IMAGE SEGMENTATION BY AREA DECOMPOSITION OF HSV COMPONENTS

Stephen J. Impey and J. Andrew Bangham

School of Information Systems, University of East Anglia, Norwich NR4 7TJ, UK

Tel +44 (0) 1603 592245.

Email: `sji@sys.uea.ac.uk`

Abstract

Coloured images may be simplified with an area based sieve whilst preserving edges and, usually, colour up to the edges using either the hue, saturation and value (HSV) or red, blue, green (RGB) components. Furthermore, an image may be segmented by area. Applying the sieve to HSV components from a colour image appears to significantly improve the chances of finding objects in a scene, particularly when the objects have different colours. An example of finding cars in a car park scene is presented.

1. INTRODUCTION

It has been shown that image shape can be characterised by granulometry [1] shapes [2, 3]. The asymmetry introduced by using sets of increasing scale openings or closings can sometimes be an advantage but often it renders the decomposition sensitive to noise. An alternative is to use alternating sequential filters [4, 5] but these do not produce an invertible transform to a "granularity domain" and so filters implemented in the granularity domain are not idempotent. Furthermore, the rigid structuring elements cause shape to be distorted at larger scales [6].

Two different refinements can solve these problems, the first is to implement the alternating sequential filter with reconstruction [6, 7], experiment suggests they are invertible [8] and probably preserve scale-space causality. The second is to operate on *graphs* and use *area* based operators [9, 10]. The latter have been shown to both preserve scale-space causality and produce invertible transforms [11]. They are also reject noise robustly [see elsewhere in this volume] when compared to classical scale-space preserving linear filters [12, 13, 14, 15].

We call the alternating sequential filter with these properties, sieves. Not all alternating sequential filters are sieves and not all sieves are alternating sequential filters, e.g. [16]. Here, area based sieves are used to simplify and segment colour images in the hue, saturation, value (HSV) representation.

2. DEFINITIONS:

Figure 1. Shows a greyscale representation of a colour carpark image.

A connected set, or area, sieve in N dimensions $\Phi_m : Z^N \rightarrow Z^N$ operating on graphs [17] is used. An image can be described by a connected graph, *H=(V,E)* where *V* is the set of vertices and *E* the set of pairs that describe the edges. Let $C_r(H)$ denote the set of connected subsets, when N=2 these represent areas and when N=3, volumes.

$$\Phi_m(X) = \phi(\Phi_{m-1}(X)), \quad \text{where } \Phi_0(X) = X \qquad 1$$

The operator ϕ_m may be an open/close (*M-sieve*) *max(min(min(max($C_r(H)$))))* or area close/open (*N-sieve*) *min(max(max(min($C_r(H)$))))* or a recursive equivalent, each operating on graphs with a set of connected subsets with *m* elements [11]. It is emphasised that ϕ_m operates on connected sets without regard to shape, there is no structuring element. The granularity of a signal is obtained from

$$Gran_R(X)(m) = (R_m(X)) - (R_{m+1}(X)) \qquad 4$$

The set of granules *{G}* represent the non-zero sets in the granule functions and the sieve transform maps the signal into a set of granules

$$S : Z \rightarrow \{G\} \qquad 5$$

The inverse is S^{-1}, $S^{-1}S=1$, can be obtained by summing the re-expanded granules. Area channels are obtained by keeping just those granules in a range of scales. In an analogous fashion, area opening channels are obtained by keeping just those "grains", obtained from an area granulometry [9]. The method is similar to that of the sieve except that the operator ϕ_m is area open to a particular scale *m*.

Figure 2. The RGB images associated with Figure 1.

Figure 3. Images in Figure 2 area sieved to remove all extrema of area less than 500.

Figure 4. HSV from Figure 1 area sieved to remove all extrema of area less than 500 and rebuilt to form colour image shown in greyscale

Suitable algorithms for computing area sieves and openings can be based on those described in [9, 17].

3. RESULTS:

First the problem of simplifying coloured images is addressed by using area sieves. Figure 1 shows the intensity of an image of a number of coloured cars and Figures 2 and 3 show the associated red, blue, green (RGB) and the result of area sieving to remove all regional extrema of area less than 500. The recombined image (not shown) shows good colour, shape and edge preservation. Similar results are obtained by sieving the HSV images. Figure 4, shows the rebuilt image in greyscale. This is unlike Gaussian filtering of HSV images where, not only are edges smoothed, but the colours can change markedly at the boundaries of objects (not shown).

Secondly the problem of segmenting the image to locate objects of interest, such as the cars in Figure 1. These are

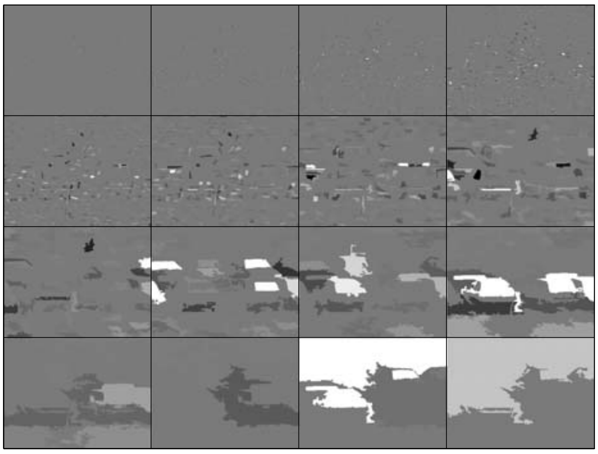

Figure 5. Area granularity channels obtained by sieving the intensity (V) from Figure 1. Small scale areas forming local extrema are in the small scale channels, top left and largest scale extrema are shown in the bottom right. Maxima are paler than the background and minima darker.

often distinguished by a combination of colour and size. The cars in Figure 1 are represented by between about 1000 and 5000 pixels. The area sieve can be used to decompose an image into a set of increasing scale area

Figure 6. Left panel, V from Figure 1. Right panel, area sieved to keep just those local extrema of areas between 1000 and 4000 pixels.

channels, equation 4. By combining, equation 5, the granularity functions in this range of scales it should be possible to find objects of the correct size.

Figure 5 shows granule functions in the ranges, area=1, 2 to 3, 4 to 5, 6 to 9, 10 to 17, 18 to 35, ..., obtained by sieving the V channel from Figure 1. The intensity of the images has been adjusted to make as many granules as possible visible. What were maxima are shown paler than the background and minima darker than the background. Large area scale extrema are to be found in larger scale channels.

On this segmentation criterion alone, it is not possible to segment out the cars perfectly, however, by just concentrating on the appropriate scale ranges (channels 10 upwards) the process can limit the number of candidate segments. Figure 6 right panel, shows just those objects that form extrema with areas in the range 1000 to 4000 pixels. Clearly large chunks of some of the cars are well segmented. In the end it will be necessary to identify the cars by generating a "signature" or feature vector for each

Figure 7. Left panel, H from Figure 1 that has been rotated before taking the cosine to emphasise blues. Right panel, the result of taking just the maxima of scales 1000 to 4000 pixels obtained from an area opening decomposition.

candidate segment and use that to perform the classification. However, it is possible to further reduce the

Figure 8. Left panel, H from Figure 1 that has been rotated before taking the cosine to emphasise reds. Right panel, the result of taking just the maxima of scales 1000 to 4000 pixels obtained from an area opening decomposition.

number of candidate signatures by area sieving the colour information in H.

The colour information, H component, is an angle so red (Matlab image processing toolbox) falls on the 0 to 360 degrees divide. As result the colour that forms a regional maximum or minimum is arbitrary. Here, we choose a colour of interest, blue for example, then rotate the H values such that the colour of interest is centred on 0 degrees then window the values by taking the cosine. This causes features of the required colour to become regional maxima. These can then be segmented out by picking the appropriate scale outputs from an area opening decomposition. Figure 7 right panel, shows the result of selecting blue features in the appropriate area range. No attempt has been made to selected on the intensity of blue, so all blue and bluish segments are shown. The meaning associated with this is left to the classifier to interpret. Likewise, Figure 8 shows the result of picking out red objects.

4. DISCUSSION:

Diffusion based filters preserve scale-space causality but when used to simplify images they simultaneously blur edges and round corners. When filtering in the HSV domain blurring of the colour component leads to colour fringing, the arithmetic “synthesis” of colours not present in the original image. Certain morphological filters also preserve scale-space but with rigid structuring elements they too round corners and the results are prone to distortion due to noise and clutter. Area-based sieves appear to reduce these problems. They can be used to satisfactorily simplify colour images by separately filtering either the RGB or the HSV components. Since every pixel in a sieve output is an element of the input signal, the output colours must exist in the input.

The problem of selecting features by colour is addressed by rotating the colours so that those of interest have the largest numerical values. This means that a decomposition using openings, which selectively separate out those regions that form regional maxima, can localise the relevant areas. Here, the colours were selected by analysing the colours in the image. However, it would be reasonable to tune the colour maps to match those of, say, the cones in a human vision system. A similar approach to the saturation component was not helpful in distinguishing cars but might well be useful for distinguishing textured colours such as balls of wool. This coupled with the area openings and sieves might prove effective for a more general image analysis system.

5. REFERENCES:

1. Matheron, G. 1975. Random Sets and Integral Geometry. London: Wiley.
2. Pitas, I., and A. N. Venetsanopoulos. 1990. Morphological shape decomposition. IEEE Trans. on Pattern Anal. and Machine Intell. vol. 12: pp 38-45.
3. Maragos, P. 1989. Pattern spectrum and multiscale shape representation. IEEE Trans. on Pattern Anal. and Machine Intell. vol. 11: pp 701-716.
4. Sternberg S R. 1986. Grayscale Morphology. Comp. vision, graphics and image proc. vol. 35: pp 333-355.
5. Serra, J. 1988. Image Analysis and Mathematical Morphology vol 2. Academic Press.
6. Salembier, P., and M. Kunt. 1992. Size sensitive multiresolution decomposition of images with rank order based filters. Signal Processing vol. 27: pp 205-241.
7. Serra, J., and P. Salembier. 1993. Connected operators and pyramids. SPIE Proceedings on Image Algebra and Mathematical Morphology 2030: 65-76.
8. Bangham, J. A., and T. G. Campbell. 1993. Sieves and wavelets: multiscale transforms for pattern recognition. Proceedings IEEE Workshop on Nonlinear Signal Processing, pp 1.1-4.1 -1.1-4.6. Tampere Finland:
9. Vincent, L. 3 1992. Morphological area openings and closings of greyscale images. Workshop on Shape in Picture, NATO.
10. Vachier, C., and F. Meyer. 1995. Extinction value: a new measurement of persistence. IEEE Workshop on Nonlinear and Image Signal Processing, I. Pitas, 254-257. June. IEEE.
11. Bangham, J. A., R. W. Harvey, P. D. Ling, and R. V. Aldridge. 2 1996. Nonlinear scale-space in many dimensions. European Conference on Computer Vision, IEEE.
12. Witkin, A. P. 1983. Scale-space filtering. 8th Int. Joint Conf. Artificial Intelligence, 1019-1022. IEEE.
13. Koenderink, J. J. 1984. The structure of Images. Biological Cybernetics 50: 363-370.
14. Babaud, J., A. P. Witkin, M. Baudin, and R. O. Duda. 1986. Uniquenss of the Gaussian Kernal for Scale-Space Filtering. IEEE Transactions on Patern Analysis and Machine Intelligence Vol. 8: pp 26-33.
15. Perona, P., and J. Malik. 1990. Scale-space and edge detection using anisotropic diffusion. IEEE Trans. Pattern Analysis and Machine Intelligence 12: 629-639.
16. Bangham, J. A., P. Ling, and R. Young. 1 1996. Multiscale recursive medians, scale-space and sieve transforms with an inverse. IEEE Transactions on Image Processing Aug.
17. Vincent, L. 1989. Graphs and mathematical morphology. Signal Processing 16: 365:388.