

Denoising in cryoET

Jose-Jesus Fernandez

Spanish National Research Council (CINN-CSIC),
Health Research Institute of Asturias (ISPA),
Oviedo, Spain.

JJ.Fernandez@csic.es

To appear as Chapter 7 in the book 'Cryo-electron Tomography' edited by Dorit Hanein and Niels Volkmann, Academic Press, Elsevier 2025, ISBN 9780443188299. <https://doi.org/10.1016/C2021-0-02831-5>

Table of contents

- 1. Introduction.**
- 2. Basic concepts of filtering in the real and Fourier space.**
- 3. Handling noise during tomographic reconstruction.**
- 4. Denoising tomograms.**
 - 4.1 Linear methods.**

Gaussian filter.
 - 4.2 Nonlinear methods.**

Median filter.

Beltrami flow.

Other nonlinear methods.
 - 4.3 Anisotropic methods.**

Bilateral filtering.

Nonlocal means.

Anisotropic nonlinear diffusion.
- 5. Deep-learning approaches to denoising tomograms.**
- 6. Contrast improvement through CTF deconvolution.**
- 7. Software.**
- 8. Concluding remarks.**

1. Introduction

Cryo-electron tomography (cryoET) combines the power of 3D imaging and the optimal structural preservation to enable in situ visualization of the subcellular architecture and molecular organization of cells and tissues at the nanoscale. In cryoET a series of projection images (so-called tilt-series) is acquired from the frozen-hydrated sample by tilting it around an axis, usually in the range $\pm 60^\circ$ at intervals of $1-3^\circ$. The acquired images are then mutually aligned and, afterwards, combined through tomographic reconstruction methods to yield a 3D volume or tomogram.

Tomograms contain a wealth of information but their analysis is not an easy task. Apart from the inherent biological complexity, there are a number of factors that make tomogram interpretation tough. The tomograms in cryoET are characterized by an extremely low contrast and signal-to-noise ratio (SNR). This is mainly caused by the use of minimal electron dose during data acquisition (typically an accumulated dose around $100 \text{ e}^-/\text{\AA}^2$) due to the sensitivity of the frozen-hydrated sample to radiation damage. Another cause is the thickness of the sample, typically thicker than 100 nm (and increasingly thicker with the tilt angle), which leads to events of inelastic and/or multiple elastic electron scattering that contribute only noise to the images. In addition, the low contrast of tomograms also derives from the contrast transfer function of the microscope, which substantially attenuates the low-resolution components. Along with the low SNR and contrast, another factor that makes interpretation of the tomograms difficult is the limited tilt range (i.e. not covering the full range $\pm 90^\circ$), which translates into blurring along the electron beam direction (so-called "missing wedge" artefacts).

To increase the SNR and facilitate the analysis of tomograms, denoising techniques are thus required. In cases where repetitive motifs are found in the tomograms, aligning and averaging the motif instances increases the SNR of the specific motif. This is the subtomogram averaging approach to high-resolution structure determination of macromolecules in situ ([Wan and Briggs, 2016](#)). However, this approach is not applicable in a general cryoET study, where the focus may be on the whole cellular environment or pleomorphic structures (e.g. organelles) within the field of view. Denoising may also be an important pre-processing step for the downstream computational workflow (e.g. Segmentation, Feature detection).

There are several stages of the computational workflow in cryoET where the noise can be handled to provide tomograms with reduced noise and higher contrast. First, before reconstructing the tomograms, the acquired projection images may be filtered according to the accumulated electron dose to account for the loss of structural information caused by the radiation damage ([Grant and Grigorieff, 2015](#); [Wan and Briggs, 2016](#)), along with other possible filters. Second, the tomographic reconstruction methods and their parameters can significantly influence on the noise transferred to the tomogram. Finally, after reconstruction, tomogram denoising is a common post-processing stage in the cryoET workflow. Several techniques to denoise tomograms become well

established in the first decade of this century and are still commonly used. Lately, novel strategies based on deep-learning methods have emerged.

This chapter describes standard and emerging approaches used in the reconstruction and in the post-processing stages to reduce noise and increase the contrast in tomograms. The reader will find a discussion of their advantages and limitations, recommendations and a list of available software tools.

2. Basic concepts of filtering in the real and Fourier space.

Transformation of an image or tomogram into the Fourier domain allows straightforward separation of coarse and slowly varying features (low-resolution components) from finer details, edges and sharp transitions (represented by medium and high-resolution components). Since noise is mainly predominant over signal at the medium and high frequencies, attenuating or filtering out the amplitude at those frequencies reduces the contribution of noise. A filter of this kind that preserves the low frequency components is called a low-pass filter. By contrast, a high-pass filter preserves the high frequency components and attenuates or eliminates the others. So-called band-pass filters retain a specified range of medium frequencies.

Filtering in the Fourier space consists in computing the Fourier transform (FT) of the image or tomogram, multiplying the Fourier filter with the original Fourier components, and computing the inverse FT to obtain the filtered image in the real space or spatial domain. The Fourier filters are usually implemented through smooth functions, thus avoiding abrupt cut-offs that would translate into ringing artefacts in real space. Through the convolution theorem of the FT, it is possible to find an equivalent linear, position invariant filter that operates in the real space ([Gonzalez and Woods, 2018](#)).

Filtering in the real space is a procedure that operates directly on the pixels/voxels of the image or tomogram. The spatial filter, also known as mask or kernel, is an operator whose coefficients determine the nature of the filter and is defined over some neighbourhood. The approach is indeed a convolution and performs a sum-of-products operation between the kernel coefficients and the density of the pixels/voxels under the kernel at a specific location in the image/tomogram. In other words, the density of a pixel/voxel of the filtered tomogram is determined by a weighted average of itself and the pixel/voxels in its neighbourhood. This operation is applied for all pixels/voxels.

[Figure 1](#) illustrates these basic concepts. An image (A) and its FT (B) are shown. The low frequencies are around the origin of the FT whereas the medium and high frequencies spread farther throughout the FT. The maximum frequency (also known as Nyquist frequency), given by $0.5/P$ with P being the pixel size of the image or just 0.5 expressed as a normalized frequency, is marked in yellow. The cyan circumference indicates an example frequency for filtering. A low-pass filtered version of the image is shown in (C), which looks blurred and where the lack of details is apparent. As seen in its FT (C, inset) the amplitude of the medium and high frequencies are cancelled out. In (D) the high-pass filtered image is presented, where the fine details (e.g. edges) are preserved and there is loss of global contrast and coarse features. The FT (D, inset) shows that the low frequencies are removed. Panels (E) and (F) show the profile of the filters in the Fourier space as well as the corresponding spatial kernels to be convolved with the original image (sketched by the orange grid in (A)).

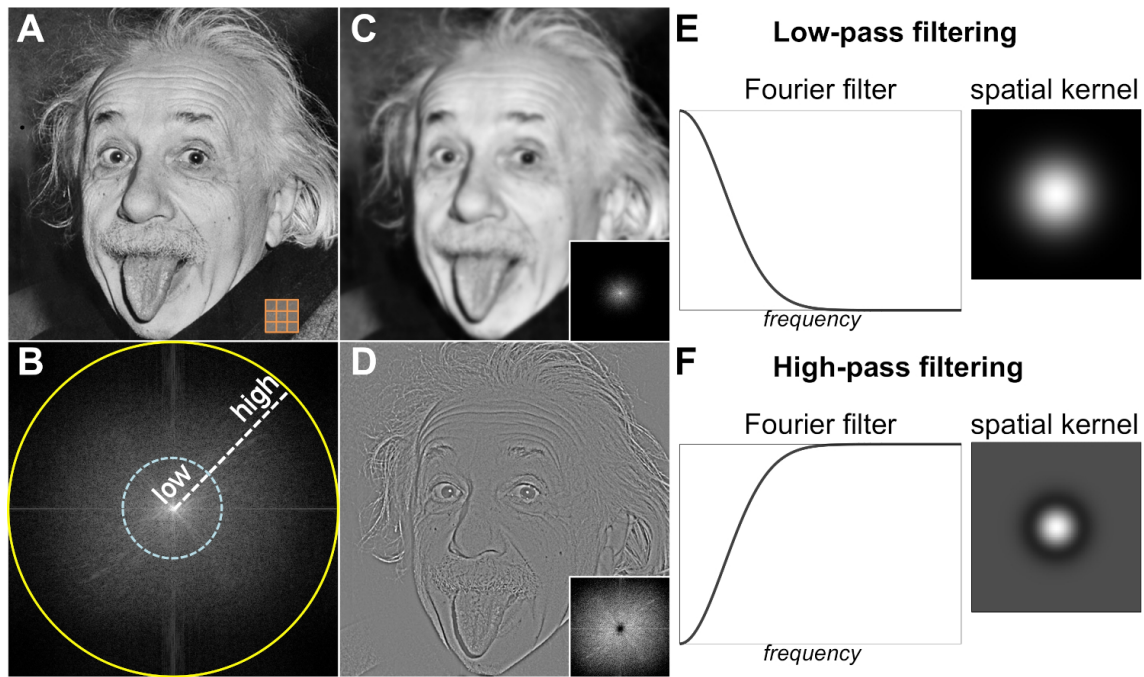


Figure 1. Basic concepts of filtering.

Original image (A), its Fourier transform (B) with the Nyquist frequency highlighted in yellow and a specific example frequency for filtering represented in cyan. Low-pass (C) and high-pass (D) filtered versions of the image, with the FTs shown in the insets. Low-pass (E) and high-pass (F) filters in the Fourier space, as well as equivalent spatial kernels to convolve with the original image (convolution sketched by the orange grid in (A)), to obtain the filtered images shown in (C) and (D).

3. Handling noise during tomographic reconstruction

The tomographic reconstruction problem in cryoET is to compute the tomogram containing the 3D structure of the specimen from the set of aligned projection images. There are a number of reconstruction methods with their own parameters, and the selection of the method and its tuning have an important influence on the noise present in the tomogram.

The mathematical principles of tomographic reconstruction are based upon the central section theorem, which states that the FT of a 2D projection of a 3D object is a central section of the 3D FT of the object ([Gonzalez and Woods, 2018](#)). Therefore, the 3D FT of the specimen can be computed by assembling the 2D FTs of the projection images, which yields the 3D structure of the specimen by an inverse FT. The drawbacks of this approach are related to the non-trivial interpolation in Fourier space and to the computational burden. As a consequence, it is not often used to compute tomograms.

There exist real-space iterative reconstruction algorithms that progressively refine the tomogram by minimizing the error between the experimental projection images and the equivalent projections calculated from the reconstructed tomogram. These methods are robust to face the particularities found in cryoET, namely the limited tilt range, high noise and low contrast. However, the main disadvantage is their high computational expensiveness. SIRT (Simultaneous Iterative Reconstruction Technique) ([Gilbert, 1972](#)) is one such method, very commonly used in cryoET to obtain tomograms with high contrast ([Fernandez, 2012](#)).

The standard method for tomographic reconstruction in cryoET is Weighted Backprojection (WBP), which essentially is equivalent to the Fourier approach above described but working in the real space ([Frank, 2006](#)). The relevance of WBP in ET mainly stems from its computational simplicity, though it is particularly sensitive to noise.

Backprojection distributes the specimen mass present in the projection images evenly over computed back-projection rays according to the experimental tilt angles ([Figure 2A](#)). Backprojection rays from the different images intersect and reinforce each other at the points where mass is found in the original structure, resulting in the reconstructed tomogram. The backprojection process involves an implicit low-pass filtering that makes reconstructions strongly blurred. A high-pass filtering (i.e., weighting), hence the term "Weighted Backprojection (WBP)", is necessary to compensate for this blurring and properly represent the high frequency information in the reconstruction ([Figure 2A](#)). The weighting is to be applied in the directions perpendicular to the tilt axis (planes XZ of the tomogram, with the tilt axis running along the Y axis, as is standard in cryoET). However, for computational reasons, it is more efficient to apply the weighting directly to the aligned projection images (along the X axis) prior to the reconstruction.

The weighting filter in WBP plays a paramount role in the noise and contrast of the reconstructed tomogram. There is a trade-off between the level of detail and

noise present in the tomogram that has to be considered when selecting and tuning the filter. A description of the most common filters used in cryoET follows.

The high-pass weighting filter used in WBP to compensate the blurring involved by the backprojection process is a linear ramp filter ([Frank, 2006](#); [Gonzalez and Woods, 2018](#)):

$$W_{\text{ramp}}(f) = f$$

where f denotes the normalized frequency (i.e. in $[0,0.5]$, with 0.5 denoting the Nyquist frequency). This filter works well under ideal, noise-free conditions ([Figure 2A, 2B](#)), but its susceptibility to transfer and amplify the noise from the projection images to the tomogram is a severe disadvantage, particularly under the highly noisy conditions in cryoET.

In practice, the actual weighting filter combines the ramp filter with a low-pass filter so as to attenuate or filter out the high frequency noise ([Figure 2B](#)). One common weighting filter is based on the Hamming window:

$$W_{\text{Hamming}}(f) = \begin{cases} W_{\text{ramp}}(f), & 0 \leq f \leq f_t \\ W_{\text{ramp}}(f) * (0.54 - 0.46 * \cos(\pi (0.5 - f)/(0.5 - f_t))), & f_t < f \leq 0.5 \end{cases}$$

where f_t denotes the normalized frequency from which the Hamming window is applied and 0.5 represents the Nyquist frequency. A value of $f_t = 0.5$ means that the ramp filter is used throughout the range, without low-pass filtering at all. The most common value for f_t is 0, which provides the highest level of smoothing by having an effect on the whole frequency range. This is the default behaviour in Tomo3D ([Agulleiro and Fernandez, 2015, 2011](#)), one standard software tool for tomographic reconstruction in cryoET. IMOD ([Kremer et al., 1996](#)), the standard package in the cryoET field, also includes this weighting filter as an option.

Another common weighting filter is based on the Gaussian function. It behaves similarly to the above Hamming-based filter: the ramp filter is used for all frequencies up to f_t , from which a Gaussian fall-off is applied. This is the default option in IMOD, where the default value for f_t is 0.35.

A new weighting filter that provides visual effects similar to the tomographic reconstructions with SIRT has emerged in the last few years. The great advantage with respect to the real SIRT method is that the processing time is the same as WBP. And the advantage with regard to the other filters is that it provides significantly cleaner tomograms with higher contrast ([Figure 2C](#)). The filter was originally developed in the context of medical computerized tomography ([Zeng, 2012](#)), and it is rapidly becoming standard in cryoET, especially for visual inspection of tomograms containing pleomorphic structures or for selection of particles of interest for later subtomogram averaging. The weighting filter is commonly referred to as "Fake SIRT" or "SIRT-like filter" and is defined as:

$$W_{\text{SIRT}}(f) = W_{\text{ramp}}(f) * (1 - (1 - \alpha/f)^k)$$

where k denotes the number of SIRT iterations to simulate and α is a constant. These parameters were slightly tweaked in IMOD to fairly match the SIRT reconstructions obtained with this package. This filter is now available in IMOD and in Tomo3D.

[Figure 2](#) illustrates the process of tomographic reconstruction with WBP, presents the weighting filters described here, and shows the effects on an experimental tomogram of ribosomes ([Bharat and Scheres, 2016](#)) (EMPIAR-10045). While the ramp filtering is required to compensate the global blurring and properly weight the structural information of the tomogram, the importance of strong low-pass filtering to substantially attenuate the medium and high frequency noise is clearly evident. A Hamming-based filter throughout the entire range (i.e. $f_t = 0$) and, remarkably, SIRT-like filters provide the highest level of cleanliness and contrast.

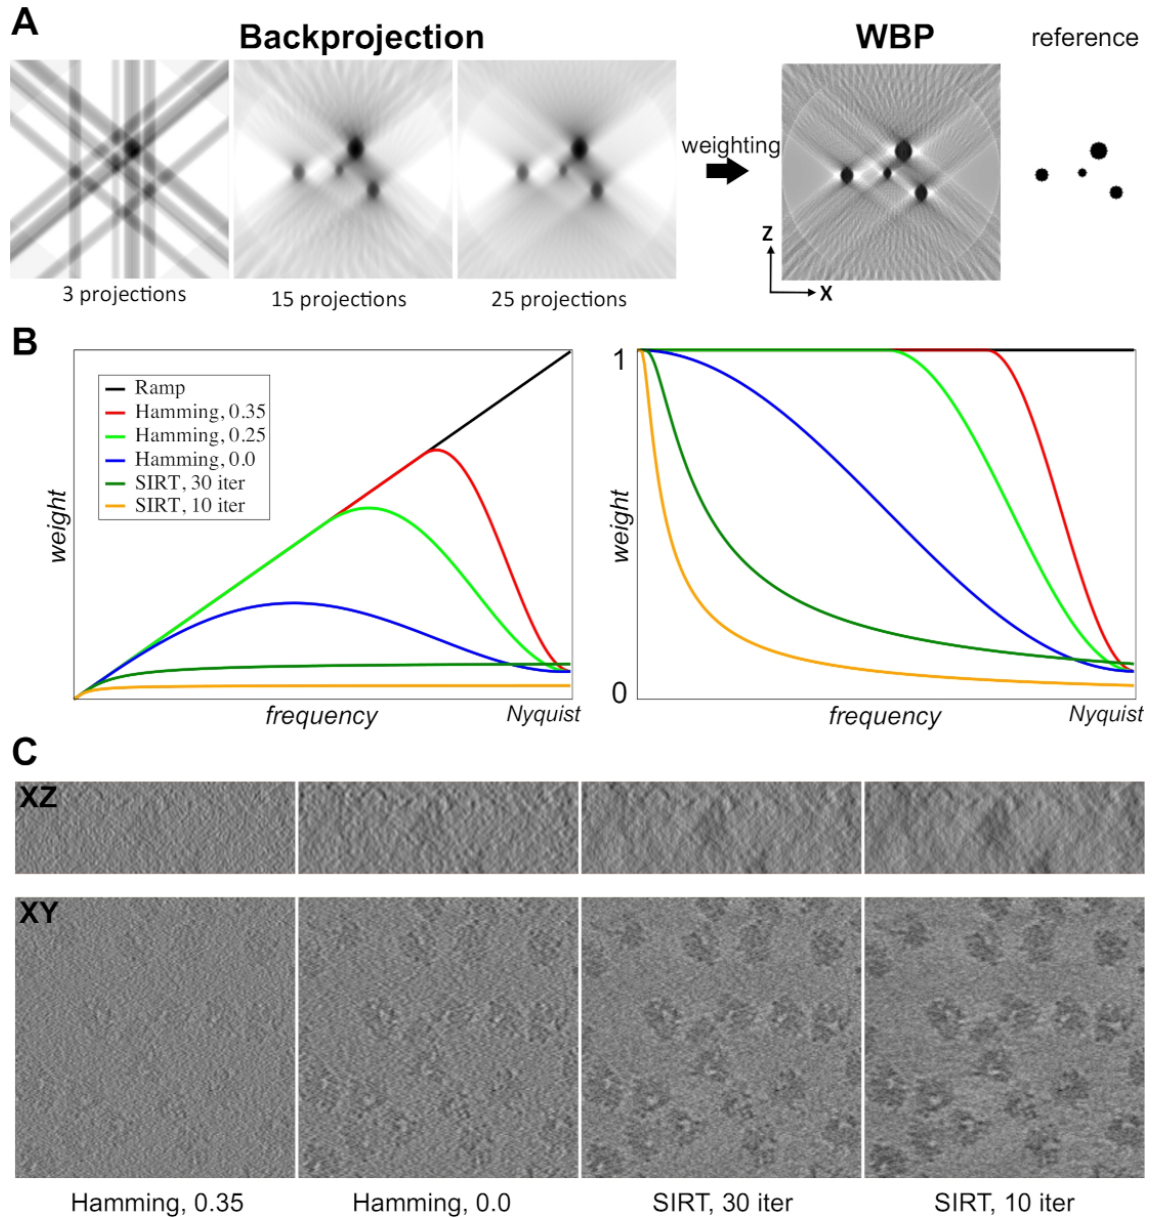


Figure 2. Handling noise during tomographic reconstruction with WBP.

(A) Tomographic reconstruction with WBP in 2D. The projection data are backprojected towards the reconstruction space, and their sum provides the reconstruction. To compensate for the blurring caused by the backprojection process, a weighting is required, hence Weighted Backprojection (WBP). The blur and elongation caused by the missing wedge (off $\pm 50^\circ$) is observable in the reconstruction when compared to the reference data (right).

(B) Weighting filters actually used in WBP to enhance the high frequency information and control the noise, formed by a multiplication of the ramp and a low-pass filter (left). On the right, only the low-pass filters are shown.

(C) Effect of the weighting filters in WBP with a ribosome tomogram from EMPIAR-10045. Planes XY and XZ are shown.

4. Denoising tomograms

Tomograms in cryoET may be degraded by substantial noise, despite the strategies to improve the SNR and contrast during tomographic reconstruction. Denoising intends to filter out noise but trying to preserve the structural details as much as possible with the ultimate aim of facilitating visualization and interpretation of tomograms. Several techniques originated in the image processing and computer vision fields were explored in cryoET early this century, some of which become standard tools and are extensively used. This section reviews these techniques.

In general, denoising techniques can be grouped into three wide categories: linear, non-linear and anisotropic. [Figure 3](#) illustrates them in 2D. Linear techniques are based on low-pass filters or, equivalently, on local averages using uniform or Gaussian-like kernels. Here, all voxels are substituted by a weighted average of the voxels in the local neighbourhood, with the weights given by the kernel. The term "linear" indicates that all voxels in the tomogram are subjected to the same kernel regardless of the underlying structural detail. These techniques thus succeed in reducing the noise at the expense of blurring edges and features. Non-linear techniques then overcome this limitation by tuning the strength of the filtering (i.e. the parameters/weights of the kernel) to the specific detail the voxel belongs to ([Figure 3C](#)). This way, the filtering is strong in homogeneous areas whereas it is greatly attenuated or cancelled in voxels with a high gradient, hence indicating that there is an edge or a potentially important detail. The unwanted effect of this strategy is that edges might remain somewhat noisy. Anisotropic non-linear (or just anisotropic for short) methods tune not only the strength, but also the direction of the filtering. There exist different ways to do it ([Figure 3D](#)). The outcome is that the edges are subjected to a filtering process that runs parallel to them, or at least not across them, thereby cleaning and enhancing them.

Linear methods have their equivalent in Fourier space, as described in Section 2, so they could be implemented in either domain. However, non-linear and anisotropic methods only can be implemented in real space. Linear methods, particularly those based on spherically symmetric kernels, can be applied to the tomogram in 3D or, equivalently, in 2D to the tilt-series of acquired images from which the tomogram is calculated, which is a faster process. Non-linear or anisotropic methods could also be applied to the tilt-series (e.g. [Maiorca et al., 2012](#)), though it must be done with caution since the linear relationship between the projection images and tomogram might break, with artefacts potentially arising in 3D.

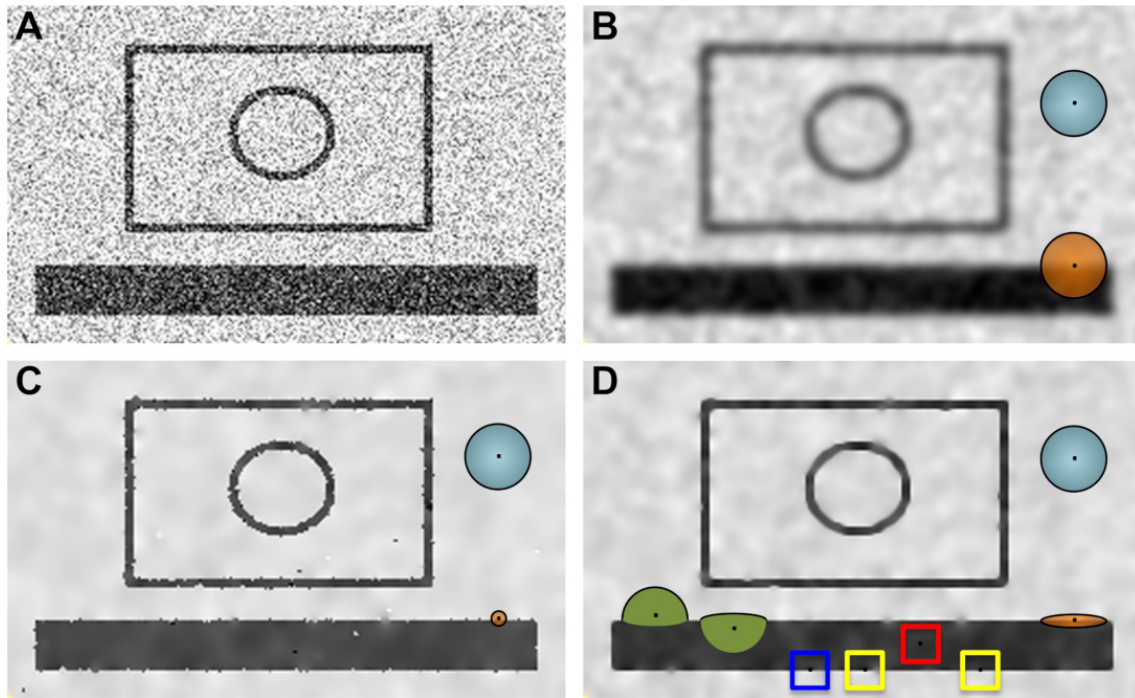


Figure 3. Categories of denoising techniques.

(A) Example of a noisy image. (B) In linear methods, the same kernel (denoted by colored circles whose radii intend to resemble the standard deviation of a Gaussian function) is applied to all voxels, regardless of the existence of edges or not. As a result, the noise present in the image is reduced, but the edges look blurred. (C) In non-linear methods, the strength of the filtering is tuned according to an edge indicator. The filtering is attenuated or cancelled at edges (orange kernel with small width) whereas relatively homogeneous areas are strongly smoothed (blue kernel). Though preserved, the edges may thus remain partially noisy. (D) Anisotropic methods tune the strength and direction of the filtering. There exist different strategies. First, at edges the filtering is applied only parallel to them, but not across, as denoted by the elongated orange kernel. Second, the kernel is adjusted so that only voxels with similar density to the voxel at the center of the kernel are involved in the filtering. This is sketched by the green kernels centered on a couple of points at the two sides of an edge. Another option is based on averaging pixels, not necessarily close, whose surrounding area looks similar. The colored boxes represent this strategy. To filter the voxel at the center of the blue box, itself and those at the yellow boxes are considered, but not the red box because it is clearly different. As a result of these anisotropic strategies, the edges look clean and enhanced.

4.1 Linear methods

Gaussian filter

The most common linear filter is the Gaussian filter, which essentially replaces any voxel by a weighted average of its neighbour voxels with decreasing weights for those at larger distances ([Gonzalez and Woods, 2018](#)). Several features make it particularly attractive. First, it is spherically symmetric, meaning that its response is independent of orientation. Other linear filters, for instance box or uniform kernels, are not and thus might favour smoothing in specific directions. Second, the FT of a Gaussian is also a Gaussian function, so the Gaussian filter yields a low-pass filter with smooth behaviour in both real and Fourier spaces. Third, there is only one parameter, the standard deviation, that is very well known and intuitive to be tuned according to the filtering needs. Finally, the Gaussian filter is separable, meaning that it can be implemented as a succession of convolutions along the three axes of the tomogram. In other words, instead of a 3D convolution with a 3D Gaussian kernel, it can be implemented as three 1D convolutions, which is significantly faster from the computational point of view. The Gaussian kernel and its separable form can be expressed as:

$$G(\mathbf{v}) = G(x, y, z) = e^{-\frac{x^2+y^2+z^2}{2\sigma^2}} = e^{-\frac{x^2}{2\sigma^2}} \cdot e^{-\frac{y^2}{2\sigma^2}} \cdot e^{-\frac{z^2}{2\sigma^2}}$$

with $\mathbf{v} = (x, y, z)$ denoting the coordinates of a voxel. The standard deviation controls the extent of the spatial filtering, with larger values causing stronger smoothing. The filtering in real space can be expressed as

$$I_{\text{denoised}}(\mathbf{v}) = C \cdot \sum_{\mathbf{u} \in K} I(\mathbf{u}) \cdot G(\mathbf{v} - \mathbf{u})$$

where I denotes the original tomogram, G the Gaussian kernel, \mathbf{v} a voxel and $\mathbf{u} \in K$ its neighbourhood within the window K defined by the kernel. C denotes a normalization factor (i.e. inverse of the sum of all the weights involved in the kernel) that prevents a bias during filtering (i.e. the average density in the original and filtered version will be the same). As a linear filter, the Gaussian filter has the drawback of blurring edges and features of interest. However, it is still useful for scientists to inspect the tomograms and for manual segmentation ([Danita et al., 2022](#)). Moreover, an initial Gaussian filtering is essential as a pre-processing step for other more sophisticated denoising filters ([Fernandez, 2009](#); [Fernández and Li, 2003](#)) or for segmentation algorithms ([Martinez-Sanchez et al., 2014](#)) because it ensures more reliable computation of gradients, which is important in those algorithms.

4.2 Non-linear methods

Median filter

The median filter is the best-known non-linear denoising method (Gonzalez and Woods, 2018). An iterative version of the median filter in 3D was introduced and successfully used in cryoET (van der Heide et al., 2007). It consists in substituting a voxel by the median of its neighbours: the density values of the voxels encompassed by the 3D spatial kernel are sorted and the value placed in the middle is picked up as the new density value of the center voxel. Using the same notation as above, the filtering can be expressed as follows:

$$I_{\text{denoised}}(\mathbf{v}) = \underset{\mathbf{u} \in K}{\text{median}}\{I(\mathbf{u})\}$$

The method thus eliminates extreme density values and forces voxels to be more like their neighbours. The process is repeated for a number of cycles, and three iterations were observed to be optimal in terms of performance vs. processing time (van der Heide et al., 2007). The great advantage of the method is the lack of parameters difficult to tune, which makes it user-friendly.

Beltrami flow

The Beltrami flow is a non-linear method that tunes the strength of the filtering according to an edge indicator based on geometry operators, hence preserving structures of interest (Fernandez, 2009). It relies upon a geometric diffusion approach (Kimmel et al., 1997) and considers the tomogram I as a 3-manifold embedded into a 4D space, with the fourth dimension being the density $S = (x, y, z, I(x, y, z))$. The Beltrami flow aims to minimize the area of the manifold while maintaining edges. It is formulated as a partial differential equation:

$$I_t = \frac{1}{\sqrt{g}} \text{div} \left(\frac{\nabla I}{\sqrt{g}} \right)$$

where I_t denotes the derivative of I with respect to the time, ∇I is the gradient vector whose components are the partial derivatives of I with respect to x , y and z . g denotes the determinant of the first fundamental form of the surface S , which is $g = 1 + |\nabla I|^2$. Finally, div is the divergence operator (i.e. the sum of the partial derivatives of the argument with respect to x , y and z). The term $\frac{1}{\sqrt{g}}$ in the previous equation acts as an edge indicator since it is proven to be the projection of the normal-to-the-surface to the vector representing the 4th dimension (Kimmel et al., 1997) (Figure 4A). Therefore, the Beltrami flow is a selective denoising method that minimizes the filtering (i.e. diffusion) at and across edges whereas it applies extensive filtering elsewhere. The method has no complicated parameters to tune, as the estimation of the edges and their strength is performed directly from the gradient. Nonetheless, the method is solved in an iterative way and hence a number of iterations have to be specified. A value around 100 cycles has been shown to provide good solutions in cryoET in terms of noise filtering and structure preservation (Fernandez, 2009).

Other non-linear methods

There have also been several other approaches within this category of non-linear methods, but they have not been extensively used in the field. Among them are those based on multiscale Wavelet transformation ([Frangakis et al., 2001](#); [Huang et al., 2018](#); [Stoschek and Hegerl, 1997](#)).

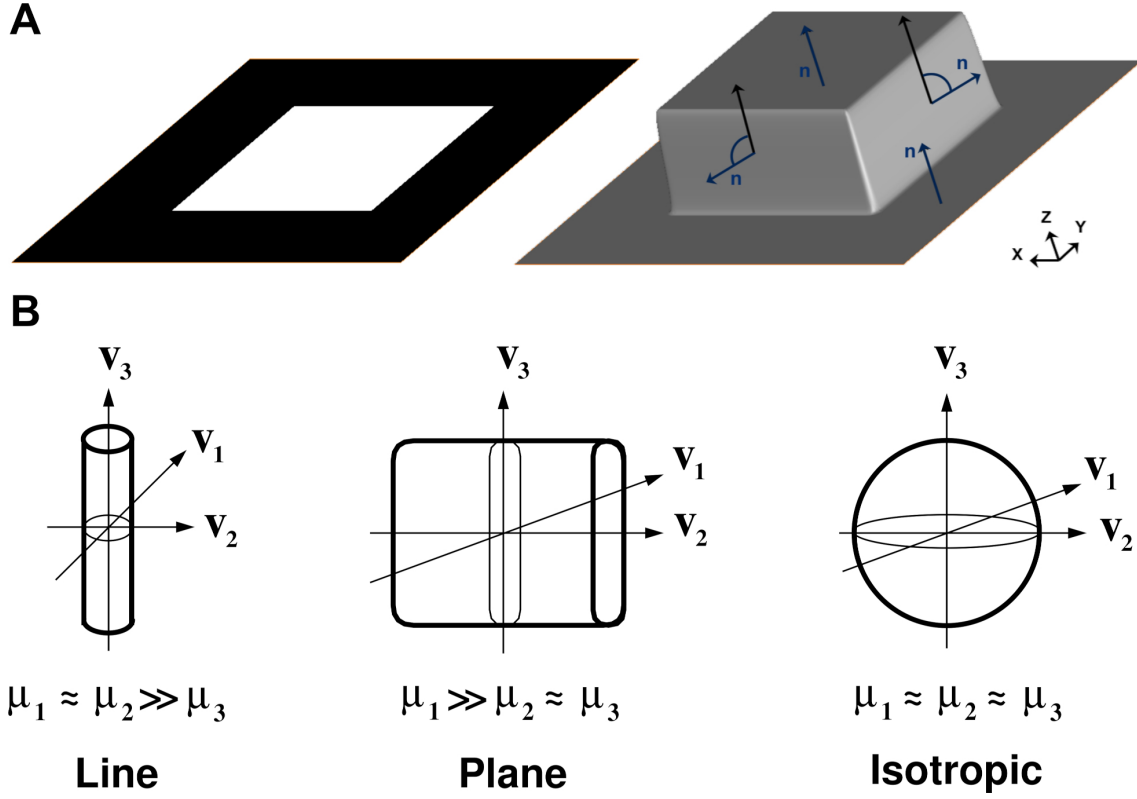


Figure 4. Fundamentals of diffusion-based denoising methods.

(A) Beltrami Flow. A 2D image $I(x, y)$ made up of a white square over black background (left) is viewed as a surface $S = (x, y, I(x, y))$ in a 3D space (right). The edges are seen as cliffs in the Z direction. At each point of the surface, the projection of the normal \mathbf{n} (arrows in blue) to the Z direction (arrows in black) acts as an edge indicator, yielding little value at sharp edges. In uniform areas, however, the normal to the surface \mathbf{n} runs parallel to Z and the projection thus yields maximum value. The Beltrami flow uses this information to minimize diffusion at edges whereas it applies extensive diffusion elsewhere.

(B) Anisotropic non-linear diffusion (AND). The local structure around each voxel is determined based on an eigen-analysis of the structure tensor that provides three orthogonal eigenvectors $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ and their corresponding eigenvalues $\mu_1 \geq \mu_2 \geq \mu_3$. This analysis allows identification of these basic local structures and the strength and direction of the smoothing are adaptively tuned so that the edges of these structures are preserved and enhanced. Therefore, in linear structures the filtering should be applied along the major direction \mathbf{v}_3 while in planar ones should be across the plane defined by \mathbf{v}_2 and \mathbf{v}_3 .

4.3 Anisotropic methods

Bilateral filtering

Bilateral filtering is an anisotropic non-linear technique that was introduced in the computer vision field with great success (Tomasi and Manduchi, 1998) and has been often used in cryoET (Jiang et al., 2003). It reduces noise by weighted averages of neighbour voxels with weights reflecting both the spatial distance and the similarity in density (Figure 3D, green kernels). The filter is composed of two Gaussian kernels, one on the spatial domain and other on the density domain. The former is the standard Gaussian low-pass filter. The second Gaussian kernel is the key of the method and it accounts for the difference in density with respect to the voxel at the center of the kernel: the larger difference, the smaller the contribution in the weighted average. This second kernel has potential not only to avoid edge blurring but also to enhance edges by smoothing along them. The filtering is expressed as

$$I_{\text{denoised}}(\mathbf{v}) = C \cdot \sum_{\mathbf{u} \in K} I(\mathbf{u}) \cdot e^{-\frac{\|\mathbf{v}-\mathbf{u}\|^2}{2\sigma_s^2}} \cdot e^{-\frac{|I(\mathbf{v})-I(\mathbf{u})|^2}{2\sigma_d^2}}$$

where the bilateral kernel is given by the product of the two Gaussian functions whose width is defined by the two parameters of the method, the standard deviations σ_s and σ_d . C represents a normalization factor (i.e. inverse of the sum of the weights) with the same aim as above. σ_s specifies the strength of the standard spatial Gaussian filtering. σ_d controls the discrimination between true structural details and noise and it needs to be finely tuned. It must be high enough to smooth noise but, at the same time, lower than the density variations across the edges to preserve (i.e. between different structural features). If σ_d is too high, the bilateral filter will behave as a standard Gaussian low-pass filter. If it is too low, there will be no filtering and the noise will remain.

Non-local means

The principle of non-local means (NLM) relies on the redundancy present in any image or tomogram. The original method (Buades et al., 2005) revolutionized the image processing and computer vision fields because its abilities to attenuate noise without destroying textures and fine details. The impact in cryoET, where several improvements were presented (Darbon et al., 2008; Wei and Yin, 2010), has been relatively modest, though. While most denoising methods work by averaging voxels in local neighbourhoods, NLM exploits the existence of regions, or patches, spread throughout the tomogram that show similar patterns, and then reduces noise by a weighted average of the pixels in those similar regions, regardless of the spatial distance (Figure 3D, colored boxes). Thus, compared to the bilateral filtering, the similarity in the density domain, but not the spatial neighbourhood, is the central basis of the NLM. The filtering process can be expressed as:

$$I_{\text{denoised}}(\mathbf{v}) = C \cdot \sum_{\mathbf{u} \in W(\mathbf{v})} I(\mathbf{u}) \cdot e^{-\frac{\|V(N_{\mathbf{v}}) - V(N_{\mathbf{u}})\|^2}{h^2}}$$

where $V(N_{\mathbf{v}})$ and $V(N_{\mathbf{u}})$ denote vectors with the density of the voxels in the square neighbourhood windows (patches) $N_{\mathbf{v}}$ and $N_{\mathbf{u}}$ centered at the voxels \mathbf{v} and \mathbf{u} , respectively. $W(\mathbf{v})$ is the search area, which may be the whole tomogram or a area around the voxel \mathbf{v} , where patches $N_{\mathbf{u}}$ similar to $N_{\mathbf{v}}$ are to be found. The similarity between the patches $N_{\mathbf{v}}$ and $N_{\mathbf{u}}$ is determined by the Euclidean distance between the vectors $V(N_{\mathbf{v}})$ and $V(N_{\mathbf{u}})$, and the final weight is given by the Gaussian function whose standard deviation h acts as a filtering parameter and C is a normalization factor. Apart from the patch size and search area, the important parameter to control the degree of filtering is h . Values too high of h will not discriminate patches, resulting in strong smoothing and blurring. By contrast, too low values will provide a stringent setting of weights, with little denoising effects.

Anisotropic non-linear diffusion

Anisotropic non-linear diffusion (AND) was successfully introduced in the computer vision field (Perona and Malik, 1990; Weickert, 1998) and by far has been the predominant denoising method in cryoET (Fernandez and Li, 2005; Fernández and Li, 2003; Frangakis and Hegerl, 2001), where it was proved to be superior to other methods described in this section (Narasimha et al., 2008). AND is inspired on diffusion, the physical process whereby concentration differences are equilibrated as a function of time without creating or destroying mass. In AND the density values play the role of concentration, and the aim is to diffuse (smooth) the density but preventing it from crossing edges (Figure 3D, orange kernel). AND achieves feature preservation and enhancement as the strength and direction of the filtering are adaptively tuned to the local structure around each voxel. This local structure is estimated by eigen-analysis of the structure tensor:

$$J(I) = \nabla I \cdot \nabla I^T = \begin{bmatrix} I_x^2 & I_x I_y & I_x I_z \\ I_x I_y & I_y^2 & I_y I_z \\ I_x I_z & I_y I_z & I_z^2 \end{bmatrix} = \mathbf{V} \mathbf{Q} \mathbf{V}^T$$

with $\nabla I = (I_x, I_y, I_z)$ being the gradient vector of the tomogram I . $\mathbf{V} \mathbf{Q} \mathbf{V}^T$ denotes the eigen-decomposition of $J(I)$, which yields three orthogonal eigenvectors $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ and their corresponding eigenvalues $\mu_1 \geq \mu_2 \geq \mu_3$. The first eigenvector points to the direction of the maximum density variation whereas the third one to the direction of the minimum one. Based on the relative values of the eigenvalues, basic local structures can be recognized (Figure 4B), which are then used to adapt the smoothing, as described next.

AND follows the diffusion equation $I_t = \text{div}(\mathbf{D} \cdot \nabla I)$, with I_t being the derivative of I with respect to the time and div the divergence operator, as already

mentioned above. The 3x3 matrix \mathbf{D} is the diffusion tensor and tunes the filtering according to the local structure. \mathbf{D} is built from the eigenvectors $\{\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3\}$ of the structure tensor and its eigenvalues λ_i (ranking in $[0,1]$) define the strength of the smoothing along the corresponding direction \mathbf{v}_i :

$$\mathbf{D} = \mathbf{V}\mathbf{L}\mathbf{V}^T = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3] \begin{bmatrix} \lambda_1 & 0 & 0 \\ 0 & \lambda_2 & 0 \\ 0 & 0 & \lambda_3 \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3]^T$$

The standard mode to set the diffusion tensor \mathbf{D} is called Edge-Enhancing Diffusion (EED), though other modes exist (Fernández and Li, 2003; Frangakis and Hegerl, 2001). For edge preservation and enhancement in locally planar features, the smoothing must run mainly along the major directions of the plane (\mathbf{v}_2 and \mathbf{v}_3), hence $\lambda_2 = \lambda_3 = 1.0$, whereas the strength along the normal to the plane (\mathbf{v}_1) is set as a monotonically decreasing function of the gradient (the higher the gradient, the lower the strength): $\lambda_1 = 1.0 - \exp\left(\frac{-3.31488}{(|\nabla I|/K)^8}\right)$. Here, the parameter K acts as a gradient threshold that defines edges. This strategy is also valid for isotropic structures. In linear features, however, the filtering is applied along their major direction ($\lambda_3 = 1.0$) while along the other directions it depends inversely on the gradient, as above.

The method has two important parameters. First, a number of iterations have to be specified since it is solved in an iterative way, and a number around 10-20 usually works well. Second, the most important parameter is the gradient threshold K . Voxels with higher gradient are considered edges to be preserved, thus decreasing the filtering strength along the corresponding eigen-direction(s). This parameter has to be set precisely because too high values provide solutions similar to the Gaussian low-pass filtering whereas too low values may leave the tomogram noisy. It is dataset-dependent, its tuning is not trivial and it is usually set by trial-and-error. A strategy was devised to set the parameter K automatically based on the average gradient of the whole 3D volume or a noise sub-region, and in time-varying fashion to minimize the effects of excessive iterations. This strategy facilitates user operation by providing acceptable denoised solutions from which manual refinement could follow (Fernández et al., 2007).

Figure 5 shows the performance of the most common denoising methods used in cryoET with a tomogram of HIV virions (Briggs et al., 2006) (EMD-1155) that has traditionally served as a kind of benchmark.

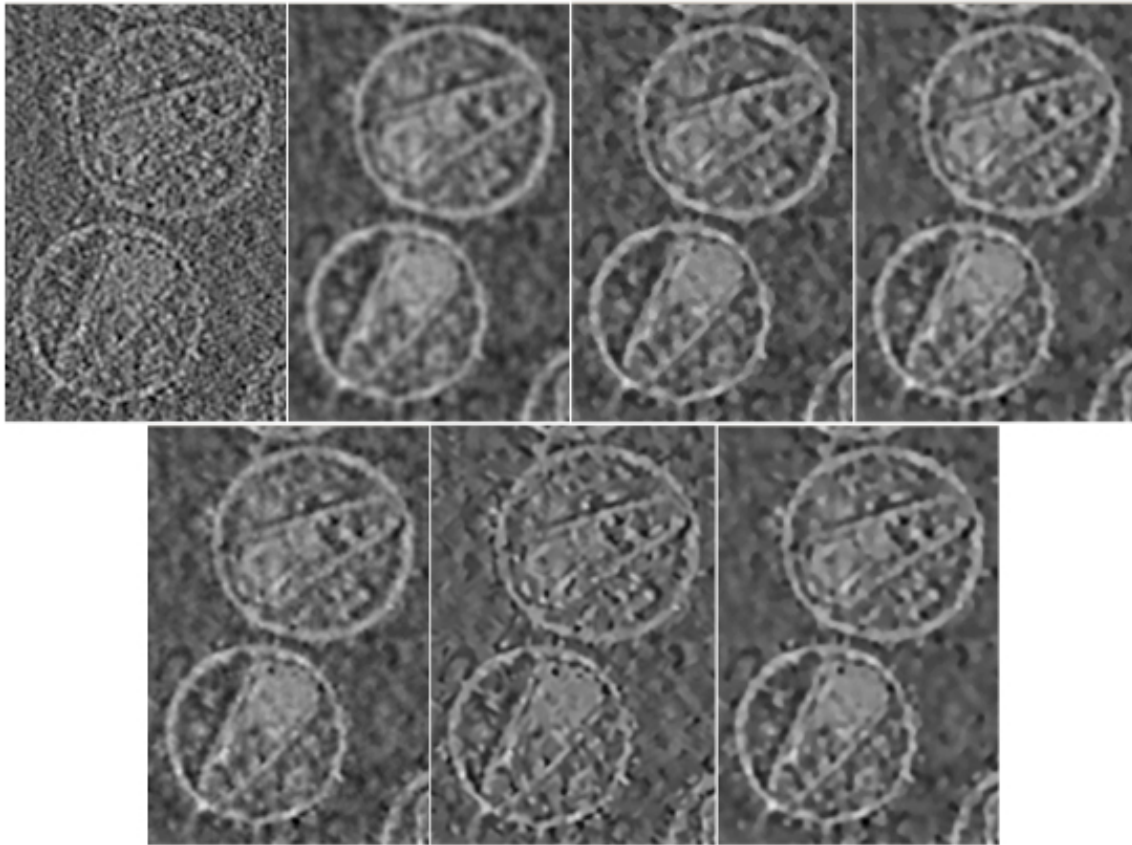


Figure 5. Standard denoising methods applied to a tomogram of HIV virions.

Top row: from left to right, original tomogram, denoised with a Gaussian filter, with three iterations of median filter and with Beltrami Flow. Bottom row, anisotropic methods: from left to right, denoised with bilateral filtering, non-local means and anisotropic non-linear diffusion. Note how the features look sharper in anisotropic methods.

5. Deep-learning approaches to denoising tomograms

The recent revolutionary advances in deep learning (LeCun et al., 2015) have brought about novel and remarkable approaches to denoising in the image processing and computer vision fields, which have also been adapted to the cryoET field. A neural network (NN) can be interpreted as a highly parameterized function f_{θ} that maps an input x to an output y , with θ denoting the parameters of the NN. The parameters θ are determined by a training process that, using a large set of pairs input-target (x_i, y_i) , aims at minimizing a loss function L (e.g., the square error) between the targets y_i and the predictions from the NN $f_{\theta}(x_i)$, that is:

$$\operatorname{argmin}_{\theta} \sum_i L(f_{\theta}(x_i), y_i)$$

Convolutional NNs (CNNs or ConvNets, for short) are a particular type of NN especially designed to process 2D images and 3D volumes that achieve spectacular results in many image processing and computer vision tasks (LeCun et al., 2015). U-nets are a specific kind of CNN very well suited for segmentation and regression.

For the denoising task, the traditional deep-learning approach so far consisted in mapping noise-corrupted inputs (x_i) to clean, noise-free targets (y_i) . However, the requirement of clean targets for training prevents application of this approach in many domains, including cryoET, where acquisition of noise-free signals is unfeasible. To overcome this limitation, the Noise2Noise approach was introduced (Lehtinen et al., 2018), where the authors demonstrated that a signal can be recovered from noisy observations without knowledge of the clean signal and without explicit characterization of the noise. Here, the NN (actually, a CNN with U-net architecture) is trained using pairs of independent noisy instances of the same signal acting as inputs (x_i) and targets (y_i) , and the NN model learns the expectation value, actually the underlying noise-free signal.

Application of the Noise2Noise approach is nearly straightforward in cryoET thanks to the data acquisition based on dose-fractionation (Bepler et al., 2020; Buchholz et al., 2019a, 2019b; Tegunov and Cramer, 2019). Every image of the tilt-series is acquired by fractionating the electron dose into multiple movie frames, which are then aligned and summed to yield the actual acquired image. By splitting the collected frames into two sets, even and odd frames, it is possible to render two independent noisy observations of the image with similar accumulated electron dose. This is repeated for all images of the tilt-series and, as a result, two independent half tilt-series (from the even and the odd frames) are obtained, which are then subjected to 3D reconstruction to yield two independent half tomograms. An alternative option to produce the pair of independent tomograms is by reconstructing them from all odd and from all even tilt images (with each image involving all its movie frames), as in Warp (Tegunov and Cramer, 2019).

The Noise2Noise NN is then trained with the two independent tomograms, which are used as the input and target indistinctly (Figure 6A). For the training, the tomograms are actually split into 3D patches (with size typically in the range

64x64x64 to 128x128x128) that are randomly selected to serve as the pairs input-target (x_i, y_i) . Once the training is finished, the NN can infer the denoised version of the all patches and, upon their assembly, the denoised tomogram.

There are two strategies to obtain the definite denoised tomogram. First, the trained denoiser NN can be applied to the two independent half tomograms, and the two results are then voxel-wise averaged, as in Cryo-Care ([Buchholz et al., 2019a](#)). Alternatively, the original whole tomogram (i.e. reconstructed with all acquired frames and tilts) can be fed directly to the denoiser NN, as in Topaz-Denoise ([Bepler et al., 2020](#)) and Warp ([Tegunov and Cramer, 2019](#)).

The performance of deep-learning denoising approaches applied to tomograms in cryoET can be outstanding. Some denoised results show remarkable noise reduction with excellent preservation of structural features. A unique feature is their exceptionally good abilities to flatten the background, as observed in some cases. [Figure 6](#) shows an illustrative example with a tomogram from *Thermoanaerobacter kivui* (EMPIAR-11058) ([Dietrich et al., 2022](#)) processed with Cryo-Care ([Buchholz et al., 2019a](#)).

One disadvantage of these approaches is the computational requirements to train the NNs. Fortunately, there is availability of NNs that have already been pre-trained with numerous representative tomograms, thus providing a general denoising model in cryoET, as specifically in Topaz-Denoise ([Bepler et al., 2020](#)). This makes application of this deep-learning denoising approach straightforward, without the need for dataset-specific training.

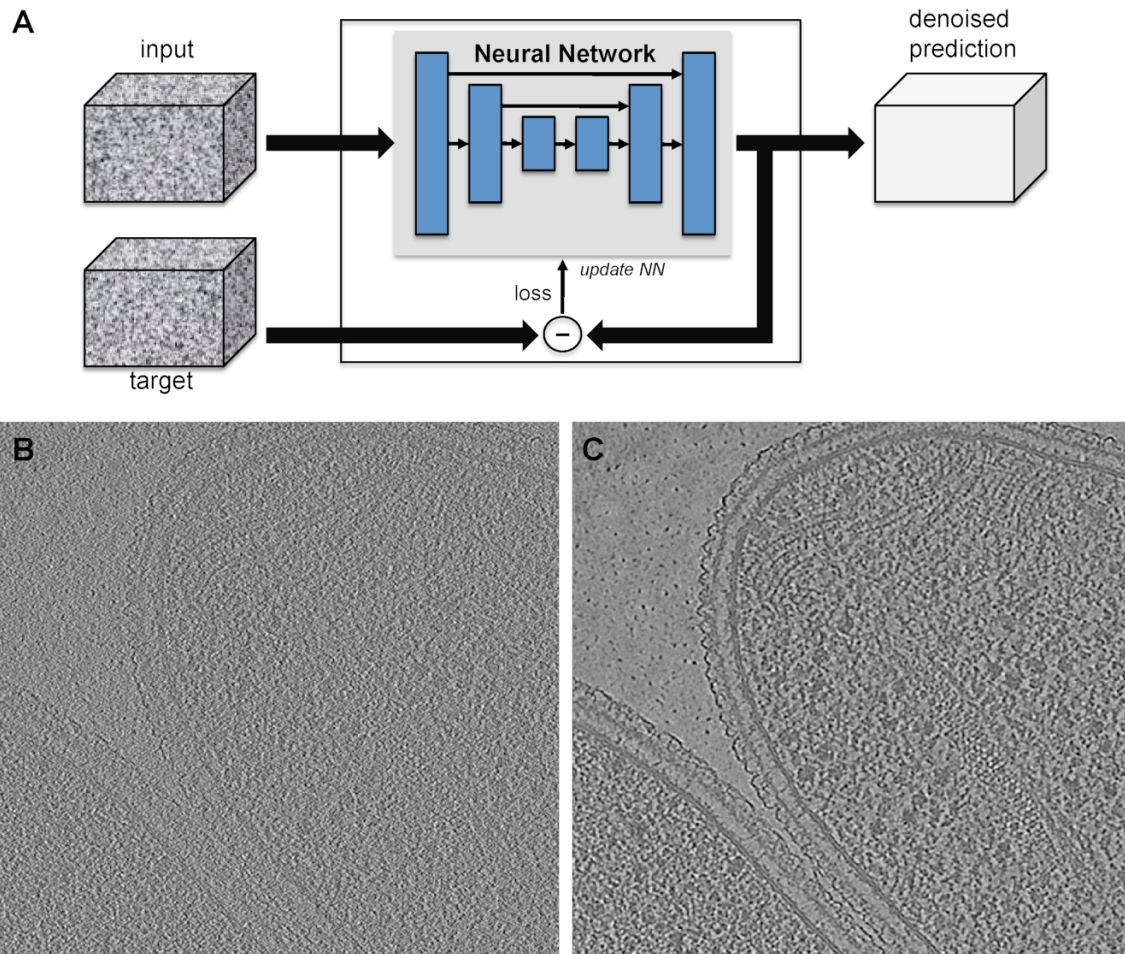


Figure 6. Deep-learning approaches to denoising.

(A) Training of a Noise2Noise neural network (NN, actually a U-net). Two independent tomograms, computed from the acquired even/odd frames throughout the tilt-series (or alternatively, even/odd tilt images), are used as input and target indistinctly. The tomograms are split into 3D patches to be fed to the NN. An input patch passes through the NN and a prediction of the denoised version of the patch is obtained, which is then compared to the target to obtain a loss measure. Based on the loss, the NN is then updated to make the prediction more accurate. This process is repeated for all patches and for a number of times (also known as epochs). Once the NN is fully trained, the NN will infer the denoised version of all patches, which will be assembled to yield the denoised tomogram.

(B) Original tomogram from *Thermoanaerobacter kivui* (EMPIAR-11058).

(C) Denoised result with the Noise2Noise approach.

6. Contrast improvement through CTF deconvolution

Tomograms in cryoET have very poor low-resolution contrast that hinders visual inspection of their contents. Apart from the low electron density of the unstained frozen biological samples, one major cause of this low contrast is the Contrast Transfer Function (CTF) of the electron microscope (Fernandez, 2012; Wan and Briggs, 2016). The CTF arises from the aberrations of the lenses and from the defocus used in imaging, and has an oscillatory nature that produces amplitude modulation and phase reversals in the Fourier components of the images. At the lowest frequencies, the CTF exhibits very low values, thereby significantly attenuating the amplitude of those Fourier components and thus reducing the low-resolution contrast. CTF correction is an essential, well-established step in subtomogram averaging studies aiming at high resolution (Wan and Briggs, 2016). However, compensation for the CTF effects is not usually applied at the level of raw tomograms thus far.

Recently, the application of CTF deconvolution to tomograms has been shown to be advantageous, providing a striking increase of the contrast and enabling visual inspection of the tomograms and identification of individual biological components (Tegunov and Cramer, 2019). The technique relies on a Wiener-like filter:

$$\mathfrak{I}\{I_{\text{deconv}}\}(f) = \mathfrak{I}\{I\}(f) \cdot \frac{\text{CTF}(f)}{\text{CTF}^2(f) + \text{SNR}^{-1}(f)}$$

where I and I_{deconv} are the tomogram and its CTF-deconvolved version, respectively, \mathfrak{I} denotes Fourier transformation and f is the spatial frequency. The CTF term in the equation only depends on a variable parameter, the defocus used during the tilt-series acquisition (normally either the defocus of the untilted image or the average defocus of the images of the tilt-series), with other parameters being given by the specific microscope (voltage, spherical aberration). SNR represents the spectral signal-to-noise ratio, estimated as a combination of an exponential decay curve and an optional raised-cosine high-pass filter H :

$$\text{SNR}(f) = e^{-f \cdot 100F} \cdot 10^{3S} \cdot H(f)$$

with S and F being custom parameters controlling the strength and decay of the filter, whose default unit values provide good results in general.

The Wiener-like filter mainly boosts the Fourier components up to the first CTF peak while preventing over-amplification of those components dominated by noise (at the mid- and high-frequencies and, also, at the lowest end, i.e. close to the Fourier origin). As a consequence, the filter restores the low-resolution contrast and, at the same time, also acts as a low-pass denoising filter, with dramatic benefits for the visual interpretation of tomograms. Figure 7 illustrates the deconvolution filter and its performance on several experimental tomograms of ribosomes (Bharat and Scheres, 2016) (EMPIAR-10045), HIV-1 virus like particles (Schur et al., 2016) (EMPIAR-10164) and *Thermoanaerobacter kivui* (Dietrich et al., 2022) (EMPIAR-11058).

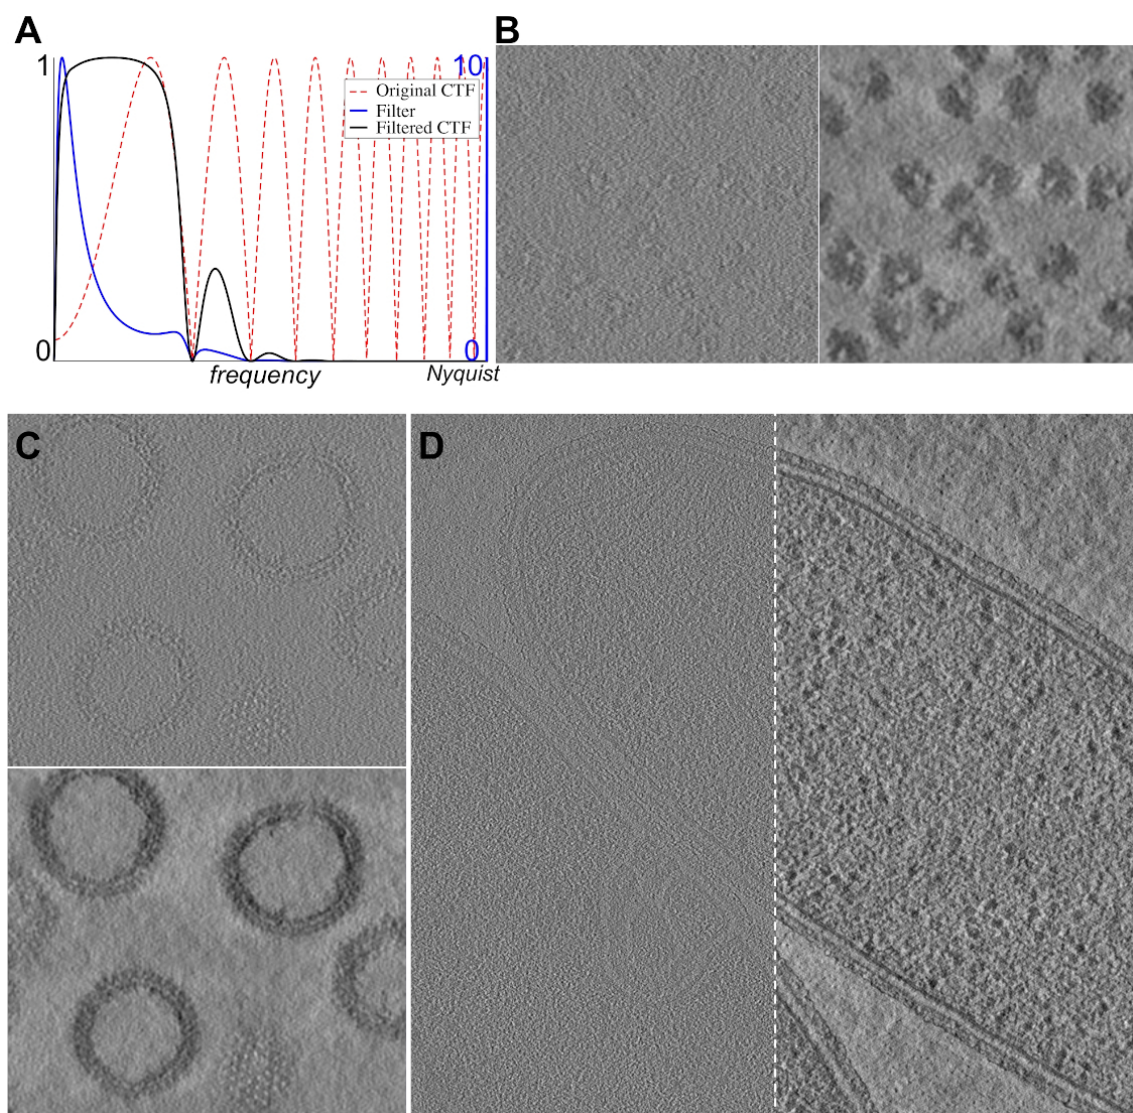


Figure 7. Contrast improvement with CTF deconvolution.

(A) CTF deconvolution filter. The CTF (in absolute value), the Wiener-like filter and the resulting restoration are shown in dashed red, solid blue and solid black lines, respectively. The CTF corresponds to a tomogram acquired at an average defocus of 3.7 microns on a 300 kV microscope and the Nyquist frequency here is 8.68 Angstroms. The filter restores the amplitude of the low-resolution components up to the first CTF peak, followed by a rapid decay acting as a low-pass filter.

(B) Result of the CTF deconvolution applied to a ribosome tomogram (EMPIAR-10045) using the filter shown in (A). Original (left) and deconvolved (right) tomogram.

(C) Result of the CTF deconvolution on a tomogram of HIV-1 virus like particles (EMPIAR-10164) acquired on a 300 kV microscope with an average defocus of 3.9 microns. Original (top) and deconvolved (bottom) tomogram.

(D) Result of the CTF deconvolution on a tomogram of *Thermoanaerobacter kivui* (EMPIAR-11058) acquired on a 300 kV microscope with an average defocus of 4.4 microns. Original (left) and deconvolved (right) tomogram.

7. Software

Table 1 presents a list of software packages or programs available in cryoET to reduce noise in tomograms at the different steps of the computational workflow using the methods described along this chapter. A companion website has also been prepared with links to access the software and associated publications. The website is:

<https://sites.google.com/view/Denoising-cryoET>

Regarding the control of noise during tomographic reconstruction, two major software tools in the cryoET field, namely IMOD (Kremer et al., 1996) and Tomo3D (Agulleiro and Fernandez, 2015, 2011), provide essentially the same mechanisms. Both include the iterative reconstruction method SIRT as well as WBP equipped with the weighting filters described in Section 3: the Gaussian and/or Hamming filter and the 'SIRT-like' filter. The latter is a particularly interesting option to obtain reconstructions with good contrast, similar to SIRT, at a speed of WBP.

The urgent needs to cope with the extremely low SNR in cryoET stimulated the development of advanced denoising methods in the 2000's (labeled as classical methods in Table 1). As a result, there are numerous tools available in standard software packages. Two important ones are IMOD (Kremer et al., 1996) and Bsoft (Heymann, 2021), both of them including Gaussian filtering, median filtering and anisotropic non-linear diffusion (AND). Bsoft also includes an implementation of Bilateral filtering.

The computational burden in terms of processing time and memory consumption of these methods or their difficult parameter tuning led us to develop specialized, easy-to-use, efficient and parallel implementations. As a consequence, there are fast, memory-efficient and user-friendly implementations of the Gaussian filtering (Martinez-Sanchez et al., 2014, 2011), Beltrami flow (Fernandez, 2009; Fernández and Martínez, 2010) and AND (Moreno et al., 2018). These optimized programs have been integrated into a single package named TomoDenoise (<https://tiny.cc/tomodenoise>).

The emergence of deep learning and the excellent performance of the Noise2Noise denoising approach have driven adaptations to cryoET, with several software tools now available. Cryo-CARE (Buchholz et al., 2019a), Topaz-Denoise (Bepler et al., 2020) and Warp (Tegunov and Cramer, 2019) take advantage of these denoising neural networks to work directly on tomograms in 3D.

CTF deconvolution is rapidly getting interest in the field owing to its striking performance to improve the contrast and enable visual analysis of the tomograms. Initially provided by the main author as a MATLAB script and later in Warp, implementations of the method are also being included within several software tools, such as IsoNet (Liu et al., 2022), pyCoan (Gaietta et al., 2021), TomoDenoise, IMOD and Tomo3D.

Apart from those tools, there is a number of free general bioimaging software suites with denoising methods that can be used for cryoET as well, for instance ImageJ/Fiji ([Schindelin et al., 2012](#)).

Table 1. Software to reduce noise in cryoET

Method	Software
<i>Classical denoising methods</i>	
Gaussian filtering	IMOD, Bsoft, TomoDenoise
Median filtering	IMOD, Bsoft
Beltrami flow	TomoDenoise
Bilateral filtering	Bsoft
Anisotropic diffusion (AND)	IMOD, Bsoft, TomoDenoise
<i>Deep-learning based denoising methods</i>	
Noise2Nose	Warp, Cryo-CARE, Topaz-Denoise
<i>CTF Deconvolution</i>	
Wiener-like filter	Warp, Isonet, pyCoan, TomoDenoise, IMOD, Tomo3D https://github.com/dtegunov/tom_deconv
<i>Handling noise during tomographic reconstruction</i>	
Iterative reconstruction (SIRT)	IMOD, Tomo3D
SIRT-like filter in WBP	IMOD, Tomo3D
Hamming/Gaussian filter in WBP	IMOD, Tomo3D

8. Concluding remarks

Denoising plays an important role within the computational workflow in cryoET, where the SNR and contrast are very low. Denoising intends to facilitate the interpretation of tomograms by reducing noise and increasing the contrast, making them better suited for the downstream processing (e.g. segmentation and delineation of structures, feature detection, 3D visualization). This is essential not only for cryoET studies focused on cellular landscapes or pleomorphic structures, but also for those aimed at *in situ* high-resolution structure determination where the features to average first have to be identified.

Several stages of the computational workflow may have an influence on the SNR and contrast of the tomograms. During tomographic reconstruction, the use of weighting filters in WBP is an important decision to control the amount of noise transferred to the tomogram. In the last few years, the use of WBP with the SIRT-like filter has got increasing interest thanks to the quality of the resulting tomograms obtained in short time.

There exist several denoising options as a post-processing stage. Out of all the methods developed during the 2000's, Anisotropic Non-linear Diffusion (AND) is by far the predominant one and its availability in several well-known software packages makes it easily applicable. Lately, deep-learning methods have arisen as powerful alternatives thanks to their ability to reduce noise with significant flattening of the background. The Noise2Noise approach was easily adapted to cryoET, with several packages now available. Thus, application of these methods to tomograms in experimental cryoET studies is progressively increasing.

CTF deconvolution has recently emerged as a new approach to improve the quality of tomograms. Equipped with an integrated low-pass filter, it manages to significantly increase the contrast and reduce the noise, resulting in an impressive enhancement of the tomograms that are ready for analysis.

Depending on the needs for contrast improvement and noise reduction, these methods can be combined. Tomographic reconstruction with the SIRT-like filter along with CTF deconvolution should already make tomograms amenable to visual interpretation. Should further polishing happen to be required, subsequent application of denoising techniques (e.g. AND or a deep-learning-based one) is possible.

Acknowledgements

Manuscript finished in December 2022. Our work and the writing of this chapter are possible thanks to support from the Spanish Agencia Estatal de Investigacion (grants TED2021-132020B-I00 and PID2022-139071NB-I00 funded by MICIU/AEI/10.13039/501100011033, NextGenerationEU/PRTR and ERDF/EU) and a 2022 HDSA HD Human Biology Project Fellowship.

References

- Agulleiro, J.-I., Fernandez, J.-J., 2015. Tomo3D 2.0--exploitation of advanced vector extensions (AVX) for 3D reconstruction. *J. Struct. Biol.* 189, 147–152. <https://doi.org/10.1016/j.jsb.2014.11.009>
- Agulleiro, J.I., Fernandez, J.J., 2011. Fast tomographic reconstruction on multicore computers. *Bioinformatics* 27, 582–583. <https://doi.org/10.1093/bioinformatics/btq692>
- Bepler, T., Kelley, K., Noble, A.J., Berger, B., 2020. Topaz-Denoise: general deep denoising models for cryoEM and cryoET. *Nat. Commun.* 11, 5208. <https://doi.org/10.1038/s41467-020-18952-1>
- Bharat, T.A.M., Scheres, S.H.W., 2016. Resolving macromolecular structures from electron cryo-tomography data using subtomogram averaging in RELION. *Nat. Protoc.* 11, 2054–2065. <https://doi.org/10.1038/nprot.2016.124>
- Briggs, J.A.G., Grünewald, K., Glass, B., Förster, F., Kräusslich, H.-G., Fuller, S.D., 2006. The Mechanism of HIV-1 Core Assembly: Insights from Three-Dimensional Reconstructions of Authentic Virions. *Structure* 14, 15–20. <https://doi.org/10.1016/j.str.2005.09.010>
- Buades, A., Coll, B., Morel, J.-M., 2005. A Non-Local Algorithm for Image Denoising, in: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), IEEE, San Diego, CA, USA, pp. 60–65. <https://doi.org/10.1109/CVPR.2005.38>
- Buchholz, T.-O., Jordan, M., Pigino, G., Jug, F., 2019a. Cryo-CARE: Content-Aware Image Restoration for Cryo-Transmission Electron Microscopy Data, in: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI), IEEE, Venice, Italy, pp. 502–506. <https://doi.org/10.1109/ISBI.2019.8759519>
- Buchholz, T.-O., Krull, A., Shahidi, R., Pigino, G., Jékely, G., Jug, F., 2019b. Content-aware image restoration for electron microscopy, in: *Methods in Cell Biology* 152. Elsevier, pp. 277–289.
- Danita, C., Chiu, W., Galaz-Montoya, J.G., 2022. Efficient manual annotation of cryogenic electron tomograms using IMOD. *STAR Protoc.* 3, 101658. <https://doi.org/10.1016/j.xpro.2022.101658>
- Darbon, J., Cunha, A., Chan, T.F., Osher, S., Jensen, G.J., 2008. Fast nonlocal filtering applied to electron cryomicroscopy, in: 2008 5th IEEE International Symposium on Biomedical Imaging (ISBI 2008), IEEE, Paris, France, pp. 1331–1334. <https://doi.org/10.1109/ISBI.2008.4541250>
- Dietrich, H.M., Righetto, R.D., Kumar, A., Wietrzynski, W., Trischler, R., Schuller, S.K., Wagner, J., Schwarz, F.M., Engel, B.D., Müller, V., Schuller, J.M., 2022. Membrane-anchored HDCR nanowires drive hydrogen-powered CO₂ fixation. *Nature* 607, 823–830. <https://doi.org/10.1038/s41586-022-04971-z>
- Fernandez, J.J., Li, S., 2005. Anisotropic Nonlinear Filtering of Cellular Structures in Cryoelectron Tomography. *Comput. Sci. Eng.* 7, 54–61. <https://doi.org/10.1109/MCSE.2005.89>
- Fernandez, J.J., 2012. Computational methods for electron tomography. *Micron* 43, 1010–1030. <https://doi.org/10.1016/j.micron.2012.05.003>

- Fernandez, J.J., 2009. TOMOBFLOW: feature-preserving noise filtering for electron tomography. *BMC Bioinformatics* 10, 178. <https://doi.org/10.1186/1471-2105-10-178>
- Fernández, J.J., Li, S., 2003. An improved algorithm for anisotropic nonlinear diffusion for denoising cryo-tomograms. *J. Struct. Biol.* 144, 152–161. <https://doi.org/10.1016/j.jsb.2003.09.010>
- Fernández, J.J., Li, S., Lucic, V., 2007. Three-Dimensional Anisotropic Noise Reduction with Automated Parameter Tuning: Application to Electron Cryotomography, in: *Current Topics in Artificial Intelligence*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 60–69.
- Fernández, J.J., Martínez, J.A., 2010. Three-dimensional feature-preserving noise reduction for real-time electron tomography. *Digit. Signal Process.* 20, 1162–1172. <https://doi.org/10.1016/j.dsp.2009.12.010>
- Frangakis, A.S., Hegerl, R., 2001. Noise Reduction in Electron Tomographic Reconstructions Using Nonlinear Anisotropic Diffusion. *J. Struct. Biol.* 135, 239–250. <https://doi.org/10.1006/jsbi.2001.4406>
- 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, 213–222. <https://doi.org/10.1109/10.909642>
- Frank, J. (Ed.), 2006. *Electron Tomography*. Springer New York, New York, NY.
- Gaietta, G., Swift, M.F., Volkmann, N., Hanein, D., 2021. Rapid tool for cell nanoarchitecture integrity assessment. *J. Struct. Biol.* 213, 107801. <https://doi.org/10.1016/j.jsb.2021.107801>
- Gilbert, P., 1972. Iterative methods for the three-dimensional reconstruction of an object from projections. *J. Theor. Biol.* 36, 105–117. [https://doi.org/10.1016/0022-5193\(72\)90180-4](https://doi.org/10.1016/0022-5193(72)90180-4)
- Gonzalez, R.C., Woods, R.E., 2018. *Digital image processing*. Pearson, New York, NY.
- 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. <https://doi.org/10.7554/eLife.06980>
- Heymann, J.B., 2021. High resolution electron tomography and segmentation-by-modeling interpretation in Bsoft. *Protein Sci.* 30, 44–59. <https://doi.org/10.1002/pro.3938>
- Huang, X., Li, S., Gao, S., 2018. Exploring an optimal wavelet-based filter for cryo-ET imaging. *Sci. Rep.* 8, 2582. <https://doi.org/10.1038/s41598-018-20945-6>
- 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. <https://doi.org/10.1016/j.jsb.2003.09.028>
- Kimmel, R., Sochen, N., Malladi, R., 1997. From high energy physics to low level vision, in: Haar Romeny, B., Florack, L., Koenderink, J., Viergever, M. (Eds.), *Scale-Space Theory in Computer Vision*. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 236–247.
- Kremer, J.R., Mastronarde, D.N., McIntosh, J.R., 1996. Computer Visualization of Three-Dimensional Image Data Using IMOD. *J. Struct. Biol.* 116, 71–76. <https://doi.org/10.1006/jsbi.1996.0013>
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521, 436–444. <https://doi.org/10.1038/nature14539>

- Lehtinen, J., Munkberg, J., Hasselgren, J., Laine, S., Karras, T., Aittala, M., Aila, T., 2018. Noise2Noise: Learning Image Restoration without Clean Data, in: Proceedings of the 35th International Conference on Machine Learning Research. PMLR, pp. 2965–2974.
- Liu, Y.-T., Zhang, H., Wang, H., Tao, C.-L., Bi, G.-Q., Zhou, Z.H., 2022. Isotropic reconstruction for electron tomography with deep learning. *Nat. Commun.* 13, 6482. <https://doi.org/10.1038/s41467-022-33957-8>
- 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, 132–142. <https://doi.org/10.1016/j.jsb.2012.05.019>
- Martinez-Sanchez, A., Garcia, I., Asano, S., Lucic, V., Fernandez, J.J., 2014. Robust membrane detection based on tensor voting for electron tomography. *J. Struct. Biol.* 186, 49–61. <https://doi.org/10.1016/j.jsb.2014.02.015>
- Martinez-Sanchez, A., Garcia, I., Fernandez, J.J., 2011. A differential structure approach to membrane segmentation in electron tomography. *J. Struct. Biol.* 175, 372–383. <https://doi.org/10.1016/j.jsb.2011.05.010>
- Moreno, J.J., Martínez-Sánchez, A., Martínez, J.A., Garzón, E.M., Fernández, J.J., 2018. TomoEED: fast edge-enhancing denoising of tomographic volumes. *Bioinformatics* 34, 3776–3778. <https://doi.org/10.1093/bioinformatics/bty435>
- 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, 7–17. <https://doi.org/10.1016/j.jsb.2008.04.006>
- Perona, P., Malik, J., 1990. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 12, 629–639. <https://doi.org/10.1109/34.56205>
- Schindelin, J., Arganda-Carreras, I., Frise, E., Kaynig, V., Longair, M., Pietzsch, T., Preibisch, S., Rueden, C., Saalfeld, S., Schmid, B., Tinevez, J.-Y., White, D.J., Hartenstein, V., Eliceiri, K., Tomancak, P., Cardona, A., 2012. Fiji: an open-source platform for biological-image analysis. *Nat. Methods* 9, 676–682. <https://doi.org/10.1038/nmeth.2019>
- Schur, F.K.M., Obr, M., Hagen, W.J.H., Wan, W., Jakobi, A.J., Kirkpatrick, J.M., Sachse, C., Kräusslich, H.-G., Briggs, J.A.G., 2016. An atomic model of HIV-1 capsid-SP1 reveals structures regulating assembly and maturation. *Science* 353, 506–508. <https://doi.org/10.1126/science.aaf9620>
- Stoschek, A., Hegerl, R., 1997. Denoising of Electron Tomographic Reconstructions Using Multiscale Transformations. *J. Struct. Biol.* 120, 257–265. <https://doi.org/10.1006/jsbi.1997.3925>
- Tegunov, D., Cramer, P., 2019. Real-time cryo-electron microscopy data preprocessing with Warp. *Nat. Methods* 16, 1146–1152. <https://doi.org/10.1038/s41592-019-0580-y>
- Tomasi, C., Manduchi, R., 1998. Bilateral filtering for gray and color images, in: IEEE 6th International Conference on Computer Vision, pp. 839–846. <https://doi.org/10.1109/ICCV.1998.710815>
- van der Heide, P., Xu, X.-P., Marsh, B.J., Hanein, D., Volkmann, N., 2007. Efficient automatic noise reduction of electron tomographic reconstructions based

- on iterative median filtering. *J. Struct. Biol.* 158, 196–204.
<https://doi.org/10.1016/j.jsb.2006.10.030>
- Wan, W., Briggs, J.A.G., 2016. Cryo-Electron Tomography and Subtomogram Averaging, in: *Methods in Enzymology*. Elsevier, pp. 329–367.
- 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. <https://doi.org/10.1016/j.jsb.2010.06.021>
- Weickert, J., 1998. Anisotropic diffusion in image processing, European Consortium for Mathematics in Industry. Teubner, Stuttgart.
- Zeng, G.L., 2012. A filtered backprojection algorithm with characteristics of the iterative landweber algorithm. *Med. Phys.* 39, 603–607.
<https://doi.org/10.1118/1.3673956>