Citation for published version:

Evans, AN 2010, 'Color area morphology scale-spaces', *Advances in Imaging and Electron Physics*, vol. 160, pp. 35-74. [https://doi.org/10.1016/S1076-5670\(10\)60002-X](https://doi.org/10.1016/S1076-5670(10)60002-X)

DOI:

[10.1016/S1076-5670\(10\)60002-X](https://doi.org/10.1016/S1076-5670(10)60002-X)

Publication date:

2010

[Link to publication](#)

©Elsevier.

Evans, A. N., 2010. Color area morphology scale-spaces. *Advances in Imaging and Electron Physics*, 160, pp. 35-74. The definitive version is available from: [http://dx.doi.org/10.1016/S1076-5670\(10\)60002-X](http://dx.doi.org/10.1016/S1076-5670(10)60002-X)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Colour Area Morphology Scale-Spaces

Adrian N. Evans

*Department of Electronic and Electrical Engineering, University of Bath,
BA2 7AY, United Kingdom*

Contents

1	Introduction	2
2	Area openings and closings	5
3	Area morphology scale-spaces	8
4	Colour connected filters	11
4.1	Colour area morphology scale-spaces	12
4.1.1	Colour Extrema	13
4.1.2	Distance metrics	16
4.2	Convex Colour Sieve	17
4.3	Vector area morphology sieve (VAMS)	19
4.3.1	VAMS parameter evaluation	20
4.4	Vector area morphology open-close sieve (VAMOCs)	26
5	Performance evaluation	29
5.1	Implementation and timings	30
5.2	Application to image segmentation	32
5.3	Robustness to noise	38
6	Conclusions	40
7	Acknowledgements	41

1. Introduction

Multiscale image analysis using morphological scale-spaces has many important applications in computer vision, for example segmentation (Gauch, 1999) and classification (Acton and Mukherjee, 2000). The concept was originally introduced using linear scale-spaces (Witkin, 1983; Koenderink, 1984) in which a succession of images are obtained by blurring the original using Gaussian filters with increasing scale-space parameters. One drawback with this approach is that the locations of image boundaries tend to drift with increasing scale, an effect that can have the undesired result of merging distinct objects and destroying corners (Perona and Malik, 1990). The anisotropic diffusion algorithm of Perona and Malik was an attempt to address this problem by reducing the blurring in the presence of edges (Perona and Malik, 1990).

An alternative approach to Gaussian filters is to employ morphological operators in the process of generating the scale-spaces (Jackway and Deriche, 1996; van den Boomgaard and Smeulders, 1994). Morphological scale-spaces that are based on successive openings and closings using structuring elements of increasing size can suffer from the same drifting of edges through scale that characterises linear scale-spaces (Maragos, 1989; Park and Lee, 1996). In contrast, scale-spaces constructed using area morphology operators do not suffer from this problem. This is because area openings and closings are types of connected operators, that work by merging the regions of constant signal. In common with scale-spaces based on other connected operators such as opening by reconstruction and levelings (Meyer and Maragos, 2000), area morphology scale-spaces do not introduce any new edges with increasing scale. Further, they do not alter any edges present at finer scales until the scale where they are completely removed is reached (Acton and Mukherjee, 2000).

Area openings and closings belong to a class of morphological techniques called connected operators (Salembier and Serra, 1995). A full discussion and introduction to the subject of mathematical morphology is beyond the remit of the chapter and, indeed, is well covered elsewhere. An authoritative description of the mathematical basis of morphology is contained in the books by Serra (Serra, 1982, 1988) and an excellent practical treatment, including example applications, is provided by Soille (Soille, 2003). For a recent update on the advances in mathematical morphology, and in particular those pertaining to the theme of connective seg-

mentation, the reader is referred to a chapter from volume 150 of this publication and the references therein (Serra, 2008).

In common with many other morphological techniques, problems exist when trying to extend connected operators to colour and other multichannel images due to the absence of an unambiguous ordering. Colour morphology has received much attention from the research community and was the subject of a recent review paper by Aptoula and Lefèvre (Aptoula and Lefèvre, 2007). However, for the most part researchers have concentrated on the development of multichannel approaches for structural morphology rather than connected operators, as evidenced by the fact that only three of the 98 references in (Aptoula and Lefèvre, 2007) consider the particular problem of extending connected operators to multichannel images.

The conventional approach to extended morphological operators to colour and other multichannel images is to employ marginal ordering, in which each channel is filtered independently. As marginal processing completely disregards any inter-channel correlation, when the image to be processed is highly correlated it is often necessary to change the colour space or apply a decorrelating transform before the ordering operation. In addition, as marginal processing can alter the spectral composition of an image it can result in edge jitter and create colours that were not in the original image. Despite these problems, marginal ordering does have the advantage of allowing any of the greyscale morphological techniques to be directly applied to multichannel imagery.

In addition to marginal ordering, there are a number of other vectorial ordering techniques that can be employed. These have the advantage of eliminating the possibility of creating new colours, by virtue of treating each pixel as a vector. The ordering schemes can be classified according to the taxonomy of sub-ordering schemes detailed in the classic paper on the ordering of multivariate data by Barnett (Barnett, 1976). In addition to marginal ordering, reduced (or aggregate) ordering, partial ordering and conditional ordering can be employed and the use of all of these for colour morphology has been proposed at some time or other. In particular, there has been much interest in the use of lexicographical ordering, the best known type of conditional ordering, for defining colour morphology operations in hue-based colour spaces (Hanbury and Serra, 2001a,b; Louverdis et al., 2002). Reduced ordering sorts vectors according to scalar values

that are a function of the observations in each channel. Functions that have been investigated include luminance and individual colour channels (Comer and Delp, 1999). A reduced ordering scheme based on principal component analysis has also been proposed (Li and Li, 2004).

This chapter discusses the development of area morphology scale-spaces for colour images. Unlike morphological techniques based on the use of structuring elements, area morphology does not employ erosions and dilations but instead works by using openings and closings. Therefore, the requirements for extending area scale-spaces to colour images are subtly different to those of structural scale-spaces. The specific approach to be detailed is the vector area morphology sieve (VAMS) that first appeared in the literature in 2003 (Evans, 2003a,b). An approach known as the convex colour sieve (CCS) (Gibson et al., 2003b) that is similar in spirit to the VAMS was developed independently around the same time and both approaches were first presented in June 2003 at the IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing (Evans, 2003a) and Scale Space 2003, respectively. Subsequent work has shown that the two techniques have a common algorithm and only differ in the specific implementation of some of the algorithmic steps (Gimenez and Evans, 2005, 2008). The VAMS was first proposed for analysing motion vector fields and has also been applied to colour images. However, as its approach is generic in nature it can, in theory, be extended to images with an arbitrary number of channels.

The format of this chapter is as follows. Greyscale area operations are introduced in section 2 and their use in scale-spaces is described in section 3. Some recently proposed colour connected filters that have been used for colour area morphology scale-spaces are described in more detail in section 4. Section 5 contains an evaluation of the performance of the colour scale-spaces, with particular emphasis on their application to image segmentation. Finally, conclusions are drawn in section 6. The chapter gathers together work from a number of sources and for a bibliography the reader is referred to references (Evans, 2003a,b; Gimenez and Evans, 2005, 2008; Evans and Gimenez, 2008; Gimenez, 2007).

2. Area openings and closings

Openings and closings are two of the fundamental operations in mathematical morphology and, as such, are the building blocks for many other morphological operations (Serra, 1982). The performance of the opening and closing operations depends in a large part on the structuring element used. If some a priori knowledge of the size, shape and orientation of objects within the image is available, then this can be used to select the structuring elements for the filtering operation. An alternative approach that has application when little or no a priori shape information is available is to use an area morphology opening or closing, which remove light or dark structures of a given number of pixels regardless of their shape or orientation.

Area openings and closings belong to the class of morphological operations known as connected operators. The first connected operator to appear in the literature was openings by reconstruction (Klein, 1976). The generalisation of openings by reconstruction to greyscale images and the development of efficient algorithms was proposed by Vincent (Vincent, 1993c). The use of reconstruction criteria to constrain the reconstruction process has also been proposed and the development of a fast approximation recently reported (Wilkinson, 2008). Also belonging to the general class of connected operators are levelings, which filter an image by removing or attenuating its contours (Meyer, 2004; Serra, 2008).

Area openings and closings work by processing (merging or removing) the flat zones in an image. The flat zones of a signal are the largest connected components where the signal is constant (Salembier and Serra, 1995). If S is a subset of image pixels, then two pixels p and q are said to be connected if there exists a path between them consisting entirely of pixels in S . Typically the path is defined using 4 or 8 connectivity. For any pixel p in S , a connected component of S is the set of pixels that are connected to p in S . Therefore any pixel p belongs to a flat zone that consists of the set of connected pixels that have the same signal value. For greyscale images the flat zones are connected components with the same intensity while for colour images they are connected components of constant colour.

The term area opening was introduced by Vincent (Vincent, 1993a) although, as pointed out in (Meijster and Wilkinson, 2002), area openings and closings first appeared in the literature a year earlier as “a new type of opening operators (NOP) and closing operators (NCP)” (Cheng

and Venetsanopoulos, 1992). For a greyscale image I , the area opening (NOP) and area closing (NCP) operations from (Cheng and Venetsanopoulos, 1992) can be denoted γ_λ^a and φ_λ^a , respectively, and defined as

$$\gamma_\lambda^a(I) = \bigvee_{B \in A_\lambda} (I \circ B) \quad (1)$$

and

$$\varphi_\lambda^a(I) = \bigwedge_{B \in A_\lambda} (I \bullet B) \quad (2)$$

where A_λ is the set of connected subsets with area $\geq \lambda$. Examination of (1) reveals that an area opening is the maximum of openings with all possible connected structuring elements of a minimal size (Vincent, 1993b; Cheng and Venetsanopoulos, 1992). It can be considered to be adaptive in that the structuring element adapts itself to the image structure at every location so that the image is changed as little as possible.

Subsequent area opening and closing definitions use a threshold decomposition (Acton and Mukherjee, 2000; Meijster and Wilkinson, 2002) to produce a set of binary images from a greyscale input by applying a threshold t for $t = 0, 1, \dots, L - 1$, where L is the number of discrete grey levels. An area opening can be applied to each binary image by removing all connected components whose area is less than λ . The resulting greyscale image is reconstructed by a stacking operation, such that the output at each pixel position is the maximum threshold at which the area opened binary image is true. Area closings can be implemented by applying the same operations to the complement image.

It is immediately apparent that a naïve implementation of (1) or (2) requires an image to be opened (or closed) using all possible connected structuring elements consisting of λ pixels. The number of structuring elements involved rapidly becomes unfeasible and although a more efficient algorithm was proposed in (Cheng and Venetsanopoulos, 1992) it was not until the pixel-queue algorithm was proposed by Vincent (Vincent, 1993a,b) that the application of area openings and closings to greyscale images became a practical proposition.

Figure 1 compares the operation of an area closing with that of a conventional closing using a structuring element. When the image is closed using a 9×9 structuring element [Figure 1(b)] many of the darker details, for example the boats' masts, are removed. In addition, the size

Figure 1: Boats. (a) 576×720 pixels original image, (b) closing with a 9×9 structuring element, (c) and (d) area opening and closing to area $\lambda = 81$, respectively.

and shape of the structuring element used to probe the image can clearly be seen in the result and this results in the distortion of edge positions. Area closing to an area $\lambda = 81$ [Figure 1(d)] removes the smaller, darker connected components, for example from the beach, without altering the positions of any of the other image boundaries.

As illustrated by Figure 1(c) and (d), area openings and closings work by removing light and dark structures with area less than λ , respectively. This observation provides the basis for Vincent's pixel-queue algorithm, where the light and dark structures to be removed are synonymous with regional maxima and minima respectively. Regional maxima (resp. minima) are connected sets of pixels whose intensity is greater (resp. less) than that of their connected

neighbours. Rather than processing every pixel directly, to perform an area opening the pixel-queue algorithm identifies all regional maxima and places them in a list. Each maximum is then processed by adding the neighbour with the highest grey level and the process repeated until the desired area size has been reached or the region becomes non-maximal. Full details of this algorithm are provided in (Vincent, 1993a,b) and a convenient summary is given in (Meijster and Wilkinson, 2002).

Other efficient algorithms have been proposed, including the Max-tree of Salembier et al. (Salembier et al., 1998) and Tarjan’s union find algorithm (Wilkinson and Roerdink, 2000). A full comparison of the performance of these two fast algorithms with the pixel-queue has demonstrated their advantages (Meijster and Wilkinson, 2002). In particular, it was demonstrated that the processing times for area closings using the max-tree and union find algorithms do not exhibit the strong dependence of λ and image content that is found with the pixel-queue algorithm. The processing speed for all methods is linear with image size, with the union find generally being the best performing method.

Attribute openings and closings are a more generalised approach to area morphology and use attributes other than area to control the filtering action (Breen and Jones, 1996). Both increasing and nonincreasing criteria can be accommodated using attribute openings and thinnings respectively. Examples of increasing criteria include contrast, volume (Salembier et al., 1996), moments of inertia and power (Young and Evans, 2003). Nonincreasing attributes allow use of shape criteria.

3. Area morphology scale-spaces

All nonlinear scale-spaces based on morphological operations have at their heart the application of successive morphological operations of increasing size. Using standard openings and closings, Maragos proposed a morphological scale-space that has the important property of not creating any spurious extrema with increasing scale (Maragos, 1989). However, in common with other approaches using structuring elements, problems with causality exist when they are applied to more than one dimension (Park and Lee, 1996; Jackway and Deriche, 1996). In essence, this means that the positions of image edges are not fixed as scale increases and can

drift through the scale-space which is the same problem found with linear scale-spaces based on Gaussian filters.

A way to overcome this problem is to replace the structural openings and closings with the equivalent area morphology operations and to combine them in an alternating sequential filter (ASF) structure (Bangham et al., 1996). Bangham et al. termed these filters $\mathcal{M}$ - or $\mathcal{N}$ -sieves, respectively, according to whether the area opening or closing operation was performed first. Acton and Mukherjee refer to these as area open-close (AOC) and area close-open (ACO) scale-spaces and have used them for scale-space classification (Acton and Mukherjee, 2000). Given an image I , the AOC and ACO image representation at scale λ are defined by

$$AOC_\lambda(I) = \varphi_\lambda^a \gamma_\lambda^a (\varphi_{\lambda-1}^a \gamma_{\lambda-1}^a (\dots (\varphi_2^a \gamma_2^a (\varphi_1^a \gamma_1^a (I)))))) \quad (3)$$

and

$$ACO_\lambda(I) = \gamma_\lambda^a \varphi_\lambda^a (\gamma_{\lambda-1}^a \varphi_{\lambda-1}^a (\dots (\gamma_2^a \varphi_2^a (\gamma_1^a \varphi_1^a (I))))). \quad (4)$$

where γ_λ^a and φ_λ^a are the area opening and closing operations to an area λ , given by (1) and (2) respectively.

These scale-spaces possess some important properties for multiscale image analysis. Firstly, combining area openings and closings in this manner results in a sieve structure that is guaranteed not to produce any new extrema as the scale increases. Secondly, the ACO and AOC sieves possess the property of *strong causality* (Acton and Mukherjee, 2000). This means that not only will no new edges be created as scale increases but the position of existing edges will not drift through scale space. The ACO and AOC are both sieves and ASF, although not all ASF have the properties of sieves (Bangham et al., 1996). As openings and closings do not commute, although the results of sieving an image to area λ using AOC and ACO sieves are often very similar they are not guaranteed to be identical.

Conceptually, the algorithms for the AOC and ACO area morphology sieves described by equations (3) and (4) have two main steps: (1) the identification of extrema regions and (2) the merging of the extrema with the neighbouring region with the closest greyscale value. These steps are repeated with increasing scale giving rise to the algorithm shown below:

1. Identify all regional extrema;

2. Merge all scale 1 regional extrema with their nearest neighbour;
3. Repeat step 2 with increasing scale, up to scale = λ .

This algorithm clearly shows that it is possible to have a clear separation between the stages of identifying and merging the extrema. For greyscale images the mechanisms used to perform these steps are clear and unambiguous: the extrema are the regional maxima and minima and are merged with the region with the closest intensity value. For colour images, there are different choices that can be made for both these stages and these are explored in more detail in Section 4.1 below.

An efficient and convenient way to interact with the scale-space representations derived using area operators is to use a tree-based representation. If either image maxima and minima are to be processed independently the max-tree or the min-tree of Salembier and Garrido can be used to represent the connected components of the space (Salembier and Garrido, 2000). When it is desirable to represent and process both maxima and minima simultaneously, the inclusion tree can be built from the level line image representation (Monasse and Guichard, 2000), where the level lines are the boundaries of the level sets of the image. The inclusion tree induces a scale-space image structure in which scale is in terms of the number of pixels, i.e. area for two-dimensional image data. A review of region tree representations, namely the max-tree, min-tree, inclusion tree and binary partition tree, that have been used to create connected operators is provided in (Salembier, 2008).

Other connected operators can be used instead of area openings and closings within an ASF structure to generate a scale-space representation. Openings and closings by reconstruction are one possibility but these are far more computationally expensive as, unlike the scale-spaces derived using fast area morphological operators (Meijster and Wilkinson, 2002), the algorithms used in their computation are not nearly so efficient. As an alternative to the use of ASF, levelings can be used to provide a symmetric treatment of the image peaks and valleys (Meyer and Maragos, 2000).

4. Colour connected filters

The problem of extending connected operators to multichannel images has only received limited attention. Weber and Acton proposed a colour connected filter based on applying marginal ordering in the hue, saturation and value (HSV) colour space that overcome the lack of ordering in the hue by applying a rotational shift (Weber and Acton, 2004). The filter was shown to outperform RGB marginal connected filtering for impulsive noise reduction but, like all marginal filters, is not vector preserving.

Several approaches to vector levelings have also been proposed including separable and non-separable levelings [(Meyer, 2000) and (Zanoguera and Meyer, 2002) respectively]. A comparison of the performance of vector levelings using lexicographical total orderings for image segmentation has been performed by Angulo and Serra (Angulo and Serra, 2003)

Related work includes the application of seeded region growing to the quasi-flat zones of multichannel images, as proposed by Brunner and Soille (Brunner and Soille, 2007). More recently an approach based on constrained connectivity, that considers the greyscale variations both along connected paths and within connected components, has been described and results for its extension to colour images presented (Soille, 2008). Alternatively, a connective criterion can be used to aggregate regions around a set of seeds. Jump connections provide one such criterion that forms connected regions in which each pixel is less than a jump connection k above a minimum (Serra, 2008). Iterated jumps have been applied used for colour image segmentation by separately segmenting the hue, luminance and saturation and then using the saturation to choose between the hue and luminance partitions (Angulo and Serra, 2007). Serra has recently extended this approach to images with an arbitrary number of channels using the unit sphere (Serra, 2009). An elegant solution for 4 principal components is also detailed.

The remainder of this section is devoted to the extension of area morphological filters to colour images and the development of colour area morphology scale-spaces. The three scale-spaces described in detail are those based on the CCS, the VAMS and the vector area morphology open-close sieve (VAMOCS).

4.1. Colour area morphology scale-spaces

To extend the greyscale area morphology scale-space algorithm described in section 3 to colour images there are a number of key challenges that must be overcome. The CCS of Gibson et al. (Gibson et al., 2003b,a) and the VAMS of (Evans, 2003a,b) are two approaches to overcoming these challenges. Although these sieves were developed independently, subsequent analysis has shown that they share a common algorithm and only differ in their choices for implementing some of the key choices that must be made at various stages in the algorithm. The algorithm for the filters has the following steps:

1. Identify all regional extrema;
2. Merge all scale 1 regional extrema with their nearest neighbour;
3. Repeat steps 1 and 2 until no extrema exist at current scale;
4. Repeat steps 1 to 3 with increasing scale, up to scale = λ .

Comparison with the greyscale area morphology scale-space algorithm presented in section 3 shows that the major difference is the inclusion of an additional step, namely, *Repeat steps 1 and 2 until no extrema exist at current scale*. This step is required because, unlike the greyscale case, with colour imagery the merging process can create new extrema in the vicinity of the merged regions. Therefore, to ensure that the property of idempotence is preserved, any new extrema with area less than the current value of λ must be processed before moving to the next scale.

Several other significant differences between the colour and greyscale sieves exist. Firstly, in greyscale area morphology extrema are unambiguously identified as either regional maxima or minima regions. In the colour equivalent, extrema are not further classified as maxima or minima and there is no unique method for their identification. Furthermore, as the additional step in the colour algorithm requires the extrema identification and merging steps to be repeated at each scale the performance of the colour area morphology sieves is largely determined on the extrema definition used. The approach used by the VAMS to identify colour extrema is discussed in detail in section 4.1.1 below.

Likewise, there are a number of choices for the mechanism for selecting the neighbouring region with the closest colour to merge with in step 2 of the algorithm. Unlike the greyscale case where the closest intensity is given by the difference in greyscale, for colour images there are a number of norms and distance measures that can be employed. A brief discussion of some of the distance metrics that can be used are discussed in section 4.1.2.

Finally, the second step of the above algorithm contains a further subtlety, which is the value to assign to the merged region. The conventional approach for greyscale images is to set the merged region to the intensity of the non-extreme region. Salembier and Garrido have proposed other possibilities that also take the area of the two merging regions into account (Salembier and Garrido, 2000) and for colour images the freedom to explore various other options may be advantageous.

4.1.1. Colour Extrema

The problem of vector ordering is found in all areas of multi-channel morphology. Likewise, in colour area morphological sieves vector ordering provides the mechanism for identifying the extremal pixels. In addition to identifying extrema in its initial step, the colour area morphology scale-space sieve algorithm also creates new extrema during the merging process of step 2, which must be subsequently processed. Therefore, the proportion of regions that are classified as extrema is of critical importance to the sieve's performance. Extrema definitions that result in a high proportion of image regions being classified as extreme produce sieves that are characterised by an extremely aggressive sieving action, as few regions smaller than the current area will survive at each scale. Another view of extrema is that they are the seeds from which the image simplification process commences, in which case an extrema definition that identifies a relatively low proportion of extreme regions will not significantly alter the image until larger scales. The other significance of the proportion of extreme regions is that it directly relates to processing time at each scale.

Some of the various approaches to this problem can be considered with reference to the example set of vectors shown in figure 2. Applying lexicographical ordering to this vector set results in very different vectors being selected as the maximum and minimum, depending on which channel is given priority. Furthermore, the categorisation of vectors as either maxima

Figure 2: Example vector set $(2, 3)$, $(4, 1)$, $(4, 7)$, $(6, 6)$ and $(8, 2)$. The vectors' aggregate distances using equation (5) with the L_2 norm are 18.38, 18.34, 19.11, 17.09 and 21.08 respectively. ©IEEE 2008.

or minima is not as intuitive as it is in the scalar case. This problem is addressed by an alternative vector ordering scheme based on reduced ordering, that has been used in the design of many multichannel filters including the vector median (Astola et al., 1990) and colour edge detectors (Trahanias and Venetsanopoulos, 1993a, 1996). For a set of n vectors, $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n$, the aggregate distance of $\vec{x}_i$ to the other vectors in the set is denoted d_i and given by

$$d_i = \sum_{k=1}^n \|\vec{x}_i - \vec{x}_k\|_p, \quad i = 1, 2, \dots, n, \quad (5)$$

where $\|\cdot\|_p$ is a vector norm. When the vectors are ordered according to their d_i s such that $d_{(1)} \leq d_{(2)} \leq \dots \leq d_{(n)}$ the result is the ordered sequence

$$\vec{x}_{(1)} \leq \vec{x}_{(2)} \leq \dots \leq \vec{x}_{(n)}. \quad (6)$$

In this ordering $\vec{x}_{(1)}$ is the vector median ($\vec{x}_{vm}$), defined by

$$\sum_{k=1}^n \|\vec{x}_{vm} - \vec{x}_k\|_p \leq \sum_{k=1}^n \|\vec{x}_i - \vec{x}_k\|_p, \quad i = 1, 2, \dots, n$$

and is the vector whose aggregate distance to the other vectors is less than or equal to the

aggregate distances of all other vectors. This definition of the vector median forms the basis of the widely used vector median filter (Astola et al., 1990).

The highest ranked vector is known as the vector extremum $\vec{x}_{ve}$ (Evans, 2003a), such that

$$\sum_{k=1}^n \|\vec{x}_{ve} - \vec{x}_k\|_p \geq \sum_{k=1}^n \|\vec{x}_i - \vec{x}_k\|_p \quad i = 1, 2, \dots, n \quad (7)$$

This definition of the vector extremum has also been used by other authors, where it has been termed the vector outlier (Zhu et al., 1999) or the most spectrally singular pixel (Plaza et al., 2002). For the set of vectors shown in figure 2 the aggregate distances, given in the figure's caption, show that the vector median and extremum are (6,6) and (8,2) respectively. Although in this example the vector median and extremum are uniquely defined, it should be noted that there is no guarantee that this is always the case. As an extreme example, consider any set of only two vectors, where this ordering results in two vectors being both the median and extremum.

To provide an insight into the vector ordering of equations (5) and (6), it is interesting to compare the results it produces when applied to scalar data with those of a simple sort operation. For example, consider the sorted set of numbers shown below:

$$25, 32, 71, 139, 153, 208, 231, 233, 244.$$

A simple inspection reveals that the minimum = 25, the median = 153 and the maximum = 244. If the values are then sorted according to their aggregate distances using equations (5) and (6) then their order is:

$$153, 139, 208, 71, 231, 233, 244, 32, 25.$$

Comparing these two orderings, it can be seen that the (vector) median provides a fulcrum about which the other values are pivoted and, in the scalar case, the highest ranked value will be either the maximum or the minimum. Therefore, with this ordering, instead of two extrema (the maximum and the minimum) being identified only a single extremum exists, as given by the highest ranked vector. For vector data, such as that in figure 2, the notion of a vector minimum or maximum is not consistent with an intuitive interpretation and therefore

the identification of a single extremum is more rational. For example, in figure 2 the vector extremum is (8,2) and is synonymous with an intuitive interpretation of an outlier but can not be easily classified as a maximum or minimum.

Other, similar vector orderings that can be used to identify vector extrema have been proposed, such as ranking vectors according to their distance from the vector median and using the convex hull (Gibson et al., 2003b). In terms of area morphology scale-space sieves for colour images, the existence of only a single class of extrema is very helpful, as the problems of accommodating both maxima and minima within the filter structure are avoided.

4.1.2. Distance metrics

L_p norms are widely used metrics and therefore are an obvious choice for determining the aggregate distances in equation (7). However, there are many alternative distance metrics that can be used both to determine the vector extrema and to identify the neighbouring region with the closest colour to merge with in step 2 of the algorithm.

In other areas of colour image processing the use of the angular difference between two vectors as the distance metric has been proposed, for example in the vector directional filters of Trahanias and Venetsanopoulos (Trahanias and Venetsanopoulos, 1993b). In this approach, for two vectors $\vec{x}_i$ and $\vec{x}_j$ the angular difference L_θ is given by:

$$L_\theta = \cos^{-1} \left(\frac{\vec{x}_i \cdot \vec{x}_j}{\|\vec{x}_i\|_p \|\vec{x}_j\|_p} \right) \quad (8)$$

Combined metrics consider both magnitude and direction, thus capturing the advantages of both properties. For example, Androutsos et al. (Androutsos et al., 1998) normalise the direction and magnitude differences by their maximum possible values to give the combined metric:

$$L_{ae1} = 1 - \left[1 - \frac{2}{\pi} L_\theta \right] \left[1 - \frac{\|\vec{x}_i - \vec{x}_j\|_p}{\sqrt{3 \cdot 255^2}} \right] \quad (9)$$

Equation (9) is advantageous in that only the magnitude difference is used when the two vectors under consideration are collinear. When the magnitudes of both vectors is low, the magnitude term on the right will be close to unity and the metric is dominated by the angular difference. However, in these cases a small change in the response in an individual colour channel can have significant influence on the angular difference and, as low intensity is often associated

Figure 3: Local convex hull example. (a) Two-dimensional vector image and (b) convex hull for $(-6, -5)$ vector and eight-neighbour connectivity.

with random hue, an alternative combined metric, denoted L_{ae2} , takes the angle from the point $\text{RGB} = [255 \ 255 \ 255]$, if this is smaller than L_θ (Gimenez, 2007).

4.2. Convex Colour Sieve

The CCS was proposed by Gibson et al. (Gibson et al., 2003b) and was an attempt to extend the greyscale area morphology sieve of Bangham et al. (Bangham et al., 1996) to colour images. As its name implies, the mechanism used by the CCS to determine image extrema is based upon a convex hull. The method proposed was a local convex hull, consisting of a pixel and its connected neighbours. For images the local neighbourhood is given by either four or eight nearest neighbours connectivity and a pixel is defined as extreme if it lies on the exterior of the local convex hull. Figure 3 gives an example of the local convex hull for the $(6, -5)$ pixel from a two channel vector image. In this example all of the pixels in the set are extreme, as they are all on the exterior of the hull. In fact, this approach results in all except one of the vectors in Figure 3(a) being classed as extreme.

For colour images using the RGB colour space, with a maximum of five or nine vertices for four- or eight-connectivity respectively, degenerate cases result, where the hull is reduced to either a line or a plane. To accommodate this, the approach in (Gibson et al., 2003b) was to define a point as extreme if it lies of the exterior of the reduced dimensionality convex hull.

Defining extrema using the local convex hull approach has the advantage that the hull topology is invariant to monotonic scaling and linear axes transformations. However, it typically results in a large proportion of regions being defined as extreme. For example, the colour test

Figure 4: CCS extrema for the Lily test image using 8nn connectivity. (a) Original image and (b) CCS extrema (shown in white).

image Lily shown in figure 4, has in excess of 80% of its pixels classified as extrema. The high proportion of extrema is at odds with an intuitive interpretation of extrema as outliers. There are also many instances of different pixels that are neighbours being labelled as extreme. Although this can occur in the greyscale case, where maxima and minima regions can be neighbours, this does not provide a satisfactory explanation of the many adjacent extrema.

To determine the neighbour to merge with in the Step 2 of the colour area sieve algorithm, the CCS uses the Euclidean distance. However, as there is a possibility that two or more different colours can be equidistant from a reference colour a further ordering using the difference in luminance was proposed. Any remaining ties are then ordered by their G, R and B values respectively.

Although not explicitly stated in (Gibson et al., 2003b), the convention used to select the colour to assign to the combined pixels, after the merging process, follows that of standard greyscale area morphological sieves, i.e. the colour of the pixel selected to be merged with. Therefore, for the case where two single-pixel regions are neighbours and are both classified as extreme, the colour assigned to the resultant merged region will arbitrarily depend on which extrema is processed first. New extrema that are created by the merging process are identified and processed by the recursion introduced by Step 3 of the algorithm.

As with other sieves, the set of scale-space images produced by the CCS can be presented in a tree-based representation. An evaluation of how semantically meaningful the segmenta-

tions from the CCS are was undertaken by Gibson et al. (Gibson et al., 2003a). In this work the approximately 50 regions were extracted from the scale-space tree, starting from (but excluding) the root. Comparison with sets of human-segmented ground truths from a set of 50 images showed some agreement, although there was better agreement amongst the humans who performed the segmentations. A comprehensive comparison of segmentation performance produced by the CCS and other colour area-morphology scale-spaces is given in section 5.2 below.

4.3. *Vector area morphology sieve (VAMS)*

The VAMS was developed independently to the CCS and was first presented in 2003 (Evans, 2003a,b). Although its algorithmic steps are identical to the CCS, it differs in how the stages are implemented. In particular, the mechanism used to identify the extrema is significantly different. As discussed in section 4.1 above, as the extrema definition has a major influence on the sieve’s performance the VAMS and CCS can produced substantially different results.

The starting point for the VAMS’s extrema definition is the vector extrema ($\vec{x}_{ve}$) of equation (7). However, unlike the CCS that applies a binary extreme/non-extreme decision at each individual pixel, the VAMS uses slightly different approach. Firstly, a scalar image is derived from the vector data by calculating the aggregate difference from each pixel to its neighbours using equation (5). Then, for each flat zone containing more than 1 pixel, the aggregate distances are averaged and the mean value assigned to all its pixels. Finally, the extrema are identified as the maxima in the scalar surface of all the aggregate distances. In this approach, the extrema are those regions whose sum of distances to their connected neighbours is greater than that of any neighbouring region. This process is illustrated in Figure 5(a) that shows the example motion field from figure 3, with the flat zones of area > 1 clearly identified, and the corresponding scalar surface. Note that in comparison with the CCS, where all but one flat zone was classified as extreme, the VAMS only produces four initial extrema.

The VAMS uses the Euclidean distance to identify the region to merge with in Step 2 of the algorithm. This approach is very similar to the CCS, the only difference being in how ties are resolved, where the VAMS simply uses scan order. New extrema that are created by the merging process are themselves merged and the scalar image updated. The sieving process

is shown in figure 5, where (b)-(d) show the vector field, its flat zones and the corresponding scalar images up to scale $\lambda = 2$ when all the extrema have an area ≥ 2 . In this example the merged regions are assigned the vector of the non-extreme region, as is conventional.

This results of applying the VAMS to the Lily test image of figure 4(a) are shown in figure 6. As the scale increases [figure 6(b) to (d)] patches of extreme colour are removed from the image, for example from the lily's petals, without altering the position of any other image boundaries. Figure 6(a) shows the initial extrema, which constitute approximately 6% of image regions. Comparing the initial extrema for the VAMS to those produced by the CCS for the same image in figure 4(b), which has over 80% of the regions classified as extreme, it can be seen that the VAMS has a much less aggressive sieving action.

As discussed in section 4.1, there are other choices that can be made in the implementation of the various steps of the VAMS algorithm. Some of the different possibilities for these are examined with respect to the noise reduction performance of the VAMS in the following section.

4.3.1. VAMS parameter evaluation

This section looks at the effect that some of the different choices that can be made at the VAMS algorithm's steps has on the sieve's performance. In particular the influence of these on the noise reduction performance of the VAMS are considered. The first choice to be considered is the distance metric used to form scalar images from the aggregate distances, which consequently affects the proportion of regions classified as image extrema. Figure 7 shows the aggregate distances for the Lily test image using the L_2 and L_{ae2} norms, as described in section 4.1.2. In these scalar images there are 2118 and 2209 initial extrema for the L_2 and L_{ae2} norms respectively.

To evaluate the noise reduction performance both the normalised mean square error (NMSE) and mean chromaticity error (MCRE) were used, as these are widely used metrics for determining objective error measures for colour images (Trahalias and Venetsanopoulos, 1993b). The NMSE is given by

$$NMSE = \frac{\sum_{i=0}^M \sum_{j=0}^N \left\| \vec{f}(i, j) - \vec{f}'(i, j) \right\|^2}{\sum_{i=0}^M \sum_{j=0}^N \left\| \vec{f}(i, j) \right\|^2} \quad (10)$$

where $\vec{f}(i, j)$ and $\vec{f}'(i, j)$ are the original and filtered vectors at pixel (i, j) and M and N are

1,0	3,0	8,7	2,0	2,0	2,0	10.1	12.0	28.5	4.0	4.0	4.0
8,-4	2,-1	1,-1	2,0	2,0	2,0	21.2	10.1	14.5	4.0	4.0	4.0
3,0	1,-1	2,0	2,0	2,0	2,0	10.9	5.7	4.0	4.0	4.0	4.0
1,-1	2,-1	2,0	2,0	-6,-6	2,0	2.2	3.0	4.0	4.0	33.6	4.0
1,-1	2,-1	2,0	-6,-5	-9,-8	1,-1	2.2	3.0	4.0	33.4	33.7	15.0
1,-1	2,0	2,0	3,0	2,0	2,0	2.2	4.0	4.0	14.5	11.9	11.9
2,-1	2,0	2,0	1,-1	1,0	-4,-3	2.0	4.0	4.0	4.7	7.8	12.5

(a)

1,0	3,0	3,0	2,0	2,0	2,0	4.0	3.3	3.3	3.5	3.5	3.5
3,0	2,-1	1,-1	2,0	2,0	2,0	3.9	4.8	6.1	3.5	3.5	3.5
3,0	1,-1	2,0	2,0	2,0	2,0	3.9	5.7	3.5	3.5	3.5	3.5
1,-1	2,-1	2,0	2,0	-6,-6	2,0	2.2	3.0	3.5	3.5	24.8	3.5
1,-1	2,-1	2,0	-6,-5	-6,-6	1,-1	2.2	3.0	3.5	30.2	24.8	11.4
1,-1	2,0	2,0	3,0	2,0	2,0	2.2	3.5	3.5	14.5	7.2	7.2
2,-1	2,0	2,0	1,-1	1,0	1,0	2.0	3.5	3.5	4.7	1.5	1.5

(b)

3,0	3,0	3,0	2,0	2,0	2,0	1.9	1.9	1.9	3.5	3.5	3.5
3,0	2,-1	2,-1	2,0	2,0	2,0	1.9	3.6	3.6	3.5	3.5	3.5
3,0	1,-1	2,0	2,0	2,0	2,0	1.9	5.7	3.5	3.5	3.5	3.5
1,-1	2,-1	2,0	2,0	-6,-6	2,0	2.2	3.0	3.5	3.5	26.5	3.5
1,-1	2,-1	2,0	-6,-6	-6,-6	1,-1	2.2	3.0	3.5	26.5	26.5	11.4
1,-1	2,0	2,0	3,0	2,0	2,0	2.2	3.5	3.5	15.1	7.2	7.2
2,-1	2,0	2,0	1,-1	1,0	1,0	2.0	3.5	3.5	4.7	1.5	1.5

(c)

3,0	3,0	3,0	2,0	2,0	2,0	1.8	1.8	1.8	3.5	3.5	3.5
3,0	2,-1	2,-1	2,0	2,0	2,0	1.8	2.7	2.7	3.5	3.5	3.5
3,0	2,-1	2,0	2,0	2,0	2,0	1.8	2.7	3.5	3.5	3.5	3.5
1,-1	2,-1	2,0	2,0	-6,-6	2,0	2.2	2.7	3.5	3.5	26.5	3.5
1,-1	2,-1	2,0	-6,-6	-6,-6	1,-1	2.2	2.7	3.5	26.5	26.5	11.4
1,-1	2,0	2,0	3,0	2,0	2,0	2.2	3.5	3.5	15.1	7.2	7.2
2,-1	2,0	2,0	1,-1	1,0	1,0	2.0	3.5	3.5	4.7	1.5	1.5

(d)

Figure 5: VAMS operation using 4-neighbour connectivity and the L_2 norm. (a) Original vector image with flat zones of area > 1 identified using colour and corresponding scalar image, with extrema coloured red. (b)-(d) Corresponding vector and scalar images as extrema of area ≤ 2 are merged.

Figure 6: VAMS example for Lily test image using 8-connectivity. (a) Initial extrema (black), (b) area=9, (c) area=49 and (d) area=81.

Figure 7: VAMS aggregate distance images using (a) L_2 norm and (b) L_{ae2} metric.

the image dimensions. The MCRE is defined as the distance between the intersection points of $\vec{f}(i, j)$ and $\vec{f}'(i, j)$ and the Maxwell triangle. When the triangle is defined on the unit plane this gives

$$MCRE = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N \left[(\vec{r}(i, j) - \vec{r}'(i, j))^2 + (\vec{g}(i, j) - \vec{g}'(i, j))^2 + \vec{b}(i, j) - \vec{b}'(i, j))^2 \right]^{1/2} \quad (11)$$

where r , g and b are the normalised RGB values (Trahanias and Venetsanopoulos, 1993b).

A set of six test images (Lily, Autumn, Sample1, Lenna, Boats and Baboon) were corrupted with uncorrelated Gaussian ($\sigma^2 = 1000$), 10% impulsive and mixed noise using the noise model of Viero et al. (Viero et al., 1994). The images were sieved using the VAMS with the extrema identified using the L_1 and L_2 norms, the differences in luminance (L_Y), angle (L_θ) and the combined distance and angle differences L_{ae1} and L_{ae2} . All images were sieved up to area $\lambda = 12$ and, for each image, the minimum NMSE between the filtered and original, noise-free image found. For all aggregate distances used to identify the extrema, the L_2 norm was used to select the region to merge with in step 2 of the algorithm and the merged region was assigned the colour of the non-extreme region.

Figure 8 gives the average NMSE results for the six test images. The MCRE at the scale that produced the minimum NMSE was also recorded and these results are also shown. All choices for the extrema definition result in a VAMS that reduces the image noise, with the reduction being greater for impulsive noise than for Gaussian noise. Using the luminance difference L_Y produced the worst noise reduction performance for all noise types for both the NMSE and the MCRE. The performance of the L_1 and L_2 norms were broadly similar, with the L_2 norm having a slight advantage for impulsive noise. Both combined distance metrics performed well and had a slightly lower NMSE and MCRE than the L_1 and L_2 norms for Gaussian and mixed noise. Their results were comparable for impulsive noise. The angular difference L_θ produced the lowest MCRE for Gaussian and mixed noise. Also shown are the results for the VAMS applied to the original image, which give a measure of the minimum distortion the sieve will introduce at small scales.

The six distance metrics described above were also used to investigate the influence of the

Figure 8: VAMS results using various distance metrics to calculate the aggregate distances used for extrema calculation. ©IEEE 2008.

distance metric in selecting the region to merge with on the noise reduction performance. As before, the noise reduction performance was evaluated using the NMSE and MCRE, see Fig. 9. Here, regardless of the method used to find the closest region, the L_2 norm was used to identify the extrema and distance ties were resolved using luminance.

The results show that the angular difference introduced the greatest distortion on the original, noise free image and also performed poorly for all noise types. For example, with Gaussian noise it does not produce any improvement on the noisy image. The L_1 and L_2 norms produce slightly lower NMSE than the combined metrics, with the L_1 norm performing particularly well for impulsive noise. Overall the noise reductions produced by the VAMS are greater for impulsive noise than for the other noise types.

The final VAMS choice to be considered is the colour to assign to the merged regions. With greyscale connected sieves, the convention is to assign the merged region the value of the non-extreme region. Some other alternatives have been put forward, for example Salembier and Garrido proposed using the value of the largest of the two regions being merged (Salembier and Garrido, 2000). The colour equivalent of this is simply choosing the colour of the largest region. There are also other options that can be considered, such as selecting the mean or the marginal median of the regions being merged. Although this will not preserve the original vectors, in

colour morphology this approach has been found to be advantageous for noise reduction (Aptoula and Lefèvre, 2007). If vector preservation is desired an intermediate approach can be used in which the central vector is used as a “guide” and the merged region is then assigned the closest colour to the guide from the colours of the input regions.

Figure 9: VAMS results using various distance metrics to select the region for the extrema to merge with. ©IEEE 2008.

Figure 10 presents the NMSE and MCRE results using the different mechanisms to select the colour to assign the merged region. In particular, the merged region is assigned the colour of the non-extreme region, the mean colour, the marginal median and the mean and median guided colours, as described above. For all noise types, selecting the colour of the non-extreme region results in the lowest NMSE and MCRE. For impulse noise this can be explained by the fact that the noise gives rise to extrema that should be completely removed from the image without affecting the surrounding values. For Gaussian and mixed noise the desideratum to completely remove extrema applies to a lesser degree but the non-extreme region’s colour still proves to be the best choice. This evaluation has also been performed with other colour area morphology scale-space sieves and the non-extreme region proves to be the best choice in nearly all cases (Gimenez, 2007). For the other choices, the mean and mean-guided colour choices perform badly for impulsive and mixed noise while the marginal median performs better than the median-guided.

Figure 10: VAMS results produced using different mechanisms to determine the colours of the merged regions. ©IEEE 2008.

4.4. Vector area morphology open-close sieve (VAMOCS)

The VAMOCS is a development of the VAMS that aims to improve the performance of the colour area morphology scale-space sieve by increasing the number of image extrema (Gimenez and Evans, 2005, 2008). As the extrema act as seeds for the process of image simplification through the merging of flat zones, increasing their numbers should make the filter more aggressive and consequently has the potential for improving the segmentation performance of the sieve. Unlike the CCS, where the use of the local convex hull results in a binary extreme/non-extreme decision for each flat zone, the scalar surface employed by the VAMS allows more flexibility to be introduced to the sieve structure.

To introduce more extrema, the VAMOCS classifies both the minima and maxima of the scalar surface as extreme. Unlike the maxima, which are true outliers, the minima are those flat zones with a difference in colour to their neighbouring regions that is less than those of their connected neighbours. For the example vector image of figure 2(a) there are two initial minima with values $(1, 0)$ and $(2, -1)$, in the top left and bottom left corners of the image respectively. Although the maxima and minima are different conceptually and are processed by openings and closings respectively, in the VAMOCS structure both types of extrema are treated identically and are merged with their closest neighbour in step 2 of the algorithm. However, the choice of colour to assign the merged region may differ according to whether the extremum is a maximum

Figure 11: Maxima (black) and minima (white) for Lily test image. (a) Greyscale extrema and (b) VAMOCS extrema.

or a minimum.

Processing the minima in this manner merges those flat zones whose colour is similar to the surrounding colours and will create larger flat zones in regions that have little variation in colour. This approach is similar in spirit to that of Salembier et al. (Salembier et al., 1998), who proposed the use of a bound on the maximum allowable greyscale fluctuations in response to the observation that visual entities may not be strictly flat. When merging the minima regions particular care must be taken with the selection of the colour to assign to the merged region. To try and ensure that the choice reflects a colour that is representative of the region to be merged the colour of the flat zone with the largest area was selected, an approach previously proposed for greyscale images (Salembier and Garrido, 2000).

The VAMOCS also has a further refinement, introduced to improve the segmentation performance. This is to normalise the distance metric for flat zones by their perimeters rather than their areas (Gimenez and Evans, 2005). The VAMOCS extrema using the normalisation for the Lily test image are shown in 11. For comparison, the greyscale maxima and minima for the luminance image are also shown. Both the VAMOCS and the greyscale sieve define approximately 10% of the pixels as extreme, although the distributions of these extrema differ.

The sieve results for increasing the number of image extrema by including the minima of the scalar surface are presented in figure 12 for the Lily test image. Here, the VAMS, VAMOCS and CCS were used to sieve the image to scales 10, 100, 1000 and 1000. For comparison the

VAMOCS closings are also shown, where only the minima selected as extrema. The figure illustrates that the actions of the VAMOCS openings and closings are complimentary: processing the maxima removes outliers whilst using closings to remove minima extends the image regions that are relatively flat, leaving isolated extrema unaffected. The VAMOCS (column 3 of figure 12) combines the effects of openings and closings to produce an area morphology colour sieve whose filtering action is similar in aggressiveness to the CCS. However, this aggressiveness is achieved with far fewer extrema, as can be seen by comparing figure 11(b) with the CCS extrema in figure 4(b).

The proportion of VAMOCS flat zones defined as extreme is further characterised in Table 1 where, for comparison, results for the VAMS, CCS and greyscale area open-close (GS-AOC) sieves are also shown. The CCS has 82% of its flat zones initially classified as extreme and the number of regions rapidly reduces with increasing area until there are only 25 regions when the scale reaches 1000. The VAMOCS only has 10.6% of regions initially defined as extrema but its aggressive action rapidly reduces the total number of flat zones until only 132 remain at scale 1000. In contrast the GS-AOC sieve starts with 9.9% of initial extrema but still has over 600 flat zones at scale 1000. The VAMS has the fewest initial extrema and still has nearly 7000 extrema at scale 1000, confirming its comparatively non-aggressive action.

scale	1	2	10	50	100	500	1000
VAMS	<u>2209</u> 41207	<u>1054</u> 37199	<u>151</u> 27547	<u>21</u> 19660	<u>8</u> 16102	<u>1</u> 8024	<u>1</u> 6871
CCS	<u>33616</u> 41207	<u>12852</u> 16402	<u>2211</u> 2878	<u>456</u> 572	<u>223</u> 279	<u>44</u> 53	<u>23</u> 25
VAMOCS	<u>4388</u> 41207	<u>2353</u> 33882	<u>476</u> 14326	<u>112</u> 3856	<u>16</u> 1845	<u>8</u> 298	<u>9</u> 132
GS-AOC	<u>4062</u> 41207	<u>1973</u> 36553	<u>424</u> 28457	<u>110</u> 20697	<u>63</u> 17040	<u>11</u> 9114	<u>7</u> 6150

Table 1: Area morphology scale-space sieves extrema for Lily image (expressed as fractions): number of flat zones classified as extreme to the total number of flat zones.

(a) Area = 10.

(b) Area = 100.

(c) Area = 1000.

(d) Area = 10000.

Figure 12: Colour sieve results for Lily test image using 8-connectivity. Columns 1 to 4: VAMS, VAMOCs (closings only) and VAMOCs (combined openings and closings) and CCS.

5. Performance evaluation

The aim of this section is twofold. Firstly, the implementation used for the colour sieves is described and the relationship between the extrema definition and the computational complexity investigated. Secondly, the suitability of the colour area-morphology scale-spaces for image

segmentation is investigated. Segmentation is one of the most important problems of colour image analysis and has been widely considered elsewhere and has been the subject of several survey papers (Cheng et al., 2001; Lucchese and Mitra, 2000). Colour image segmentation was also the focus of (Gimenez and Evans, 2005) and (Gimenez and Evans, 2008), from which material in this section is drawn. Although only the RGB colour space is used in this work the relationship between colour image segmentation and the colour space used is a topic in its own right and has been addressed at length by Busin et al. (Busin et al., 2008). The use of the colour area morphology scale-spaces in other colour spaces has been investigated (Gimenez, 2007), as has their use for image noise reduction (Gimenez, 2007; Gimenez and Evans, 2008).

5.1. Implementation and timings

To implement the colour area-morphology scale-space sieves an approach based on Vincent’s pixel-queue algorithm (Vincent, 1993b) has been proposed (Gimenez, 2007; Gimenez and Evans, 2008). Although other, more efficient algorithms exist for greyscale sieves the pixel-queue lends itself to a straightforward adaption to the colour sieve algorithm presented in section 4.1.

Figure 13 presents the algorithm used to implement the colour sieves and a brief explanation of its steps follows. In line 1, all flat zones identified as extreme are placed in a list according to its area, each list being a first-in-first-out (FIFO) queue. The lists are processed from area 1 to λ (line 2) until no extrema of area $\leq \lambda$ remain. The use of the *while* loop in line 3 ensures that any new extrema created by the merging process (line 6) that are smaller than the current area are processed, thus ensuring idempotency. New extrema can only be created in the vicinity of the merging regions and therefore this check in line 7 can be performed locally. In addition to creating new extrema, the merging process can also result in regions that were classified as extrema becoming non-extreme. These cases could be detected and removed from the extrema list. Alternatively, line 5 performs a simple test on each previously identified extrema, to confirm that it is still extreme, before the merging process.

The CCS and VAMS implementations both follow this algorithm, only differing in how they identify the extrema. The VAMOCs requires a slight modification as the choice of colour to assign to the merged region is different for maxima and minima. Therefore the extrema must either be placed in separate list or distinguished between before the merging process, for

1. Extract all extremum flat zones and place in list *extrema(area)*, according to their area
2. For *area* = 1 to λ :
3. While *extrema(i)* $\forall i = 1, 2, \dots, area$ are not empty
4. For each flat zone in *extrema(area)*
5. If flat zone is still an extremum
6. Merge region with its closest neighbour;
7. Append any new extrema created by merging process to
 appropriate list;
8. Else remove flat zone from *extrema(area)*

Figure 13: Colour area morphology scale-space sieve algorithm.

example by determining the extrema type in line 5.

The colour area morphology scale-space sieve algorithm has several properties in common with its greyscale counterpart. These include idempotence, strong causality and a reduction in the total number of flat zones with increasing scale. However, the outputs of the colour sieves are not invariant to the order in which the extrema are merged. A similar situation is found with greyscale sieves, where different results are obtained according to whether the maxima or minima are processed first i.e. whether the sieve is an AOC or an ACO. For the colour sieves there are typically many adjacent extrema and the order in which they are processed can have a more significant influence on the final result.

Using the above algorithm, the colour area morphology sieves were implemented in C++, then compiled and run as mex files within Matlab. The processing times and proportion of regions classified as extreme using eight-connectivity and the L_2 norm are shown in figure 14. Results for the GS-AOC sieve applied to the luminance component are also included. To provide a meaningful comparison, results for a pixel-queue implementation of a GS-AOC sieve are also included, although it should be noted that more computationally efficient algorithms have been developed for the greyscale case (Meijster and Wilkinson, 2002).

The processing times can, for a large part, be explained by considering the proportion of extreme regions. The CCS has a high proportion of extrema and the calculation of the

Figure 14: Colour sieve performance for Lily image. (a) Proportion of regions classified as extreme and (b) processing times.

convex hull is also relatively expensive. These factors are reflected in a high processing time for low values of area. However, as the total number of regions rapidly reduces with scale (see table 1) the processing time is relatively constant above an area of 100. The VAMS has a lower proportion of extrema than the GS-AOC sieve but has a slightly higher processing time below an area of 1000 due to its extrema definition being more computationally expensive to calculate. Compared to the VAMS, the more complex structure of the VAMOCS is reflected by its higher processing times. However, its total number of regions reduces more rapidly because of its more aggressive action. Therefore the computational advantage of the VAMOCS increases with increasing scale and it is the quickest colour sieves for areas > 1000 .

5.2. Application to image segmentation

The ability of connected filters to simplify images by either completely removing or perfectly preserving edges means that they are well suited for image segmentation and their application to

Figure 15: Example trees for Lily test image [figure 4(a)], sieved to 500 regions.

this problem has been proposed by several authors, for example see (Crespo et al., 1997; Soille, 2008; Serra, 2008). For greyscale images, a formal relationship between connected operators and segmentation algorithms based on region merging has been established (Gatica-Perez et al., 2001). One useful approach is to derive a tree-based image representation from the scale-space and apply an appropriate pruning technique. This strategy has been demonstrated for scale-spaces from area morphology and other connected operators (Salembier et al., 1998; Salembier and Garrido, 2000).

This section explores the potential of the colour area morphology scale-spaces for image segmentation and contains some material drawn from (Gimenez and Evans, 2008). It is not the intention to propose a pruning strategy that is optimal in a segmentation sense. Instead, the

Figure 16: Example image from the Berkeley Segmentation Dataset. (a) Koala image and (b) collection of human segmentations.

ability of the colour sieves' tree representations to contain semantically meaningful segmentation is investigated by applying the sieves over a range of scales, such that the sieved images have a predetermined target number of regions remaining, and quantitatively assessing the resulting segmentations. All experiments were performed in the RGB colour space and the L_2 norm was used both to determine the extrema and to select the regions to merge with. Figure 15 shows the trees for 500 regions produced by the GS-AOC and the colour sieves on the Lily test image. Here, it can be seen that although the target number of regions is the same, there are marked differences in the resulting tree structures.

To provide an objective measure of the segmentation performance the Berkeley Segmentation Dataset and Benchmark is used (available from <http://www.cs.berkeley.edu/projects/vision/grouping/segbench>) (Martin et al., 2001, 2004). The public dataset contains 300 (100 test and 200 training) images and the segmentation results produced by a number of human observers, using both colour and greyscale versions of the images (Martin et al., 2001). An example colour image and its human segmentations are shown in figure 16. Using each collection of human segmentations to provide ground truth,

a methodology for quantifying the segmentation performance is also proposed (Martin et al., 2004). This methodology uses precision-recall (P-R) curves to characterise the segmentation performance, where precision measures the probability that a detected boundary pixel is contained in the ground truth and recall is the probability of detecting a true boundary pixel. Although other quantitative evaluation methods exist, for example those described in (Busin et al., 2008), P-R curves are particularly well adapted for colour image segmentation.

P-R curves for each of the 100 colour test images were found by applying the colour area morphology scale-space sieves with increasing area until the number of regions was less than a predetermined threshold. For each target number of regions the precision and recall values were recorded. Predetermined thresholds for the number of regions were used in preference to explicitly varying the area parameter as it is less dependent on image content and also allows direct comparison with sieves results using other attributes. The approach is also compatible with the dataset ground truth images which contain a small number of equally important regions.

To enable a comparison to be made between different P-R curves the F-measure can be calculated for all points on the curves. The F-measure is given by the harmonic mean of precision and recall which can be simply calculated by

$$F - \text{measure} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (12)$$

and the maximum value of the F-measure along a P-R curve provides a single number characterising the segmentation performance (Martin et al., 2004).

Figure 17 presents the P-R curves for the Koala test image shown in figure 16(a) produced by the GS-AOC sieve, the CCS, the VAMS and the VAMOCs. The position on each curve at which the maximum F-measure occurs is also shown and the corresponding values given in the figure's caption. The F-measures show that all that colour sieves produce a better segmentation performance than the greyscale sieve, with the CCS and VAMOCs producing the best F-measures of 0.66 and 0.65 respectively. The segmentation results for the maximum F-measures are shown in figure 18. These clearly show the improved performance of the colour sieves over their greyscale counterpart. In addition to improved segmentation performances, the colour

Figure 17: P-R curves for Koala test image. The maximum F-measure is marked in bold with exact value given in brackets (a) GS-AOC Sieve (0.54), (b) CCS (0.66), (c) VAMS (0.60) and (d) VAMOCS (0.65).

sieves typically achieve their best results with far fewer regions. For example, the CCS and VAMOCS results in figure 18(b) and (d) have between 30 and 40 regions whereas the GS-AOC segmentation has several thousand regions. The number of regions is important as, although it does not directly contribute to the segmentation performance, when two segmentations have the same maximum F-measure, the segmentation with fewer regions is generally preferable.

To provide a more comprehensive evaluation the maximum F-measures from the P-R curves for all 100 colour test images were found and the average values produced by each sieve calculated; these are presented in table 2. Also given are the number of regions at which the maximum F-measures occurred. These results confirm the initial impressions gained from the

Figure 18: Sieve segmentation results corresponding to the maximum F-measures in figure 17 produced by (a) GS-AOC Sieve, (b) CCS, (c) VAMS and (d) VAMOCs.

Koala image; compared with the greyscale sieve, all the colour area-morphology scale-spaces produce an improved segmentations with fewer regions. The maximum F-measures for the colour sieves are $0.03 - 0.14$ higher than the greyscale case and the number of regions is at least an order of magnitude lower.

Comparing the performance of the colour sieves, the VAMS has the lowest F-measure and this is achieved with, on average, the highest number of regions. The CCS has a significantly better segmentation performance than the VAMS, with an F-measure that is improved by 0.09. However, the inclusion of area closings in addition to openings by the VAMOCs results in the highest F-measure of 0.51 which is 0.02 above that of the CCS, although this is achieved with a slightly increased number of regions. Results using contrast as the attribute instead of area have also been reported (Gimenez and Evans, 2008) and show the same general trends.

Morphological Sieve	Maximum F-measure	Number of Regions
GS-AOC	0.37	3000
CCS	0.49	30
VAMS	0.40	60
VAMOCs	0.51	70

Table 2: Average F-measure for 100 images from the Berkeley dataset achieved by the CCS, area-normalised VAMS, perimeter-normalised VAMOCs and a greyscale AOC ASF applied to the luminance component.

5.3. Robustness to noise

One important aspect of segmentation algorithms is how sensitive they are to perturbations of the input image, for example by noise. To this end, the robustness of the colour area morphology scale-spaces was investigated by applying the colour sieves to original and noise-corrupted images. The resulting segmentations were compared, again using the quantitative evaluation methodology provided by the Berkeley Segmentation Dataset and Benchmark. For each noise-free image the segmentation that produced the maximum F-measure on the P-R curve was used as the best segmentation result. The images were then corrupted by different levels of impulsive and Gaussian noise and the resulting segmentations compared to the best noise-free results produced by the same sieves. In practice this was accomplished by generating P-R curves for the noisy images, using the noise-free segmentation results as the ground truths. The maximum F-measure then provides a measure of how closely the segmentation results of the noisy images match those of the noise-free cases.

The results for one image from the dataset are shown in figure 19. In the figure, the overlaps between the original and noise-corrupted segmentations are shown in black and the noise-free

(a) Original image and collection of human segmentations.

(b) CCS results for impulsive noise (left) and Gaussian noise (right).

(c) VAMS results for impulsive noise (left) and Gaussian noise (right).

(d) VAMOCS results for impulsive noise (left) and Gaussian noise (right).

Figure 19: Segmentation results for image corrupted by 10% impulsive noise and Gaussian noise ($\sigma^2 = 10^3$). Segmentation results for noise-corrupted images shown in blue, noise-free segmentation results shown in yellow and overlaps shown in black.

and noisy segmentations in yellow and blue respectively. If the introduction of image noise caused no change in the segmentation result then the segmentations will be congruent and only the black overlap would be visible. In practice the segmentations differ but are very hard to evaluate subjectively, confirming the need for a quantitative evaluation methodology to assess the segmentation performances.

The 100 colour test images from the dataset were corrupted by 3 levels of independent impulsive and Gaussian noise and the average F-measure used to quantify the robustness of the results, see Table 3. All colour sieves show a decreasing F-measure as the level of noise increases, for both impulses and Gaussian noise. The VAMS has the lowest average F-measure for all noise types and levels, implying that is the least robust. For a low and median level of impulsive noise the CCS is the most robust. For a high level of impulsive noise and all levels of Gaussian noise the VAMOCS is the most robust.

Noise Type	CCS	VAMS	VAMOCS
0.1% Impulsive	0.670	0.478	0.610
1% Impulsive	0.590	0.452	0.577
10% Impulsive	0.543	0.456	0.568
Gaussian ($\sigma^2 = 10^1$)	0.537	0.436	0.558
Gaussian ($\sigma^2 = 10^3$)	0.490	0.452	0.499
Gaussian ($\sigma^2 = 10^4$)	0.431	0.415	0.436

Table 3: Robustness of colour sieve segmentations to (independent) image noise. Average F-measures between noisy and noise-free segmentations for 100 images from the Berkeley dataset.

6. Conclusions

By far the majority of the colour morphology work to date has been concerned with structural techniques. In comparison, the development of multichannel connected filters is an area that is somewhat more inchoate though recent contributions, for example those of Serra (Serra, 2008) and Soille (Soille, 2008), show that it is gathering momentum. Area morphology scale-space sieves have proved a useful tool for image analysis and have been widely applied to image segmentation. Their extension to the colour domain has some attractions as well as several problems that must be overcome. Here, a vectorial approach is preferred as simply applying a

greyscale sieve to individual channels can result in colour bleeding and perceived edge shifts. Although there is an ordering problem to be addressed, in multichannel images the distinction between maxima and minima that must be accommodated in the greyscale case vanishes, resulting in an inherently balanced treatment.

A detailed description of the underlying algorithm of three such colour area morphology sieves, the CCS, the VAMS and the VAMOCs, has been presented and their different mechanisms for determining extrema described. For the CCS, the extrema definition is binary, resulting in a high proportion of extrema and an aggressive sieving action. However, some discretion in the merging process is possible. The VAMS has additional flexibility resulting from its capacity to accommodate different distance metrics in its extrema definition. Its scalar surface can also be processed using closings, giving rise to the VAMOCs. Although it has an increased proportion of extrema and hence a more aggressive action, the complexity of the VAMOCs is still comparable to that of the VAMS. The colour area morphology filters are idempotent by virtue of the fact that they process any new extrema created by the merging process at each scale. However, they do lack invariance to the order in which the extrema are processed, which may be an issue in some applications.

The application of the colour area morphology scale-spaces for segmentation has been evaluated using the Berkeley Segmentation Dataset and Benchmark and an improvement on the GS-AOC sieve shown. The VAMOCs produced the best overall segmentation performance and is also the most robust to image noise. Although not competitive with state-of-the-art colour segmentation techniques, the results show the potential of the scale-spaces to generate trees from which semantically meaningful segmentations can be extracted. The scale-spaces are also suitable for use with other multivariate images, for example from remote sensing applications.

7. Acknowledgements

The contribution made to the material on which this chapter is based by D. Gimenez during the course of his PhD is gratefully acknowledged.

References

- Acton, S., Mukherjee, D., April 2000. Scale-space classification using area morphology. *IEEE Trans. Image Processing* 9 (4), 623–635.
- Androutsos, D., Plataniotis, K., Venetsanopoulos, A., 1998. Distance measures for color image retrieval. In: *IEEE International Conference on Image Processing*. pp. Vol II: 770–774.
- Angulo, J., Serra, J., 2003. Color segmentation by ordered mergings. In: *IEEE International Conference on Image Processing*. pp. II: 125–128.
- Angulo, J., Serra, J., April 2007. Modelling and segmentation of colour images in polar representations. *Image and Vision Computing* 25 (4), 475–495.
- Aptoula, E., Lefèvre, S., November 2007. A comparative study on multivariate mathematical morphology. *Pattern Recognition* 40 (11), 2914–2929.
- Astola, J., Haavisto, P., Neuvo, Y., 1990. Vector median filters. *Proceedings of IEEE* 78, 678–689.
- Bangham, J., Harvey, R., Ling, P., Aldridge, R., July 1996. Morphological scale-space preserving transforms in many dimensions. *Journal of Electronic Imaging* 5 (3), 283–299.
- Barnett, V., 1976. The ordering of multivariate data. *J. Royal Statistical Society A* 139 (Part 3), 318–343.
- Breen, E., Jones, R., 1996. Attribute openings, thinnings, and granulometries. *Computer Vision and Image Understanding* 64, 377–389.
- Brunner, D., Soille, P., Aug. 2007. Iterative area filtering of multichannel images. *Image and Vision Computing* 25 (8), 1352–1364.
- Busin, L., Vandenbroucke, N., Macaire, L., 2008. Color spaces and image segmentation. Vol. 151 of *Advances in Imaging and Electron Physics*. Elsevier, pp. 65 – 168.

- Cheng, F., Venetsanopoulos, A., 1992. An adaptive morphological filter for image processing. *IEEE Trans. Image Processing* 1, 533–539.
- Cheng, H., Jiang, X., Sun, Y., Wang, J., December 2001. Color image segmentation: advances and prospects. *Pattern Recognition* 34 (12), 2259–2281.
- Comer, M., Delp, E., 1999. Morphological operations for color image processing. *Journal of Electronic Imaging* 8 (3), 279–289.
- Crespo, J., Schafer, R., Serra, J., Gratin, C., Meyer, F., October 1997. The flat zone approach: A general low-level region merging segmentation method. *Signal Processing* 62 (1), 37–60.
- Evans, A. N., June 2003a. Extending area morphology to multivariate images. In: *Proc. 6th IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing*.
- Evans, A. N., August 2003b. Vector area morphology for motion field smoothing and interpretation. *IEE Proc. Vision, Image and Signal Processing* 150 (4), 219–226.
- Evans, A. N., Gimenez, D., 2008. Extending connected operators to colour images. In: *International Conference on Image Processing*. pp. 2184–2187.
- Gatica-Perez, D., Gu, C., Sun, M., Ruiz-Correa, S., September 2001. Extensive partition operators, gray-level connected operators, and region merging/classification segmentation algorithms: theoretical links. *IEEE Trans. Image Processing* 10 (9), 1332–1345.
- Gauch, J., January 1999. Image segmentation and analysis via multiscale gradient watershed hierarchies. *IEEE Trans. Image Processing* 8 (1), 69–79.
- Gibson, S., Bangham, J., Harvey, R., June 2003a. Evaluating a colour scale-space. In: *Proc. British Machine Vision Conference*.
- Gibson, S., Harvey, R., Finlayson, G., 2003b. Convex colour sieves. In: *Proc. 4th International Conference on Scale Space Methods in Computer Vision*. Vol. LNCS 2695. pp. 550–563.
- Gimenez, D., 2007. Colour morphological sieves for scale-space image processing. Ph.D. thesis, University of Bath.

- Gimenez, D., Evans, A. N., 2005. Colour morphological scale-spaces for image segmentation. In: Proc. British Machine Vision Conference. pp. 909–918.
- Gimenez, D., Evans, A. N., April 2008. An evaluation of area morphology scale-spaces for colour images. *Computer Vision and Image Understanding* 110, 32–42.
- Hanbury, A., Serra, J., 2001a. Mathematical morphology in the HLS colour space. In: 12th British Machine Vision Conference.
- Hanbury, A., Serra, J., December 2001b. Morphological operators on the unit circle. *IEEE Trans. Image Processing* 10 (12), 1842–1850.
- Jackway, P., Deriche, M., January 1996. Scale-space properties of the multiscale morphological dilation erosion. *IEEE Trans. Pattern Analysis and Machine Intelligence* 18 (1), 38–51.
- Klein, J. C., 1976. Conception et réalisation d’une unité logique pour l’analyse quantitative d’images. Ph.D. thesis, Nancy University, France.
- Koenderink, J. J., 1984. The structure of images. *Biol. Cybern.* 50, 363370.
- Li, J., Li, Y., 2004. Multivariate mathematical morphology based on principal component analysis: initial results in building extraction. *International Archives for Photogrammetry* 35 (B7), 1168–1173.
- Louverdis, G., Vardavoulia, M., Andreadis, I., Tsalides, P., August 2002. A new approach to morphological color image processing. *Pattern Recognition* 35 (8), 1733–1741.
- Lucchese, L., Mitra, S., 2000. Filtering color images in the xyY color space. In: *IEEE International Conference on Image Processing*. pp. Vol III: 501–503.
- Maragos, P., July 1989. Pattern spectrum and multiscale shape representation. *IEEE Trans. Pattern Analysis and Machine Intelligence* 11 (7), 701–716.
- Martin, D., Fowlkes, C., Malik, J., May 2004. Learning to detect natural image boundaries using local brightness, colour and texture cues. *IEEE Trans. Pattern Analysis and Machine Intelligence* 26 (5), 530–549.

- Martin, D., Fowlkes, C., Tal, D., Malik, J., July 2001. 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. Vol. 2. pp. 416–423.
- Meijster, A., Wilkinson, M., April 2002. A comparison of algorithms for connected set openings and closings. *IEEE Trans. Pattern Analysis and Machine Intelligence* 24 (4), 484–494.
- Meyer, F., 2000. Vector levelings and flattenings. In: J. Goutsias, L. V., Bloomberg, D. (Eds.), *Mathematical Morphology and Its Applications to Image Processing*. Kluwer, pp. 51–60.
- Meyer, F., January 2004. Levelings, image simplification filters for segmentation. *Journal of Mathematical Imaging and Vision* 20 (1-2), 59–72.
- Meyer, F., Maragos, P., June 2000. Nonlinear scale-space representation with morphological levelings. *Journal of Visual Communication and Image Representation* 11 (2), 245–265.
- Monasse, P., Guichard, F., May 2000. Fast computation of a contrast-invariant image representation. *IEEE Trans. Image Processing* 9 (5), 860–872.
- Park, K., Lee, C., November 1996. Scale-space using mathematical morphology. *IEEE Trans. Pattern Analysis and Machine Intelligence* 18 (11), 1121–1126.
- Perona, P., Malik, J., July 1990. Scale space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Analysis and Machine Intelligence* 12 (7), 629–639.
- Plaza, A., Martínez, P., Pérez, R., Plaza, J., 2002. Spatial/spectral endmember extraction by multidimensional morphological operations. *IEEE Transactions on Geosciences and Remote Sensing* 40 (9), 2025–2041.
- Salembier, P., Oct. 2008. Connected operators based on region-trees. pp. 2176–2179.
- Salembier, P., Garrido, L., April 2000. Binary partition tree as an efficient representation for image processing, segmentation, and information retrieval. *IEEE Trans. Image Processing* 9 (4), 561–576.

- Salembier, P., Meyer, F., Brigger, P., Bouchard, L., 1996. Morphological operators for very low bit rate video coding. In: IEEE International Conference on Image Processing. pp. III: 659–662.
- Salembier, P., Oliveras, A., Garrido, L., April 1998. Antiextensive connected operators for image and sequence processing. *IEEE Trans. Image Processing* 7 (4), 555–570.
- Salembier, P., Serra, J., August 1995. Flat zones filtering, connected operators, and filters by reconstruction. *IEEE Trans. Image Processing* 4 (8), 1153–1160.
- Serra, J., 1982. *Image Analysis and Mathematical Morphology*. Academic Press.
- Serra, J., 1988. *Image Analysis and Mathematical Morphology*. Vol. II: Theoretical Advances. Academic Press.
- Serra, J., 2008. Advances in mathematical morphology: Segmentation. Vol. 150 of *Advances in Imaging and Electron Physics*. Elsevier, pp. 185 – 219.
- Serra, J., 2009. Pilot lattices. In: to appear in *Proc. International Symposium on Mathematical Morphology*, Groningen, The Netherlands.
- Soille, P., 2003. *Morphological Image Analysis: Principles and Applications*, 2nd Edition. Springer-Verlag.
- Soille, P., July 2008. Constrained connectivity for hierarchical image decomposition and simplification. *IEEE Trans. Pattern Analysis and Machine Intelligence* 30 (7), 1132–1145.
- Trahanias, P., Venetsanopoulos, A., 1993a. Color edge detection using vector order statistics. *IEEE Trans. Image Processing* 2 (2), 259–264.
- Trahanias, P., Venetsanopoulos, A., February 1996. Vector order statistics operators as color edge detectors. *IEEE Trans. Systems, Man and Cybernetics* 26 (1), 135–143.
- Trahanias, P. E., Venetsanopoulos, A. N., 1993b. Vector directional filters - a new class of multichannel image processing filters. *IEEE Trans. Image Processing* 2 (4), 528–534.

- van den Boomgaard, R., Smeulders, A., November 1994. The morphological structure of images: The differential equations of morphological scale-space. *IEEE Trans. Pattern Analysis and Machine Intelligence* 16 (11), 1101–1113.
- Viero, T., Oistamo, K., Neuvo, Y., 1994. Three-dimensional median-related filters for colour image sequence filtering. *IEEE Trans. Circuits and Systems for Video Technology* 4 (2), 129–142.
- Vincent, L., 1993a. Morphological area openings and closings for grey-scale images. In: *Shape in Picture: Mathematical Description of Shape in Grey-level Images*. pp. 196–208.
- Vincent, L., 1993b. Morphological area openings and closings, their efficient implementation and applications. In: *Proceedings of the “EURASIP Workshop on Mathematical Morphological and its Application to Signal Processing”*, May 1993, Barcelona, Spain. pp. 22–27.
- Vincent, L., 1993c. Morphological gray scale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Trans. Image Processing* 2, 176–201.
- Weber, K., Acton, S., 2004. On connected filters in color image processing. *Journal of Electronic Imaging* 13 (3), 619–629.
- Wilkinson, M., 2008. Connected filtering by reconstruction: Basis and new advances. In: *IEEE International Conference on Image Processing*. pp. 2180–2183.
- Wilkinson, M. H. F., Roerdink, J. B. T. M., 2000. Fast morphological attribute operations using Tarjan’s union-find algorithm. In: *Mathematical Morphology and its Applications to Image and Signal Processing*. Kluwer, pp. 311–320.
- Witkin, A. P., 1983. Scale-space filtering. In: *Proc. 8th Int. Joint Conf. Artificial Intelligence*. pp. 1019–1022.
- Young, N., Evans, A., October 2003. Psychovisually tuned attribute operators for pre-processing digital video. *IEE Proc. Vision, Image and Signal Processing* 150 (5), 277–286.

- Zanoguera, F., Meyer, F., 2002. On the implementation of non-separable vector levelings. In: Talbot, H and Beare, R (Ed.), MATHEMATICAL MORPHOLOGY, PROCEEDINGS. pp. 369–377.
- Zhu, S.-Y., Plataniotis, K., Venetsanopoulos, A., April 1999. Comprehensive analysis of edge detection in color image processing. Optical Engineering 38 (4), 612–625.