

Comments on the paper “A novel 3D wavelet-based filter for visualizing features in noisy biological data”, by Moss *et al.*

C. L. Luengo Hendriks * and D. W. Knowles

Life Science Division, Lawrence Berkeley National Laboratory,
One Cyclotron Road, mailstop 84R171, Berkeley, CA 94720, USA.

Summary

Moss *et al.* (2005) describe, in a recent paper, a filter that they use to detect lines. We noticed that the wavelet on which this filter is based is a difference of uniform filters. This filter is an approximation to the second derivative operator, which is commonly implemented as the Laplace of Gaussian (or Marr-Hildreth) operator (Marr & Hildreth, 1980; Jähne, 2002), Figure 1. We have compared Moss’ filter with 1) the Laplace of Gaussian operator, 2) an approximation of the Laplace of Gaussian using uniform filters, and 3) a few common noise reduction filters. The Laplace-like operators detect lines by suppressing image features both larger and smaller than the filter size. The noise reduction filters only suppress image features smaller than the filter size. By estimating the signal to noise ratio (SNR) and mean square difference (MSD) of the filtered results, we found that the filter proposed by Moss *et al.* does not outperform the Laplace of Gaussian operator. We also found that for images with extreme noise content, line detection filters perform better than the noise reduction filters when trying to enhance line structures. In less extreme cases of noise, the standard noise reduction filters perform significantly better than both the Laplace of Gaussian and Moss’ filter.




Figure 1: Wavelet used by Moss *et al.* to construct their filter (—), compared to the second derivative of a Gaussian function (– –).

*Tel: +1-510-4865359; fax: +1-510-4865730; e-mail: clluengo@lbl.gov

Method

The comparisons here were done in two dimensions. Moss' wavelet-based filter is described for three-dimensional images, but is dimensionality independent. A 2D test image was created following the recipe given by Moss: We drew three 1-pixel thick lines of intensity 10 on a background of intensity 0, and dilated them with decreasing intensity to 5 pixels width. To this test image we then added Gaussian distributed noise with a mean of 0 and a standard deviation of $10/\text{SNR}$ (SNR being the chosen value for the signal to noise ratio as defined by Moss *et al.*), and linearly stretched the result to fit the $[0,255]$ range. This stretching did not change the content of the image, since we used 32-bit floating point representation for the image data. The parameter a from Moss' filter was set to 6 pixels. Each of the filters, described below, was then applied to test images with a SNR of 8, 2 and 0.5.

To estimate the SNR in the filtered results we used the following procedure: 1) Find the multiplication factor that minimizes the mean square difference between the noiseless input image and the filter output, taking into account only the pixels that fall within the test image lines (i.e. those pixels that are non-zero in the noiseless input image). 2) Multiply the filtering output with this value. This normalizes the "signal level" to 10. 3) Compute the standard deviation s over a set of background pixels in the normalized filter output (we exclude all pixels within a distance of $3a$ of the lines and the edge, see the mask image in Figure 2). 4) The SNR is given by $10/s$. The value of the minimized mean square difference gives an indication of how well the filter was able to preserve the line.

The scripts we used for this comparison have been made available online at <http://clluengo.lbl.gov/mossfilter.html>

Moss' Wavelet-Based Filter

For 2D, Moss' filter requires eight 1D filter passes, eight clips (which sets all negative values to 0), three image averagings and two rotations. The rotations are the most expensive components of this filter, and are used to improve the rotation-invariance of the filter (applying 1D uniform filters, which is what the wavelet is built on, sequentially in each dimension yields a rectangular filter kernel).

The Laplace Operator

The Laplace operator ∇^2 is constructed by addition of the second order derivative in each dimension (Jähne, 2002), that is

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2} \quad (1)$$

The second order derivative is best implemented as a Gaussian derivative (Marr & Hildreth, 1980; Jähne, 2002), to avoid enhancing the noise. This is often

referred to as the Laplace of Gaussian (LoG) or the Marr-Hildreth operator, and can be implemented as a separable filter (Huertas & Medioni, 1986). The Gaussian smoothing gives the LoG operator a parameter, which has to be tuned to the thickness of the lines expected in the image.

For this comparison we choose the parameter for the Gaussian to be $\sigma = 0.6a$, with a being the parameter for Moss' filter. This makes the two filters quite comparable, as can be seen in Figure 1. Also, we clipped the result of the operator, setting all negative values to 0, to mimic the behavior of Moss' filter.

The Difference of Uniform Filters

The LoG operator can be approximated by a Difference of Gaussians (DoG) (Jähne, 2002). The image is smoothed at two scales, and the difference taken. In 1D, if the Gaussian smoothing is replaced by a uniform smoothing kernel, we obtain an operator identical to the wavelet used by Moss. We implemented this in 2D (and called it “difference of uniform filters”, DoU) by using a disk of diameter a as the smaller smoothing kernel, and a disk of diameter $3a$ as the larger. As with the LoG, we clip negative values to 0.

There are other filters that improve upon the LoG, such as the non-linear Laplace operator described by van Vliet *et al.* (1989). Another nice work on line detection is by Danielsson *et al.* (2001). We did not consider it necessary to include any of these here.

Noise Reduction filters

Reducing noise must be one of the topics most studied in signal and image processing. We added several common methods to this comparison: the median filter (Jähne, 2002), the Kuwahara-Nagao filter (Kuwahara *et al.*, 1976; Nagao & Matsuyama, 1979), and an alternating sequential filter (Sternberg, 1986; Serra, 1988). These filters do not set the background value to 0, but rather to the noise mean. Thus, to better compare the result of these filters with the line detectors we subtract the median of the noisy input image, and then clip all negative values to 0. For both the median and the Kuwahara-Nagao operator we used a disk with diameter $a - 1$ as the filter kernel. The alternating sequential filter is implemented as an open-close filter applied recursively with disks of increasing size between 2 and $a - 3$ pixels.

Results and Discussion

Figure 2 shows the input images and the result of each of the filters described, for different setting of the input SNR. As can be seen none of the methods work particularly well for very high noise levels. Moss *et al.* show a better result at $\text{SNR} = 0.5$ on their 3D test image than the result the same filter obtains here in 2D. In 3D a higher noise level is allowable because of the increased number of pixels within a neighborhood, which leads to better statistics on the noise.

| SNR | 8 | | 2 | | 0.5 | |
|-----------------|--------|---------|-------|---------|-------|----------|
| Input image | 8.02 | (0.163) | 2.00 | (2.439) | 0.50 | (39.712) |
| Moss' filter | 37.57 | (0.971) | 9.24 | (1.323) | 3.39 | (2.933) |
| LoG | 47.16 | (0.292) | 11.35 | (0.543) | 3.82 | (2.167) |
| DoU | 34.73 | (0.604) | 8.81 | (0.989) | 3.34 | (2.991) |
| Median filter | 54.10 | (0.076) | 12.44 | (0.436) | 3.75 | (2.599) |
| Kuwahara filter | 67.50 | (0.073) | 15.61 | (0.761) | 4.04 | (2.619) |
| Seq. open-close | 566.94 | (0.062) | 72.06 | (1.241) | 14.34 | (6.851) |

Table 1: SNR estimated for the filter results with, between brackets, the MSD obtained after fitting the signal (see text for details).

This advantage would equally improve the results of the other filters in this comparison. For all but the lowest SNR levels, standard noise reduction filters do a really good job of recovering the original noiseless image.

Table 1 shows the SNR and MSD estimated for the images discussed above. High values of SNR are desired, but not at the expense of increased MSD, which would indicate that the signal is not being recovered well. As can be seen, Moss' filter and the DoU obtain a similar performance (with this measure), whereas the LoG performs slightly better. We attribute this difference to the optimality of the Gaussian as a low-pass filter (Marr & Hildreth, 1980). The noise reduction filters break down for extremely low SNR, but in less severe cases perform much better than any of the line detecting filters.

We cannot compare the execution time of these filters because our implementation of Moss' filter is not optimized. However, we can compare the computational complexity of this filter with the LoG. As explained above, Moss' filter is implemented by adding four filtering results (or 24 in 3D), obtained by a separable filtering kernel; some of these filtering passes must be applied to a rotated version of the input image, the output of which must be rotated back (two rotations in 2D or six in 3D). In contrast, the LoG can be computed by adding the output of two separable filters (or three in 3D). On our AMD Opteron (2.2 GHz) processor, the LoG filter with $\sigma = 0.6a$ ($a = 6$) takes 13 seconds to process a $300 \times 300 \times 300$ voxel image. In comparison, a single rotation of this image around the z-axis takes 15 seconds. Moss' filter requires 6 rotations, 24 filter passes and some additional arithmetic, which Moss *et al.* reported taking 3.6 minutes on their Intel Xeon (2.8 GHz) processor, on an image of the same size.

Conclusion

The wavelet-based filter presented by Moss *et al.* is very similar to a Laplace-type filter. The complex implementation of Moss' filter makes it computationally expensive and, in our study using 2D synthetic data, does not perform better than other established techniques, such as the Laplace of Gaussian operator.




Figure 2: The noisy input images and the results for the various filters in our test. On the top right corner is the mask image used to identify background pixels for the SNR measure.

References

- Danielsson, P.-E., Lin, Q., and Ye, Q.-Z. (2001). Efficient detection of second-degree variations in 2D and 3D images. *Journal of Visual Communication and Image Representation*, 12:255–305.
- Huertas, A. and Medioni, G. (1986). Detection of intensity changes with sub-pixel accuracy using Laplacian-Gaussian masks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8(5):651–664.
- Jähne, B. (2002). *Digital Image Processing*. Springer, Berlin, 5th edition.
- Kuwahara, M., Hachimura, K., Eiho, S., and Kinoshita, M. (1976). Processing of RI-angiographic images. *Digital Processing of Biomedical Images* (ed. by K. Preston and M. Onoe), pp. 187–203. Plenum Press, New York.
- Marr, D. and Hildreth, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London B*, 207:187–217.
- Moss, W. C., Haase, S., Lyle, J. M., Agard, D. A., and Sedat, J. W. (2005). A novel 3D wavelet-based filter for visualizing features in noisy biological data. *Journal of Microscopy*, 219(2):43–49.
- Nagao, M. and Matsuyama, T. (1979). Edge preserving smoothing. *Computer Graphics and Image Processing*, 9:394–407.
- Serra, J. (1988). Alternating sequential filters. *Image Analysis and Mathematical Morphology, Volume 2: Theoretical Advances* (ed. by J. Serra), pp. 203–216. Academic Press, London.
- Sternberg, S. R. (1986). Grayscale morphology. *Computer Vision, Graphics and Image Processing*, 35:333–355.
- van Vliet, L. J., Young, I. T., and Beckers, A. L. D. (1989). A nonlinear laplace operator as edge detector in noisy images. *Computer Vision, Graphics and Image Processing*, 45(2):167–195.