

Supplementary Material for
“Developing a Denoising Filter for Electron Microscopy and
Tomography Data in the Cloud” by
Zbigniew Starosolski, Marek Szczepanski, Manuel Wahle,
Mirabela Rusu, and Willy Wriggers

1. On the Treatment of Noise in Electron Microscopy and Tomography

As described in Section 1 of the main text, recorded cryo-EM and tomography images show a very low SNR due to the low radiation dose that is necessary to protect the biological specimen.

In the past few decades, the influence of environmental factors has been carefully studied, and a consensus has emerged that specimens at lower temperature are more robust to high-energy electrons in terms of specimen degradation (Wright et al., 2006). However, high-energy electrons also cause changes in ice density and induce displacements of the specimen. Recent studies have shown that for high resolution (4-10 Å) cryo-EM reconstructions, there is little benefit of a temperature below liquid nitrogen (Bammes et al., 2010), and only the lower resolution (20-60 Å) maps tend to benefit from a lower temperature. Therefore, noisy micrographs will remain an important problem in the foreseeable future.

Single-particle and tomographic reconstruction overcome these limitations by processing vast numbers of individual particle projections. The particles are identified on 2D micrographs, classified or aligned according to their rotation, and averaged before (cryo-EM) or after 3D reconstruction (tomography) to reduce noise and strengthen the true signal. Diverse processing techniques for digitization, particle picking, and classification are well established Frank (2006a,b).

The limitations of 3D data mentioned above pose a formidable challenge to denoising methods. Unfortunately, linear filtering methods such as Gaussian averages are not able to effectively reduce the excessive amount of noise while preserving the detail of edges (Gonzales and Woods, 2002). One solution developed by us and many others (see below) is to allow for non-linearities in the filtering.

We acknowledge that cryo-EM and tomography data are often collected and processed in a way that preserves linearity. Ideally, one wishes for 3D maps to be interpretable in terms of macromolecular mass density. The introduction of non-linearities distorts this relation and, therefore, must be applied with caution. The main text shows that non-linear filters can be advantageous, but we believe that they should be limited to cases where linear filters are clearly insufficient.

Non-linear methods also require a tradeoff between noise reduction and precision based on empirical criteria that are unrelated to the first principles of image formation or 3D reconstruction. In fact, the most useful denoising methods are those that adaptively change their performance based on the local conditions of the specific data and use intricate noise models (Baxter et al., 2009). For example, the widely used bilateral filtering method employs two competing low-pass filters working together (Jiang et al., 2003). One low-pass filter is applied to intensities while the other is applied to distances. The filter parameters are set for all the input data, causing a loss of robustness to different noise values. An extension of this method is a discriminative bilateral filter in which an additional photometric discriminant function is embedded (Pantelic et al., 2006). The additional discriminant distinguishes between edges and impulse noise. With properly chosen parameters, high frequency noise pixels are smoothed while the object edge detail is preserved.

An alternative approach builds on anisotropic nonlinear diffusion (Fernandez and Li, 2003) by adapting the parameters of a Gaussian-based filter to local features. These local features are determined by eigenanalysis of the averages of the image tensors.

More recently, Wei and Yin (2010) introduced a locally optimized adaptive non-local means method, which improves the approach originally developed by Buades et al. (2005). This method uses a local and global standard deviation of the image data and a parameter-sensitive decision rule. The adjustment of these parameters is a challenging task that requires knowledge and experience to achieve the desired effect. The DPSV filter developed by us and described in the following section is an alternative to these local variance- (or variability-) based filters.

2. The DPSV Filter Algorithm as Applied to Cryo-EM and Tomography Data

A schematic overview of the algorithm is presented in Fig. 1 of the main text. To assist users, we will provide here, for the first time, a complete review of the DPSV theory as applied to cryo-EM and tomography maps.

2.1. Local Search Strategy and the Digital Path Approach

This section describes a local search strategy based on digital paths generated by a self-avoiding walk (Havlin and Ben-Avraham, 1982). Two neighborhood models can be defined in 2D data: a 4-neighborhood and a 8-neighborhood. In the case of 3D data, a 6-neighborhood and a 26-neighborhood can be defined (see Szczepanski,

$P \setminus M$	3	5	7	9	3	5	7	9
2D	4-neighborhoods				8-neighborhoods			
1	4	4	4	4	8	8	8	8
2	8	12	12	12	24	56	56	56
3	8	32	36	36	56	264	368	368
4		64	96	96		1056	2016	2336
5			240	280			9888	13608
6				718				73632
3D	6-neighborhoods				26-neighborhoods			
1	6	6	6	6	26	26	26	26
2	24	30	30	30		650	650	650
3		144	150	150		11928	15986	15986
4			720	726			350592	388994
5				3528				9051144

Supplementary Table 1: Cardinality of a set of paths for a 2D and 3D filtering scheme as a function of the mask size M and path length P

2008). The chosen model of the neighborhood affects the cardinality of the path set (see Suppl. Table 1), the computation time, and the quality of the results. The method is based on a local filtering approach where the local (2D or 3D) space is limited by the dimension, M , of a (square or cubic) mask, and the reach and curvature of the digital walk within the mask is set by its length, P . An increase in these parameters typically results in a more precise exploration of local space. Small mask sizes and a long paths ($P > (M - 1)/2$) cause paths to fold inside the mask, which allows them to sample curved or bent features. Short walk lengths $P \leq (M - 1)/2$ are unrestricted by the mask and are better suited for straight edges in the data. The mask size and path length should be adjusted to the structure found in the data upon examination.

Fig. 1b in the main text shows a 2D example of a digital walk through one of four nearest neighbors. The virtual particle (black circle) is initially at position p_i , then takes a step through the first neighbor $p_{i(1),*,1}$ (gray circle), and then continues through its second neighbor from the set: $N_3 = \{p_{i(1),1,2}, p_{i(1),2,2}, p_{i(1),3,2}\}$ (white circles). This walk will produce three different digital paths of length $P = 2$: $p_i \rightarrow p_{i(1),l,1} \rightarrow p_{i(1),l,2}$ where $l = 1 \dots 3$ for the nearest neighbor $p_{i(1),*,1}$. For all four neighbors, there will be a total of 12 paths. During each step along the path, the virtual particle is defined by its spatial position and the value of the intensity $I(p_i)$, which is the intensity of a pixel normalized across the image.

2.2. Spatial Intensity of Digital Paths

The spatial intensity of a path is defined as a connection cost. This is the maximum cost observed among pixels that are linked by one path. The individual connection cost is defined as the absolute difference of (normalized) intensities between the center pixel p_i and a linked pixel $p_{i(n),l,k}$, divided by their Euclidean distance. The connection cost of the l -th path passing through the n -th closest neighbor is thus defined as

$$\Lambda_{n,l}^P \{p_i, p_{i(n),l,1}, p_{i(n),l,2}, \dots, p_{i(n),l,P}\} = \max \left(\frac{|I(p_i) - I(p_{i(n),l,k})|}{\text{dist}(p_i, p_{i(n),l,k})} \right), \quad (1)$$

where the maximum is computed among $k = 1, \dots, P$ pixels along the l -th path, $|I(p_i) - I(p_{i(n),l,k})|$ is an absolute difference of intensities between linked pixels p_i and $p_{i(n),l,k}$ along the path, and $\text{dist}(\cdot)$ denotes the Euclidean distance.

2.3. Supervised Classification

Ideally, we wish to distinguish between paths that include noise and ones that do not, and then to recover pixel values using only the paths that include the true signal. Here, we assume that the digital paths that explore a smooth intensity landscape correspond to a noise-free neighborhood. Paths with a relatively large connection cost are assumed to hold high-frequency noise or cross the edge of an object. We note that any object edge detail is preserved by those paths that follow the edge, thus preventing a softening of edges.

At this point, we require a classification procedure to group the paths into two classes, which should be robust under different background levels and should preserve the paths with a relatively low connection cost. Following Smolka (2008), we adapt Fisher's discriminant analysis (FDA) using a maximized Fisher functional in Eqn. 2 below, which separates the set of paths into two classes that (ideally) correspond to signal and noise. The paths are first sorted in descending order based on the value of $r_k = \Lambda_{n,l}^P \{.\}$. The Fisher discriminant (Smolka, 2008) is then defined as

$$F(k) = \frac{[m_1(k) - m_2(k)]^2}{v_1(k) + v_2(k)}, \quad k = 1, \dots, u - 1, \quad (2)$$

where m_1, m_2 and v_1, v_2 are the means and the variances, respectively, as calculated for two classes:

$$m_1(k) = \frac{1}{k} \sum_{i=1}^k r_i, \quad m_2(k) = \frac{1}{u-k} \sum_{i=k+1}^u r_i, \quad (3)$$

$$v_1(k) = \sum_{i=1}^k (r_i - m_1)^2, \quad v_2(k) = \sum_{i=k+1}^u (r_i - m_2)^2, k = 1, \dots, u - 1. \quad (4)$$

Here, u is the cardinality of the set of paths. The cardinality depends on mask size M , path length P , and the neighborhood model as listed in Suppl. Table 1.

The discriminant analysis succeeds if one can find a k^* that fulfills

$$k^* = \arg \max F(k). \quad (5)$$

If FDA cannot divide the set of paths into two sets, the first $r_1 = \min(r_k)$, $k = 1, \dots, u - 1$ will be chosen for further processing (see Eqn. 6 below), i.e. neighboring pixels will always contribute to the filter response.

Fig. 1c in the main text provides an illustrative example of FDA for a set of four paths. The input set of the four paths labeled 1, 2, 3, and 4 is sorted with respect to the connection cost. After applying FDA, the labeled paths (3, 2, 4) remain for further analysis. For a more detailed description of FDA see Smolka (2008) and Kenney et al. (2001).

2.4. Similarity Function

The output intensity of the central pixel is calculated as the cost-weighted mean of the surviving paths through the neighborhood (similarity function):

$$I(p_i) = \sum_{n=1}^N \frac{I(p_{i(n),*,1}) \Xi(p_{i(n)})}{\sum_{n=1}^N \Xi(p_{i(n)})}, \quad (6)$$

where $I(p_{i(n),*,1})$ is the (normalized) intensity associated with the closest neighbor from the focal pixel p_i and size N is the number of neighbors in the neighborhood model. The cumulative cost is defined as

$$\Xi(p_{i(n)}) = \sum_{l=1}^v \left[K \left(\beta, \Lambda_{n,l}^P \left\{ p_i, p_{i(n),l,1}, p_{i(n),l,2}, \dots, p_{i(1),l,P} \right\} \right) \right], \quad (7)$$

where the kernel function $K(\beta, \Lambda)$ has to fulfill the imposed conditions: smooth, convex, and decreasing function on R_+ . After empirically exploring a number of possible functions (data not shown) we chose the exponential function $K(\beta, \Lambda) = e^{-\beta\Lambda}$ as a kernel. Section 4 explores the sharpness parameter, β .

As the mask (see diagram in Fig. 1a in the main text) is moved along pixels, i , the weighting process depends on the shared relations between the intensity associated

with the closest neighbor $p_{i(n),*,1}$ and the cumulative cost calculated in the mask described by $\Xi(p_{i(n)})$.

2.5. Usage and 3D Applications

In *Sculptor* version 2.1 (<http://sculptor.biomachina.org>), the DPSV filter can be applied to a map via the menus “Volume” → “DPSV Filter” (entering parameters M , P , and β in the pop-up dialog box). In *Situs* version 2.7 (<http://situs.biomachina.org>), a separate command-line utility *volfltr* was implemented for DPSV (see online user guide).

The applications in the main text were selected as qualitative examples of the use of DPSV for 3D filtering. To gain a quantitative understanding of the filter performance in future work, one could compare raw and filtered versions of noisy cryo-EM maps with known atomic resolution structures. There are several known cryo-EM structures that were also solved by X-ray crystallography and would lend themselves as test candidates for a quantitative (real-space or Fourier space) validation of filter robustness.

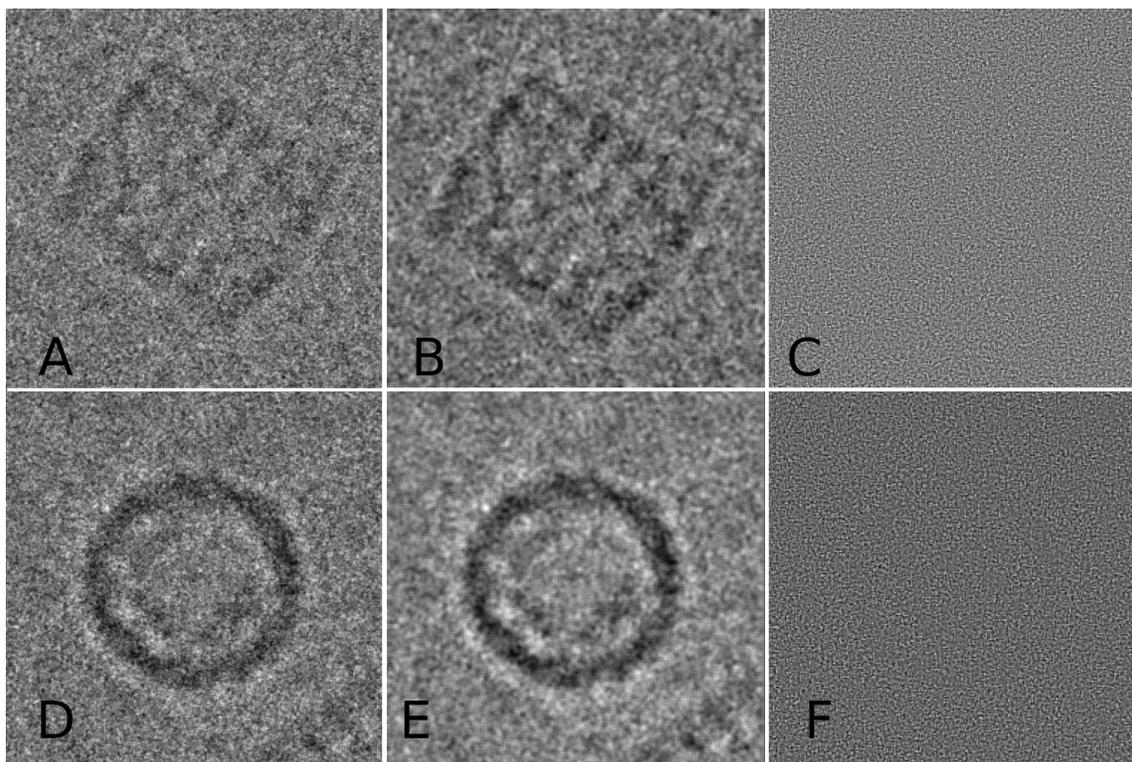
3. Denoising Experimental 2D Micrographs

In addition to the 3D applications in the main text, we investigated here the effect of filtration on experimental 2D micrographs. In the first 2D test (Suppl. Fig. 1), we examined the difference images of data before and after filtration (Russ, 2002). If there are any shadows resembling the original shapes in the difference images, then it is an indication that the filter greatly influences the structure hidden under noise. On the other hand, if the difference image shows only random and uniform noise, then the result indicates that the filtration preserves the features.

For this test, we chose a representative projection of the Keyhole Limpet Hemocyanin (KHL) protein (pixel size 2.22 Å at the specimen scale), see Zhu et al. (2003). Suppl. Fig. 1A and D present selected micrographs of the side and top view. Suppl. Fig. 1B and E show the results after filtering (with mask size $M = 7$, path length $P = 3$, 8-neighborhood model, and filter parameter $\beta = 0.0005$). The FDA reduced the number of paths entering Eqn. 6 by an average of 12% (Suppl. Fig. 1B) or 28% (Suppl. Fig. 1E). As can be seen in the differences between images Suppl. Fig. 1C and F, there are no noticeable shapes or structures in the noise. This result suggests that DPSV behaves properly for these conditions.

4. Denoising Simulated 2D Image Stacks and Setting β

This section demonstrates the effect of DPSV on class averages derived from the simulated projection of image stacks. The results are dependent on the noise level

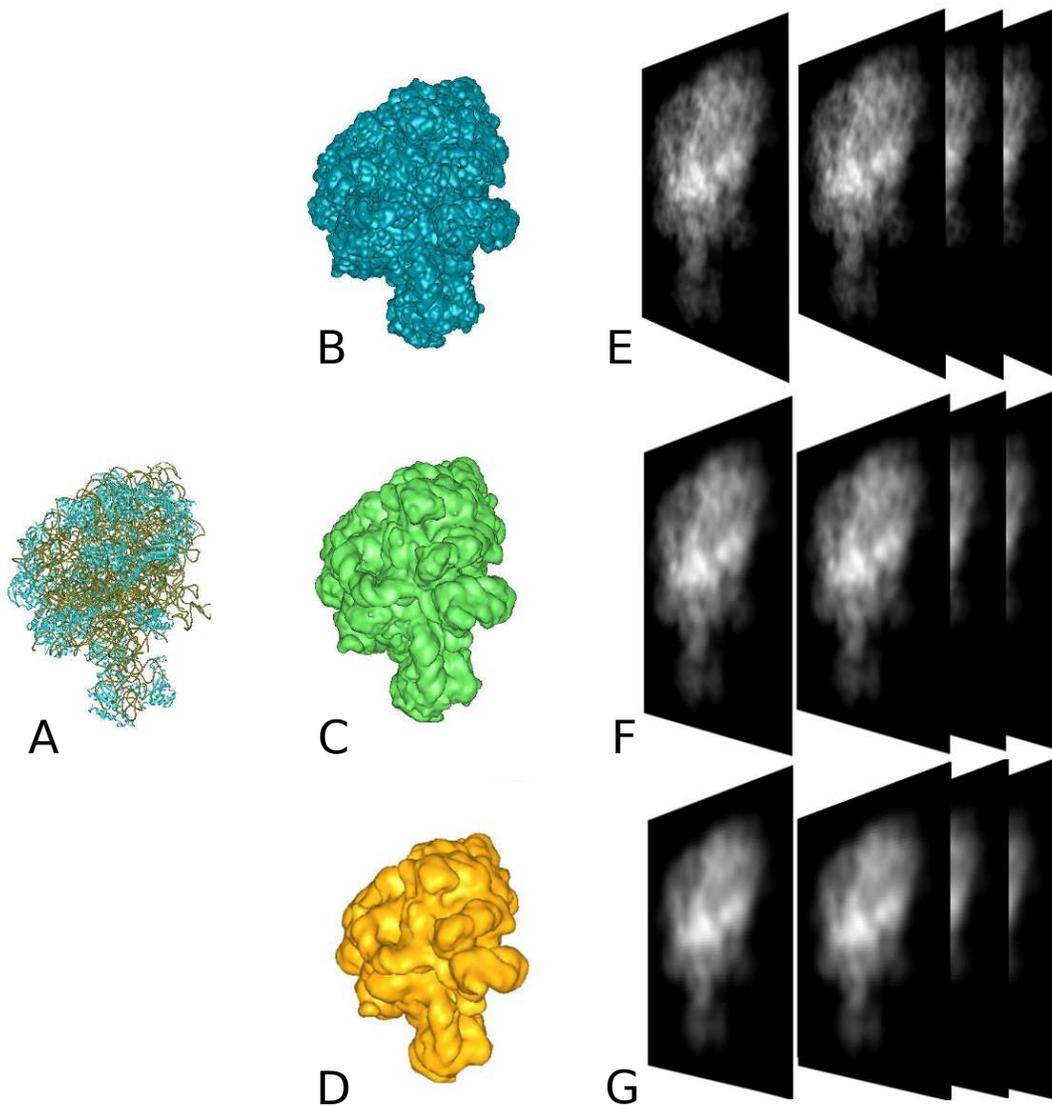


Supplementary Figure 1: Filtering results of representative projections of the Keyhole Limpet Hemocyanin protein obtained by cryo-EM (pixel size 2.22 Å; top row = side view; bottom row = top view): (A) Original data; (B) data after filtering; (C) difference between (A) and (B); (D) original data; (E) data after filtering; (F) difference between (D) and (E).

and the resolution of the simulated projection. As shown in Suppl. Fig. 2, we created three data sets derived from a single molecule using different resolutions and pixel sizes. To facilitate the calculation of class averages, we collected data from separate stacks of images. Each stack contained projections taken from the same angle subject to different realizations of simulated noise.

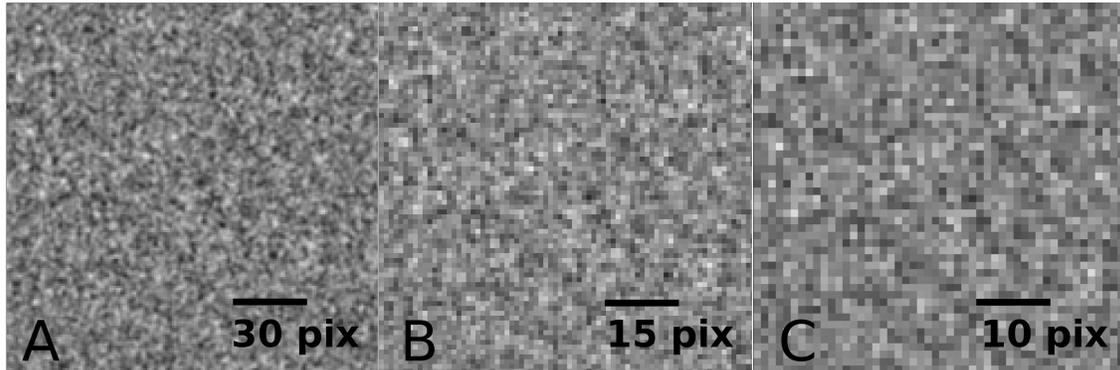
For the validation (Suppl. Fig. 2), we chose the ribosomal subunit of *Haloarcula marismortui* solved by Klein et al. (2001) (PDB ID 1jj2). We simulated a density map from atomic coordinates with different resolutions 6 Å, 10 Å, 15 Å and voxel spacings 1 Å, 2 Å, 3 Å, respectively.

To simulate noise in the stacks of images (Suppl. Fig. 2), we did not simply apply plain Gaussian white noise or impulse noise. Such simple noise models are easy to eliminate due to the lack of spatial correlation. In micrographs obtained experimentally from cryo-EM, the noise is intricate (Baxter et al., 2009) and does



Supplementary Figure 2: (A) Atomic structure of ribosomal unit (PDB ID 1jj2). (B-D) Simulated density maps with resolution 6 Å, 10 Å, 15 Å at voxel spacings 1 Å, 2 Å, 3 Å respectively. (E-G) Stacks of simulated 2D projections.

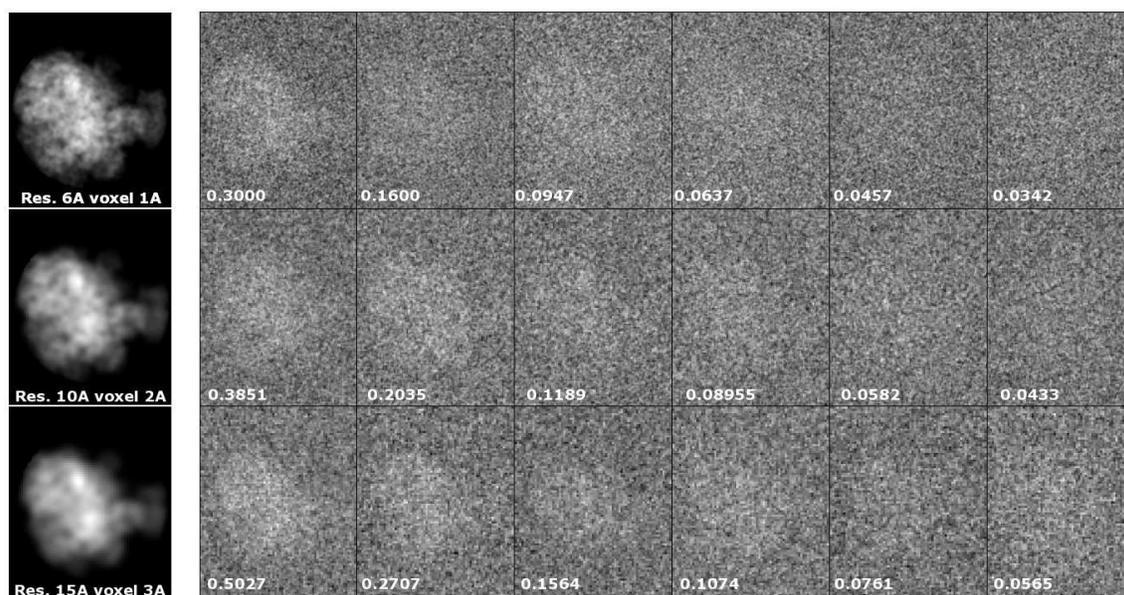
not follow a simple model. Therefore, we mimicked this complexity using a procedure that generates stacks of 2D images that are subject to Gaussian white noise with different distribution parameters for each stack. Next, we applied a Gaussian low-pass filter to each 2D image, which caused the spatial dependence of noise (i.e.,



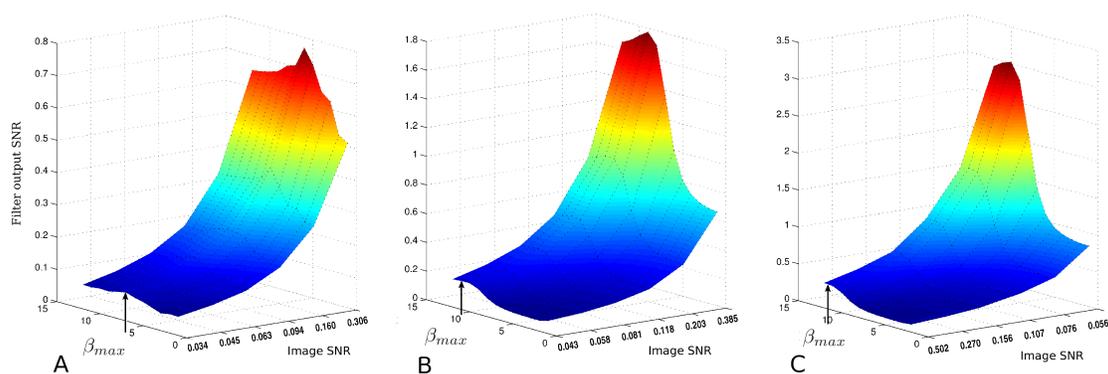
Supplementary Figure 3: Example of one realization of colored noise: (A) Pixel size 1 Å; (B) projection to 2 Å pixel size; (C) projection to 3 Å pixel size. The length of the displayed bar is 30 Å, i.e. 30 pixels for 1 Å, 15 pixels for 2 Å and 10 pixels for 3 Å pixel size.

color noise) as is typical in experimental cryo-EM data that is filtered to a specific resolution. To show that the behavior of the filter depends not only on the noise level but also on the pixel size, we projected the noise onto a larger 2 Å and 3 Å pixel spacing by local averaging of 2×2 and 3×3 pixels, respectively. Then, we added it to the stack of 2D particle projections of the corresponding pixel spacing with a proper SNR ratio in the range of 0.5 to 0.03.

Suppl. Fig. 3 provides an example of digitized noise with different pixel spacing. The described procedure yields three stacks of 2D images with different pixel spacings 1 Å, 2 Å, 3 Å, and at different noise levels. Each stack of images presents the same projection but with different realizations of noise (128 realizations in total). To show how the algorithm influences the SNR for the class averages, we calculated the averages of images from the (sub-) sets of $\{1, 2, 4, 8, 16, 32, 64, 128\}$ images. A simulated noise-free projection and a series of projections corrupted by simulated noise are shown for each of the three cases in Suppl. Fig. 4. The results show that a larger pixel size leads to an increase in observed SNR values.



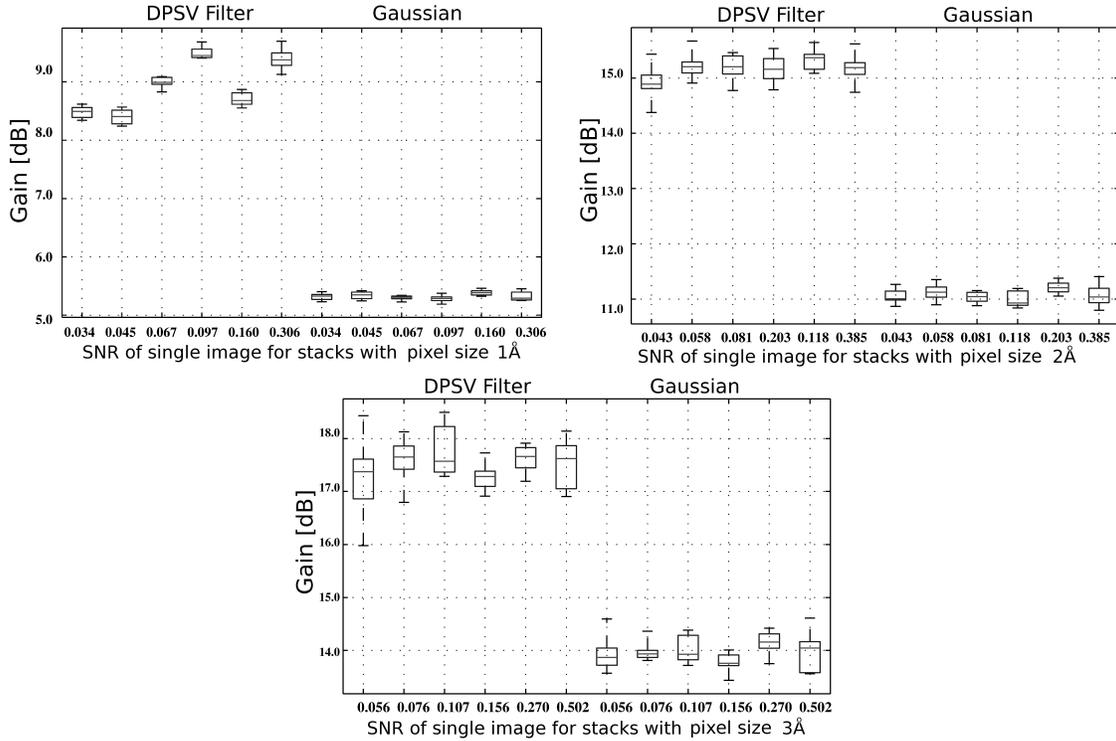
Supplementary Figure 4: Example of noise-free projections and different noise levels. The first column shows simulated projections, and the following columns show projections corrupted with simulated noise at increasing SNR values in white.



Supplementary Figure 5: The influence of parameter, β , on filter performance as judged by SNR values after filtering. (A) Density map with resolution 6 Å and voxel size 1 Å; (B) density map with resolution 10 Å and voxel size 2 Å; (C) density map with resolution 15 Å and voxel size 3 Å. Units for axis β represent the exponent i in the form: $\beta = 0.5^{(i)}$. Values of β_{max} indicate the value of the filter parameter that maximizes filter gain for each resolution.

We have optimized the filter parameters based on the size of the particle and pixel size. Local parameters such as mask size M and path length P are intuitive, but the selection of the filtering kernel parameter β requires more experience. To investigate the influence of β on values of SNR after filtering, we chose $\beta \in \{0.5^i, i = 1, \dots, 12\}$ for all three data sets. Suppl. Fig. 5 presents the results. It is clear that the optimal value of parameter β depends on the pixel spacings 1 \AA , 2 \AA , and 3 \AA : The SNR is maximized for $\beta_{max}(1 \text{ \AA}) = 0.5^7$, $\beta_{max}(2 \text{ \AA}) = 0.5^{10}$, and $\beta_{max}(3 \text{ \AA}) = 0.5^{11}$, respectively.

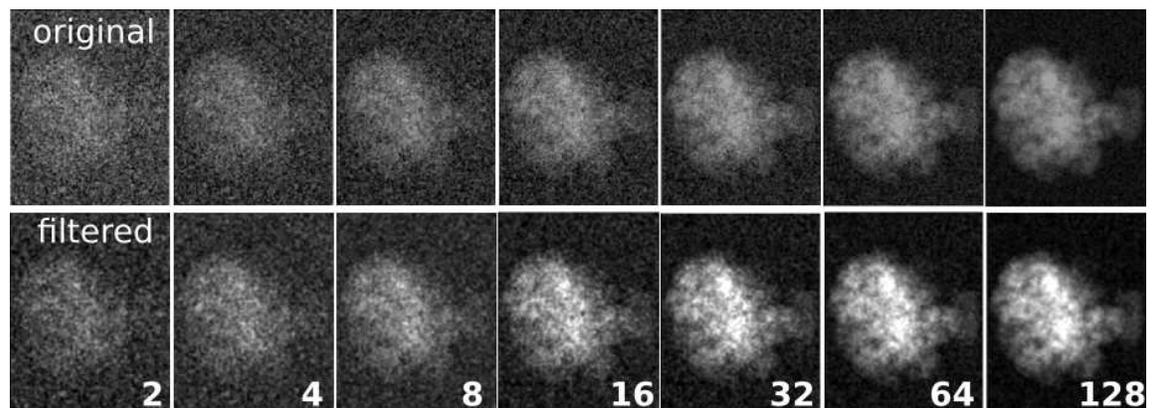
Next, we investigated the performance of the DPSV filter under the realistic condition of averaging the subsets of the image stack. We also compared the effect of filtration on the averaging process to the effect of a standard Gaussian filter using a sigma (standard deviation of the 1D Gaussian envelope) of two pixels. Suppl. Fig. 6



Supplementary Figure 6: Descriptive statistics of filter performance for each of the investigated image stacks. The box plots show the distribution parameters of DPSV and Gaussian filter gain (log of difference between averages of image stack subsets $\{1, 2, 4, 8, 16, 32, 64, 128\}$ after filtration and the averages of subsets before filtration), calculated as a function of the SNR values of individual images in the stack. Each box shows the mean, maximum, minimum, lower quartile, and upper quartile. Optimum Gaussian parameters were selected for the Gaussian control calculations.

presents the results. Each box plot in Suppl. Fig. 6 shows the statistical distribution of gain values observed in the averaged subsets (from $\{1, 2, 4, 8, 16, 32, 64, 128\}$ stack images) for a given stack image SNR level. The descriptive statistics show that the filtering technique is stable in the tested SNR range of 0.03 to 0.5. We also note that the DPSV filter does not behave linearly (the results depend on supervised classification), as was observed in the Gaussian control calculations. However, the obtained results vary within reasonable limits. In all of the tested cases, the filter gain is at least 3 dB higher compared to the Gaussian filter.

Suppl. Fig. 7 provides a comparison of class averages of filtered and unfiltered images (with mask size $M = 7$, path length $P = 3$, 8-neighborhood, and $\beta = \beta_{max}(1\text{\AA})$). As can be clearly seen, beginning from a stack of 32 images, filtering and averaging both greatly improve the image details and resulting SNR. The gain of ~ 9 dB achieved by filtering is largely independent of the stack size (c.f. Suppl. Fig. 6 upper left).



Supplementary Figure 7: Example of images before (top) and after (bottom) filtering, simulated for 6\AA resolution, 1\AA pixel spacing, averages from (sub-) sets $\{1, 2, 4, 8, 16, 32, 64, 128\}$, filtered with mask size $M = 7$, path length $P = 3$, and $\beta_{max}(1\text{\AA}) = 0.5^7$.

We recognize that any filter with low-pass properties will improve the overall SNR of the image, simply by virtue of its ability to up-weight the contribution of the high SNR information at a low frequency. Therefore, more extreme low-pass filtration will improve SNR, but the loss of details make the results unattractive. We have offered a simple test of the preservation of high resolution features in Section 3. We have also explored the trade-off between low-pass filtration and feature preservation in Supplementary Data 1 of (Rusu and Wriggers, 2012) for DPSV and the Gaussian filter: The results demonstrate better feature preservation properties in DPSV at a comparable level of low-pass filtration.

References

- Bammes, B., Jakana, J., Schmid, M., Chiu, W., 2010. Radiation damage effects at four specimen temperatures from 4 to 100 K. *J. Struct. Biol.* 169, 331–334.
- Baxter, W., Grassucci, R., Gao, H., Frank, J., 2009. Determination of signal-to-noise ratios and spectral SNRs in cryo-EM low-dose imaging of molecules. *J. Struct. Biol.* 166, 126–132.
- Buades, A., Coll, B., Morel, J., 2005. A review of image denoising algorithms, with a new one. *Multiscale Model. Simul.* 4, 490–530.
- Fernandez, J., Li, S., 2003. An improved algorithm for anisotropic nonlinear diffusion for denoising cryo-tomograms. *J. Struct. Biol.* 144, 152–161.
- Frank, J., 2006a. *Electron Tomography*. Springer Verlag, New York.
- Frank, J., 2006b. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies*. Oxford University Press, New York.
- Gonzales, R., Woods, R., 2002. *Digital Image Processing, 2nd Edition*. Prentice Hall.
- Havlin, S., Ben-Avraham, D., 1982. New approach to self-avoiding walks as a critical phenomenon. *Journal of Physics* 15, 321–328.
- Jiang, W., Baker, M. L., Wu, Q., Bajaj, C., Chiu, W., 2003. Applications of a bilateral denoising filter in biological electron microscopy. *J. Struct. Biol.* 144, 114–122.
- Kenney, C., Deng, Y., Manjunath, B., Hower, G., 2001. Peer group image enhancement. *IEEE Trans. on Image Processing* 10, 326–334.
- Klein, D. J., Schmeing, T. M., Moore, P. B., Steitz, T. A., 2001. The kink-turn: a new RNA secondary structure motif. *EMBO J.* 20, 4214–4221.
- Pantelic, R. S., Rothnagel, R., Huang, C.-Y., Muller, D., Woolford, D., Landsberg, M. J., McDowall, A., Pailthorpe, B., Young, P. R., Banks, J., Hankamer, B., Ericksson, G., 2006. The discriminative bilateral filter: An enhanced denoising filter for electron microscopy data. *J. Struct. Biol.* 155, 395–408.
- Russ, J. C., 2002. *The Image Processing Handbook, 4th Edition*. CRC Press.

- Rusu, M., Wriggers, W., 2012. Evolutionary bidirectional expansion for the annotation of alpha helices in cryo-electron microscopy reconstructions. *J. Struct. Biol.* 177, 410–419.
- Smolka, B., 2008. Peer group filter for impulsive noise removal in color images. *IEEE Trans. Med. Imaging* 27, 699–707.
- Szczepanski, M., 2008. Fast digital approach spatio-temporal filter. *Zeszyty Naukowe Politechniki Slaskiej, seria Automatyka* 150, 207–222.
- Wei, D. Y., Yin, C. C., 2010. An optimized locally adaptive non-local means denoising filter for cryo-electron microscopy data. *J. Struct. Biol.* 172, 211–218.
- Wright, E., Iancu, C., Tivol, W., Jensen, G., 2006. Observations on the behavior of vitreous ice at approximately 82 and approximately 12 K. *J. Struct. Biol.* 153, 241–252.
- Zhu, Y., Carragher, B., Potter, C., 2003. Automatic particle detection through efficient Hough transforms. *IEEE Trans. Med. Imaging* 22, 1053–1062.