Some morphological scale-space filters and their properties

Richard Harvey, J. Andrew Bangham, Alison Bosson
School of Information Systems, University of East Anglia,
Norwich, NR4 7TJ, UK.
Tel: +44 1603 593257; fax: +44 1603 593345
e-mail: {rwh,ab,bosson}@sys.uea.ac.uk

November 20, 1996

Abstract

We discuss some properties of then class of scale-space processor called the *sieve*. Sieves are useful because, like a diffusion processor they have an increasing support region, but, unlike a diffusion processor, the region follows extremal regions in the image. We describe the poperties of sieve processors and show how they may reduce some of the problems associated with diffusion systems.

1 Introduction

This paper describes a class of systems known as *sieves* that derive from mathematical morphology and rank filtering. Section 3 is a brief review of the relevant literature in these fields. Section 2 illustrates some well known problems with scale-spaces created from diffusion systems and shows how some of these may be solved using scale-spaces from sieves. Section 4 describes some of the properties of sieves[1]. They differ from the more widely known mathematical morphology filters in their ability to reject noise robustly and so experiments to test the robustness of diffusion- and sieve- based processors are described in Section 5. Section 6 shows how some of these properties are beginning to produce useful systems.

2 Scale-space processors

Systems that preserve scale-space causality are usually associated with Gaussian filters [3, 4, 5] and diffusion [6] in which the image forms the initial conditions for a discretization of the continuous diffusion equation:

$$\nabla \cdot (c \nabla f) = \frac{\partial f}{\partial s} \tag{1}$$

If the diffusivity is a constant this becomes the linear diffusion equation, $\nabla^2 f = f_s$ which may be implemented by convolving the image with the Green's function of diffusion equation: a Gaussian filter. Of course, care is needed when discretising (1) [7] but, if it is done correctly, a scale-space with discrete space and continuous scale may be formed[2]. It is possible to obtain separable filters ($\gamma = 0$) where the image at scale s is computed as

$$f^{(s)}(x, y) = \sum_{m=-\infty}^{\infty} T(m; s) \sum_{n=-\infty}^{\infty} T(n; s) f(x - m, y - n) \tag{2}$$

where $f(x, y)$ is the pixel value at position (x, y) and $f^{(s)}(x, y)$ is the pixel value after smoothing to scale s. $T(n; s)$ is the discrete approximation to the Gaussian kernel

$$T(n; s) = e^{-s} I_n(s) \tag{3}$$

and $I_n(s)$ is the modified Bessel function of the first kind [8].

However, such scale-space systems have several problems:

1. The edges of sharp-edged objects become blurred at large scales. This leads to poor localisation and a necessity to track edges back through scale-space.
2. They introduce extrema, albeit in rather pathological cases.
3. They are not scale calibrated. An image filtered to scale s contains features at many scales.
4. Multiple convolutions may require significant amounts of computation using real arithmetic.

Problem 1 may be overcome reverting to (1) and allowing the conduction coefficent, c, to be a function of $\|\nabla f\|$ thus $c = g(\|\nabla f\|)$. If this function is carefully chosen (Perona and Malik suggest $g(\|\nabla f\|) = \exp\left(-\left(\|\nabla f\| / K\right)^2\right)$ or $g(\|\nabla f\|) = \left(1 + \left(\frac{\|\nabla f\|}{K}\right)^2\right)^{-1}$) then the effect is to allow diffusion in low contrast regions but not at sharp-edges. Unfortunately anisotropic diffusion requires even more

[1]The mathematical properties have been described elsewhere [1, 2]. We include this last reference for the convenience of the reviewers.

[2]Or with a discrete scale parameter if preferred.

computation than linear diffusion and problem 3 is exacerbated – sharp-edged small scale objects persist to large scales.

The situation is summarised in Figure 1. In the left-hand column is a photograph of Norwich cathedral. The top row shows this image progressively simplified using the separable discrete Gaussian kernel of (3). The right-hand image is considerably simpler than the original but barely recognisable. The centre row shows the same image simplified using anisotropic diffusion. The boundaries between sharped-edged objects are preserved but the large-scale picture contains objects at a variety of scales. The bottom row shows the same image simplified with one of the sieve processors.

Figure 1: Scale-space filters

Problem 2 is illustrated in Figure 2. Each panel shows an image and a three-dimensional representation of its greyscale structure. The original image (two discs connected by thin isthmus form its main feature) is shown on the top left. The top right shows the result after processing with a discretised Gaussian filter. The single maximum associated with the two discs and isthmus has now become two maxima. This example, due to Pizer [9], is well known and leads to a new definition of scale-space causality namely the "non-enhancement of exisiting regional extrema" principle which is satisfied by diffusion systems. The lower diagrams in Figure 2 show the result of applying an M-sieve. In the left-hand diagram all small-scale objects have been removed (the small disc has gone) and, on the right, small and medium-scale objects (the large intense disc on the left has also disappeared). In both cases the large connected discs are untouched. Not only do sieves not enhance existing regional extrema, they do not introduce new ones [10, 11, 2], a property shared with multidimensional erosions and dilations [12, 13] and one-dimensional openings and closings [14, 15].

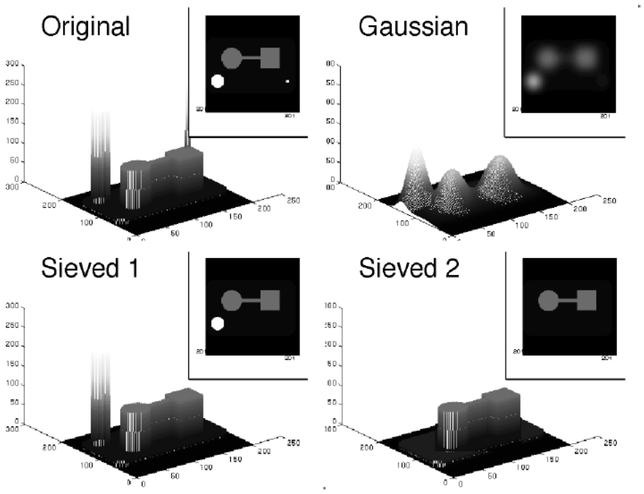

Figure 2: Diffusion systems introduce extrema

Figure 2 also illustrates problem 3. In the Gaussian processor the intensity of the output at scale s is proportional to the scale and intensity of the original feature. So to recover the parameters of the original signal implies either deconvolution or scale-space tracking. Both of these are difficult tasks and, as far as image interpretation goes, having peaks representing extrema from objects of a variety of scales is inconvenient. As will be seen, morphological filters produce at scale s only objects of size s.

3 Morphological processors

Mathematical morphology is the analysis of signals, particularly images, by shape and has been developed by Serra [16, 17, 18] from work by Matheron [19] but it also derives from work of others (Blum [20] for example). Morphological methods are widely used in commercial image processing packages where they appear as greyscale, or more usually, binary operations, such as erosion, dilation, opening, closing, top-hat or watershed [21, 22]. Several of these functions are available in hardware (DataCube, Imaging Technology, Matrox) and custom hardware appears to be used in several industrial systems.

Aspects of mathematical morphology relevant to sieves include: multiple scale erosions [12], multiple scale openings or "granulometry" [19], increasing scale sequences of open/closings or "alternating sequential filters" [23] but only with flat structuring elements [24] and alternating sequential filters with reconstruction, "reconstruction filters", again with flat structuring elements [25, 26, 27].

Granulometry has been used to characterise size distributions of rock particles [19] and shapes [28, 29] and there are methods for choosing the relevant granulometry in a principled fashion [30]. Granulometry algorithms are particularly fast in one dimension [31, 32, 33]. Such multiscale openings or closings preserve scale-space causality in one-dimension [14] but not higher [15]. There appear to be two ways to preserve scale-space causality in higher dimensions: by using simpler granulometry analogues based on erosions and dilations [12]; or by using *connected set* granulometries that do not use rigid structuring elements [34, 35, 36].

Alternating sequential filters with greyscale structuring elements that match features [23, 17] are said to be a powerful analogue of linear matched filters and there are strategies for designing optimal filters with flat, rigid, two dimensional flat structuring elements [24]. More recently, alternating sequential filters, that do not impose a rigid geometry on objects, and connected sets have been described[34, 37, 36, 35, 27], that fall within our definition of sieves [38, 2].

A separate stream of development has been that of rank filters, including median [39], root median [40], recursive median [41] and, more generally, stack filters [42]. Such filters have been developed primarily to remove random noise from signals and images, although there are suggestions for using them for tasks such as shape recognition [43]. There is a close relationship between rank filters and morphological filters [44]. Rank (median) filters are usually self-dual and robustly reject random noise. The advantages of these self-dual filters are incorporated into median and root-median sieves [32, 33, 45] and recursive median sieves [11]. There are other, related, nonlinear multiscale approaches to object location and recognition. For example, "SUSAN" which uses rigid structuring elements with median-like operators acting as "highpass" filters [46]. Of the mathematical morphology based filters that preserve scale-space causality, we are concerned with the category we call sieves. These are now considered in relation to the way they address the problems outlined in Section 1.

4 Sieves and their properties

The basis of the sieve [1, 2] is the representation of an image as a graph [47, 48] $G = (V, E)$. The set of edges E describes the adjacency of the pixels (which are the vertices V). For a one-dimensional image the graph is just a list [49, 11, 10] but for a multidimensional image the graph defines the neighbourhood of a particular pixel. For example, in Figure 3 which shows a four-connected two-dimensional image the graph would be $G = (V, E)$ with $V = \{1, 2, 3, \ldots, 16\}$ and $G = \{\{1, 2\}, \{2, 3\}, \{3, 4\}, \{1, 5\}, \ldots\}$. In Figure 3 the numbers are labels not intensities.

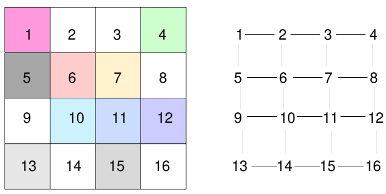

Figure 3: An image represented as a graph.

The algorithm proceeds by defining a region, $C_r(G, x)$ over the graph that encloses the pixel (vertex) x,

$$C_r(G, x) = \{\xi \in C_r(G) | x \in \xi\} \tag{4}$$

where $C_r(G)$ is the set of connected subsets of G with r elements. Thus $C_r(G, x)$ is the set of connected subsets of r elements that contain x. In Figure 3, for example, $C_2(G, 6) = \{\{5, 6\}, \{2, 6\}, \{6, 10\}, \{6, 7\}\}$. For each integer $r \geq 1$ the operators ψ_r, γ_r, $\mathcal{M}^r$, $\mathcal{N}^r : \boldsymbol{Z}^V \to \boldsymbol{Z}^V$ are defined as

$$\psi_r f(x) = \max_{\xi \in C_r(G,x)} \min_{u \in \xi} f(u), \tag{5}$$

$$\gamma_r f(x) = \min_{\xi \in C_r(G,x)} \max_{u \in \xi} f(u), \tag{6}$$

$$\mathcal{M}^r = \gamma_r \psi_r, \qquad \mathcal{N}^r = \psi_r \gamma_r. \tag{7}$$

$\mathcal{M}^r$ is a greyscale opening followed by a closing defined over a region of size r and $\mathcal{N}^r$ is a greyscale closing followed by an opening over the same region.

The types of sieve known as M- or N-sieve are formed by repeated operation of the $\mathcal{M}$ or $\mathcal{N}$ operators. They are also known as connected alternating sequential filters. An M-sieve of f is the sequence $(f^{(r)})_{r=1}^{\infty}$ given by

$$f^{(1)} = \mathcal{M}^1 f, \quad f^{(r+1)} = \mathcal{M}^{r+1} f^{(r)}, \quad r \geq 1 \tag{8}$$

The N-sieve is defined similarly. The output of an area sieve is usually taken to be the set of *granule functions*

$$d^{(r)} = f^{(r)} - f^{(r+1)} \text{ for each integer } \quad r \geq 1 \tag{9}$$

These form the scale selection surface and non-zero connected regions within granule functions are called granules. Each granule has sharp edges and, at a particular scale, all granules have the same area. In this sense the sieve is *scale calibrated*, a characteristic illustrated in Figure 2, that addresses Problem 3 in Section 1.

An example of the operation of the area sieve is shown in Figure 4 and 5. Figure 4 shows a twenty-five pixel image with one maximum in the centre and one minimum in the top left.

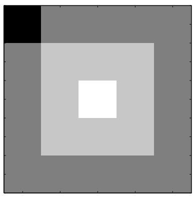

Figure 4: An example image, f, with two unit-area extrema and one of area nine

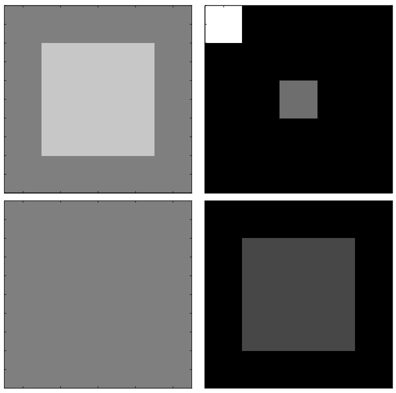

Figure 5: $f^{(2)}$ (top left), $d^{(1)}$ (top right), $f^{(10)}$ (bottom left), $d^{(9)}$ (bottom right)

The first application of the sieve yields $f^{(1)} = f$ and so is not shown. The top left of Figure 5 shows $f^{(2)}$ which is $f^{(1)}$ with all the extrema of unit area removed. The top right of Figure 5 shows $|d^{(1)}|$. There are granules where there was a maximum and a minimum and area one. The remaining feature has an area of nine pixels so $f^{(2)} \cdots f^{(9)}$ are identical and are not shown. The bottom left of Figure 5 shows $f^{(10)}$ in which central nine-pixel maximum has been removed and on the bottom right of Figure 5 is $|d^{(9)}|$ which shows the corresponding granule.

The M operation depicted in Figure 5 is the same as a removal maxima followed by the removal of minima. In an N-sieve these operations are reversed. With this in mind it is possible to define an

m-sieve in which the extrema are processed in the order in which they occur as the graph is parsed. In one-dimension such an algorithm is equivalent to the recursive median operator. The m-sieve has properties similar to M- and N-sieves and so not all sieves are alternating sequential filters and not all alternating sequential filters are sieves. There is a higher dimensional equivalent of the m-sieve (unpublished).

It might appear that the sieving operation could be computationally expensive. In practice, however, the algorithms referred to here, in which the image graph is rewritten as flat-zones are merged, produce scale-spaces quicker than linear diffusion. Typically a 512 by 512 image is processed in seconds [50, 1]. This addresses Problem 4 in Section 1.

Figure 2 suggests that area sieves preserve scale-space causality and this has been proved formally (Theorem 6.36 in [2]) together with a number of other properties of area sieves (Theorem 6.49 in [2]). This addresses Problem 2 in Section 1.

For the simple image shown in Figure 4 the sieve decomposition yields only a few images but for a 100 by 100 pixel image there are 10^4 possible scale-space filtered images. So, for the purposes of display and for approximate scale determination, the outputs over a range of scales are summed to form a *channel*. Figure 6 shows the channels from an image of a telephone. Features of different sizes appear in different spatial channels and those with sharp edges retain them. Likewise the bottom row of Figure 1 shows that large scale objects retain their sharp edges. This addresses Problem 1 in Section 1.

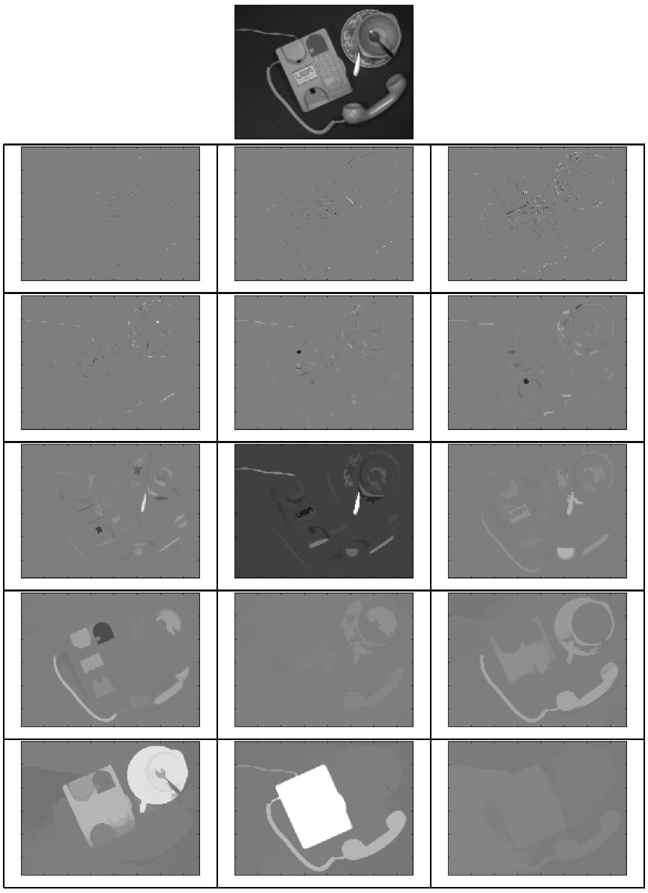

Figure 6: A decomposition into channels. The original image is shown at the top.

The properties described so far would appear to be desirable. However, it is well known that morphological operators such as erosion and dilation are sensitive to noise, more complicated morphological filters are less so, and Gaussian filters, having the form of matched filters, are fairly insensitive to noise. Whether sieves have practical application depends on whether they reject image corruptions robustly.

5 Robustness

Noise interacts with CCD images and photographs in a complicated manner [51, 52]. Furthermore the underlying noiseless image distribution is also extremely complicated. For this reason it is usual to study the performance of filters using rather stylised models of images and noise. Here we shall consider the underlying image to be a disc or a square which we shall assume is corrupted with either Gaussian noise (a model of noise in the imaging system) or impulsive noise (a model for 'glint' or pixel drop-out).

Figure 7 shows examples of the target images. The target image is a disc or square of amplitude 144 in the centre of a 100 by 100 pixel image with background 112. To this, is added either uncorrelated Gaussian noise ($\mu = 0, \sigma = 24$), or alternatively, pixels are replaced with a random value in the range $(-6, 6)$ with a probability of 0.2. The resulting image is clipped into the range (0,255).

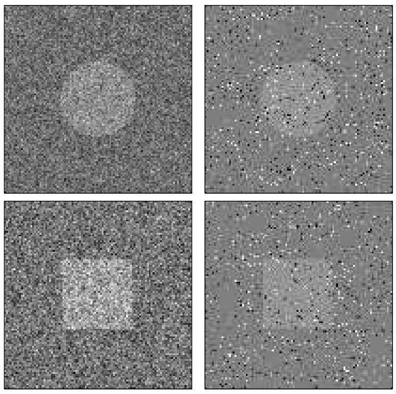

Figure 7: Examples of noise-corrupted images: Gaussian noise (top) and Impulsive noise (bottom)

The performance of sieves was compared to Gaussian processors in which the scale-space was generated using separable filters ($\gamma = 0$) where the image at scale s is computed as in (2). The scale-selection surface ([7] p. 323)

$$d = s^2 \left(f_{xx}^{(s)} + f_{yy}^{(s)} \right)^2 \tag{10}$$

is used to locate the scale-space estimate of the target parameters. For the sieve the scale-selection surface is formed from the granule functions. In the absence of noise both scale-selection surfaces will have a peak corresponding to the target.

For both systems an iterative search is conducted over all scales to find this maximum. Scattergrams of results from 150 trials of the image corrupted by Gaussian noise for the two systems are shown in Figure 8. Each graph shows the estimates as circular markers with their co-ordinates being the estimates of x, y and s. The light-grey dots are the projection of the points in three-space onto each pair of axes. The linear scale-space processor is scale-calibrated by normalising the mean of the cluster to equal the true area of the disc.

In the linear scale-space processor the peak is the maximum value of the normalised Hessian (10). To improve the precision of the position estimates a quadratic surface is least-squares fitted to the peak and its four-connected neighbours. The position of this peak gives the (x, y) position estimate. Since the area sieve operates by “slicing off” the peaks and troughs, at a particular scale, the extrema are large and flat. To obtain the position of the disc the centroid of all the pixels that form the maximum, at a particular scale, is computed.

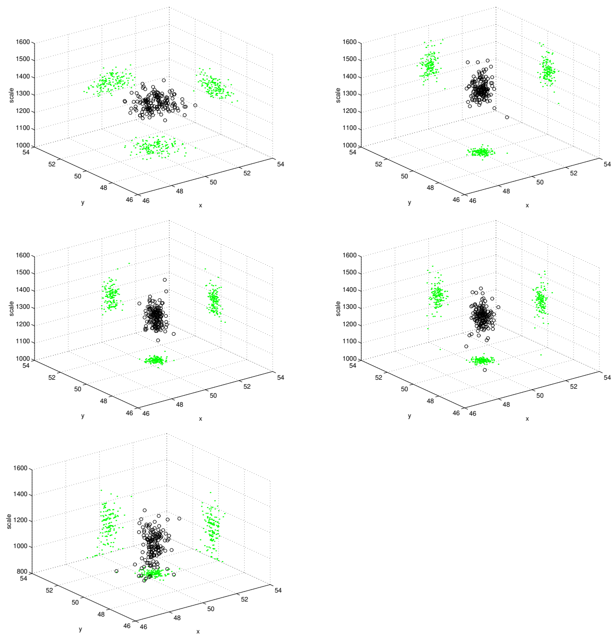

Figure 8: Estimates with Gaussian noise, disc target. Top row shows diffusion and m-sieve, middle row shows N-sieve and M-sieve and lower row shows an opening sieve

Examining Figure 8 shows that, in Gaussian noise, the diffusion system, M-, $N-$ and m-sieves have comparable performance. Multiscale openings (opening sieve) have worse erformance in scale and the multiscale closing cannot resolve a positive disc under these conditions[3]. Studies with the equivalent one-dimensional filters [53] show that the linear diffusion system is more sensitive to Gaussian noise than the sieve. In two-dimensions this is not the case, probably because the scale-selection surface (10) in two-dimensions performs more spatial averaging.

Figure 9 shows that in impulsive noise, M-, $N-$, m- and opening- sieves can produce better scale and position estimates than those from the diffusion system. This is an important property since in real vision systems impulses, glint and occlusion are commonplace. Again, the asymmetry of multiscale closings means that the algorithm would have to be specially tuned to handle positive targets.

The performance of sieves in replacement noise is not surprising since it is known that, in one dimension, $\mathcal{M}$ and $\mathcal{N}$ filters behave similarly to median filters which have good performance in impulsive noise. The sieve retains its performance in Gaussian noise because the cascaded operation described in (8) means that, like wavelets, the output at large scales has a large support.

A key question is why is the performance of image processors in noise important? Studies of imaging noise in general purpose CCD images and photographs show that the noise model is complicated but has a Gaussian pdf in which the parameters of the pdf depend on the image intensity. However for most images the noise level is low. We believe it is the impulsive noise model that is the most interesting since it forms the basis for a study of occlusion – an important reality that is considered now.

[3]The closing removes minima, leaving maxima intact. Since, in this simple experiment, we search for maxima, this is a problem.

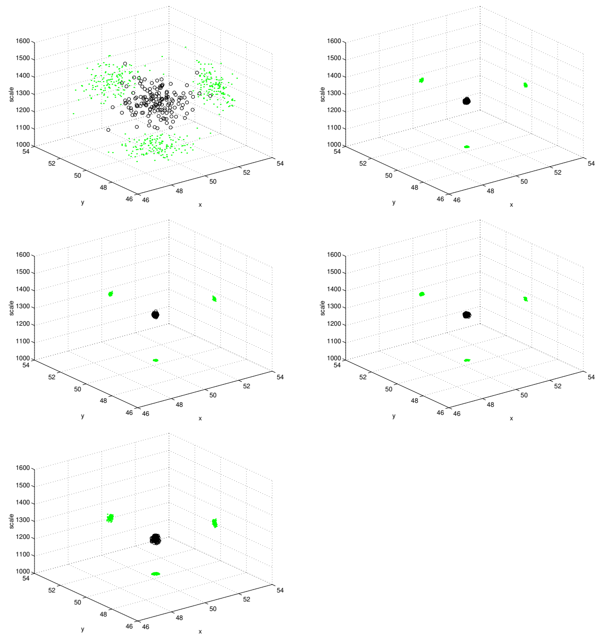

Figure 9: Estimates with replacement noise, disc target. Top row shows diffusion and m-sieve, middle row shows N-sieve and M-sieve and lower row shows an opening sieve

Test images, for example Figure 10, are generated by placing discs of differing area onto a background of amplitude 32. Each 100 by 100 pixel image contains five discs. The target is the only constant image feature. It has an area of 901 pixels and an amplitude of 47. The four cluttering discs, areas 49, 253, 481 and 709 pixels, are selected at random (with replacement). They are placed at random, have random intensity (uniform range 0 to 255) and randomly occlude one another.

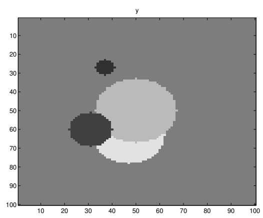

Figure 10: Synthetic occlusion image. In this particular example, the central target partially occludes and is partially occluded by other objects.

Were the six objects not to have occluded one another, the sieve would produce six granules at the true positions of the objects in scale and space. With occlusion the estimates are degraded. One may judge the degradation by repeatedly generating test images and plotting the granule positions, as was done to produce Figures 8, 9. This is shown in Figure 11

In Figure 11 the estimates are plotted as circles and their projections on the position and scale planes are plotted as grey dots. The grey dots show a clustering. There are bands corresponding to the scales of the clutter objects and there is a cluster at the centre corresponding to the 901 pixel target.

Repeating this experiment using a diffusion system is more complicated since one has to track extrema through scale-space as shown in Figure 12. Scale-space extrema tracking in two-dimensions needs careful

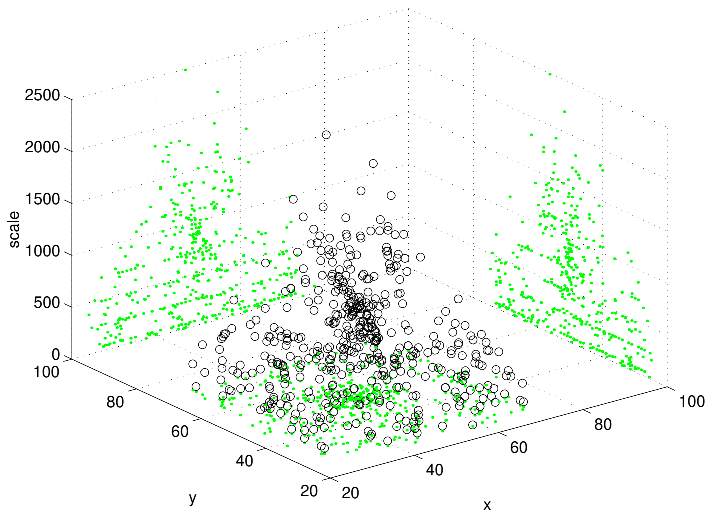

Figure 11: Estimates, obtained from an m- sieve, of the scale and position of objects in 100 test images.

treatment as insignificant extrema due to finite precision effects need to be distinguished from genuine ones. Objects are denoted by maxima along the scale-space tracks.

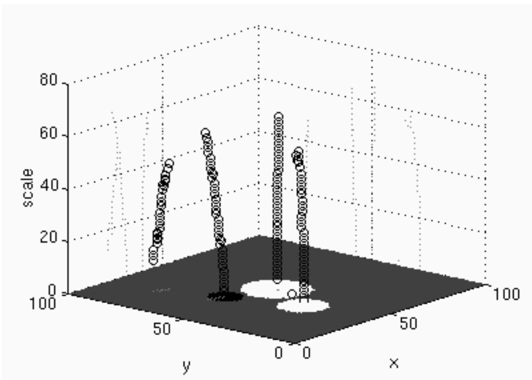

Figure 12: Scale-space extrema tracks from Figure 10. The tracks are sensitive to the boundary conditions.

Figure 13 shows the result of performing this scale-space tracking on 100 images and plotting the maxima on three-dimensional plot as before.

The results from the diffusion system are very confused. The detail at fine scale represents extrema associated with the edges of the disc (no anti-aliasing was used when generating these images) and is to be expected. However, at larger scale very little structure is apparent despite the structure of the input image.

This is a stylised problem which could be solved easily using a processor that has more prior information. For example, if it is assumed that the shape of the target image is known. That is not the point. What is required is a primary vision system that obeys Marr's "principle of least committment" [54] but preserves as much structure as possible. It appears that the sieve preserves structure and does so more robustly than the diffusion system.

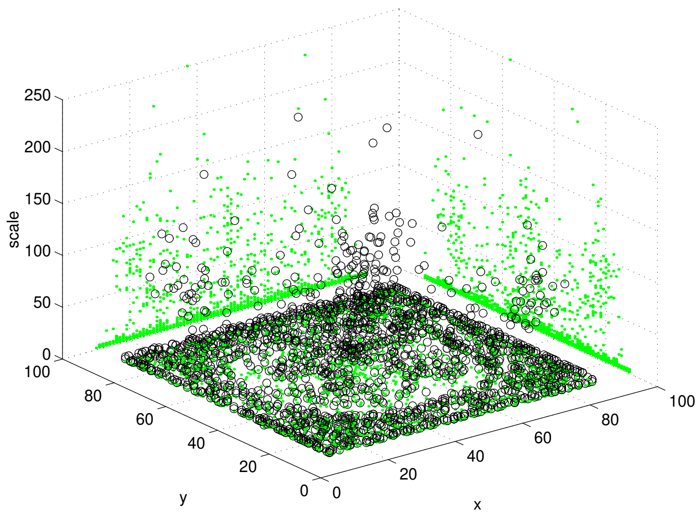

Figure 13: Scale-space extrema tracks from Figure 10. The tracks are sensitive to the boundary conditions.

6 An application

The sieve is a primary vision system and so has potential in many areas. Initial studies have shown that sieves and related filters have value in image coding [55], object recognition [56], motion [57], image understanding [58] and noise reduction [59].

A recent application has been the use of a one-dimensional sieve to the problem of lip-reading as an element of an audio-visual speech recognition system. Here we illustrate the use of an area sieve to address the same problem. The left panel of Figure 14 shows the results. At the bottom of the figure is the audio signal for the spoken letters 'D', 'G' and 'M' and at the top are samples from the corresponding 25 frames-per-second greyscale image sequence. These images show the lips in their quiescient state and in the middle of enunciating the three letters. Each image in the 140 image sequence has been processed with four-connected opening, closing, m-, M- and N-sieves (panels two to six) and the number of granules in each channel is plotted as a greyscale value. We call these area-sieve spectrograms. They form the raw data from which feature vectors are drawn for subsequent analysis. In the top plot (due to openings) one sees that the dark line, associated with the dark area surrounded by the lips, increses scale as the letters are spoken – the dark areas increase as the mouth opens. The same effect can be seen, but to a much lesser degree in the area-sieve spectrogram for closings shown underneath. The pattern is much less pronounced for m-, M- and N-sieves since these processors track the area of both dark and light regions. For all of these letters the mouth opens (an increased number of dark regions) to reveal teeth (an increased number of light regions).

The right panel of Figure **??** indicates the interchangeability of opening and closing sieves. In Figure **??** the colourmap of the images has been inverted and the outputs of the open and close sieves have been interchanged compared to that of the left panel of Figure 14. The outputs from the m-, M- and N-filter are little affected. It is this robustness that makes the m-, M- and N-sieves preferable in most applications and initial experiments suggest that this is indeed the case for this application.

It is worth emphasising that the area-sieve spectrograms are insensitive to shape. As a result, these data complement shape-sensitive length-sieve spectrograms that we have in previous work [58]. We expect that practical systems will use a combination of the two.

7 Concluding remarks

As contenders for primary vision systems, sieves have some advantages. They use some of the recent results from mathematical morphology to form a processor that satisfies many of the desirable scale-space axioms and robustly separates noise and clutter according to scale.

References

[1] J. A. Bangham, R. W. Harvey, P. D. Ling, and R. V. Aldridge. Nonlinear scale-space in many dimensions. In *European Conference on Computer Vision*, volume 1, page 186:192. Kluwer, 1996.

[2] J. Andrew Bangham, Richard Harvey, and Paul D. Ling. Morphological scale-space preserving transforms in many dimensions. *Journal of ELectronic Imaging*, 5(3):283–299, July 1996.

[3] A. P. Witkin. Scale-space filtering. In *8th Int. Joint Conf. Artificial Intelligence*, pages 1019–1022. IEEE, 1983.

[4] J. J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.

[5] J. Babaud, A. P. Witkin, M. Baudin, and R. O. Duda. Uniquenss of the gaussian kernal for scale-space filtering. *IEEE Transactions on Patern Analysis and Machine Intelligence*, Vol. 8:pp 26–33, 1986.

[6] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, July 1990.

[7] Lindeberg Tony. *Scale-space theory in computer vision*. Kluwer, 1994. ISBN 9-7923-9418-6.

[8] G.N.Watson. *A treatise on the Theory of Bessel Functions (2nd Ed)*. Cambridge University Press, Cambridge, 1958.

[9] L. M. Lifshitz and S. M. Pizer. Multi-resolution hierarchical approach to image segmentation based on intensity extrema. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(6):529–540, June 1990.

[10] J. Andrew Bangham, Paul W. Ling, and Richard Harvey. Nonlinear scale-space causality preserving filters. *IEEE Transactions in Pattern Analysis and Machine Intelligence*, 18:520–528, 1996.

[11] J. A. Bangham, P. Ling, and R. Young. Multiscale recursive medians, scale-space and sieve transforms with an inverse. *IEEE Transactions on Image Processing*, 5:1043–1047, 1996.

[12] P. T. Jackway. Morphological scale-space. In *Proc. 11th IAPR International conference on pattern recognition*, pages C252–255. IEEE, 1992.

[13] Paul T. Jackway and Mohamed Deriche. Scale-space properties of the multiscale morphological dilation-erosion. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18:38:51, 1996.

[14] M. H. Chen and P. F. Yan. A multiscale approach based upon morphological filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:694–700, 1989.

[15] P. F. M. Nacken. Openings can introduce zero crossings in boudary curvature. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16:656–658, 1994.

[16] J. Serra. *Image Analysis and Mathematical Morphology*, volume 1. Academic Press, 1982.

[17] J. Serra. *Image Analysis and Mathematical Morphology*, volume 2. Academic Press, 1988.

[18] J.Serra. Introduction to mathematical morphology. *Computer vision, graphics and image processing*, 35:283–305, 1986.

[19] G.Matheron. *Random sets and integral geometry*. Wiley, 1975.

[20] H. Blum. Biological shape and visual science (part 1). *J. Theoret. Biol.*, pages pp 205–287, 1967.

[21] P. Soille and L. Vincent. Determining watersheds in digital pictures via flooding simulations. In M. Kunt, editor, *Visual Communications and Image Processing '90*, volume SPIE-1360, pages 240–250, 1990.

[22] F. Meyer and S. Beucher. Morphological segmentation. *J. Visual Communication and Image Representation*, 1:21–46, 1990.

[23] S. R. Sternberg. Grayscale morphology. *Computer Vision, Graphics and Image Processing*, 35:333–335, 1986.

[24] D. Schonfeld and J. Goutsias. Optimal morphological filters for pattern restoration. In *Proc. SPIE*, volume 1199, pages 158–169, 1989.

[25] P. Salembier and M. Kunt. Size sensitive multiresolution decomposition of images with rank order based filters. *Signal Processing*, 27:205–241, 1992.

[26] J. Serra and P. Salembier. Connected operators and pyramids. In *Proceedings of SPIE Conference on Image Algebra and Mathematical Morphology, Volume 2030*, pages 65–67, 1993.

[27] P.Salembier and J.Serra. Flat zones filtering, connected operators and filters by reconstruction. *IEEE Transactions on Image Processing*, 8(4):1153–1160, AUGUST 1995.

[28] I. Pitas and A. N. Venetsanopoulos. Morphological shape decomposition. *IEEE Trans. on Pattern Anal. and Machine Intell.*, vol. 12:pp 38–45, 1990.

[29] P. Maragos. Pattern spectrum and multiscale shape representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:701–716, 1989.

[30] R. M. Haralick, P. L. Katz, and E. R. Dougherty. Model-based morphology: the opening spectrum. *Graphical models and image processing*, 57:1–12, 1995.

[31] Luc Vincent. Fast opening functions and morphological granulometries. In *SPIE*, volume 2300, page 253:267, 1994.

[32] J. A. Bangham. Data-sieving hydrophobicity plots. *Anal. Biochem*, 174:694–700, 1988.

[33] J. A. Bangham. Properties of an improved variant on the median filter. In *in: E. Arikan, ed., Communication, Control and Signal Proc.*, volume Elsevier press, pages pp 1591–1597, 1990.

[34] L. Vincent. Morphological area openings and closings of greyscale images. In *Workshop on Shape in Picture*. NATO, 1992.

[35] Corinne Vachier and Fernand Meyer. Extinction value: a new measurement of persistence. In Ionas Pitas, editor, *Proceedings of 1995 IEEE Workshop on nonlinear signal and image processing*, volume 1, pages 254–257, JUNE 1995.

[36] J.Serra and P.Salembier. Connected operators and pyramids. In *Proceedings SPIE*, volume 2030, pages 65–76, 1994.

[37] Luc Vincent. Grayscale area openings and closings, their efficent implementation and applications. In Jean Serra and Phillipe Salembier, editors, *Proceedings of the international workshop on mathematical morphology and its applications to signal processing*, pages 22–27, May 1993.

[38] J. A. Bangham. Median and morphological scale space filtering and zero-crossings. In *Proc. SPIE Nonlinear signal processing symp.*, volume 2180, page 90:98, 1994.

[39] J. W. Tukey. *Exploratory Data Analysis.* Addison-Wesley, London, 1971.

[40] Gallagher N C. and Wise G N. A theoretical analysis of the properties of median filters. *IEEE Trans. Acoust. Speech, Signal Processing*, 21:pp 1136–1141, 1981.

[41] M. P. McLoughlin and G. R. Arce. Deterministic properties of the recursive separable median filter. *IEEE Acoustics, Speech and Signal Processing*, 35:98–108, 1987.

[42] E. J. Coyle, J. H. Lin, and M. Gabbouj. Optimal stack filtering and the estimation and structural approaches to image processing. *IEEE Trans. Acoust. Speech, Signal Processing*, 37:2037–2066, 1989.

[43] Jack M. Sacks. Target discrimination utilizing median filters, 1984. US Patent number, 4,603,430.

[44] P. Maragos and R. W. Schafer. Morphological filters - part 1: Their set theoretic analysis and relations to linear shift-invariant filters. *IEEE Trans. Acoust, Speech, Signal Processing*, 35:1153–1169, 1987.

[45] J. A. Bangham. Properties of a series of nested median filters, namely the datasieve. *IEEE Trans. Signal Processing*, vol. 41:pp 31–42, 1993.

[46] S.M. Smith. Flexible filter neighbourhood designation. In *Proc. 13th Int. Conf. on Pattern Recognition*, volume 1, pages 206–212, 1996.

[47] L. Vincent. Graphs and mathematical morphology. *Signal Processing*, 16:365:388, 1989.

[48] H.J.A.M.Heijmans, P.Nacken, A.Toet, and L.Vincent. Graph morphology. *Journal of Visual Computing and Image Representation*, 3(1):24–38, March 1992.

[49] J. A. Bangham. Properties of a series of nested median filters, namely the datasieve. *IEEE Trans. Signal Processing*, 41:pp 31–42, 1993.

[50] J. A. Bangham, S. J. Impey, and F. W. D. Woodhams. A fast 1d sieve transform for multiscale signal decomposition. In *EUSIPCO*, 1994.

[51] Robert A. Boie and Ingemar J. Cox. An analysis of camera noise. *IEEE Trans. Pattern analysis and machine intelligence*, 14:671:674, 1992.

[52] Glenn E. Healey and Raghava Kondepudy. Radiometric ccd camera calibration and noise estimation. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 16:267:275, 1994.

[53] J. Andrew Bangham, Paul Ling, and Richard Harvey. Scale-space from nonlinear filters. In *Proceedings of the First International Conference on Computer Vision*, pages 163–168, 1995.

[54] D. Marr. *Vision*. W. H. Freeman and Company, New York, 1982.

[55] Olivier Egger, Wei Li, and M. Kunt. High compression image coding using an adaptive morphogical subband decomposition. *Proc. IEEE*, 83:272:287, 1995.

[56] J. A. Bangham, T. G. Campbell, and R. V. Aldridge. Multiscale median and morphological filters used for 2d pattern recognition. *Signal Processing*, 38:387–415, 1994.

[57] P. Salembier, A. Olivieras, and L. Garrido. Motion connected operators for image sequences. In *Signal Processing VIII Theories and applications*, volume II, pages 1083–1086, 1996. ISBN 88-86179-83-9.

[58] I. Matthews, J.A. Bangham, and S. Cox. Audio-visual speech recognition using multiscale nonlinear image decomposition. In *International Conference on Spoken Language Processing, ICSLP'96*, pages 38–41, Philadelphia, USA, October 1996.

[59] S. J. Impey and J. A. Bangham. Variations in the morphological 'top hat' transform to improve the statistical robustness for analysing pigmented patches on flower petals. In *Proc. SPIE Nonlinear signal processing symp.*, volume 2180, pages 262–267, 1994.

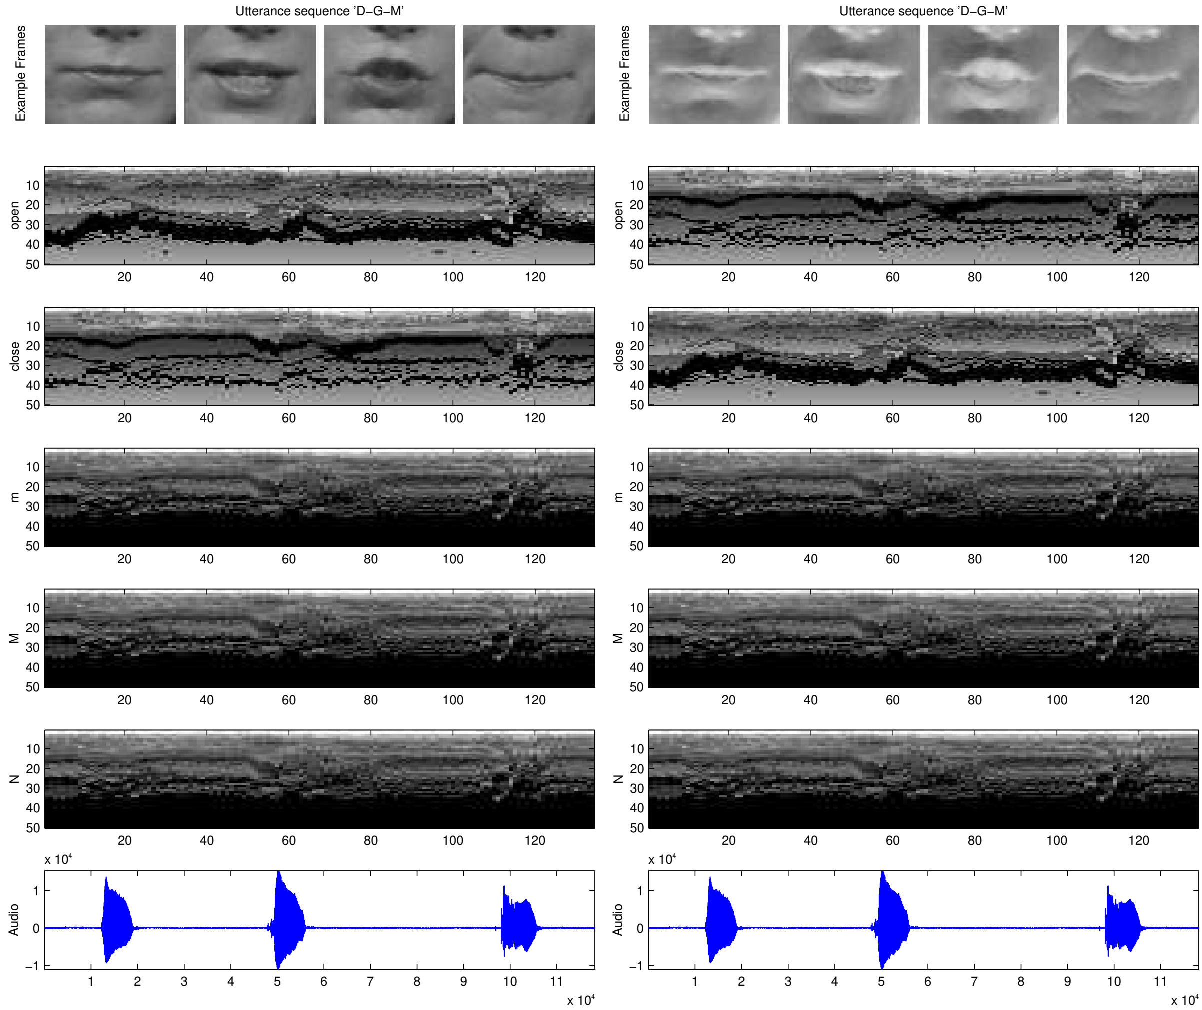

Figure 14: Left panel: tracking moving features using, from panels two to six, sieves of openings, closings, m-, M- and N-filters. The ordinate is frames and the abscissa is channel number. Intensity is granule-count. The bottom panel shows the audio signal. The right shows the inverse.