It's noisy out there! A review of denoising techniques in cryo-electron tomography

Achilleas S. Frangakis

Buchmann Institute for Molecular Life Sciences and Institute for Biophysics, Goethe University Frankfurt Max-von-Laue-Str. 15, Frankfurt am Main, D-60438, Germany

ABSTRACT

Cryo-electron tomography is the only technique that can provide sub-nanometer resolved images of cell regions or even whole cells, without the need of labeling or staining methods. Technological advances over the past decade in electron microscope stability, cameras, stage precision and software have resulted in faster acquisition speeds and considerably improved resolution. In pursuit of even better image resolution, researchers seek to reduce noise – a crucial factor affecting the reliability of the tomogram interpretation and ultimately limiting the achieved resolution. Sub-tomogram averaging is the method of choice for reducing noise in repetitive objects. However, when averaging is not applicable, a trade-off between reducing noise and conserving genuine image details must be achieved. Thus, denoising is an important process that improves the interpretability of the tomogram not only directly but also by facilitating other downstream tasks, such as segmentation and 3D visualization. Here, I review contemporary denoising techniques for cryo-electron tomography by taking into account noise-specific properties of both reconstruction and detector noise. The outcomes of different techniques are compared, in order to help researchers select the most appropriate for each dataset and to achieve better and more reliable interpretation of the tomograms.

1. Introduction

Denoising is an image processing procedure that aims to reduce the noise in two-dimensional (2D) or three-dimensional (3D) images (Russ and Neal, 2018). In electron microscopy, many identical macromolecules (often numbering in the thousands) are captured in a single image. This repetitive signal is averaged again and again, thus achieving denoising, which is reflected by the gain in resolution (Frank, 2006). On the other hand, when the signal of non-repetitive objects needs to be denoised, very sophisticated techniques – other than averaging – are required. These techniques typically explore features in the image such as edges or patterns that are to be retained and try to remove everything else which is considered to be noise. Denoising techniques other than averaging, tend to be dismissed as cosmetic operations. This is because the gain in resolution and signal-to-noise ratio (SNR) is not yet sufficient to allow for a mechanistic understanding of the objects under scrutiny. Most biological objects such as organelles, macromolecules and polymers are non-repetitive, very dynamic and cannot be captured at discrete states, as required by conventional structural techniques. This pleiomorphism ultimately requires a solitary analysis (Sikora et al., 2020; Turoňová et al., 2020) and prohibits the use of averaging techniques, where the SNR improves in proportion to the number of entities used in the average (Papoulis, 1984). Thus, it is foreseeable that with the further development of electron tomography that opens up new avenues for the analysis of pleiomorphic objects, denoising methods will become even more essential.

Cryo-electron tomography (cryoET) is currently the only technique capable of recording 3D images of biological objects with a resolution in the nanometer range (Sali et al., 2003). The biological specimen is typically preserved at liquid nitrogen temperatures in a close-to-native state. Unfortunately, cryo-electron tomograms have a very low SNR at higher resolutions due to the high radiation sensitivity of biological specimens (Hattne et al., 2018; Sali et al., 2003). The use of higher doses of radiation increases the SNR but also destroys or falsifies high-frequency information due to alterations to the specimens. When seeking increasingly better resolutions, noise is a major limiting factor and improving the SNR is essential. There are many sources of noise, and noise can occur before and after image reconstruction. Before the reconstruction, the noise can be modeled by a Poisson or Gaussian distribution (Baumann and Reimer, 1981) depending on the radiation dose, whereas after the reconstruction, the noise modeling becomes more complex. Tomograms are prone to artifacts caused by samples changing during the recording process, which leads to differences between the individual images (Veesler et al., 2013), while 3D reconstructions are affected most significantly by the limited angular range of the tilt series (typically from −65 to +65 degrees), which results in an incomplete data set that is visible as an empty wedge in the Fourier domain, as predicted by the Fourier slice theorem (Mastronarde, 1997). In addition, the contrast transfer function (CTF) of the microscope is not trivial to be estimated in images of thick specimens recorded at higher tilt angles

(Kunz and Frangakis, 2017). Each of these issues alone is difficult to overcome, their combination poses a major challenge to any denoising technique.

For the reconstruction, the well-established filtered back projection (FBP) algorithm is typically used (Mastronarde, 1997). Despite the acceptable performance of the FBP (Heumann et al., 2011; Hrabe et al., 2012), in particular with respect to sub-tomogram averaging, the ramp filtering amplifies high-frequency noise and generally results in low-contrast images (Kunz and Frangakis, 2014; Smith et al., 1973). This limitation of FBP has probably motivated the parallel use of iterative reconstruction algorithms (Gilbert, 1972). Thus, although image reconstruction techniques also provide some level of denoising, this review will address denoising techniques independent of the reconstruction method.

Here, I discuss different denoising techniques used in cryo-electron tomography and compare their outcomes in particular regarding the visual quality and ability to preserve small details, which are particularly important for biological applications. Lastly, I provide examples of each method's performance to illustrate their advantages.

2. Evaluation measures

Evaluating the performance of denoising algorithms is challenging, because in most real-life cases the ground truth is not known (Frank, 2008). The evaluation criteria are broad and lastly depend on the result for each individual application (Narasimha et al., 2008). The typical measurement of quality for a denoising algorithm is reflected in the gain of the SNR. The SNR is defined as:

$$SNR = \frac{E[S^2]}{E[N^2]}$$

where E is the expectation value. If the noise N and signal S are zero-mean, the SNR is the ratio of the variances of the signal to noise.

In electron microscopy, resolution is typically estimated by the Fourier shell correlation function (FSC); a better (closer to the value one) FSC also reflects an increase in the SNR (Böttcher et al., 1997; Rosenthal and Henderson, 2003; van Heel, 1987; van Heel and Schatz, 2005). Thus, the SNR is related to the achieved resolution (Radermacher and Ruiz, 2019):

$$SNR = \frac{CC}{(1 - CC)}$$

where CC is the cross-correlation coefficient between two images. The FSC(s) between two images I_1 und I_2 is given by the following formula:

$$FSC(s) = \frac{\sum_{|\mathbf{f}| \in s} \left(I_1(f) \overline{I_2(\mathbf{f})} \right)}{\sqrt{\left(\sum_{|\mathbf{f}| \in s} |I_1(f)|^2 \right) \left(\sum_{|\mathbf{f}| \in s} |I_2(f)|^2 \right)}}$$

where s is the Fourier shell index, $\mathbf{f}$ is the spatial frequency, and – denotes complex conjugation.

Other metrics for estimating the quality of denoising include the peak SNR (Zhang et al., 2012), the spectral SNR (Penczek, 2002), and different interpretation of the resolution in the FSC such as the bit-FSC (van Heel and Schatz, 2005), but these performance measurements are not discussed in this review.

3. Denoising methods

The human eye is extremely powerful and can identify features within images even when noise is present. However, this ability is significantly impaired when viewing 3D images such as tomograms or when the noise level is very high. In these situations, denoising methods become particularly useful, as they aim to preserve important image details that are essential for further processing. There are essentially two points in the processing pipeline when denoising methods can be used: pre-reconstruction, applied to individual 2D images; or post-reconstruction, applied to tomograms (Fig. 1). There are advantages and disadvantages to both approaches, although the risks are significant for the former, as denoising can ambiguously affect the proper downstream interpretation of the data (Pruggnaller et al., 2008) (Heymann et al., 2008) (Roels et al., 2020).

Pre-reconstruction denoising reduces noise in individual 2D images by directly addressing the statistical attributes of noise from the microscope and the camera. Various denoising methods have been used for computed tomography, such as domain filtering techniques like Wavelet and Fourier Transformation-based filtering (FT), nonlinear and edge-keeping noise filters (e.g. nonlinear anisotropic diffusion (NAD) and non-local (NL) means) and, more recently, denoising techniques based on neural networks (Noise2Noise) (Fig. 2).

Post-reconstruction denoising reduces noise in 3D images by exploiting structural features in the objects – primarily edges. In tomogram denoising, the filter parameters are typically tuned to retain the proper appearance of the sample – in general a subjective judgement. A statistical model of the noise in tomograms is difficult to generate and depends on various effects.

Post-reconstruction denoising is typically preferred for several reasons: (i) it does not disturb the linearity between the projection images and the reconstruction; (ii) it avoids different treatments or even removal of features from the tilt-series due to a non-preferential viewing orientation; and (iii) it minimizes the risk of introduction of new artifacts that appeared during the denoising procedure.

Most importantly the SNR, but also specific artifacts and properties related to electron tomography may affect denoising performance and ultimately the choice of the best technique for a specific application (Fig. 3a). They include how the techniques perform regarding: (i) CTF corrected tomograms (Zanetti et al., 2009),(Bhamre et al., 2016), since the CTF flips the contrast of features depending on the spatial frequency and effects the contrast at the edges. (ii) the amount of missing wedge, i. e. how the elongation resulting from the missing wedge is dealt with (Mastronarde, 1997). Increasing the size of the wedge decreases the dimensionality from 3D towards 2D. Consequently, algorithms performing better in 2D may be appropriate for denoising tomograms with a very pronounced missing wedge.

3.1. Linear filters applied in either the spatial or frequency domain

Linear filters are routinely applied to image data individually or in combination, usually with the aim of removing either slow-changing background intensities (mild high-pass filter) or high frequencies (low-pass filter) that are thought to contain significantly more noise than signal. Linear filters typically denoted with $h(j)$ in real space and their Fourier transform $H(f) = \{ \mathscr{F}(h(j)) \}$ are mostly applied in the Fourier domain, due to speed considerations, although they can also be applied in the spatial domain when the kernels are very small.
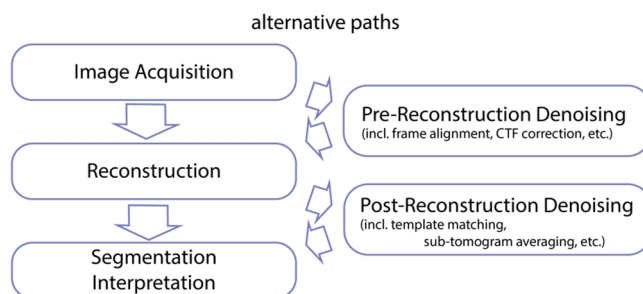

Fig. 1. Denoising can be used at two points during analysis in electron tomography.
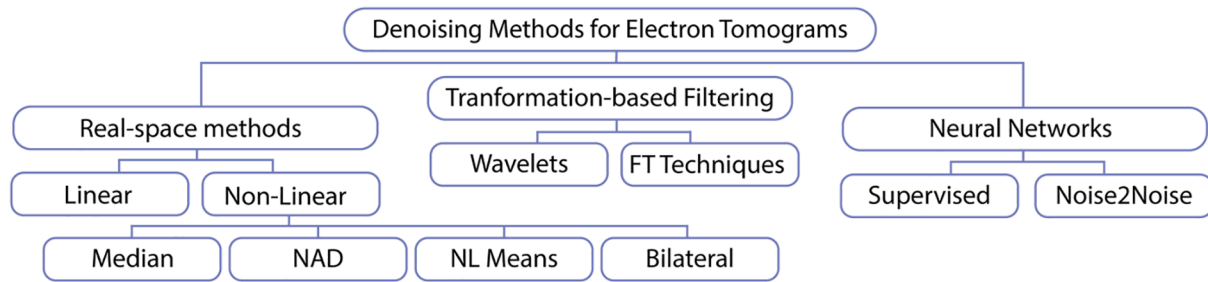
Fig. 2. Overview and classification of denoising methods in electron tomography.

$$I'(j) = (I*h)(j) = \int_{-\infty}^{\infty} I(k) \bullet h(j-k)\mathrm{d}k \rightleftharpoons \{\mathscr{T}(I(j))\} \bullet H(f)$$

Linear filters can be empirically derived and easily adapted. Characteristic examples include: (i) the low-pass filter which either downweights or sets to zero all Fourier components above the maximum achievable frequencies or resolution, (ii) the Gaussian-filter, which is also a lowpass filter, that multiplies all Fourier components with a Gaussian function $H(f) = e^{-f^2/2\sigma^2}$ with a set width σ, (iii) the "dose-filtering" to eliminate or down-weight higher frequencies that are thought not to contain any more useful information because of the dose that has already been applied to the sample (Grant and Grigorieff, 2015; Scheres, 2014; Sigworth, 2016), and others. Linear filters improve the visual quality of images and the overall SNR may also improve, because the energies of the noise and the signal are not equally distributed over all frequencies – that is, more noise is present at higher frequencies, which, for instance, in the case of low-pass filtering are removed (Fig. 3b). Instead, the FSC remains the same, because the signal and the noise are minimized in each frequency band by the same amount. In addition, there is a tradeoff between the reduction in noise at the higher frequencies and the spatial resolution of the image.

3.2. Nonlinear filters

3.2.1. Median filter

The most popular nonlinear filter is perhaps the median filter (Pitas and Venetsanopoulos, 1990). The median filter replaces the value of the central kernel pixel $I(j)$ with the median value of the kernel W

$$I'(j) \leftarrow \underset{v\in W_j}{\mathrm{median}}\{I(v+j)|v\in R\}$$

The median filter is known for preserving the edges. It works particularly well for salt and pepper noise. In contrast to linear filters, the median filter improves both the SNR and the FSC. In general, it is a very versatile filter that is routinely used for many applications (Omer et al., 2018).

3.2.2. Nonlinear anisotropic diffusion

Nonlinear anisotropic diffusion (NAD) is based on the original proposal of Perona and Malik (Perona and Malik, 1990) and further mathematical improvements and the introduction of coherence enhancing diffusion (Weickert, 1999). It has been a popular technique for denoising electron tomograms, because the diffusion filters are designed to preserve the edges while smoothing along them. NAD filters considerably improve the SNR and the FSC, because they effectively remove noise in smooth regions and ensure edge preservation. Image smoothening is designed as a diffusion process that is equivalent to Gaussian filtering, but it can be controlled along edges as a function of the local gradient. Thus, the image I_t, with respect to the processing time t, is a function of the gradient of the image ∇I and a diffusivity matrix G:

$$I_t = \mathrm{div}(\mathbf{G} \cdot \nabla I)$$

The properties of G determine the development of the image denoising. It is composed of the three eigenvectors v_i of the structure tensor $\mathbf{J}_\sigma = \nabla I \cdot \nabla I^\mathrm{T} * \mathbf{K}_\sigma$ which allow for an averaged direction of the gradient ($\mathbf{K}_\sigma$ is a Gaussian kernel of width σ and * denotes the convolution) and the corresponding eigenvalues μ_i.

The diffusivity G can be structured in the following way:

$$\mathbf{G} = \begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix} \cdot \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} \cdot \begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix}^\mathrm{T}$$

Eigenvectors and eigenvalues characterize the local structural features of the image within a neighborhood defined by $\mathbf{K}_\sigma$.

For edge-enhancing diffusion, $\lambda_1 = \lambda_2 = 1/(1+\mu_1^2/\kappa^2)$ and $\lambda_3 = 1$ are chosen; for coherence enhancing diffusion, $\lambda_1 = \lambda_2 = \alpha$ and $\lambda_3 = \alpha + (1-\alpha) \cdot \exp\left(-C/(\mu_1 - \mu_3)^2\right)$ are chosen (κ, α and C are user-defined, although α and C are typically close to zero). Hybrid approaches have been shown to perform well for electron microscopy (Frangakis and Hegerl, 2001). NAD is mathematically well-posed and can be used as a pre- or post-reconstruction filter (Fig. 3c) (Frangakis and Hegerl, 2001; Maiorca et al., 2012).

3.2.3. Non-local means filter

Averaging reduces the noise by summing many identical structures. A conceptually similar filter is the non-local means (NLM) filter, which removes noise by averaging similar patterns within repeating information in the image (Buades et al., 2005). As a result, the new pixel intensity is determined as the weighted mean of similar patterns inside the image. For a given noisy image I, a denoised image $I'(j)$ at pixel j is computed as a weighted mean of all pixels in an image or within a neighborhood IN,

$$I'(j) = \sum_{k \in IN} we(j,k)I(k)$$

where the weights $we(j,k) = \frac{1}{NC(j)} e^{-\left(\|I(V_j) - I(V_k)\|_{2,b}^2\right)/h^2}$ indicate the similarity between two pixels or a small region around those pixels. b is the Gaussian kernel standard deviation, V_j and V_k are the intensity grey level vectors, and NC is a normalizing constant, such that the overall sum of the weights is equal to one (Zhang et al., 2013). NLM is also used in electron microscopy (Wei and Yin, 2010). The Gaussian kernel with standard deviation h acts as a filtering parameter, controlling the degree of filtering. When h is very small, most weights of the pixels in the neighborhood IN can be ignored and only a few pixels whose neighborhood is very similar to that of the target pixels contribute to the average. Therefore, a very small h leads to weak smoothening of the image. For large h values, the opposite is true. For electron microscopy, the choice of an appropriate h was proposed as an automatic and locally adaptive method to estimate the optimal h value (Wei and Yin, 2010). While admittedly in 2D this method is among the best denoising techniques, its proper implementation in a 3D setting with 6 degrees of freedom is still lacking, as it can become very computationally expensive (Fig. 3d).
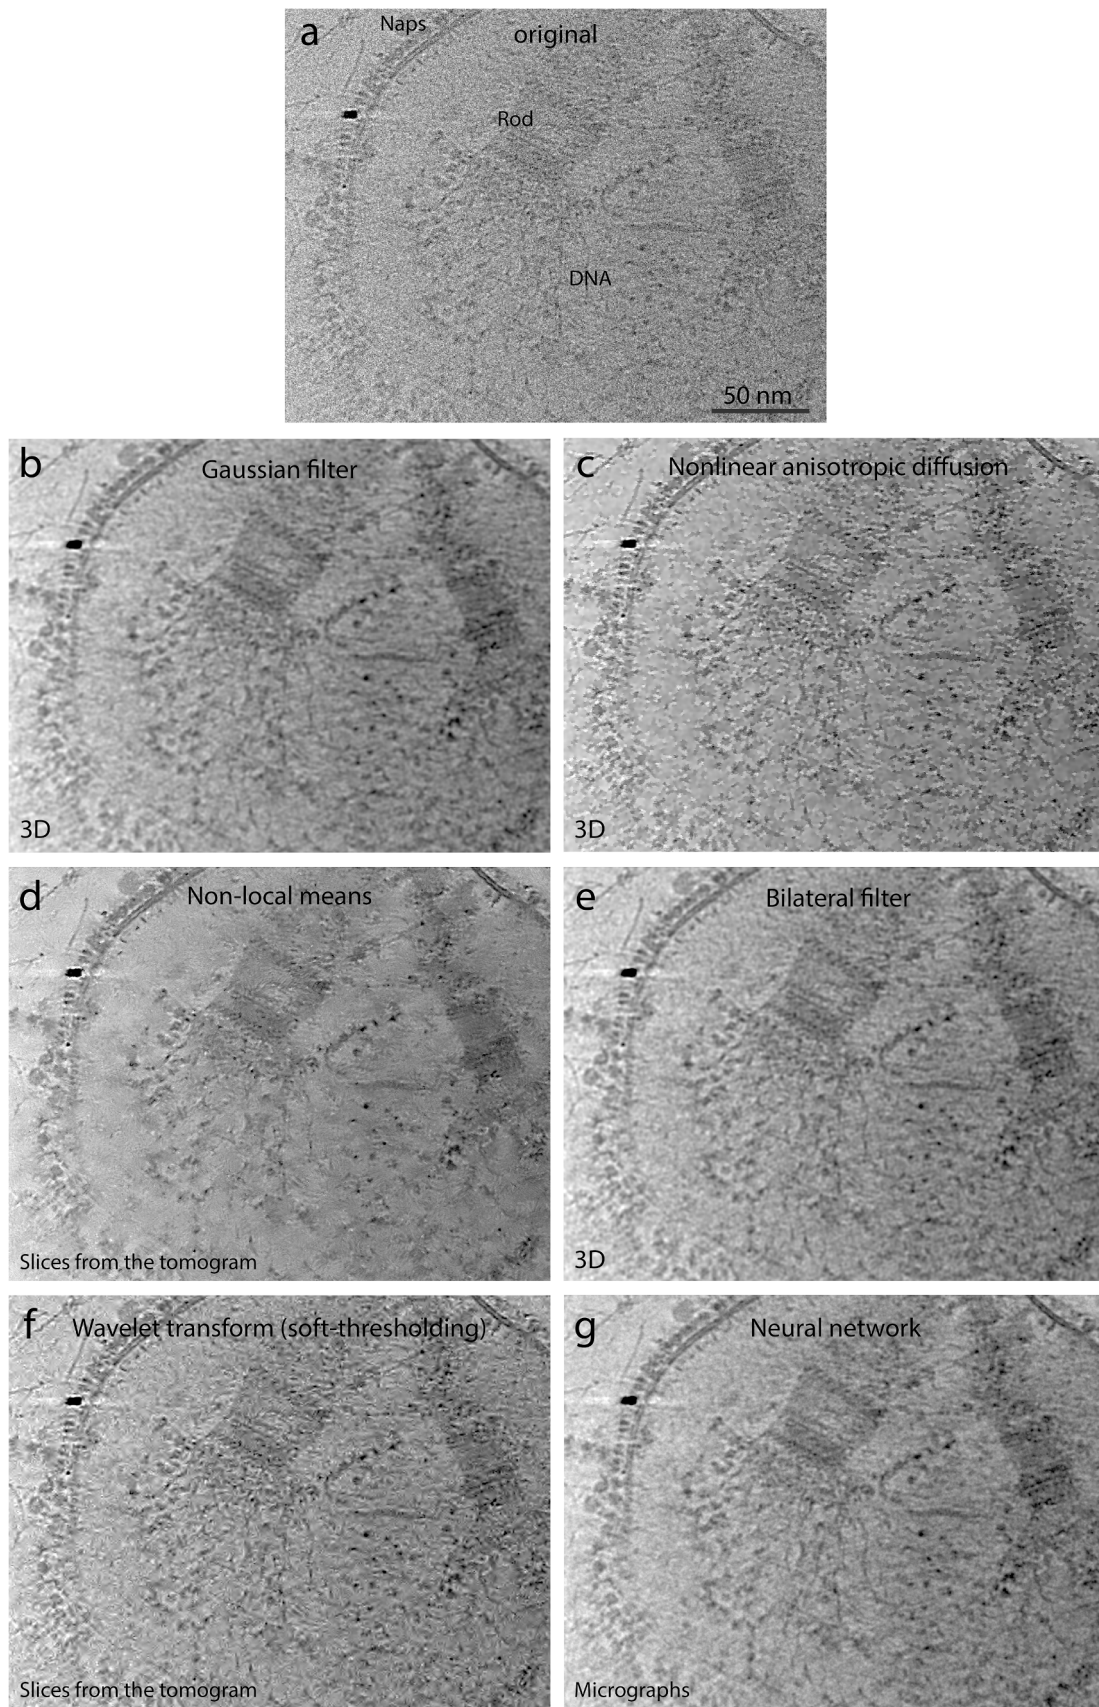
(caption on next page)

Fig. 3. Visual comparison of denoising results of different techniques. All images were set to be mean-value free and were scaled to ±3 the variance. Parameters were selected as such, in order to show the individual effects of each denoising technique on the image (i.e. not as a quantitative performance comparison). Denoted at the bottom left corner of each image is the level at which the tomograms were denoised, whether in 3D, as individual tomographic slices or whether the individual micrographs were denoised prior to reconstruction. This depended on the limitations of individual techniques, but also implementation constraints. (a) 1.6 nm thick tomographic slice through a *Mycoplasma pneumoniae* cell (Scale bar is 50 nm for all subsequent images). The cell membrane, adhesion machinery (NAPs), two rods (R) in a top and side view, and DNA with various complexes attached can be seen. This slice was selected in order to compare the denoising properties on high spatial frequencies at the area of the rod, the continuity of the membrane and preservation of the lipid bilayer, the contrast enhancement on individual features attached to the DNA. (b) Gaussian filter with $\sigma = 3$ pixels. (c) Nonlinear anisotropic diffusion performed for 50 iterations with $\kappa = 0.001$. (d) Non-local means applied on individual slices with h = 10. (e) Bilateral filter with $\sigma_r = 3$ pixels and $\sigma_d = 3$ pixels. (f) Wavelet transformation applied on individual slices with soft-thresholding and a threshold value at 10% of the highest intensity coefficients. (g) Neural network denoising with the pretrained network of Topaz (Bepler et al., 2020), applied on the individual movie frames of the direct detector and subsequent reconstruction.

4. Bilateral and trilateral filters (or intensity and range filters, respectively)

Bilateral and trilateral filters are local nonlinear filters (Tomasi and Manduchi, 1998). Their shape depends on the local properties of the image, such as intensity, distance and similarity. The shape of the filter resembles nonlinear anisotropic diffusion in the sense that filtering over edges is prevented, while smoothing along the edge is possible, typically with a Gaussian function. The bilateral filter is defined for each pixel individually and does not need to be applied iteratively like NAD. Several implementations in 3D also exist for electron microscopy (Jiang et al., 2003), (Pantelic et al., 2007).

The filtered image *I'* subjected to the bilateral filter is defined as:

$$I'(x) = \frac{1}{W_p}\sum_{x_i \in N} I(x_i) f_r(\|I(x_i) - I(x)\|) g_s(\|x_i - x\|)$$

with a normalization term W_p so that the energy of the filter is 1; *x* are the coordinates; *N* is the window centered in *x*, so $x_i \in N$ is another pixel.

In the case of the bilateral (trilateral) filter, there are two (three) kernels, f_r is the range kernel for smoothing intensities and g_s is the spatial kernel that penalizes pixels that are further away.

Assuming that the range and spatial kernels are Gaussian functions, the weight assigned for pixel (*j*) to denoise pixel (*k*) is given by $w(j,k) = \exp\left(-\frac{(i-k)^2}{2\sigma_d^2} - \frac{\|I(j) - I(k)\|^2}{2\sigma_r^2}\right)$ where σ_d and σ_r are smoothing parameters (Fig. 3e).

5. Transform domain techniques

5.1. Fourier transform (also pin-filtering)

Elaborate linear filters can be implemented in the Fourier domain. For electron microscopy of (mini)-crystals either in 2D or in 3D that have a discrete Fourier transform, pin-filtering can be applied as a very effective method for reducing the noise while retaining the signal that is present only at discrete positions. Pin-filtering can also be considered a basic form of geometrical constraint filtering, assuming that all unit-cells are packed in a crystal. Pre-orientation and constraining of unit-cells, which is also a form of a packing, is used in many sub-tomogram averaging applications, at least for the initial model. Pin-filtering generates a hole mask in the Fourier domain after the crystal has been indexed. Areas inside the mask at the positions of the indexes are set to 1 and outside the mask are set to 0. The inverse transformation results in a noise-reduced version of the image. Of course, local variations and sample deformations are not considered, so averaging of the individual unit cells, by releasing the constrains, typically leads to better results.

5.2. Wavelet transform

There are several approaches to denoise using the wavelet transform (Mallat, 1989) (Münch et al., 2009; Starck and Bijaoui, 1994; Zhao, 1999). For electron microscopy, the steerable pyramid, translation

invariant wavelets, or translation-invariant wavelet packets have been suggested (Donoho, 1995). In the wavelet transform, the signal is divided into different sub-bands.

Many image features consist of high-frequency components with small spatial extension and low-frequency components with large spatial extension. The main idea of the wavelet transform is to provide a variable description, which has both a good spatial resolution for high-frequency components and a good frequency resolution for low-frequency components. The (highly redundant) continuous wavelet transform of the 1D signal $I(j)$ is:

$$WT(a,\tau) = \int_{-\infty}^{\infty} I(j) \bullet \psi_{a,\tau}(j - \tau)dj$$

as the inner product of the signal and a variable basis function

$$\psi_{a,\tau}(j) = \frac{1}{\sqrt{a}} \bullet \psi\left(\frac{j - \tau}{a}\right)$$

where $a > 0$ is a scaling parameter and τ is a shift parameter for the wavelet prototype $\psi(j)$. For the wavelet prototype function (or mother wavelet), various choices of functions exist, for denoising purposes the Symmlet 8 has been often used. The wavelet transform efficiently compresses characteristic signal features into a small number of significantly large coefficients, whereas noise spreads out as small coefficients in each band. The different magnitudes of the transform coefficients with respect to noise allow for efficient thresholding of the transform coefficients and the elimination of noise. Empirically, the soft thresholding operator of Donoho showed a good performance for noisy EM data (Donoho, 1995; Frangakis et al., 2001; Moss et al., 2005; Stoschek and Hegerl, 1997). Denoising with this transform sometimes leads to visual artifacts, which can be attributed to the lack of translation- and rotation-invariance of the traditional wavelet transform (Huang et al., 2018) (Fig. 3f).

5.3. Denoising by neural networks

Denoising by neural networks has been possible for decades, but has recently become more popular as a result of increased processing speeds of the training (Liu et al., 2021). This is due to gains in computing power, which also allow for deeper (multiple layers) and more complex networks (nonpolynomial activation functions), from which the expression "deep-learning" rises. From an engineering point of view, designing a neural network (i.e. the number of hidden layers and number of neurons per layer) has always been a bit of an art, as the network architecture influences the learning speed, the information stored, and the information transfer. In a training session, the weights of the individual neurons need to be adjusted. To achieve this, input–output pairs (for denoising purposes, typically a noise-corrupted input x_i and noise-free (clean) output y_i, but others can be used) are presented to the network such that the weights of the individual neurons are progressively adapted by the backpropagation algorithm. The goal of the network is to minimize

$$\underset{L}{\mathrm{argmin}} \sum_i L(\{L(f_L(x_i), y_i)\})$$

where f_L is a parametric family of mappings.

While this feature of the supervised learning approach was previously used exclusively, it was recently shown that image restoration is possible without clean data (Lehtinen et al., 2018). Thus, in the training session, the correspondence of noisy input and noise-free output is no longer required, and an ensemble of outputs for which the expectation value is equivalent to that of the noise-free output can be used.

This is particularly interesting because most often – if not always – in electron microscopy, the noise-free output is not known. However, many noise realizations of identical objects exist. This starts at the level of the detector output: the individual movie frames (short exposures) contain many noisy realizations of the recorded signal of a less noisy object, i.e. the final accumulated exposure. Every short exposure is one noisy realization of the original signal. The task of the network is to "learn" the noise from the short exposures and produce a denoised image that is less noisy than the simple average (the expectation value) of all the short exposures (Bepler et al., 2020) (Ramírez-Aportela et al., 2019). The same approach can be applied on many applications in electron microscopy, such as the denoising of sub-tomograms, half-sets of whole tomograms and others (Buchholz et al., 2019). Thus, altogether, neural networks may produce denoised images that are better than the simple averaging of individual events (Fig. 3g).

6. Algorithm performance assessment

Here, I present the results of the above denoising methods, to help readers assess the outcomes and choose the best method for a given dataset (Table 1). When selecting a method, particular consideration should be given to the presence of artifacts and, equally important, to the preservation of texture, edges, and fine image structures.

Table 1
Electron microscopy denoising techniques with their main area of application as well as potential advantages and disadvantages of each technique.

Electron microscopy denoising technique	Advantages	Disadvantages
Linear filter	Linear filters (low-pass and band-pass filters) are easy to implement and remain the first choice when it comes to denoising the image. Can be used as a pre- and post-reconstruction filter.	Many limitations as they do not discriminate specific signal properties; most importantly, they blur the edges, as they dampen the high frequencies. The FSC remains unaltered in contrast to all other techniques, where the FSC improves.
Bilateral and trilateral filter	Very effective and does not need to be run iteratively. Only a few parameters need to be properly setup.	Very much dependent on the window size. Potentially, the coherence of the structure may be lost. Should not be used as a pre-reconstruction filter.
Wavelet transform	Wavelets define local features spectrally or spatially, which enable the method to filter noise by preserving the fine details and edges of the image. Very effective for compressing information.	Thresholding is difficult. Translational and rotational invariance is not guaranteed, thus artifacts may appear. After thresholding, the output cannot be used for reconstruction as artifacts are prevalent.
Non-local means	Maybe the best denoising method in 2D. It removes noise by exploiting redundant information in the image. Performs in a similar fashion to averaging, thus intuitive to use.	In 3D in particular, computationally very expensive because of its complexity in evaluating the pixel weights.

Table 1 (*continued*)

Electron microscopy denoising technique	Advantages	Disadvantages
Anisotropic diffusion filter	Performs smoothing by preserving important edge details of the image. Achieves probably the most improvement in FSC compared to all other techniques.	Generates areas with constant intensities in the image and may remove small details.
Convolutional neural networks	Ideally suited for direct detector recording properties. In certain cases, produces sterling results.	General properties not yet well understood. Unclear whether it is acceptable as a pre-reconstruction method, or if it introduces artifacts and/or bias.

7. Discussion

As a result of various contributions from many researchers in the field of electron microscopy, often inspired by the respective computer science developments, numerous well-performing denoising techniques are applicable to cryo-electron tomograms. These techniques are currently used systematically to improve the downstream processing of electron tomograms and to highlight information hidden in the noise. Any isosurface or volume rendering requires the application of some denoising, because otherwise a proper visualization is hindered by the noise, which in combination with the missing wedge, make a meaningful interpretation impossible. Segmentation techniques similarly profit, because the tracing of boundaries, identifying coherent regions and textures, and ultimately annotating entire objects, all profit significantly from a denoised image.

Summarizing, it is difficult to make concluding statement on which is the best denoising technique, because this rather depends on the application. Here, several denoising filters were presented and Table 1 provides a general guideline for which applications they are best suited. In a real case scenario, one would start with the most versatile filters and progress to using more specialized filters as the requirements dictate. This will involve choosing the appropriate parameters, which affects the outcome but also the quantitative evaluation of the results. Tuning of several parameters, or extensive training sessions are often necessary in order to achieve an optimal result. Measuring the performance of a denoising technique is easy, when the ground truth is known. However, when real data is evaluated this can be extremely complicated. In electron microscopy, the dataset can typically be split into two halves to allow an FSC to be calculated. This indicates the gain in SNR, which most often is also reflected in the visual appearance.

When pleiomorphic objects are analyzed, denoising becomes increasingly important, as there is no other way to extract the information from the noise since averaging is not an option. In addition, automated tomogram analysis, such as template matching (Frangakis et al., 2002), or feature detection with neural networks (Chen et al., 2017) also benefits from improved denoising, for instance after denoising with neural networks (Buchholz et al., 2019) (Tegunov et al., 2021). Unfortunately, the resolution of the tomograms does not yet allow for systematically obtaining structural information at 6–7 Å, which is necessary in order to draw mechanistic insight. While denoising significantly improves the SNR, it does not manage to tip the scale towards quantitative and mechanistic insight, yet. This may change with advances in electron microscopy hardware technology. This is where the promise lies, that with improving tomographic resolution, in combination with advanced denoising techniques, all possible structures depicted in the tomograms can be precisely analyzed.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

References

Baumann, W., Reimer, L., 1981. Comparison of the noise of different electron detection systems using a scintillator-photomultiplier combination. Scanning 4 (3), 141–151.

Bepler, T., Kelley, K., Noble, A.J., Berger, B., 2020. Topaz-Denoise: general deep denoising models for cryoEM and cryoET. Nat Commun 11, 5208.

Bhamre, T., Zhang, T., Singer, A., 2016. Denoising and covariance estimation of single particle cryo-EM images. J Struct Biol 195 (1), 72–81.

Böttcher, B., Wynne, S.A., Crowther, R.A., 1997. Determination of the fold of the core protein of hepatitis B virus by electron cryomicroscopy. Nature 386 (6620), 88–91.

Buades, A., Coll, B., Morel, J.-M., 2005. A Non-Local Algorithm for Image Denoising. IEEE Computer Society, DC, USA, Washington.

Buchholz, T.O., Jordan, M., Pigino, G., Jug, F. 2019. Cryo-CARE: Content-Aware Image Restoration for Cryo-Transmission Electron Microscopy Data, pp. 502-506 Ieee 16th International Symposium on Biomedical Imaging.

Chen, M., Dai, W., Sun, S.Y., Jonasch, D., He, C.Y., Schmid, M.F., Chiu, W., Ludtke, S.J., 2017. Convolutional neural networks for automated annotation of cellular cryo-electron tomograms. Nat Methods 14 (10), 983–985.

Donoho, D.L., 1995. De-noising by soft-thresholding. IEEE Trans. Inf. Theory 41 (3), 613–627.

Frangakis, A.S., Hegerl, R., 2001. Noise reduction in electron tomographic reconstructions using nonlinear anisotropic diffusion. J Struct Biol 135 (3), 239–250.

Frangakis, A.S., Stoschek, A., Hegerl, R., 2001. Wavelet transform filtering and nonlinear anisotropic diffusion assessed for signal reconstruction performance on multidimensional biomedical data. IEEE Trans. Biomed. Eng. 48 (2), 213–222.

Frangakis, A.S., Bohm, J., Forster, F., Nickell, S., Nicastro, D., Typke, D., Hegerl, R., Baumeister, W., 2002. Identification of macromolecular complexes in cryoelectron tomograms of phantom cells. Proc Natl Acad Sci U S A 99 (22), 14153–14158.

Frank, J., 2006. Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State: Visualization of Biological Molecules in Their Native. State Oxford University Press, USA.

J. Frank Electron Tomography: Methods for Three-Dimensional Visualization of Structures in the Cell Springer 2008 New York.

Gilbert, P., 1972. Iterative methods for the three-dimensional reconstruction of an object from projections. J Theor Biol 36 (1), 105–117.

Grant, T., Grigorieff, N., 2015. Measuring the optimal exposure for single particle cryo-EM using a 2.6 Å reconstruction of rotavirus VP6. Elife 4, e06980.

Hattne, J., Shi, D., Glynn, C., Zee, C.-T., Gallagher-Jones, M., Martynowycz, M.W., Rodriguez, J.A., Gonen, T., 2018. Analysis of Global and Site-Specific Radiation Damage in Cryo-EM. Structure 26 (5), 759–766.e4.

Heumann, J.M., Hoenger, A., Mastronarde, D.N., 2011. Clustering and variance maps for cryo-electron tomography using wedge-masked differences. J Struct Biol 175 (3), 288–299.

Heymann, J.B., Cardone, G., Winkler, D.C., Steven, A.C., 2008. Computational resources for cryo-electron tomography in Bsoft. J Struct Biol 161 (3), 232–242.

Hrabe, T., Chen, Y., Pfeffer, S., Kuhn Cuellar, L., Mangold, A.-V., Förster, F., 2012. PyTom: a python-based toolbox for localization of macromolecules in cryo-electron tomograms and subtomogram analysis. J Struct Biol 178 (2), 177–188.

Huang, X., Li, S., Gao, S., 2018. Applying a Modified Wavelet Shrinkage Filter to Improve Cryo-Electron Microscopy Imaging. J Comput Biol 25 (9), 1050–1058.

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 (1-2), 114–122.

Kunz, M., Frangakis, A.S., 2014. Super-sampling SART with ordered subsets. J Struct Biol 188 (2), 107–115.

Kunz, M., Frangakis, A.S., 2017. Three-dimensional CTF correction improves the resolution of electron tomograms. J Struct Biol 197 (2), 114–122.

Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T., 2018. Noise2Noise: Learning Image Restoration without Clean Data. ArXiv, 1803.04189v3.

Liu, Z., Jin, L., Chen, J., Fang, Q., Ablameyko, S., Yin, Z., Xu, Y., 2021. A survey on applications of deep learning in microscopy image analysis. Comput Biol Med 134, 104523. https://doi.org/10.1016/j.compbiomed.2021.104523.

Maiorca, M., Hanssen, E., Kazmierczak, E., Maco, B., Kudryashev, M., Hall, R., Quiney, H., Tilley, L., 2012. Improving the quality of electron tomography image volumes using pre-reconstruction filtering. J Struct Biol 180 (1), 132–142.

Mallat, S.G., 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Representation. IEEE Trans. Pattern Anal. Mach. Intell. 11 (7), 674–693.

Mastronarde, D.N., 1997. Dual-axis tomography: an approach with alignment methods that preserve resolution. J Struct Biol 120 (3), 343–352.

MOSS, W.C., HAASE, S., LYLE, J.M., AGARD, D.A., SEDAT, J.W., 2005. A novel 3D wavelet-based filter for visualizing features in noisy biological data. J Microsc 219 (2), 43–49.

Münch, B., Trtik, P., Marone, F., Stampanoni, M., 2009. Stripe and ring artifact removal with combined wavelet –- Fourier filtering. Opt. Express 17 (10), 8567. https://doi.org/10.1364/OE.17.008567.

Narasimha, R., Aganj, I., Bennett, A.E., Borgnia, M.J., Zabransky, D., Sapiro, G., McLaughlin, S.W., Milne, J.L.S., Subramaniam, S., 2008. Evaluation of denoising algorithms for biological electron tomography. J Struct Biol 164 (1), 7–17.

Omer, A.A., Hassan, O.I., Ahmed, A.I., Abdelrahman, A., 2018. Denoising CT Images using Median based Filters: a Review. 2018 International Conference on Computer, Control, Electrical, and Electronics Engineering (ICCCEEE), 1-6.

Pantelic, R.S., Erickson, G., Hamilton, N., Hankamer, B., 2007. Bilateral edge filter: photometrically weighted, discontinuity based edge detection. J Struct Biol 160 (1), 93–102.

Papoulis, A., 1984. Probability, Random Variables and Stochastic Processes McGraw Hill.

Penczek, P.A., 2002. Three-dimensional spectral signal-to-noise ratio for a class of reconstruction algorithms. J Struct Biol 138 (1-2), 34–46.

Perona, P., Malik, J., 1990. Scale-Space and Edge Detection Using Anisotropic Diffusion. IEEE Trans. Pattern Anal. Mach. Intell. 12 (7), 629–639.

Pitas, I., Venetsanopoulos, A.N., 1990. In: Nonlinear Digital Filters. Springer US, Boston, MA, pp. 63–116. https://doi.org/10.1007/978-1-4757-6017-0_4.

Pruggnaller, S., Mayr, M., Frangakis, A.S., 2008. A visualization and segmentation toolbox for electron microscopy. J Struct Biol 164 (1), 161–165.

Rademacher, M., Ruiz, T., 2019. On cross-correlations, averages and noise in electron microscopy. Acta Crystallogr F Struct Biol Commun 75 (1), 12–18.

Ramírez-Aportela, E., Mota, J., Conesa, P., Carazo, J.M., Sorzano, C.O.S., 2019. DeepRes: a new deep-learning- and aspect-based local resolution method for electron-microscopy maps. IUCrJ 6 (6), 1054–1063.

Roels, J., Vernaillen, F., Kremer, A., Gonçalves, A., Aelterman, J., Luong, H.Q., Goossens, B., Philips, W., Lippens, S., Saeys, Y., 2020. An interactive ImageJ plugin for semi-automated image denoising in electron microscopy. Nat Commun 11, 771.

Rosenthal, P.B., Henderson, R., 2003. Optimal determination of particle orientation, absolute hand, and contrast loss in single-particle electron cryomicroscopy. J Mol Biol 333 (4), 721–745.

Russ, J.C., Neal, F.B., 2018. The Image Processing Handbook. CRC Press.

Sali, A., Glaeser, R., Earnest, T., Baumeister, W., 2003. From words to literature in structural proteomics. Nature 422 (6928), 216–225.

S.HW. Scheres Beam-induced motion correction for sub-megadalton cryo-EM particles 3 2014 10.7554/eLife.03665 10.7554/eLife.03665.001 10.7554/eLife.03665.002 10.7554/eLife.03665.003 10.7554/eLife.03665.004 10.7554/eLife.03665.005 10.7554/eLife.03665.006 10.7554/eLife.03665.007 10.7554/eLife.03665.008 10.7554/eLife.03665.009.

Sigworth, F.J., 2016. Principles of cryo-EM single-particle image processing. Microscopy (Oxf) 65 (1), 57–67.

Sikora, M., Ermel, U.H., Seybold, A., Kunz, M., Calloni, G., Reitz, J., Vabulas, R.M., Hummer, G., Frangakis, A.S., 2020. Desmosome architecture derived from molecular dynamics simulations and cryo-electron tomography. Proc Natl Acad Sci U S A 117 (44), 27132–27140.

Smith, P.R., Peters, T.M., Bates, R.H.T., 1973. Image reconstruction from finite numbers of projections. J. Phys. A: Math. Nucl. Gen. 6 (3), 361–382.

Starck, J.-L., Bijaoui, A., 1994. Filtering and deconvolution by the wavelet transform. Signal Process. 35 (3), 195–211.

Stoschek, A., Hegerl, R., 1997. Denoising of electron tomographic reconstructions using multiscale transformations. J Struct Biol 120 (3), 257–265.

Tegunov, D., Xue, L., Dienemann, C., Cramer, P., Mahamid, J., 2021. Multi-particle cryo-EM refinement with M visualizes ribosome-antibiotic complex at 3.5 A in cells. Nat Methods 18, 186–193.

Tomasi, C., Manduchi, R., 1998. Bilateral filtering for gray and color images. Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271), 839–846.

Turoňová, B., Sikora, M., Schürmann, C., Hagen, W.J.H., Welsch, S., Blanc, F.E.C., von Bülow, S., Gecht, M., Bagola, K., Hörner, C., van Zandbergen, G., Landry, J., de Azevedo, N.T.D., Mosalaganti, S., Schwarz, A., Covino, R., Mühlebach, M.D., Hummer, G., Krijnse Locker, J., Beck, M., 2020. In situ structural analysis of SARS-CoV-2 spike reveals flexibility mediated by three hinges. Science 370 (6513), 203–208.

Van Heel, M., 1987. Similarity measures between images. Ultramicroscopy 21 (1), 95–100.

van Heel, M., Schatz, M., 2005. Fourier shell correlation threshold criteria. J Struct Biol 151 (3), 250–262.

Veesler, D., Campbell, M.G., Cheng, A., Fu, C.-y., Murez, Z., Johnson, J.E., Potter, C.S., Carragher, B., 2013. Maximizing the potential of electron cryomicroscopy data collected using direct detectors. J Struct Biol 184 (2), 193–202.

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 (3), 211–218.

Weickert, J., 1999. Coherence-Enhancing Diffusion Filtering. Int. J. Comput. Vision 31, 111–127.

Zanetti, G., Riches, J.D., Fuller, S.D., Briggs, J.A.G., 2009. Contrast transfer function correction applied to cryo-electron tomography and sub-tomogram averaging. J Struct Biol 168 (2), 305–312.

Zhang, C., Wang, T.T., Sun, D.J., 2013. Image edge detection based on the Euclidean distance graph. {font-family:"Cambria Math" 18, 176–183.

Zhang, Y., Cheng, H.D., Huang, J., Tang, X., 2012. An effective and objective criterion for evaluating the performance of denoising filters. Pattern Recogn. 45 (7), 2743–2757.

Zhao, S., 1999. Wavelet Filtering for Filtered Backprojection in Computed Tomography. Appl. Comput. Harmon. Anal. 6 (3), 346–373.