{(n)} = \mathbf{T}_{\text{geom}}^{(n)}(\mathbf{f}_{\text{win}}^{(n)})$ // Shape norm. (9)

$\mathbf{r}^{(n)} = \mathbf{P}_{\perp}(\mathbf{f}^{(n)} - \mathbf{g}'_0)$ // Projection (34)

$\mathbf{b}^{(n)} = (\mathbf{I} - \mathbf{T}_{\text{win}}^{(n)})(\mathbf{f}_{\text{norm}}^{(n)})$ // Windowing

$\tilde{\mathbf{f}}^{(n)} = \mathbf{T}_{\text{geom}}^{(n)-1}(\mathbf{r}^{(n)}) + \mathbf{b}^{(n)}$ // Superimpose

$\tilde{\mathbf{f}}_{\text{micr}} \leftarrow \mathbf{T}_{\text{int}}^{(n)-1}(\tilde{\mathbf{f}}^{(n)})$ // Reverse norm.

end for

$$\mathbf{P}_{\perp} = \mathbf{I} - \mathbf{P} = \mathbf{I} - \tilde{\mathbf{G}}' \tilde{\mathbf{G}}'^T \quad (33)$$

that yields the residual

$$\mathbf{r}^{(n)} = \mathbf{P}_{\perp}(\mathbf{f}^{(n)} - \mathbf{g}'_0). \quad (34)$$

The vesicle removal is achieved by reversing the shape normalisation for the residual and superimposing the result with the background windowed signal, or,

$$\tilde{\mathbf{f}}^{(n)} = \mathbf{T}_{\text{geom}}^{(n)-1} \left(\mathbf{r}^{(n)} \right) + \mathbf{b}^{(n)} \quad (35)$$

after which the reverse intensity normalisation follows. The complete vesicle removal procedure is described in Algorithm 2, and illustrated in Fig. 12.

V. Model Selection

The remaining task is to determine the complexity of the vesicle model, that is, to find the appropriate number of components $(\tilde{I}_1, \tilde{I}_2, \tilde{K})$. We assume i.i.d. additive noise, i.e., the observed normalised vesicle follows the model

$$\mathbf{f}^{(n)} = \hat{\mathbf{f}}^{(n)} + \epsilon \quad (36)$$

where $\hat{\mathbf{f}}^{(n)}$ is the noise-free vesicle and r the noise vector. Thus, in the *model selection* we aim at a model that has enough parameters to describe the structure variations in the vesicle population so that it does not *underfit* the data. However, the model should not be too complex so that it is able to represent the unrelated background noise, i.e., the model should not *overfit* the data.

To solve the model selection problem, we use two distinct principles: *Morozov's discrepancy principle* together with the *Occam razor* principle. According to Morozov's discrepancy principle [21], the right complexity is achieved when the residual has the identical level to the assumed noise level. Typically in ill-posed problems, the application of this principle however yields an ambiguity, since the solution yielding the desired level is generally not unique. To resolve this ambiguity, we apply the Occam razor principle that states that among multiple solutions, the simplest one is the most likely, i.e., we select the *simplest solution* at the desired residual noise level.

To use Morozov's discrepancy principle, we estimate the background noise level $\epsilon_{\text{noise}}^{(n)}$ by using (14) in the discrete form. The residual level $\epsilon_{\text{res}}^{(n)}$ is estimated similarly to (18) with the difference that the signal is replaced by only the residual part. As the criterion, derived from Morozov's discrepancy principle, we use the mean difference

$$d(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = \bar{\epsilon}_{\text{noise}} - \bar{\epsilon}_{\text{res}}. \quad (37)$$

between the estimated background and residual noise levels. As illustrated in Fig. 13, the case $d(\tilde{I}_1, \tilde{I}_2, \tilde{K}) > 0$ represents models which have been overfitted, whereas the case $d(\tilde{I}_1, \tilde{I}_2, \tilde{K}) < 0$ corresponds to underfitted models. The desired complexity would be achieved by the models on the isosurface $d(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = 0$, i.e., when $\bar{\epsilon}_{\text{noise}} = \bar{\epsilon}_{\text{res}}$, assuming that the parameters were continuous variables. In practice, since the model parameters are discrete, we associate the isosurface by the set of parameter configurations for which the

distance from the continuous, interpolated approximation of the isosurface is smaller than unity in the infinity norm.

To finish the model selection procedure, by using the Occam razor principle, we select the simplest solution from the family of solutions at the desired noise level. To select the simplest model, we need to characterise the effective number of parameters in the model. Using the two layer HOSVD model, the number of effective parameters is characterised by the following theorem.

Theorem 8: The number of parameters of the two-layer HOSVD model is

$$c(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = I_1 I_2 + \tilde{I}_1 \tilde{I}_2 + \tilde{K} + \tilde{I}_1 \tilde{I}_1 - \frac{\tilde{I}_1 + 1}{2} \tilde{I}_1 + \tilde{I}_2 \tilde{I}_2 - \frac{\tilde{I}_2 + 1}{2} \tilde{I}_2 + \tilde{I}_1 \tilde{I}_2 \tilde{K} - \frac{\tilde{K} + 1}{2} \tilde{K} \quad (38)$$

Algorithm 3 Model Selection

Input: Volume of differences $d(\tilde{I}_1, \tilde{I}_2, \tilde{K})$ evaluated at $\tilde{I}_1 = 1, 2, \dots, \tilde{I}_1^{\max}$, $\tilde{I}_2 = 1, 2, \dots, \tilde{I}_2^{\max}$, $\tilde{K} = 1, 2, \dots, \tilde{K}^{\max}$ on vesicle set $\{\mathbf{f}^{(n)}\}$; local noise and residual levels $\epsilon_{\text{noise}}^{(n)}, \epsilon_{\text{res}}^{(n)}$.

Output: Parameter estimates for $\tilde{I}_1$, $\tilde{I}_2$ and $\tilde{K}$.
 Step 1: Make the continuous extension $d_{\text{ext}}(i_1, i_2, k)$ for $d(\tilde{I}_1, \tilde{I}_2, \tilde{K})$ by interpolation.
 Step 2: Find $(\hat{i}_1, \hat{i}_2, \hat{k}) = \arg\min_{i_1, i_2, k} c(i_1, i_2, k)$ with subject to $d_{\text{ext}}(i_1, i_2, k) = 0$
 Step 3: Return $\tilde{I}_1 \leftarrow \text{round}(\hat{i}_1)$, $\tilde{I}_2 \leftarrow \text{round}(\hat{i}_2)$, $\tilde{K} \leftarrow \text{round}(\hat{k})$.

where I_1 and I_2 are the number of discrete radii and the number of discrete angles in a vesicle image, respectively; $\tilde{I}_1$ and $\tilde{I}_2$ are the corresponding numbers of retained components after truncation; and $\tilde{K}$ is the number of retained second layer basis vectors.

Proof: Starting from the first layer, the vesicle population mean has $I_1 I_2$ parameters. The 2d radial basis can be parameterised by parameterising the one-dimensional basis $\{\mathbf{u}_{j1}\}$. Since

the basis is orthonormal, it has $\sum_{i_1=1}^{\tilde{I}_1} (I_1 - i_1)$ degrees of freedom. Similarly, the 2d

angular basis has $\sum_{i_2=1}^{\tilde{I}_2} (I_2 - i_2)$ degrees of freedom. The second layer offset has $\tilde{I}_1 \tilde{I}_2$

parameters, and the second layer basis $\sum_{k=1}^{\tilde{K}} (\tilde{I}_1 \tilde{I}_2 - k)$ parameters. The retained second layer eigenvalues have $\tilde{K}$ parameters. The number of parameters, and thus the complexity

$c(\tilde{I}_1, \tilde{I}_2, \tilde{K})$ of the whole model is obtained by summing up the number of parameters above, and the claim follows.

The complete model selection procedure is described in Algorithm 3.

VI. Experiments

In the experiments, we used the transmission electron microscope training data set consisting of ten 2048×2048 micrographs with all together $N = 1143$ vesicles taken by different CTF settings. The model selection was performed with $N = 436$ vesicles extracted from five additional, independent micrographs, where some vesicles had a protein particle attached. The data sets contained the vesicle center positions and radii, automatically detected as reported in [27]. The bright aura around the vesicle signal is influenced by both the presence of vesicle neighbours and the microscope settings, see Fig. 21. The vesicle window radius was thus estimated so that it fully covers a few pixels of the bright vesicle aura, after which the transition region of the window brings the signal amplitude visually reasonably onto the background level (Fig. 6), to include the aura in the model while

limiting influence by vesicle neighbours. This was achieved by the setting $\alpha = \frac{1}{12}$, for the Tukey window (12). The normalised image pixel size was determined by the largest vesicle in the database. The vesicle models were trained as described in Alg. 1. The training and testing were implemented in Matlab. The computations were performed with a conventional laptop computer, where the training took 110s, model selection 1h 20 min, and the vesicle removal 9–12s/micrograph.

The model selection was carried out as described in Alg. 3, over the model class

$\tilde{I}_1 \in \{1, 2, \dots, 15\}, \tilde{I}_2 \in \{1, 2, \dots, 15\}, \tilde{K} \in \{1, 2, \dots, \tilde{I}_1 \tilde{I}_2\}$. As the ground truth for model selection, we manually evaluated the most relevant combinations of $(\tilde{I}_1, \tilde{I}_2, \tilde{K})$ and picked up the combination (7, 7, 40) that showed visually the best vesicle removal result. The estimated continuous extension of the mean difference $d_{\text{ext}}(i_1, i_2, k)$ between the background noise and residual noise is illustrated in Fig. 14(a). In other words, the isosurface in Fig. 14(a) intersects all zero-valued residuals, corresponding to valid vesicle models, in the volume of residuals measured for each model complexity $(\tilde{I}_1, \tilde{I}_2, \tilde{K})$ in the experiment.

Our model selection method yielded the result $(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = (6, 5, 27)$ that is illustrated in Fig. 14(b), which displays the isosurface from Fig. 14(a) in the volume of model complexities. Here, the solution is found at the point where the isosurface of constant complexity and the desired residual level isosurfaces tangent each other. The result is reasonably close to the ground truth and thus indicates that the model selection was successfully performed.

The mean vesicle $\mathbf{g}'_0$ and basis components in $\tilde{\mathbf{G}}'$ of the selected model are illustrated in Fig. 15. One can clearly see the dominance of the radial components against the angular ones, whereas some asymmetric structures are characterised by the latter components. The importance of the radial components can be likewise seen in Fig. 14(b), as many radial and few angular components yields a small $\tilde{K}$, whereas a model with many angular and few radial components yields a much larger $\tilde{K}$. In Fig. 14(a), one can also see that the parameter $\tilde{K}$ is saturated to the value $I_1 I_2$ when both I_1 and I_2 are small but the surface bends as soon

as there is redundancy in the combined bases elements that the second layer PCA is able to remove.

To understand how the vesicle removal affects the protein signal and how the vesicle removal performs under varying amounts of noise, we first generated simulated protein projections [38] and added Gaussian noise from increasing inverse signal-to-noise-ratio (SNR). A realistic amount of noise versus the membrane protein is about ten times in the inverse SNR [38]. Then we projected the simulated images, without added vesicles, onto the above trained vesicle subspace and onto its orthogonal complement. Due to the assumption of additive Gaussian noise and linearity of the subspace projection, it is equivalent to project the noise and protein components separately and superimpose the results, as Fig. 16 illustrates. It can also be seen what part of the protein signal and noise signal disappears into the vesicle subspace. From the simulated protein particle, 5% is projected onto the vesicle subspace in the square L_2 norm that illustrates the overlap of the protein and vesicle subspaces. To see the effect of noise onto the protein, we simulated the relative energy of the artifact caused by the vesicle removal as reported in Fig 17. Similarly, we evaluated the effect of noise on the removal of simulated vesicles [40]. It can be seen that the noise completely dominates the projections when the inverse SNR grows to hundred as the norm ratios tend to unity.

Finally, we compared the proposed approach against the robust principal component analysis (RPCA) [6]. However, in our application, there are three distinct components opposed to the standard setting of the RPCA since we have the low rank component (vesicle), sparse component (protein and background dirt), and noise which is not sparse. The original RPCA problem considers

$$\min_{\mathbf{L}, \mathbf{S}} \|\mathbf{L}\|_* + \lambda \|\mathbf{S}\|_1 : \mathbf{L} + \mathbf{S} = \mathbf{M} \quad (39)$$

where $\mathbf{L}$ is a low rank matrix, $\mathbf{S}$ is the sparse part of the measurement matrix $\mathbf{M} = \mathbf{A}_{(3)} - \mathbf{A}_{0,(3)}$, and λ is a parameter. To adapt the RPCA to our problem we thus modified the RPCA problem to

$$\min \|\mathbf{L}\|_* + \lambda \|\mathbf{S}\|_1 : \|\mathbf{L} + \mathbf{S} - \mathbf{M}\|_{\text{fro}} = C_\sigma, \quad (40)$$

where C_σ is a constant fixing the noise level, and solved it by Alternating Direction Method of Multipliers (ADMM) [4].

The resulting modified RPCA components are shown in Fig. 18. It can be seen that the low rank components are more noisy or less smooth than the basis elements of the proposed method (Fig. 15). It is natural since the modified RPCA model is not generated by separating the radial and angular parts that significantly reduces the dimensionality of the proposed method. The modified RPCA model likewise fails to capture the asymmetries of the vesicles. The vesicle removal results, using the learnt low rank components for a novel micrograph, are qualitatively compared in Fig. 19 and Fig. 20. The modified RPCA model works reasonably well, but the vesicle removal leaves more artifacts in the micrograph than by the proposed method. It can be seen (Fig. 20) that the noise-free estimates of the vesicles,

displayed in the center column, capture the individual structural features, yet are robust to structural noise both inside and outside the vesicles (row 3, 4), or attached protein (row 2). The vesicle removal by the two methods using a complete micrograph is furthermore compared in Fig. 21, where it can be seen that the modified RPCA model consequently leaves vesicle artifacts. We conclude that the modified RPCA method leaves more artifacts after removal than the proposed method. The latter is also robust to the global illumination changes over the micrographs, as the vesicle structures have been successfully removed without an addition of notable systematic artifacts.

VII. Conclusions and Future Work

In this paper, we have proposed a method to learn and remove objects from images, where our application is the removal of lipid vesicle structures from electron micrographs to facilitate 3d reconstruction of proteins. To model the object, we first estimated the vesicle model subspace by the Higher Order Singular Value Decomposition in the polar coordinate plane. The two dimensional HOSVD basis was then converted to one dimensional by proposing a PCA layer to the HOSVD model. The model selection was performed by applying Occam razor and Morozov's discrepancy principle for the hierarchical model. Experiments showed that the proposed approach efficiently captures the vesicle structure variations in a low dimensional subspace and therefore the removal procedure has a minimal impact on the underlying signals such as those of the interesting membrane proteins. In the future [18], we will apply the vesicle removal method to facilitate the membrane protein reconstruction in cryogenic single particle electron microscopy, where the vesicle structures deteriorate the protein signal. It is likewise possible to integrate the vesicle model directly into the reconstruction problem that we deem as a promising approach.

VIII. Software

The computer source code will become available at authors' web page.

Acknowledgments

This work was supported in part by the National Institutes of Health under Grant R01 NS021501 and Grant R21 NS085776 and in part by the Villum Foundation through the Centre for Stochastic Geometry and Advanced Bioimaging, Denmark. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jong Chul Ye.

Biographies

Katrine Hommelhoff Jensen received the M.Sc. degree from the Department of Computer Science, University of Copenhagen, Denmark, in 2006, with a specialization in image analysis and 3d computer graphics. She is currently pursuing the Ph.D. degree with the Image Section, Department of Computer Science, University of Copenhagen. She was a 3d Software Developer at the company 3Shape A/S, Denmark, from 2004 to 2007, and at the company Follow-Me! Technology Systems GmbH, Germany, from 2007 to 2011. Her work is focused on statistical methods for single-particle cryo-EM. Her research interests include bio-/medical image analysis and computer graphics, in particular 3d imaging and reconstruction, shape analysis, mathematical modeling, and statistical inversion.

Fred J. Sigworth was born in Berkeley, CA, in 1951. He received the B.S. degree in applied physics from the California Institute of Technology, Pasadena, in 1974, and the Ph.D. degree in physiology from Yale University, New Haven, CT, in 1979. He was a Research Fellow with the Max Planck Institute for Biophysical Chemistry, Göttingen, Germany, from 1979 to 1984, where he was with E. Neher on the development and applications of the patch-clamp technique for recording single ion-channel currents. Since 1984, he has been a Faculty Member with the Department of Cellular and Molecular Physiology, Yale University. His research interests include the structure and function of ion channel proteins and the development of techniques to study ion channels. He is a member of the Biophysical Society, the Society of Neuroscience, and the American Scientific Affiliation.

Sami Sebastian Brandt received the Ph.D. degree from the Helsinki University of Technology, Finland, in 2002, and the habilitation degree on the geometric branch of computer vision from the University of Oulu, Finland, in 2007. After the Ph.D. degree, he spent one year as a Research Scientist with the Instrumentarium Corporation Imaging Division, Finland, a couple of years in the Helsinki University of Technology, Oulu University, Finland, Malmö University, Sweden, and Nordic Bioscience Imaging/Synarc Imaging Technologies in Denmark. He is currently an Associate Professor with the Image Group, University of Copenhagen, Denmark, and a Senior Mathematical Software Developer with 3Shape. His research interests include applied mathematics, statistical

inverse problems, Bayes methods, electron tomography, single particle reconstruction, geometric computer vision, and image analysis.

REFERENCES

- [1]. Akaike H. A new look at the statistical model identification. *IEEE Trans. Autom. Control.* Dec; 1974 19(6):716–723.
- [2]. Akansu, AN.; Haddad, RA. *Multiresolution Signal Decomposition: Transforms, Subbands, and Wavelets.* Academic; San Diego, CA, USA: 1992.
- [3]. Ando, T. *Bayesian Model Selection and Statistical Modeling.* CRC Press; Boca Raton, FL, USA: 2010. *Statistics: A Series of Textbooks and Monographs*
- [4]. Boyd S, Parikh N, Chu E, Peleato B, Eckstein J. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.* Jan; 2011 3(1): 1–122.
- [5]. Brown, RG.; Hwang, PYC. *Introduction to Random Signals and Applied Kalman Filtering with MATLAB Exercises.* Wiley; New York, NY, USA: 1996.
- [6]. Candès EJ, Li X, Ma Y. Robust principal component analysis? *J. ACM.*
- [7]. Dhillon IS, Sra S. Generalized nonnegative matrix approximations with Bregman divergences, *Proc. NIPS.* 2005:283–290.
- [8]. Donoho DL. De-noising by soft-thresholding, *IEEE Trans. Inf. Theory.* May; 1995 41(3):613–627.
- [9]. Donoho DL. Compressed sensing, *IEEE Trans. Inf. Theory.* Apr; 2006 52(4):1289–1306.
- [10]. Elad, M. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing.* Springer-Verlag; New York, NY, USA: 2010.
- [11]. Elad M, Aharon M. Image denoising via sparse and redundant representations over learned dictionaries, *IEEE Trans. Image Process.* Dec; 2006 15(12):3736–3745. [PubMed: 17153947]
- [12]. Frank, J. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies.* Oxford Univ. Press; Oxford, U.K.: 2006.
- [13]. Hjort NL, Claeskens G. Frequentist model average estimators, *J. Amer. Statist. Assoc.* Dec.2003 98:879–899.
- [14]. Hoeting JA, Madigan D, Raftery AE, Volinsky CT. Bayesian model averaging: A tutorial, *Statist. Sci.* 1999; 14(4):382–401.
- [15]. Hyvärinen A. Sparse code shrinkage: Denoising of nongaussian data by maximum likelihood estimation, *Neural Comput.* 1999; 11(7):1739–1768. [PubMed: 10490945]
- [16]. Hyvärinen, A.; Karhunen, J.; Oja, E. *Independent Component Analysis.* Wiley; New York, NY, USA: 2004.
- [17]. Hyvärinen A, Köster U. FastISA: A fast fixed-point algorithm for independent subspace analysis, *Proc. ESANN.* 2006:371–376.
- [18]. Jensen KH, Brandt SS, Sigworth FJ. Statistical modeling and removal of lipid membrane projections for cryo-EM structure determination of reconstituted membrane proteins,. submitted for publication.
- [19]. Jolliffe, I. *Principal Component Analysis.* Springer-Verlag; New York, NY, USA: 2002. *Springer Series in Statistics*
- [20]. Kadane JB, Lazar NA. Methods and criteria for model selection, *J. Amer. Statist. Assoc.* Mar. 2004 99:279–290.
- [21]. Kaipio, J.; Somersalo, E. *Statistical and Computational Inverse Problems.* Springer; 2005.
- [22]. Kass RE, Raftery AE. Bayes factors, *J. Amer. Statist. Assoc.* 1995; 90(430):773–795.
- [23]. Kohavi R. A study of cross-validation and bootstrap for accuracy estimation and model selection, *Proc. IJCAI.* 1995:1137–1143.
- [24]. Kolda TG, Bader BW. Tensor decompositions and applications, *SIAM Rev.* 2009; 51(3):455–500.
- [25]. J. Koles Z, Lazar MS, Zhou SZ. Spatial patterns underlying population differences in the background EEG, *Brain Topogr.* 1990; 2(4):275–284. [PubMed: 2223384]

- [26]. De Lathauwer L, De Moor B, Vandewalle J. A multilinear singular value decomposition, *SIAM J. Matrix Anal. Appl.* 2000; 21(4):1253–1278.
- [27]. Liu Y, Sigworth FJ. Automatic cryo-em particle selection for membrane proteins in spherical liposomes, *J. Struct. Biol.* 2014; 185(3):295–302. [PubMed: 24468290]
- [28]. Lomax, RG.; Hahs-Vaughn, DL. *Statistical Concepts: A Second Course*. Lawrence Erlbaum Associates; Mahwah, NJ, USA: 2007.
- [29]. Mallat, S. *A Wavelet Tour of Signal Processing: The Sparse Way*. 3rd. Academic; San Diego, CA, USA: 2008.
- [30]. Mallows CL. Some comments on C_p Technometrics. 2000; 42(1):87–94.
- [31]. Marion, A. *An Introduction to Image Processing*. Chapman & Hall; London, U.K.: 1991.
- [32]. Nixon, M.; Aguado, AS. *Feature Extraction & Image Processing for Computer Vision*. Academic; San Diego, CA, USA: 2012.
- [33]. Raunser S, Walz T. Electron crystallography as a technique to study the structure on membrane proteins in a lipidic environment, *Annu. Rev. Biophys.* Jun.2009 38:89–105. [PubMed: 19416061]
- [34]. Rissanen J. Mdl denoising, *IEEE Trans. Inf. Theory*. Nov; 1999 46(7):2537–2543.
- [35]. Rissanen, J. *Information and Complexity in Statistical Modeling*. Springer-Verlag; New York, NY, USA: 2007. Information Science and Statistics
- [36]. Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms, *Proc. Annu. Int. Conf. Center Nonlinear Stud. Experim. Math.* 1992:259–268.
- [37]. Schwarz G. Estimating the dimension of a model, *Ann. Statist.* 1978; 6(2):461–464.
- [38]. Shang Z, Sigworth FJ. Hydration-layer models for cryo-EM image simulation, *J. Struct. Biol.* 2012; 180(1):10–16. [PubMed: 22609687]
- [39]. Strang, G. *Introduction to Linear Algebra*. Wellesley Cambridge Press; Wellesley, MA, USA: 2009.
- [40]. Wang L, S. Bose P, Sigworth FJ. Using cryo-EM to measure the dipole potential of a lipid membrane, *Proc. Nat. Acad. Sci. USA*. 2006; 103(49):18528–18533. [PubMed: 17116859]
- [41]. Wang L, Sigworth FJ. Structure of the BK potassium channel in a lipid membrane from electron cryomicroscopy, *Nature*. 2009; 461(7261):292–295. [PubMed: 19718020]
- [42]. Wang L, Sigworth FJ. Liposomes on a streptavidin crystal: A system to study membrane proteins by cryo-EM, *Methods Enzymol.* 2010; 481:147–164. [PubMed: 20887857]
- [43]. Weickert, J. *Anisotropic Diffusion in Image Processing*. Teubner Verlag; 1998.
- [44]. Zhao Z, Singer A. Fourier–Bessel rotational invariant eigenimages, *J. Opt. Soc. Amer.* 2013; 30(5):871–877.

Fig. 1. Membrane protein reconstituted into a lipid vesicle, frozen in vitreous ice to be imaged in a cryogenic transmission electron microscope.

Fig. 2.

(a) Micrograph with vesicles where some of them contain a membrane protein particle (blue rectangles), one is amplified; (b) after the vesicle signal removal. A vesicle model must capture the variation in size and shape, yet separate structure from background noise and structural noise like pieces of ice and protein particles.

Fig. 3.

The relation between the Cartesian coordinate representation $f(x, y)$ (right) and polar coordinate representation $f(\rho, \varphi)$ (center) in the generation of vesicle HOSVD basis for the normalised, windowed vesicles, as $\rho = \sqrt{x^2 + y^2}$ and $\varphi = \text{atan2}(y, x)$. Two radial profiles c_i

$(x, y) = f(\rho, \varphi = \varphi_i)$, $i = 1, 2$ evaluated at $\varphi_1 = \frac{50}{251}2\pi$, $\varphi_2 = \frac{140}{251}2\pi$ are illustrated as 2d images (left) together with two angular profiles $f(\rho = \rho_j, \varphi)$, $j = 1, 2$ evaluated by $\rho_1 = 5$ and $\rho_2 = 15$ (right). The HOSVD basis, constructed in this paper, can be seen as the basis for these kinds of radial and angular structures.

Fig. 4.

A set of size normalised vesicles from different micrographs, illustrating differences in vesicle contrast, background level, and intensity distribution.

Fig. 5.

Step-by-step illustration of the vesicle structure removal from a micrograph: (a) Vesicle in micrograph, (b) normalised vesicle, (c) windowed, normalised vesicle, (d) projection onto the windowed vesicle subspace, (e) vesicle subtraction as the projection onto the orthogonal complement of the windowed vesicle subspace.

Fig. 6.

(a) A normalised vesicle; (b) the corresponding Tukey window; (c) the windowed vesicle. The vesicle signal is captured by the window while a small overlap of other vesicles may be visible on the transition region.

Fig. 7.
An example how the background level drifts on a micrograph.

Fig. 8.

(a) A vesicle image; (b) the corresponding local background window; (c) windowed background pixels from which the background level is estimated.

Fig. 9.

The three-way array $\mathcal{A} \in \mathbb{R}^{I_1 \times I_2 \times I_3}$ illustrated with three vesicle images and their unfoldings $\mathcal{A}_{(1)}$ and $\mathcal{A}_{(2)}$. Each column in $\mathcal{A}_{(1)}$ is a radial profile evaluated at a fixed angle in a vesicle image. Each column in $\mathcal{A}_{(2)}$ is an angular profile evaluated at a fixed radius in a vesicle image.

Fig. 10.

Learnt vesicle basis visualised in the Cartesian coordinate frame sorted in the order of importance: (top row) pure radial components; (left column) pure angular components; (remaining basis elements) respective combinations of the pure radial and angular components.

Fig. 11.

Construction of the second layer for the HOSVD model. (Left) The most descriptive paired 2d bases are selected by pairing the most descriptive radial and angular bases. (Right) The new basis is generated by PCA as the linear combination of the selected paired components while the least significant new basis elements are truncated.

Fig. 12.

Illustration of the vesicle removal in the function of the number of basis components. (a) The normalised vesicle; (b) the mean subtracted. Projection onto the affine residual subspace with (c) 2, (d) 10, (e) 37, and (f) 42 basis elements.

Fig. 13.

Sketch of the behaviour of the mean difference $d = \bar{\epsilon}_{\text{noise}} - \bar{\epsilon}_{\text{res}}$ between the background noise and residual noise with different number of basis elements. The cases $d > 0$ represent overfitted models, $d < 0$ underfitted models, and, $d = 0$ the models with the desired complexity.

Fig. 14.

(a) The isosurface where the mean residual noise level equals to the mean background noise level or $d_{\text{ext}}(i_1, i_2, k) = 0$; the colours outside the isosurface indicate the residual level. (b) Illustration of the model selection result; the right complexity is achieved when the isosurfaces $d_{\text{ext}}(i_1, i_2, k) = 0$ (in colour) and $\alpha(i_1, i_2, k) = c_{\min}$ (gray) meet at the tangent point $(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = (6, 5, 27)$; the surface colour indicates the complexity.

Fig. 15.

The selected HOSVD model $(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = (6, 5, 27)$. (Top left) The mean vesicle $\mathbf{g}'_0$; (from the second in top left to right) the 27 basis images in the order of importance.

Fig. 16.

Illustration of how the projection onto the vesicle subspace (bottom row) and to its orthogonal complement (top row) affects the simulated protein signal (middle row, left), noise (middle row, middle) and the protein signal superimposed with noise (middle row, right). The amount of noise is controlled by the parameter α . Both the protein and noise have non-zero projection onto the vesicle subspace that induces the artifacts shown on the top row.

Fig. 17.

Relative artifact energy in the function of the inverse signal-to-noise-ratio. (Solid line) Induced relative artifact defined as the square L_2 norm ratio between the projection of protein plus noise onto the orthogonal complement of the vesicles minus the true protein signal and the projection onto the orthogonal complement. (Dashed line) Induced relative artifact in simulated vesicle removal defined as the square L_2 norm ratio between the projection of the simulated vesicle plus noise onto the vesicle subspace minus the true vesicle and the projection onto the vesicle subspace.

Fig. 18.

Modified RPCA model dividing the vesicles into low rank, sparse and noise components. (top) 14 of the most significant of the 78 low rank components; (middle) 14 sparse components; (bottom) 14 dense, noise components. By this decomposition the vesicle image naturally divides into the vesicle, substructure, and noise parts, respectively.

Fig. 19.

Vesicle image removal for novel vesicles by using the modified RPCA model with 78 low rank components. (Column a, b) normalised vesicle images; (c) projection onto the vesicle subspace; (d-e) the removal result. The images (b-d) are slightly blurred by a Gaussian filter to aid visual inspection, c.f. Fig. 20.

Fig. 20.

Vesicle image removal for novel vesicles by using the proposed method with the automatically selected model $(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = (6, 5, 27)$. (Column a, b) normalised vesicle images; (c) projection onto the vesicle subspace; (d-e) the removal result. The images (b-d) are slightly blurred by a Gaussian filter to aid visual inspection, c.f. Fig. 19.

Fig. 21.

Vesicle removal result for a full micrograph. (a) Original micrograph; (b) vesicles removed by the modified RPCA model with 78 low rank components; (c) vesicles removed by the proposed method with the automatically selected model $(\tilde{I}_1, \tilde{I}_2, \tilde{K}) = (6, 5, 27)$. Vesicle center positions are marked with blue dots and all images have been post-processed by a low-pass filter to enhance structural details.