Journal of Microscopy, Vol. 219, Pt 2 August 2005, pp. 43–49
Received 28 January 2005; accepted 11 May 2005

A novel 3D wavelet-based filter for visualizing features in noisy biological data

W. C. MOSS, S. HAASE,* J. M. LYLE,* D. A. AGARD*† & J. W. SEDAT*

Lawrence Livermore National Laboratory, Livermore, CA 94550, U.S.A.

**Keck Advanced Microscopy Laboratory and the Department of, Biochemistry & Biophysics, University of California San Francisco, San Francisco, CA 94143, U.S.A.*

†*Howard Hughes Medical Institute, University of California San Francisco, San Francisco, CA 94143, U.S.A.*

Key words. Cryo-EM, EM Tomography, fluorescence microscopy, wavelets, 3D image processing.

Summary

We have developed a three-dimensional (3D) wavelet-based filter for visualizing structural features in volumetric data. The only variable parameter is a characteristic linear size of the feature of interest. The filtered output contains only those regions that are correlated with the characteristic size, thus de-noising the image. We demonstrate the use of the filter by applying it to 3D data from a variety of electron microscopy samples, including low-contrast vitreous ice cryogenic preparations, as well as 3D optical microscopy specimens.

Introduction

The conventional wisdom in modern structural biology is that as the resolution increases, structures become self-evident. Recent advances in microscopy have increased the theoretical light/electron microscopy (EM) resolution to 50 nm/1.5 Å (Reimer, 1997; Westphal *et al.*, 2003), so that images of unprecedented clarity should be obtainable. However, at these high resolutions, practical limitations, such as avoiding radiation damage, place severe limitations on the data collection process. The consequence of this is that obtainable signal-to-noise ratios are often significantly less than 1.0, owing to a combination of shot noise and detector noise. This is an especially significant problem in cryo-EM tomography of unstained frozen specimens, where typically 100–200 tilted views need to be collected from the same sample and total doses need to kept below about 30 e^{-} $Å^{-2}$. The resulting three-dimensional (3D) reconstructions are quite noisy, which makes it a challenge to define accurately the shape and location of desired objects within the tomogram.

Correspondence to: Dr William Moss. Fax: +1 925 423 5804; e-mail: wmoss@llnl.gov

Non-specific staining is another source of 'noise' that further complicates interpretation of both EM and light microscopy data. For example, uranyl acetate is a popular EM stain that forms complexes primarily with phosphates on DNA, RNA and phosphoproteins. Images that are acquired using these stains are biased by the properties of the stain and thus do not necessarily represent a true picture of the underlying structure. Interpretation difficulties are compounded if non-specifically stained structures are packed densely, which is typical for many biological samples.

Light microscopy has related problems. Although the stains (especially fluorescent proteins) are very specific and provide strong contrast, experimental protocols often demand the collection of thousands of images from a single (often live) sample. In these cases, avoiding phototoxicity and bleaching of the fluorophore become paramount. As with the EM data, the result is a drastic reduction of the signal-to-noise ratio of each recorded image. Additional sources of noise, such as autofluorescence, background pools of unassembled fluorescent proteins and instrument noise, can contribute to the challenge of identifying and quantifying 3D cellular structures.

Many researchers believe that a simple solution to these problems is to construct software that can filter the image and 'bring out' the essential structure. Indeed, much effort by many groups, including ours, has been expended to develop filter methods to abstract structures and reduce noise (Böhm *et al.*, 2000; Nicholson & Glaeser, 2001; Frangakis *et al.*, 2002; Rath & Frank, 2004), with the best current method probably being anisotropic diffusion (Frangakis & Hegerl, 2001). Another approach is to locate known objects within the 3D reconstructions. Typically, such matched-filter correlation approaches can find objects under conditions of very high noise; however, there is much utility in developing hybrid methods that have the noise performance of matched filters yet do not require a priori knowledge of the search object.

Our approach is to develop a filter that preferentially highlights objects of defined size-classes within 3D volumes. This suggests using a wavelet transform. A wavelet transform is a convolution of a kernel (shape function) and the data. It differs from a Fourier transform in that the wavelet kernel is non-zero only over a finite spatial extent (chosen typically to equal the size of the feature of interest), whereas the Fourier kernel has infinite extent. It is this property of the wavelet kernel that makes it better than Fourier methods at defining frequency content as a function of spatial location. The wavelet transform shows how strongly the data are correlated to the kernel at each location in the data.

For ease of use and interpretation, we want the wavelet transform to have the following characteristics: (i) we want 2D or 3D data that are filtered with a wavelet of size n voxels to highlight preferentially those regions that have physical dimensions with a characteristic size of n voxels; and (ii) the wavelet transform should be invariant to rigid body rotations, i.e. rotating the wavelet transform of an object should be equivalent to the wavelet transform of the rotated object.

In general, wavelet methods can be quite complex, particularly for 3D data. Consequently, it would be an improvement to have a fast, efficient filter for 2D or 3D data whose only parameter is the characteristic spatial size of the structure of interest, and whose output is a spatial image of correlation strength. The 3D wavelet-based filter described here is a realization of these ideas. This filter, which highlights objects of a defined size-class, differs from a previous implementation of 3D wavelets (Stoschek & Hegerl, 1997) that used data-dependent thresholding at every wavelet size to denoise tomographic data globally.

We demonstrate the utility of our filter on synthetic data by showing how it can extract a pair of mathematically constructed helices from a noisy background, even at low signal-to-noise ratios. Four additional examples show how the filter works with typical noisy biological data: (i) EM data of positively stained microtubules, (ii) EM data of a negatively stained γ-tubulin ring complex (γTuRC), (iii) EM data of unstained microtubules preserved in vitreous ice and (iv) light microscopic images of *Caenorhabditis elegans* meiotic cells. Although this filter was developed primarily for biological applications, it is generally applicable to any 3D (or 2D) data.

Rationale for the filter and its mathematical structure

A general 1D transform can be written as

$$W(x,a) = \int \Psi\left(\frac{x-x'}{a}\right) f(x')\mathrm{d}x', \qquad (1)$$

where the 'signal' f is convolved with the function Ψ. A wavelet transform differs from a Fourier transform in the choice of Ψ. The Ψ for a Fourier transform [$\Psi = \exp(ix'/a)$] has non-zero values that extend over the entire x' domain, and thus W is only a function of a, whereas the Ψ for the wavelet transform is centred at x and is non-zero over a characteristic width a (Farge, 1992). The wavelet transform is a correlation function at each position, x. The correlation varies with the wavelet width a, and is maximal for some value of a, which explains how wavelets can determine where particular frequencies occur.

The only physical requirement for Ψ is that a uniform signal of infinite extent produces no correlations, i.e. $\int \Psi(x - x'/a)\mathrm{d}x' = 0$ (Farge, 1992). Beyond this restriction, the choice of Ψ is arbitrary, and is tailored usually to the particular problem of interest. A major exception is image compression and reconstruction, for which an orthonormal set of wavelets is desirable (Farge, 1992). Given that our focus is in identifying structures of a particular size, not in image compression, we have chosen a particularly simple 1D representation of Ψ that is easily extended to 3D. We write

$$\Psi(\zeta) = \frac{1}{a} \quad \text{for} \quad |\zeta| \le \frac{1}{2}.$$
$$\Psi(\zeta) = \frac{-1}{2a} \quad \text{for} \quad \frac{1}{2} < |\zeta| \le \frac{3}{2}, \qquad (2)$$

and

$$\Psi(\zeta) = 0 \quad \text{for} \quad |\zeta| > \frac{3}{2},$$

where $\zeta = (x - x')/a$. The normalization is chosen such that a signal of unit intensity and spatial extent a has a correlation of 1, independent of a, i.e. $\int_{-a/2}^{a/2} \Psi(x'/a)\mathrm{d}x' = 1$. This particular normalization ensures that large and small structural features are treated equivalently.

Figure 1 shows schematically how Eqs (1) and (2) are used in 1D. This figure illustrates the general principles that are extendable immediately to 3D. Figure 1(a) (black) shows a hypothetical signal along a spatial axis, where the signal intensity is either zero or non-zero along the axis. The wavelet is shown in red. As described in Eq. (2), the wavelet has a uniform positive value $(1/a)$ in a central region of extent a, and a negative value $(-1/2a)$ in two surrounding regions, each of extent a. The integral of the wavelet over the spatial domain is zero. The wavelet transform is obtained from Eq. (1) by (i) integrating the product of the wavelet and the signal over the entire spatial domain, (ii) assigning the resulting value to the spatial location of the centre of the wavelet

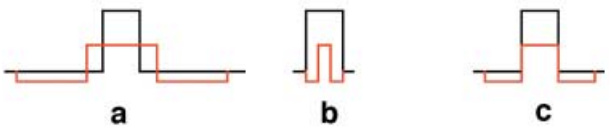

Fig. 1. Schematic 1D representation of a wavelet filter. The 1D wavelet transform is a measure of the strength of the correlation between the wavelet (red) and the data (black), at each data point. (a)–(c) show identical 1D data (black line) that underfill, overfill and match the superimposed wavelet (red line). The correlation is maximal in (c). (a) and (b) produce weaker correlations, due to the penalizing effects of the negative portions of the wavelet (b) and the wavelet normalization (a).

and (iii) translating the wavelet over the entire spatial domain and repeating the integration at each location.

It is now easy to see how this method can be used to identify particular-sized structures. Figure 1(b,c) show the same signal as in Fig. 1(a), but with different sized wavelets. Objects that just fill the positive part of the wavelet produce maximum correlations (Fig. 1c). Objects that underfill the positive part of the wavelet produce smaller correlations (Fig. 1a), because of the penalizing effects of the normalization. Objects that are larger than the wavelet also produce smaller correlations (Fig. 1b), because of the penalizing effects of the negative and zero-valued parts of the wavelet. We vary the size of the wavelet, apply it repeatedly and view the results. When the wavelet size matches the object size, large correlations are obtained at the spatial location of the object. Logical consistency requires the data to be positive definite, which can always be accomplished by adding a constant intensity to all of the data values. The two features of this procedure that define it as a wavelet method are the use of a filter that has a finite domain and varying the size of this finite domain over all length scales.

The extension of Eq. (1) to 3D is

$$W^*(x,y,z,a,b,c)=\iiint \Psi\left(\frac{x-x'}{a},\frac{y-y'}{b},\frac{z-z'}{c}\right) f(x',y',z')\,\mathrm{d}x'\mathrm{d}y'\mathrm{d}z',$$

but using such a true 3D wavelet is computationally expensive, and therefore it is desirable to seek a solution that is separable along the spatial directions, such as

$$W^*(x,y,z,a,b,c) = \iiint \Psi\left(\frac{z-z'}{c}\right)\Psi\left(\frac{y-y'}{b}\right)\Psi\left(\frac{x-x'}{a}\right) f(x',y',z')\,\mathrm{d}x'\mathrm{d}y'\mathrm{d}z'. \quad (3)$$

However, a serious problem with Eq. (3) is that two negative correlations (e.g. negative correlations in x and y) can combine to produce a positive correlation at a particular location. Consequently, Eq. (3) must be modified so that false positive correlations do not occur. In addition, the desired invariance of the wavelet transform to rigid body rotations is broken by the typical Cartesian discretization of data that is required by Eq. (3). The wavelet transform shown in Eq. (4) is an approximate and practical solution, as shown by the examples in the Results section. We write

$$\bar{W}(x,y,z,a,b,c,\phi,i) = \mathrm{H}\left[\int \Psi\left(\frac{z-z'}{c}\right)\mathrm{H}\left[\int \Psi\left(\frac{y-y'}{b}\right)\mathrm{H}\left[\int \Psi\left(\frac{x-x'}{a}\right) R_i(\phi) f(x',y',z')\,\mathrm{d}x'\right]\mathrm{d}y'\right]\mathrm{d}z'\right], \quad (4)$$

$$W_{sym}(x,y,z,a,\phi,i) = \frac{1}{6}R_i^{-1}(\phi)\sum_{\substack{x,y,z\\ perm}} \bar{W}(x,y,z,a,a,a,\phi,i), \quad \text{and}$$

$$W(x,y,z,a) = \frac{1}{4}(W_{sym}(x,y,z,a,0,x) + W_{sym}(x,y,z,a,\pi/4,x) + W_{sym}(x,y,z,a,\pi/4,y) + W_{sym}(x,y,z,a,\pi/4,z)). \quad (5)$$

H[q] in Eq. (4) is defined as follows: H[q] = q for $q > 0$ and H[q] = 0 for $q \le 0$. This produces only positive correlations, but results in $\bar{W}$ being dependent on the order of integration over x', y' and z'. W_{sym} is the average of $\bar{W}$ over all permutations of x, y and z, which symmetrizes the transform with respect to x, y and z. Rotational invariance is approximated by considering only cubical wavelets ($a = b = c$) and averaging over angles, where $R_i(\phi)$ is a rotation of ϕ radians around the i axis (0 and $\pi/4$ radians in Eq. 5). $W(x,y,z,a)$ (Eq. 5) is the resulting (approximately) symmetric and rotationally invariant 3D wavelet transform.

Results

Figures 2–7 illustrate the utility of the proposed 3D wavelet filter just discussed.

We begin with synthetic data: two 3D helices embedded in noise. The axes of the helices are displaced approximately by a helix diameter. The left column in Fig. 2 shows an analytically constructed pair of helices (strand diameter ~6 voxels) that are embedded in varying amounts of Gaussian noise (see Appendix A for details about the construction of the helices). The signal-to-noise ratios (SN) equal 8, 2, 1 and 0.5. The images in the right column are obtained by using the 3D wavelet filter, with a 6-voxel

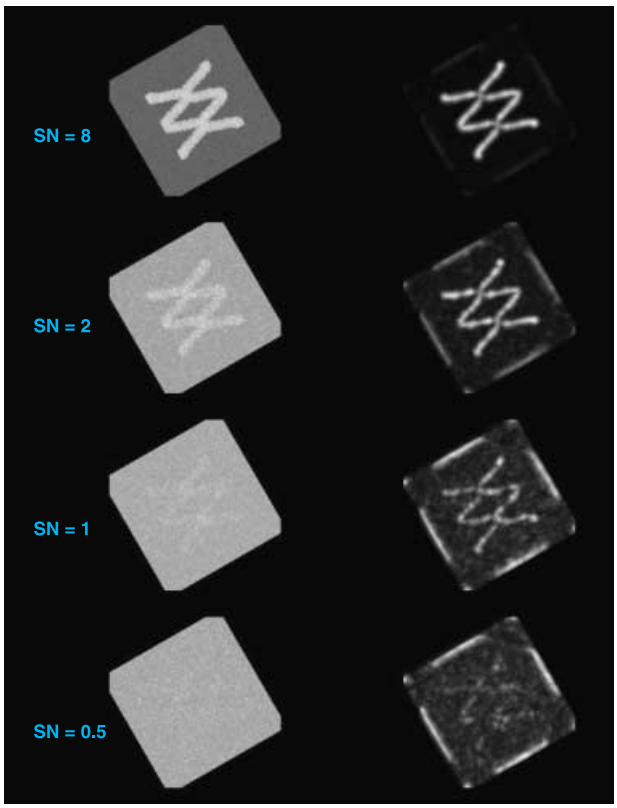

Fig. 2. 3D wavelet filtering of two helices embedded in noise. Two analytically constructed helices (strand diameter ~6 voxels) are embedded in varying amounts of Gaussian noise. Data with signal-to-noise ratios equal to 8, 2, 1 and 0.5 are shown in the left column. The 3D wavelet-filtered data (Eq. 5), with a 6-voxel spatial scale ($a = 6$) are shown in the right column. The helices can be seen even at the lowest signal-to-noise ratio.

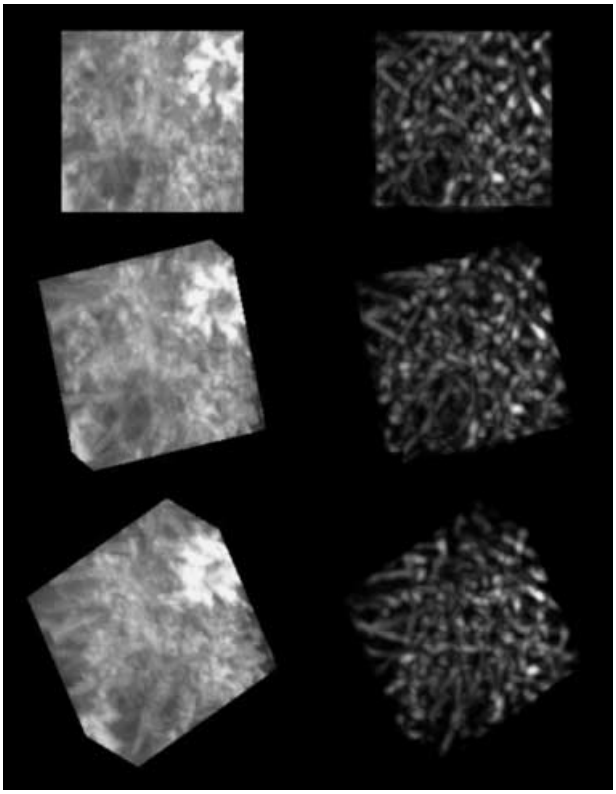

Fig. 3. 3D wavelet filtering of microtubules and centriole. The upper left volume is filtered using Eq. (5), with a 10-voxel spatial scale ($a = 10$; ~28 nm) that corresponds approximately to the diameter of a microtubule. The filtered volume is shown in the upper right. Two additional 3D-rotated volume views are also shown. The filter extracts the microtubules from the background noise, and ignores the centriole, whose substructural units have a characteristic size approximately three times the diameter of a microtubule.

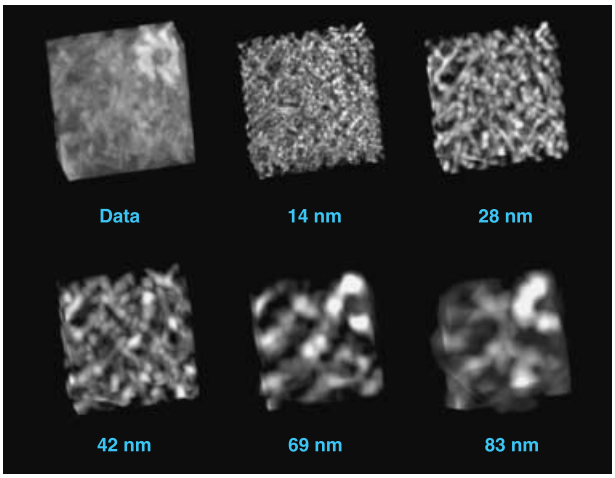

Fig. 4. 3D wavelet filtering of microtubules and centriole as a function of wavelet spatial scale. The data from Fig. 3 are filtered using spatial scales of 5, 10, 15, 25 and 30 voxels, which correspond to 14, 28, 42, 69 and 83 nm. The microtubules 'come into focus' for the 28-nm wavelet, then 'go out of focus' as the wavelet size increases. The centriole 'comes into focus' at the largest wavelet values in the figure.

spatial scale ($a = 6$ in Eq. 5), on the corresponding data volumes in the left column. Each image in the figure is a maximum intensity volume projection that was rendered using the freely available software package PRIISM (Chen *et al.*, 1992). The two helices are evident, even at the lowest SN. (Supplementary Movie 1 online shows the volumes rotating around an axis parallel to the helical axes, which emphasizes that even in the lowest SN image, elements of the helices are visible.) This example demonstrates that the filter can extract structure from a noisy environment, for constructed data. We consider next real multiscale noisy biological data.

Figure 3 shows positively stained microtubules in a typical biological environment. The upper left image (maximum intensity volume projection) in the figure shows data obtained using EM tomography of an intact centrosome embedded in plastic and stained (Moritz *et al.*, 1995). Microtubules (~25 nm characteristic diameter) scattered throughout a noisy background and a centriole in the upper right corner are visible. The two additional views (left column) were obtained by rotating the volume using PRIISM (Chen *et al.*, 1992). The upper right image shows the filtered data, using Eq. (5), with a 10-voxel spatial scale ($a = 10$, ~28 nm). Additional rotated views are shown below. The filter extracts microtubules from the noisy background, while excluding the larger centriole, whose substructural units have a characteristic size approximately three times that of a microtubule. 3D stereo overlays (not shown here) confirm that the extracted microtubules exist in the original data. Figure 4 shows the unfiltered data (upper left) and filtered data, as a function of the spatial scale that is used: wavelet sizes equal to 5, 10, 15, 25 and 30 voxels, which correspond to 14, 28, 42, 69 and 83 nm. As the spatial scale increases, the correlations change. At 14 nm [and lower (data not shown)], there are no clearly discernible features. At 28 nm, microtubules are visible. As the spatial scale continues to increase, the microtubules disappear and the centriole becomes the dominant feature, as expected.

In order to examine how wavelet filtering would enhance the contrast of a single weakly discernible structure embedded in noise, we next looked at data from a negatively stained sample of isolated *Drosophila melanogaster* γTuRC. The upper left image (maximum intensity volume projection) in Fig. 5 shows data (γTuRC preparation and imaging described by Moritz *et al.*, 2000) that were obtained using low-dose EM tomography (3.75 Å voxel size). A γTuRC is arranged much like a lock washer, with a diameter of ~25 nm and characteristic annular thickness of 8 nm, which is not apparent in the upper left image. The data were filtered with wavelet sizes equal to 5, 10 and 20 voxels, which correspond to 1.9, 3.8 and 7.5 nm. The 7.5-nm wavelet (lower right) clearly shows the lock washer shape.

By far the most challenging EM tomographic data are that collected under the lowest dose and with the least intrinsic

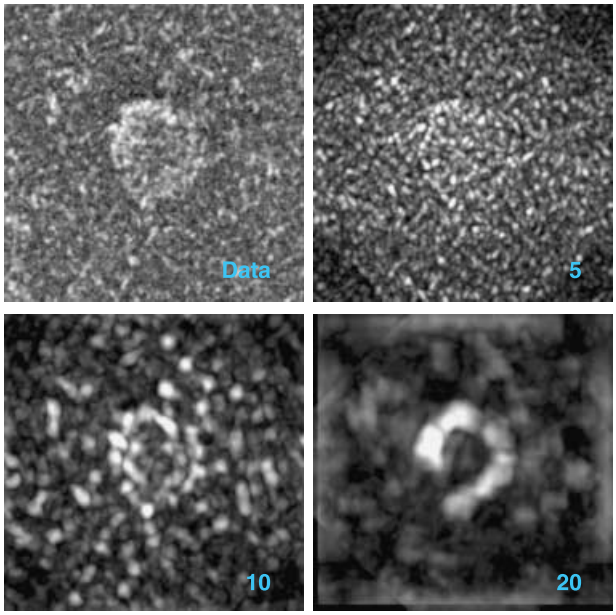

Fig. 5. 3D wavelet filtering of a γ-tubulin ring complex (γTuRC). A γTuRC has a characteristic 8-nm 'lock washer'-shaped structure, which is not apparent in a maximum-intensity projection of the data (upper left image). The images labelled 5, 10, and 20 voxels correspond to wavelet sizes equal to 1.9, 3.8 and 7.5 nm. The 7.5-nm wavelet (lower right) clearly shows the lock washer shape.

contrast. Tomographic data from unstained biological samples embedded in vitreous ice represent the current state-of-the-art in potential resolution and sample preservation, but are extremely noisy. The upper left image in Fig. 6 is a maximum-intensity volume projection of cryo-preserved (vitreous ice) microtubules (~25 nm characteristic diameter) that were imaged and reconstructed from EM tomography (Zheng *et al.*, 2004). As is readily apparent, the low contrast of cryo-EM data provides an extreme test for any visualization algorithm. In order to view the low-contrast microtubule walls (~5 nm thick) without being overwhelmed by noise, the displayed volume needs to be quite thin [5 voxels (11 nm) thick]. The full 3D volume was filtered with wavelet sizes equal to 3 and 5 voxels, which correspond to 6.6 and 11 nm. The upper middle and upper right images show the same volume slices (5 voxels thick) of the wavelet filtered data. The filtered data show an improvement in the visualization of the microtubule walls. In order to semi-quantify the improvement, a small sub volume (red box in the figure) was viewed at selected regions of the intensity histogram. The wavelet-filtered images show a greatly extended range of intensity that contains microtubule structures compared with the data images, which suggests a much enhanced signal-to-noise in the wavelet-filtered data.

The final example shows the results of wavelet filtering of optical microscope images that were obtained using a new high-resolution technique (M. G. L. Gustafson *et al.*, unpublished observations) (I5S) that is capable of producing images of approximately 0.1 μm resolution in all three dimensions by combining structured illumination (Gustafson, 2000) with interferometry from two objectives (Gustafson *et al.*, 1999). The upper left image in Fig. 7 is a maximum-intensity volume projection of oli-green-stained *C. elegans* gonad cells. [The samples were high pressure frozen and freeze-substituted, followed by embedding in epon plastic. Thick (~1 μm) sections were cut and stained with oli-green DNA-specific stain and then imaged.] The lower left image shows a side view of the data. The data were filtered with wavelet sizes equal to 3, 4 and 5 voxels, which correspond to 136, 181, and 226 nm. The wavelets enhance both clear and diffuse features. Edges are also enhanced, but appear displaced with respect to their location in the original data, e.g. compare the interstitial space between the cells with the original data. The spatial centring of the wavelet kernel (Eq. 2) is responsible for this offset of edges. We note that for accurate location of edges, the reference point of the wavelet kernel would be shifted from the centre to the leading or trailing edge of the positive valued region of the wavelet kernel. Volumes are spatially centred correctly.

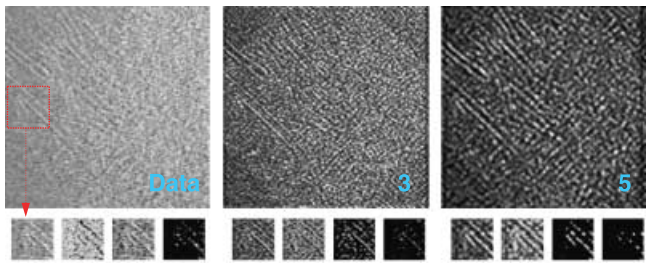

Fig. 6. 3D wavelet filtering of cryo-preserved (vitreous ice) microtubules. The upper left image shows a maximum-intensity volume projection of cryo-EM-imaged and tomographically reconstructed microtubules. The displayed volume is 5 voxels (11 nm) thick. Thicker volumes would obscure the microtubules, due to the low contrast of these data. The data (complete volume) are filtered using spatial scales of 3 and 5 voxels (upper middle and upper right), which correspond to 6.6 and 11 nm. The filtered data show an improvement in the visualization of the 5-nm-thick microtubule walls. Different regions of intensity of the same small section of each image (red square) are shown in the panels below the three upper images. Proceeding from left to right the displayed intensities are: full scale (original data), lower third, middle third, upper third.

Discussion

The need to visualize structural features in our 3D noisy biological data motivated our research into wavelet-based filters. Fourier filters, the current method of choice, are not particularly well adapted to picking out the non-repetitive, multiply orientated structures that comprise most biological data. Although wavelet theory rests on a well-established theoretical foundation (Daubechies, 1992; Farge, 1992; Laine, 2000; Boggess & Narcowich, 2001), extending such filters to 3D

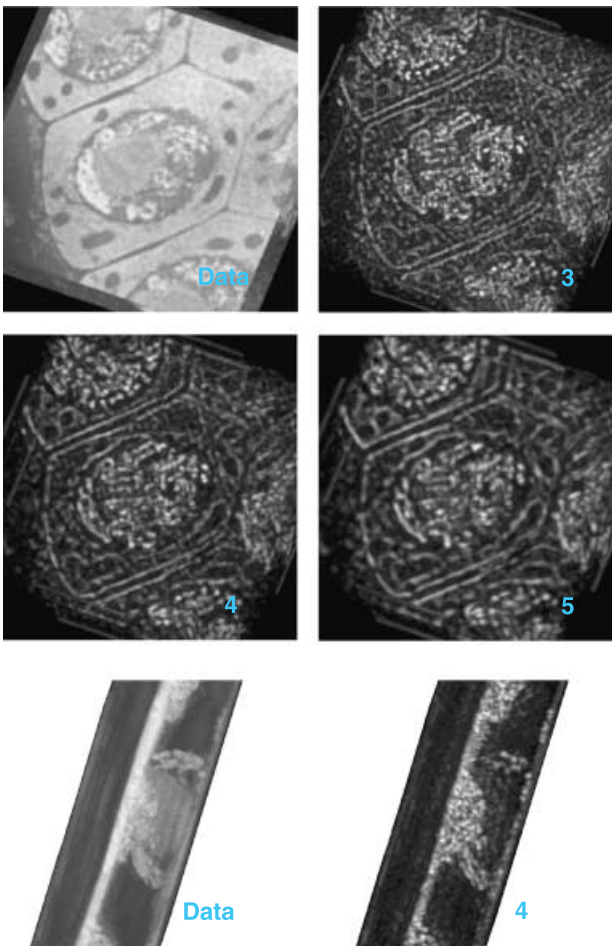

Fig. 7. 3D wavelet filtering of *C. elegans* meiotic cells. The volume data were acquired as a 3D data stack and imaged using a new microscope that employs interferometry from front and rear objectives combined with structured illumination. The upper left image shows a maximum-intensity volume projection of oli-green-stained *C. elegans* gonad cells. The voxel size is 45 nm. The lower left image shows this same volume rotated 90° around the vertical axis (side view). Diffuse structure is seen. The data are filtered using spatial scales of 3, 4 and 5 voxels, which correspond to 136, 181 and 226 nm (a side view of the 181-nm wavelet is shown in the lower right). The wavelets provide a higher contrast view of the cell structure.

results in computational complexities that for all practical purposes render the power of wavelets inaccessible to most researchers. Our goal was to produce a computationally fast wavelet-based filter that was capable of finding features of a given size in any 3D (or 2D) dataset.

The complexities in wavelet processing arise from the method that is used to evaluate Eq. (3). Most current problems to which wavelets are applied require computational speed, data compression or data reconstruction, so Eq. (3) is evaluated typically using the 'discrete wavelet transform (DWT)' (see chapters 4–6 in Laine, 2000). Our data do not require compression or reconstruction, so we have used the much simpler 'continuous wavelet transform (CWT)' (Farge, 1992), which is a direct integration of Eq. (3). Although the CWT is computationally slower than the DWT, the CWT is easier to interpret and allows very subtle information to be visualized due to its redundancy (Farge, 1992). We note that the computational speed of the CWT has not been a limiting factor in our analyses: the CWT of a $300 \times 300 \times 300$-voxel volume is computed in 3.6 min (4-voxel wavelet) and 3.7 min (12-voxel wavelet) on a PC using an Intel 2.8-GHz Xeon CPU. Computational speed results from the piecewise constant values of our wavelet kernel (Eq. 2): the value of each 1D integral at a particular x, y and z is obtained from the value at an adjacent point, by recalculating the integral at only the four locations where the kernel changes value. Consequently, the computational speed does not suffer greatly as the wavelet size increases.

The representation of the 3D transform in Eq. (3), using three 1D wavelets, simplified the development and coding of the filter and improved computational efficiency. Although alternative representations can probably be found, the simplicity and flexibility of this filter represents a significant advance. The only input parameter to the filter is a characteristic linear dimension of interest, measured in voxels, which are natural units to use for analyses of digital images. Consequently, our filter can be used to visualize structures in any 3D (or 2D) data. The filtered data show only the strength of the correlation between the data and the size of interest. Of course, the user must decide whether these mathematical correlations are meaningful scientifically, e.g. see Fig. 4.

Any type of filtering has the potential to introduce artefacts. We routinely compare our original data with our wavelet-filtered data to ensure that features that are enhanced by the filter are discernible in the original data. In addition, when it is possible, we obtain data from the same sample, using independent methods (e.g. EM, if the original data were optical), to confirm the structural identifications. Nevertheless, some artefacts are unavoidable. The mathematical structure of wavelets, in general, makes them excellent edge detectors. Our filter acts as an edge detector when the wavelet size (a in the defining equations) is small compared with the feature of interest. Although edge detection and enhancement is not necessarily bad, it could possibly result in a disc appearing as a ring, or a solid ribbon-shaped structure appearing as a pair of thin solid ribbons. (Figure 7 shows a specific example.) Filtering with multiple wavelet sizes can usually resolve these anomalies, and should be done as a matter of course if only to avoid the possibility of missing interesting previously unknown features in the data.

One of the strengths of using the CWT is that feature enhancement waxes and wanes smoothly as the wavelet size changes, so if the feature's characteristic size is not known precisely a priori, it can still be detected with a non-optimal wavelet size (although not as strongly). The wavelet size can be adjusted subsequently to optimize the visualization of the feature of interest. Our particular choice of the normalization of the filter allows the results of different wavelet sizes to be compared

directly. Larger values at a particular point mean stronger correlations with the associated wavelet size. We are currently exploring methods to determine automatically (without prior knowledge) the dominant characteristic sizes in a data volume.

What is the power of this method for biology? Simply stated, the ability to view a complex densely packed biological image and identify a backbone–ridgeline, or key features, is essential for interpreting biological structure. Once the glimpse of the essence of the structure is discerned, then other structural interpretation methodologies can be applied. Subsequent molecular and cell biological research can complement the structural studies.

What is the power of this method outside of biology? The need for feature identification is ubiquitous: medical imaging (e.g. bone structure), non-destructive evaluation (e.g. internal cracks and defects) and even the prosaic airline baggage scanner are technologies whose primary function is to identify features of interest. We are currently using our filter to examine data from all of these modalities.

Supplementary material

The following material is available for this paper online:
Movie S1.

Acknowledgements

We thank T. Gay, B. Hasegawa, Z. Kam and M. Gustafson for reviewing the manuscript. Work by W.C.M., S.H. and J.W.S. was supported by NIH grant GM25101-26 and D.A.A. was supported by NIH grant GM31627. J.M.L. is a fellow of the Jane Coffin Childs Memorial Fund for Medical Research. We thank A. Dernberg and B. Isaac for the preparation of the *C. elegans* sample. Work by W.C.M. was performed as part of a UC-DOE multilocation appointment under the auspices of the U.S. DOE by LLNL under contract W-7405-Eng-48.

References

Boggess, A. & Narcowich, F.J. (2001) *A First Course in Wavelets with Fourier Analysis*. Prentice Hall, Inc, Upper Saddle River, NJ.

Böhm, J., Frangakis, A.S., Hegerl, R., Nickell, S., Typke, D. & Baumeister, W. (2000) Toward detecting and identifying macromolecules in a cellular context: template matching applied to electron tomograms. *Proc. Natl Acad. Sci. U.S.A.* **97**, 14245–14250.

Chen, H., Clyborne, W., Sedat, J.W. & Agard, D.A. (1992) PRIISM: an integrated system for display and analysis of three-dimensional microscope images. *Proc. SPIE – Int. Soc. Opt. Engng*, **1660**, 784–790.

Daubechies, I. (1992) *Ten Lectures on Wavelets.* SIAM, Philadelphia, PA.

Farge, M. (1992) Wavelet transforms and their applications to turbulence. *Annu. Rev. Fluid Mech*. **24**, 395–457.

Frangakis, A.S., Böhm, J., Förster, 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**, 14153–14158.

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

Gustafson, M.G.L. (2000) Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy. *J. Microsc*. **198**, 82–87.

Gustafson, M.G.L., Agard, D.A. & Sedat, J.W. (1999) I5M: 3D widefield light microscopy with better than 100 nm axial resolution. *J. Microsc*. **195**, 10–16.

Laine, A.F. (2000) Wavelets in temporal and spatial processing of biomedical images. *Ann. Rev. Biomed. Eng*. **2**, 511–550.

Moritz, M., Braunfeld, M.B., Fung, J.C., Sedat, J.W., Alberts, B.M. & Agard, D.A. (1995) Three-dimensional structural characterization of centrosomes from early drosophila embryos. *J. Cell Biol*. **130**, 1149–1159.

Moritz, M., Braunfeld, M.B., Guénebaut, V., Heuser, J. & Agard, D.A. (2000) Structure of the γ-tubulin ring complex: a template for microtubule nucleation, *Nature Cell Biol*. **2**, 365–370.

Nicholson, W.V. & Glaeser, R.M. (2001) Automatic particle detection in electron microscopy. *J. Struct. Biol*. **133**, 90–101.

Rath, B.K. & Frank, J. (2004) Fast automatic particle picking from cryo-electron micrographs using a locally normalized cross-correlation function: a case study. *J. Struct. Biol*. **145**, 84–90.

Reimer, L. (1997) *Transmission Election Microscopy Physics of Image Formation and Microanalysis*. Springer-Verlag, Berlin.

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

Westphal, V., Kastrup, L. & Hell, S.W. (2003) Lateral resolution of 28nm (λ/25) in far-field fluorescence microscopy. *Appl. Phys. B*, **77**, 377–380.

Zheng, Q.S., Braunfeld, M.B., Sedat, J.W. & Agard, D.A. (2004) An improved strategy for automated electron microscopic tomography. *J. Struct. Biol*. **147**, 91–101.

Appendix A: analytical helix construction

Two single helical skeletons were constructed analytically and discretized on an $80 \times 80 \times 80$ cubic grid, where each helix had a radius equal to 10, and repeat length equal to 20. The axes of the two helices were parallel, but displaced by 18 voxels. The intensity of each point on this double helical skeleton was arbitrarily set equal to 10, while the intensity elsewhere was set equal to 0. The helices were 'grown' to their final sizes in two steps: (i) each non-zero valued voxel has 26 neighbours. The intensity of each zero valued neighbour was set equal to 9, whereas non-zero values were not changed. (ii) Repeat step (i), using a new intensity equal to 8 instead of 9. We define the SN to equal the peak signal intensity (10) divided by the standard deviation of the noise (calculated in a spatial region distant from the signal). The noise was described by a Gaussian distribution about a mean of zero, with standard deviation equal to the peak intensity (10) divided by the desired SN. The final analytical construct was obtained by summing the helix and noise values at each point. A uniform shift equal to the minimum value of the noise spectrum was added to each point so that only positive intensities were obtained. This uniform shift does not alter the SN.