

Colour Morphological Scale-spaces from the Positional Colour Sieve

David Gimenez and Adrian N. Evans
Dept of Electronic and Electrical Engineering
University of Bath
Bath, BA2 7AY
United Kingdom
{eepdg , A.N.Evans}@bath.ac.uk

Abstract

Morphological sieves are a popular tool for scale-space image analysis. Recent work has considered the development of colour and multichannel sieves, and their application to image segmentation. This paper proposes a new colour sieve based on the geometry of the local convex hull, providing a more flexible approach to extrema definition. Results show the new approach to have a similar segmentation performance to existing colour sieves but with an improved performance in terms of noise reduction.

1. Introduction

Scale-spaces formed from greyscale images are an important tool for hierarchical image analysis. Traditionally, the scale-spaces are linear and are created by employing the diffusion equation [11]. More recently, the application of a series of structuring elements of increasing scale have been used by several researchers to form morphological scale-spaces, for example see [10] and [14]. However, in common with linear scale-spaces, scale-spaces based on structural morphology do not obey the property of strong causality with the result that new boundaries can be created, and the position of edges can drift, with increasing scale [1]. Alternatively, morphological sieves based on connected operators possess the property of strong causality and, as they employ area operators, require no a priori knowledge of the shape of image objects [2, 1]. In addition, sieves have low computational complexity and are robust to noise [9].

Morphological sieves are useful for many image processing applications. At small scales they can be used to remove image noise [19, 18]. At larger scales, sieves produce regions that show correspondence with image objects and have been formally related to segmentation algorithms based on region merging/classification [6]. Consequently,

they have been employed in applications such as segmentation and classification [16, 1]. Although colour plays an important role in the segmentation process [13], the development of multichannel sieves is problematic as their underlying morphological operations require regional maxima and minima to be identified and processed. As no unique ordering for multivariate data exists, these operations cannot easily be extended to colour images.

The convex colour sieve (CCS) [7] and the vector area morphology sieve (VAMS) [5] were proposed independently in 2003 and address the problem of extending morphological sieves to colour images. These two approaches, although differing in the details, have algorithms that essentially follow the same steps and the main difference in their performance results from their approaches to defining extrema [8]. The CCS forms a convex hull from each pixel and its connected neighbours and then defines a pixel as extreme (resp. non-extreme) if it lies on the edge (resp. interior) of the hull. The VAMS first forms a scalar image in which the value of each pixel is the mean aggregate distance to its connected neighbours, assessed using a norm, and then identifies extrema as the maxima in the scalar surface. Recent analysis has shown that the binary decision used by the CCS produces a very high proportion of extrema and an aggressive filtering action. In contrast the VAMS has fewer extrema and hence a lower computational cost. However, a lower proportion of extrema also results in less image simplification for a given scale. The VAM Open-Close Sieve (VAMOCs) increases the aggressiveness of the VAMS by also sieving the minima in the scalar surface produced by the VAMS [8]. These minima correspond to “nearly flat” regions and their inclusion in the merging process was shown to produce an improved segmentation performance.

As the surface produced by the CCS is binary every pixel is either extreme or non-extreme and the inclusion of minima would result in every pixel participating in the merging process. This paper proposes a new sieve, termed the Positional Colour Sieve (PCS), which addresses this prob-

lem by deriving a scalar surface from the geometry of the local convex hull. This process reduces the proportion of pixels that are categorised as extreme and also allows the closing operation of the VAMOCS to be incorporated into the PCS structure, giving the Positional Colour Open-Close Sieve (PCOCS).

This paper is organised as follows: Section 2 reviews colour sieves and describes the PCS algorithm. A performance evaluation is presented in section 3, including noise reduction and a quantitative evaluation of the segmentation performance using the Berkeley segmentation dataset. Finally, conclusions are given in section 4.


2. Colour sieves

A main steps of a generalised colour sieve algorithm are given in [8] as:

1. Identify all extreme regions;
2. Merge all scale 1 extrema regions with their nearest neighbour;
3. Repeat steps 1 and 2 until no extrema are found at current scale;
4. Repeat steps 1 to 3 with increasing scale until only 1 region remains.

Compared with its better known greyscale counterpart that processes the maxima and minima separately, colour sieves simply process extrema as they cannot distinguish between maxima and minima. The merging process in step 2 is analogous to that of the greyscale sieve and changes the colour of each extreme region to that of its closest neighbour, as assessed using the Euclidean distance. However, the merging is not guaranteed to produce a non-extreme region and, in addition, can create new extrema close to the merged regions. As a consequence, in colour sieves the addition of step 3 is necessary to ensure idempotence. Steps 1 and 2 are therefore repeated at each scale and, as the number of extrema found depends on the definition of extrema adopted, this is a critical factor in the performance of the sieve.

The CCS defines extremeness in terms of membership of its local convex hull: if a point is on the edge of the local convex hull consisting of the point and its connected neighbours, then it is extreme. This method reduces to a combined opening and closing for greyscale images and can be extended to any number of dimensions. In addition, its extremum definition is also invariant to rotation and linear transformations of the axes. The main disadvantage of this method is the number of extrema: a n -dimensional hull needs at least $n+1$ points to be non-degenerate, which means a high proportion of the points are classified as extreme, giving an aggressive sieving action. For example, figures 1(a) and (b) show a complex image and the convex



(a) Complex image

(b) convex hull for $0 + 0i$ region

17.00	29.00	29.00	32.50	18.00
27.00	27.00	29.00	81.00	26.00
39.00	38.00	38.00	43.00	26.00
39.00	66.67	66.67	43.00	11.00
38.00	38.00	66.67	43.00	11.00

(c) VAMS scalar image using L_1 norm

5.6009	4.8543	4.8543	0.0040	3.4123
3.1419	3.1419	4.8543	5.0654	4.5705
5.0680	0.1045	0.1045	4.3594	4.5705
5.0680	5.0877	5.0877	4.3594	3.3012
5.3871	5.3871	5.0877	4.3594	3.3012

(d) PCS scalar image

Figure 1. Colour sieves example using 8nn connectivity. Extrema marked in bold in (c) and (d).

hull for the $0 + 0i$ region. As $0 + 0i$ lies inside the hull it is not extreme. However, using this definition, there is only one other non-extreme region in the image. As well as increasing the processing load, classifying the majority of pixels as extreme does not fit well with an intuitive interpretation of extrema as outliers.

An alternative definition of extrema is provided by the VAMS. Here, extrema are the maxima of a scalar surface in which each region is assigned the average aggregate distance from each pixel in the region to its neighbours. In figure 1(c) this results in only 2 extreme regions, whose values are clearly different from their neighbours. Minima in the scalar image correspond to regions that are closer to their neighbours than their connected neighbours are to theirs and the VAMOCS also identifies and processes these, producing an improved segmentation performance [8]. In figure 1(c) the 4 corners are minimum regions, with similar values to their surrounding points.

The PCS aims to derive a scalar surface from the geometry of the local convex hull. This approach has several advantages: unlike the VAMS, the extrema will not be affected by rotation of the axes, the number of extrema will be more in line with expectations, and it will also enable minima to be identified and processed. In the approach adopted here, the value of each point in the scalar image depends on the following criteria:

1. The angle with neighbouring points. For points on the exterior of the hull the angle is measured using




Figure 2. PCS example for $0+0i$ region from figure 1(a).

neighbouring exterior points. For points inside the hull the maximum angle to neighbouring exterior points is used. When all points are in a plane the maximum angle is π so, to avoid bias, all angles are scaled to a $0 - \pi$ range. As the angle reduces the point becomes more extreme, so it is subtracted from the maximum angle to give an increasing measure.

2. The normalised distance, measuring the extremeness of the position. The first step in finding this measure is to define 3 points, the original point (o), the point furthest from o (f_1) and the point furthest from f_1 (f_2). The normalised distance is then given by $(of_1 + of_2)/f_1f_2$ or $2of_1/f_1f_2 = 2$ when o and f_2 are the same point. The measure is shifted to a range of 0–1 by subtracting 1.

A measure of extremeness is then given by the product of the angle and the normalised distances measures. This measure overcomes the problem of using the angular measure in isolation, as points that are close to coplanar can be differentiated by the distance measure. Similarly, points that are located amid similar neighbours are less likely to be classified as extreme. Finally, a third criterion is added to ensure that points within the hull are not classified as extreme:

3. Whether the point is within or on the edge of the hull. As extrema are not expected to occur within the hull, the scalar product of measures 1 and 2 is augmented by π for points on the edge of the hull. Points coplanar or collinear with edges are considered to be in the hull.

This process is illustrated for the $(0 + 0i)$ region previously considered (see figure 2):

- The maximum angle is $\arccos(-14/\sqrt{5 * 49}) = 2.678$ rad (using points $(0 - 7i)$ and $(-1 + 2i)$), giving a measure of $\pi - 2.678 = 0.464$ rad.

scale	VAMS	CCS	GSAOCS	PCS	PCOCS
1	<u>2209</u>	<u>33616</u>	<u>4062</u>	<u>8889</u>	<u>17433</u>
	41207	41207	41207	41207	41207
2	<u>1054</u>	<u>12852</u>	<u>1973</u>	<u>3602</u>	<u>5695</u>
	37199	16402	36553	30775	22189
10	<u>151</u>	<u>2211</u>	<u>424</u>	<u>667</u>	<u>134</u>
	27547	2878	28457	15548	5020
50	<u>21</u>	<u>456</u>	<u>110</u>	<u>918</u>	<u>189</u>
	19660	572	20697	7546	1132
100	<u>8</u>	<u>223</u>	<u>63</u>	<u>71</u>	<u>96</u>
	16102	279	17040	5435	559
500	<u>1</u>	<u>44</u>	<u>11</u>	<u>13</u>	<u>19</u>
	8024	53	9114	2498	97
1000	<u>1</u>	<u>23</u>	<u>7</u>	<u>7</u>	<u>12</u>
	6871	25	6150	1478	40

Table 1. (Number of extreme regions)/(total number of regions) for Lily image.

- Point $f_1 = (9 + 3i)$ and is 9.49 from point o ($0+0i$), and point $f_2 = (0 - 7i)$ is 13.45 from f_1 and 7.00 from point o .
- Since $f_1f_2 \neq of_1$, points o and f_2 are not the same there is no need to scale d_2 . The normalised distance $(9.49 + 7)/13.45 - 1 = 0.225$
- The final measure is $0.464 \times 0.225 = 0.1045$.

Repeating the above process for all regions in figure 1(a) produces the PCS scalar image shown in figure 1(d) in which there are 3 extrema (maxima) and also 3 minima.

Figure 3 shows the distribution of initial extrema at selected scales for the test images Lily and Lenna. The fraction of Lily's regions that are extreme are given in table 1 and were similar to the fractions obtained for Lenna. The PCS initially has roughly four times the number of extreme regions as the VAMS but only a quarter the number of the CCS; this factor of four occurs by chance. As each extremum has to be merged with its closest neighbour, the aggressiveness of the PCS's sieving action should therefore lie somewhere between the VAMS and the CCS. The fractions for the PCOCS in table 1 also includes minima, which relate to nearly flat regions. The PCOCS also processes the minima regions in the scalar surface and, as scale increases the total number of regions reduces through the merging process; this explains why the PCOCS achieves a more rapid reduction in total regions, and hence more image simplification.

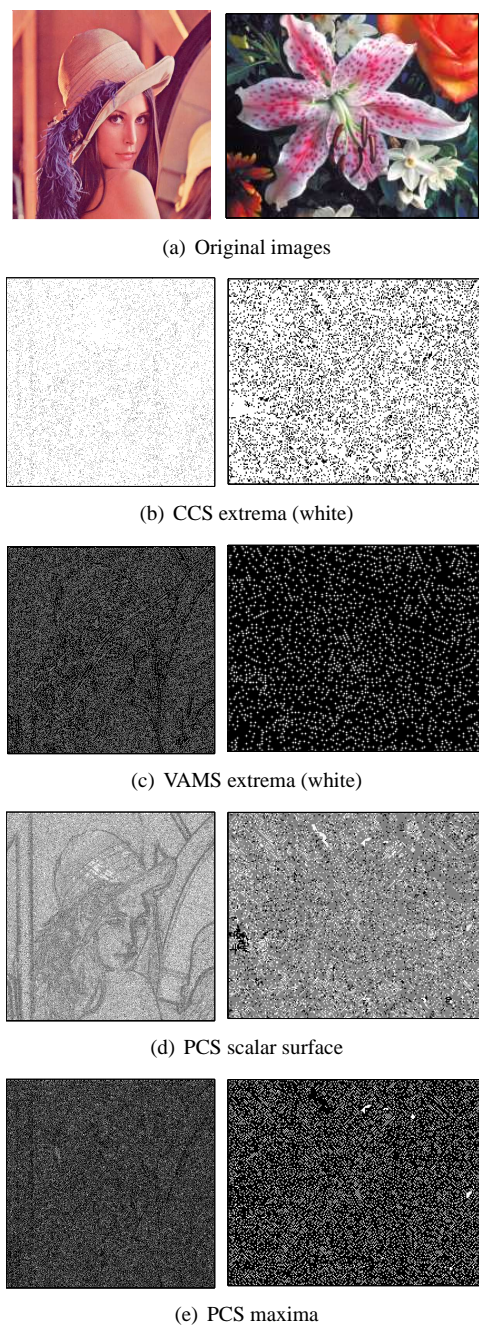


Figure 3. Colour sieve extrema for Lily and Lenna images.

3. Experimental results

The PCS was evaluated on a number of natural images and the results for the Lily (186×230 pixels) and Lenna (512×512 pixels) are representative of those obtained. For






Figure 4. Sieve results for Lily image. Area = 10 (left), 100 (centre) and 1000 (right).

comparison, the images were also sieved with the CCS, the VAMS and the VAMOCS. All colour sieves used 4 nearest neighbour connectivity and the L_2 norm. In addition, results for the greyscale area-open-close sieve (GSAOCS) were obtained, using brightness as the ordering factor.


The results of sieving the Lily and Lenna images are presented in figures 4 and 5. All sieves preserve significant boundaries through scale. An indication of the aggressiveness of the sieving action is given by the amount of image simplification at each scale and the results show broad agreement with the proportion of extrema detailed in the previous section; those sieves that have the highest proportion of extrema produce the most image simplification at a given scale. Figures 4 and 5 show the aggressiveness of the




(a) PCS




(b) PCOCS



(c) VAMS



(d) CCS



(e) GSAOCS

Figure 5. Sieve results for Lena image. Area = 100 (left), 1000 (centre) and 10000 (right).

PCS falls between the VAMS and CCS and a quantitative evaluation of the segmentation performance is presented in section 3.2.

3.1. Noise Reduction

To evaluate the ability of the sieves to remove noise, the Lily test image was corrupted with light (1%) and moderate

Method	-	Imp1	Imp2	Gaus1	Gaus2
-	0	1.17	10.77	4.56	26.13
VAMS	0.83	0.38	1.40	3.18	9.57
CCS	1.15	0.49	1.64	3.10	9.69
GSAOCS	0.40	0.21	1.80	3.45	17.95
PCS	0.52	0.25	1.19	2.44	8.52

Table 2. NMSE ($\times 10^{-2}$) for colour sieves applied to Lily image.

Method	-	Imp1	Imp2	Gaus1	Gaus2
-	0	1.05	9.95	17.56	35.12
VAMS	1.55	0.66	2.07	13.01	19.56
CCS	2.57	1.30	2.59	11.83	16.06
GSAOCS	1.03	0.49	2.49	15.01	30.30
PCS	1.64	0.66	1.86	10.20	16.75

Table 3. MCRE ($\times 10^{-2}$) for colour sieves applied to Lily image.

(10%) impulsive noise and with moderate ($\sigma^2 = 10^3$) and heavy ($\sigma^2 = 10^4$) Gaussian noise. In colour image processing, the normalised mean square error (NMSE) and the mean chromaticity error (MCRE) are two widely used metrics that provide objective error measures [17]. The NMSE is given by

$$NMSE(\mathbf{f}, \bar{\mathbf{f}}) = \frac{\sum_{x=1}^M \sum_{y=1}^N \|\mathbf{f}_{xy} - \bar{\mathbf{f}}_{xy}\|^2}{\sum_{x=1}^M \sum_{y=1}^N \|\mathbf{f}_{xy}\|^2} \quad (1)$$

where M and N are the image dimensions, and \mathbf{f}_{xy} and $\bar{\mathbf{f}}_{xy}$ are the original and processed pixels at location (x, y) respectively. The MCRE is defined as the distance between the intersection points of f_{xy} and \bar{f}_{xy} with the Maxwell triangle. When the triangle is defined on the unit plane this gives

$$MCRE(\mathbf{f}, \bar{\mathbf{f}}) = \left(\sum_{x=1}^M \sum_{y=1}^N ((r_{xy} - \bar{r}_{xy})^2 + (g_{xy} - \bar{g}_{xy})^2 + (b_{xy} - \bar{b}_{xy})^2)^{1/2} \right) / (MN) \quad (2)$$

where r , g and b are the normalised RGB values [3].

The noise-corrupted images were sieved over range of areas and, at each scale, the NMSE found. The minimum NMSE and corresponding MCRE for each sieve are given in tables 2 and 3 respectively. Results show that with no added noise the GSAOCS produces the least image distortion. When noise is added, the PCS has the lowest NMSE and MCRE measures for all noise types except light impulsive noise, where its performance is similar to the GSAOCS.

This shows that the definition of extrema used by the PCS successfully identifies noise points. It was also found that for higher levels of noise the minimum NMSE occurred at larger scales for lower noise levels. The PCOCS was not included in this evaluation as the main benefit of its use of closings is for segmentation rather than noise reduction.

3.2 Image Segmentation

An evaluation of the segmentation performance of the colour sieves was undertaken using the Berkeley database [12], a set of natural images and human ground truth segmentations. To provide a quantitative performance measure precision-recall (P-R) curves of [13] were employed. These are a variant of ROC curves where precision is the normalised number of identified boundary pixels (true positive/total positive) and recall is the proportion of identified pixels or the false-negative rate (true positive/total true positive). These metrics have the advantage of being independent of scale, unlike the ROC curve, and are an alternative to other metrics such as the global coherency error (GCE). They also enable segmentation performance to be characterised by a single figure (the F-measure), given by the harmonic mean of precision and recall.

By increasing the scale until the sieve images contained a fixed number of regions, and varying the number of regions, P-R curves for each sieve were obtained. Figure 6 presents the average P-R curves for 100 images from the dataset with the corresponding F-measures given in table 4. Results show the segmentation performances of the CCS, PCS and PCOCS to be very similar. Although the processing of minima by the PCOCS reduces the number of regions at which the peak F-measure occurs from 300 to 30, the F-measure is unchanged. This is unlike the VAMOCS where, in comparison with the VAMS, an improved F-measure occurs with a reduced number of regions. For comparison, the curve for a random segmentation is also shown. This was generated by marking all regions as extreme and randomly selecting the neighbour to merge with, an approach more appropriate than that described on the web site associated with [13].

Sieving the entire tree to the same level is known to be sub-optimal in terms of segmentation performance but has the advantage of maintaining idempotency. Segmentation can also be achieved by constructing a scale-tree from the sieve decomposition and then pruning the tree appropriately [4]. Therefore, examining the scale trees of figure 7 provides an indication of the sieves' segmentation potential. The CCS gives a balanced tree structure with a low tree depth, with similar numbers of child in each branch and few, if any, leaf nodes off the root. In contrast, the PCS tree resembles that of the VAMS, with a high tree depth and some leaves close to the root as fewer regions are removed at each

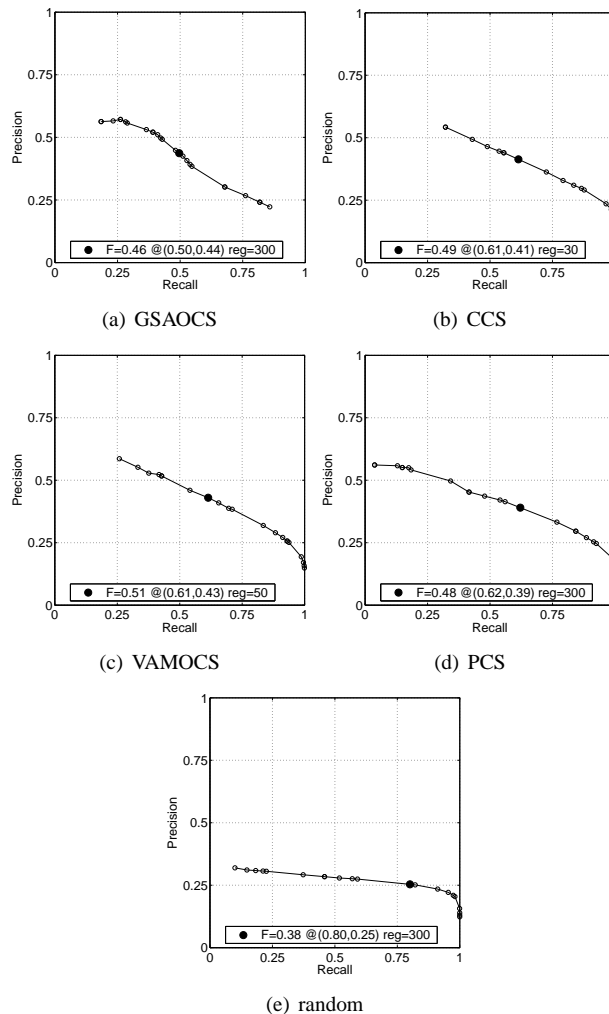


Figure 6. Average Precision-Recall curves for 100 images from the Berkeley Dataset.

step. Another factor which influences the segmentation performance is the merging rule. For the PCS and PCOCS the colour of all extrema was changed to that of their closest neighbour. However, the use of other merging rules, such as those suggested by Salembier and Garrido [15], may provide better results, especially at higher scales.

To provide an indication of the robustness of the segmentations to image noise, the Lily image was corrupted by mixed impulsive and Gaussian noise and sieved as before, see figure 8. Comparing these results with the noise-free ones of figure 4 shows that all sieves are influenced by noise to some degree, even at high scales.

The final aspect to be considered is the computational complexity, as assessed by the processing times. A comprehensive evaluation of these was undertaken in [8] for

Method	F-measure (R,P)	Regions
Human	0.79 (0.70,0.90)	-
GSAOCS	0.46 (0.50,0.44)	300
CCS	0.49 (0.61,0.41)	30
VAMS	0.40 (0.51,0.32)	100
PCS	0.48 (0.65,0.38)	300
PCOCS	0.48 (0.62,0.39)	30
VAMOCS	0.51 (0.61,0.43)	50
Random	0.38(0.80,0.25)	300

Table 4. F-measure and number of regions for P-R curves of figure 6.

VAMS, VAMOCS, CCS and GSAOCS. The PCS implementation used here is 2-3 times slower than the CCS, due to the multiple hull calculations. One reason for this is that the local convex hull is constructed at each scale rather than updated. Storing and updating the hull has the potential of reducing the complexity of the PCS, PCOCS and the CCS.

4. Conclusions

A new colour morphology sieve has been described and evaluated. The PCS uses the geometry of the local convex hull to construct a scalar surface in which extrema can be identified. Compared to the CCS, fewer extreme regions result from this approach to extrema definition. In addition, the PCS allows the use of closings to flatten homogenous regions. Result show that although the number of extrema is more realistic, this does not translate into an improved segmentation performance and this is supported by an analysis of the scale-tree. Furthermore, although the use of closings reduced the number of regions it did not result in the same benefit as for the VAMOCS.

Notwithstanding, the PCS produced the best overall performance in terms of noise removal indicating that there is some merit in its extremum definition. Areas of further work include the development of more advanced tree pruning and rules for region merging algorithms, with the aim of further improving the segmentation performance.

References

- [1] S.T. Acton and D.P. Mukherjee. Scale-space classification using area morphology. *IEEE Trans. Image Processing*, 9(4):623–635, April 2000.
- [2] J.A. Bangham, R.Harvey, P.D. Ling, and R.V. Aldridge. Morphological scale-space preserving transforms in many dimensions. *Journal of Electronic Imaging*, 5(3):283–299, July 1996.




Figure 7. Colour sieve scale-trees for simplified Lily image (top).

- [3] D.R. Charles. *Algorithmic and learning based filtering techniques with application to colour image noise suppression and enhancement*. Phd thesis, Machine Vision Group, Department of Physics, Royal Holloway, University of London, 2003.
- [4] R.J. Chen and B.C. Chieu. Three-dimensional morphological pyramid and its application to color image sequence coding. *Signal Processing*, 44(2):163–180, 1995.
- [5] A.N. Evans. Vector area morphology for motion field smoothing and interpretation. *IEE Proc. Vision, Image and Signal Processing*, 150(4):219–226, August 2003.
- [6] D. Gatica-Perez, C. Gu, M.T. Sun, and S. Ruiz-Correa. Extensive partition operators, gray-level connected operators, and region merging/classification segmentation algo-



(a) Lily image with Gaussian ($\sigma = 32$) and 10% impulsive noise.




(b) PCS



(c) PCOCS



(d) VAMS



(e) CCS



(f) GSAOCS

Figure 8. Sieve results for Lily image corrupted by mixed noise. Area = 10 (left), 100 (centre) and 1000 (right).

rithms: theoretical links. *IEEE Trans. Image Processing*, 10(9):1332–1345, September 2001.

- [7] S. Gibson, R. Harvey, and G.D. Finlayson. Convex colour sieves. In *Proc. 4th International Conference on Scale Space Methods in Computer Vision*, volume LNCS 2695, pages 550–563, 2003.
- [8] D. Gimenez and A.N. Evans. Colour morphological scale-spaces for image segmentation. In *Proc. British Machine*

Vision Conference, pages 909–918, September 2005.

- [9] R. Harvey, A. Bosson, and J.A. Bangham. The robustness of some scale-spaces. In *Proc. British Machine Vision Conference*, 1997.
- [10] P.T. Jackway and M. Deriche. Scale-space properties of the multiscale morphological dilation erosion. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18(1):38–51, January 1996.
- [11] T. Lindeberg. *Scale-space theory in computer vision*. Kluwer, 1994.
- [12] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.
- [13] D.R. Martin, C.C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, colour and texture cues. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 26(5):530–549, May 2004.
- [14] K.R. Park and C.N. Lee. Scale-space using mathematical morphology. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 18(11):1121–1126, November 1996.
- [15] P. Salembier and L. Garrido. Binary partition tree as an efficient representation for image processing, segmentation and information retrieval. *IEEE Trans. Image Processing*, 9(4):561–576, April 2000.
- [16] P. Salembier and J. Serra. Flat zones filtering, connected operators, and filters by reconstruction. *IEEE Trans. Image Processing*, 4(8):1153–1160, August 1995.
- [17] P.E. Trahanias and A.N. Venetsanopoulos. Color edge detection using vector order statistics. *IEEE Trans. Image Processing*, 2(2):259–264, Apr. 1993.
- [18] N. Young and A.N. Evans. Image noise reduction using attribute morphology filters. In *Proc. 6th IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing*, June 2003.
- [19] N. Young and A.N. Evans. Psychovisually tuned attribute operators for pre-processing digital video. *IEE Proc. Vision, Image and Signal Processing*, 150(4):277–286, October 2003.