Evaluating a colour morphological scale-space

Stuart Gibson, J. Andrew Bangham and Richard Harvey
School of Computing Sciences
University of East Anglia
Norwich, NR4 7TJ, UK
{s.e.gibson,a.bangham,r.w.harvey}@uea.ac.uk

Abstract

Scale-spaces are usually associated with the diffusion equation operating on greyscale images. This paper presents a scale-space that is based on graph morphology and works in colour. The system is shown to be a natural extension of the existing greyscale graph-morphology scale-spaces. The system is evaluated by comparing it to a specially collected database of ground-truth images created by human subjects.

1 Introduction

Scale-spaces are traditionally associated with weighted means of intensities in an image created via the diffusion equation [1–5]. However there is increasing interest in scale-space systems that form their outputs from order statistics of the image: the, so called, morphological scale-spaces. Two, rather different approaches have been described. There are those based on a series of structuring elements of increasing scale, [6] and [7] for example, which are appropriate when the structuring element shape can be tuned to the known shape of objects in the image, and graph-based techniques, [8] for example, which work best when the image objects are unknown shape which is the case in natural scenes.

Graph-morphology scale-spaces have their basis in algorithms discovered some-time ago [8–11] and, more recently, [12, 13]. Their attractiveness stems from low computational complexity, robustness to noise [14] and the observation that they produce regions that may be related to objects in the image. There are slight variations in the algorithms but they all proceed something like this:

1. Extremal regions are identified and labelled with their areas;
2. The regions of area 1 are merged to their neighbours with the closest intensity;
3. The regions of area 2 are merged to their neighbours, etc...

The algorithm terminates when there are no more nodes left to merge. In [15] the extrema are either maxima or minima forming decompositions known as the max or min tree. In [8] maxima and minima may be handled together giving a scale-tree. In [16] successive differences between simplified images are called granules[1]. The tree arises because smaller granules are contained within larger ones so the nodes represent regions and the

[1] An analogy with granulometries.

BMVC 2003 doi:10.5244/C.17.81

tree-edges represent containment. The tree forms a structure that can be used for image matching and retrieval [17], stereo matching [18] and motion detection [15].

There are several colour scale-spaces based on smoothed derivatives [19, 20]. Likewise morphological operators can, with some care, be extended to handle colour images [21] but, so far, graph-morphology scale-spaces have had only quite restricted extensions into colour [16, 22] by application to each channel for example. This paper briefly reviews a new extension to colour images [23] and compares its effectiveness to some ground-truth results.

2 Extension to colour

An effective extension of existing graph methods must tackle two problems: (i) a definition of extrema in colour and (ii) a rule for deciding on the colour similarity of two regions. In greyscale graph-morphology and image is represented by a graph $G = (V, E)$ where V is the set of vertices that index the pixels and E is a set of pairs indicating the adjacency of pixels. If $C_r(G)$ is the set of all connected-sets with r elements then $C_r(G,x) = \{\xi \in C_r(G) | x \in \xi\}$ with $x \in V$ is the set of size-r connected subsets that contain pixel x. If the image has intensity $f(x), x \in V$ then the operators $\psi_r f(x) = \max_{\xi \in C_r(G,x)} min_{u \in \xi} f(u)$ and $\gamma_r f(x) = \min_{\xi \in C_r(G,x)} max_{u \in \xi} f(u)$ have the effect of "slicing off" regional maxima (ψ_r) or minima (γ_r) or both (via combined operators $\mathscr{M}_r = \gamma_r \psi_r$ or $\mathscr{N}_r = \psi_r \gamma_r$). Thus in conventional greyscale morphology the extrema are regions that are brighter than or darker than their neighbours and the merge rule is to merge to the next most extreme grey-level.

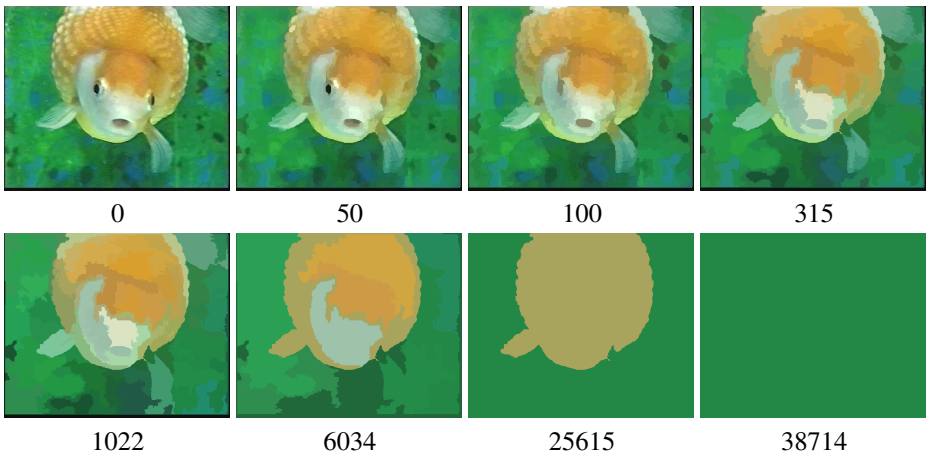

Figure 1: Example decomposition of a 288×352 image at the scales indicated. Scale is the number pixels in the region represented by a node.

In colour there is no natural ordering but an extrema definition is possible [23] either by considering points that are furthest from the vector median [24] or, as in [23], by fitting a convex-hull [2] in colour space to the colours under consideration and defining any points

[2] Fast, divide and conquer, algorithms exist that have complexity of roughly $N \log N$ where N is the number of points.

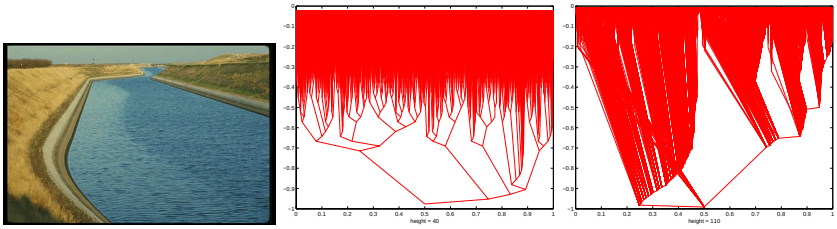

Figure 2: Example image (left) and its associated colour tree (centre) and greyscale tree (right).

on the convex hull to be extreme. This is tidy definition because extrema remain extrema under any non-degenerate linear transformation of the colour space or any monotonic transformation of the axes.

To form a scale-space processor it must merge smaller regions into larger ones without introducing additional extrema. Achieving this in greyscale is tricky and requires a two-pass algorithm using the $\mathscr{M}$ or $\mathscr{N}$ operator defined earlier in which maxima (resp. minima) are merged to their neighbour with the closest greyscale, then minima (resp. maxima) likewise. The colour system merges to the neighbour with the closest Euclidean colour distance, but because two different colours can have identical distances and there is no longer the possibility of handling maxima and minima separately, we need two modifications. Firstly, it is necessary to resolve tied distances so that neighbouring regions with identical colour distances are further ordered by computing the difference of their luminance $L = (R+G+B)/3$. If this does not resolve the tie then regions are ordered by their G, R and B values. Secondly, if, after merging, new extrema have been introduced then they are removed by reapplying the processor up to the current scale. These final iterations ensure idempotence. After this step all area-1 extrema will have been removed and we now move to area-2 and repeat the process.

An example decomposition is shown in Figure 1. Figure 2 shows an example image that has been decomposed into its colour and greyscale trees (using the $\mathscr{M}$ operator). A noticeable difference between this tree and a conventional greyscale graph-morphology example is the lack of small children off the root which avoids the need for methods that prune the tree to remove large numbers of small nodes with low tree-depth (as in [18] for example). Figure 3 expands the root region of the colour tree shown in Figure 2 and demonstrates that the nodes, which are the differences between successive simplifications, form connected sets with no holes. In this example it is a binary tree but it is not constrained to be so. One of the advantages of the colour processor derives from the fact that it identifies more extrema than the greyscale equivalent so allows the possibility of segmenting more regions. In a test of 50 images taken from the MPEG-7 Common Colour Dataset [25] we found that the probability of a pixel being a greyscale extrema was 0.17 whereas the probability of it being a colour extrema was 0.56.

The scale-space processors described so far are all examples of *primary* vision systems [26]. They may be applied to many problems and the evaluation is therefore domain-dependent. For diffusion-based systems the problem is especially acute because the outputs are not useful without further processing. The evaluation of graph-morphology based systems is easier because the outputs consist of regions so have some similarity with a

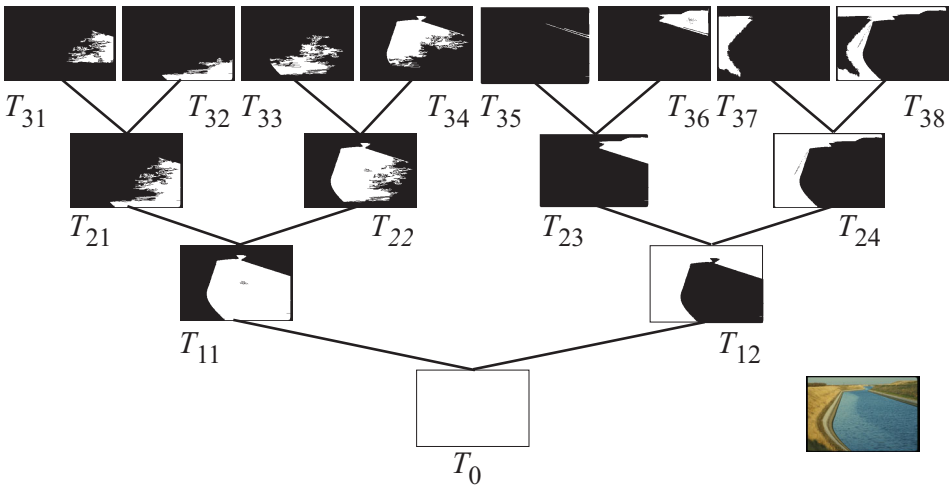

Figure 3: A schematic of the full colour tree of the colour image shown on the far right. The top fifteen nodes of the tree have been expanded are shown with their associated segments in white.

segmentation for which evaluation methods exist.

3 Evaluations

Our aim is to evaluate the extent to which semantically meaningful segments are to be found in the colour scale tree. This is similar to the evaluation of segmentation systems [27–29] but note the reservations in [30, 31]. Zhang [27] introduces a taxonomy of segmentation methods which includes methods for evaluating the segmentation when there is no ground truth available. Nowadays ground truth is available so most workers are interested in "empirical evaluation" ([31], [28] and [29] for example). The various methods consider firstly a way of measuring the error between two individual segments and secondly how these errors should be combined to give an error between two segmentations. In practice the appropriate choice for these error distances is influenced by the type of ground-truth data available[3] which we discuss now.

The Sowerby database, which is available under licence from BAE systems, contains scenes of, what appear to be, roads around Bristol. Pixels in this database are classified into one of 11 types: sky; vegetation; road surface marking; road surface; road border; building; bounding object; road sign; telegraph pole; illumination shadow; mobile object. These classes were presumably meaningful when the data were collected but they now seem domain dependent. A more serious criticism is that the hand-segmentation was obtained by grouping regions from an automatically over-segmented image which means that semantically important regions are missing and that the ground truth is actually an edited output from an automatic segmentation algorithm.

[3] In [30] for example the distance measure is based on a medical imaging problem in which the classes are known.

A more useful database for our purposes is the library of segmented Corel images that has become available during the writing of this paper [29]. These data inherit the problems of the Corel database but with 11,595 segmentations of 1020 images it represents a valuable resource that has yet to be fully explored. This paper adopts a similar

Figure 4: Some example images from the test database [25]

approach to [29] but instead uses subset of 50 images taken from the MPEG-7 common colour dataset [25]. Unlike the Corel images which are all the same size and apparently taken by professional photographers, the MPEG-7 images, Figure 4, come form a variety of sources including television broadcasts and image libraries. Some of the images are repeated (to allow analysis of the repeatability of people's segmentations) and some are inverted to test if high-level knowledge of the subject improves the segmentation. Each image has been segmented by 38 undergraduates enrolled on a final-year computer vision module. Each participant was allowed to choose up-to 16 segments that they thought represented "semantically meaningful regions". No guidance was given on what this term might mean and unlike [29] there were no strictures that the regions should have equal importance. Each segment is stored as a separate bit-plane so, unlike the data in [29] for example, there is no restriction that segments must not overlap or contain no holes. Students also constructed their own segmentation tool so there is potentially more variability than in [29] but, in practice, all students used the Matlab `roipoly` function so the segments are polygonal. Example segments are shown in Figure 5.

4 Methods

Given two segmentation, images S_1 and S_2, [29] contains some careful reasoning in favour of the *local consistency error* $\mathrm{LCE}(S_1,S_2) = \frac{1}{n}\sum_i \min\{E(S_1,S_2,p_i)E(S_2,S_1,p_i)\}$ and the *global consistency error* $\mathrm{GCE}(S_1,S_2) = \frac{1}{n}\min\{\sum_i E(S_1,S_2,p_i)E(S_2,S_1,p_i)\}$ where $E(S_1,S_2,p)$ is the *local refinement error*,

$$E(S_1,S_2,p_i) = \frac{R(S_1,p_i)\backslash R(S_2,p_i)}{|R(S_1,p_i)|} \tag{1}$$

with $R(S_j,s_i)$ denotes the set of pixels in segmentation S which are in the same segment as pixel p_i. $\backslash$ means set difference and $|\cdot|$ means the cardinality of the set. If in (1)

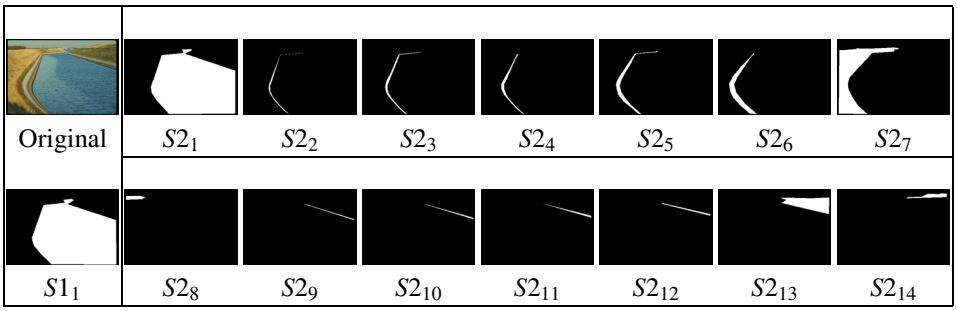

Figure 5: Showing an original image (top left). One student's segmentation consisting of one labelled region (bottom left) and another student's segmentation of the same image with fourteen regions (right).

$R(S_2, p_i)$ is a subset of $R(S_1, p_i)$ then $E(S_1, S_2, p_i)$ is zero. In a morphological scale-space it is commonplace to find that one tree node matches well with a ground-truth region and hence, because the tree encodes enclosure, all of its smaller-scale children will be subsets of the region (denote this $E(T,G) = 0$ where T is a tree-derived segment and G is the ground-truth). The larger-scale parents of the well fitting node will have $E(G,T) = 0$ which means that LCE $= \min(E(T,G), E(G,T))$ will be mostly zero except at the scale where the node fits best. GCE gives are more intuitive result with the error proportional to the ratio of the number of pixels that differ to the number of pixels in the largest region. However, in practice GCE is slow to compute so instead we adopt a faster method that gives similar results: the normalised XOR error. If a tree contains N_T regions, $T_1 ... T_{N_T}$ and the ground-truth image contains N_G regions $G_1, ..., G_{N_G}$ then the normalised XOR error is

$$E_{\oplus}(T,G) = \frac{1}{N_G} \sum_{m=1}^{N_G} \min_n \frac{|T_n \oplus G_m|}{|T_n| + |G_m|} \tag{2}$$

where $\oplus$ represents bitwise exclusive-or. Table 1 shows this error measured between some of the tree segments shown in Figure 3 and the human segments shown in Figure 5. The error measure is not a metric because $E_{\oplus}(T,G) \neq E_{\oplus}(G,T)$ which is deliberate: we do not insist that all tree nodes are meaningful rather we want to know if the tree contains semantically meaningful regions. However, if necessary the measure can be symmetrised to give $E_{\oplus\oplus}(S_1,S_2) = (E_{\oplus}(S_1,S_2) + E_{\oplus}(S_2,S_1))/2$ which is useful when measuring the error between ground-truth segmentations. Note that $E_{\oplus\oplus}$ allows the possibility of one-to-many segment matching.

5 Results

Figure 6A shows the histogram of $E_{\oplus\oplus}(G_1,G_2)$ for ground-truth segmentations of the same image. There are 38 ground-truths and 50 images so the histogram contains $38 \times 37 \times 50/2 = 35,150$ data points ignoring self-matches. The error is predominantly zero nevertheless there is significant variation between images reiterating that it is essential to make measurements over a range of images. The MPEG-7 common colour dataset contains images of different sizes so, unlike [29] comparison between images is not possible.

	$S1_1$	$S2_1$	$S2_2$	$S2_3$	$S2_4$	$S2_5$	$S2_6$	$S2_7$	$S2_8$	$S2_9$	$S2_{10}$	$S2_{11}$	$S2_{12}$	$S2_{13}$	$S2_{14}$
T_{11}	0.01	0.01	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{12}	0.99	1.00	0.98	0.95	0.95	0.91	0.87	0.53	0.97	0.98	0.98	0.98	0.97	0.79	0.95
T_{21}	0.50	0.50	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{22}	0.21	0.21	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{23}	0.98	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.95	0.94	0.93	0.92	0.47	0.87
T_{24}	1.00	1.00	0.97	0.93	0.92	0.87	0.82	0.38	0.96	1.00	1.00	1.00	1.00	1.00	1.00
T_{31}	0.73	0.73	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{32}	0.70	0.70	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{33}	0.62	0.62	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{34}	0.41	0.40	0.99	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
T_{35}	0.98	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.94	0.98	0.98	1.00	1.00	1.00
T_{36}	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.97	0.92	0.90	0.86	0.22	0.78
T_{37}	1.00	1.00	1.00	1.00	1.00	1.00	0.98	0.16	1.00	1.00	1.00	1.00	1.00	1.00	1.00
min	0.01	0.01	0.95	0.89	0.89	0.80	0.75	0.16	0.94	0.94	0.92	0.90	0.86	0.22	0.78
mean	0.01	0.71													

Table 1: Normalised XOR error measured between human segments in Figure 5 and the tree segments Figure 3

Figure 6B is the same $E_{\oplus\oplus}$ but now with the minimum operator in (2) replaced with a random selection from among the other segments. There are 20 bins [32] so $\chi^2 \approx 2600$ with 19 degrees of freedom which means we can confidently reject the hypothesis that the two histograms are identical. In other words the error measure satisfies the desideratum that it should distinguish between human segmentations of the same subject and random matches. To examine the effectiveness of the tree we measure the distance between the

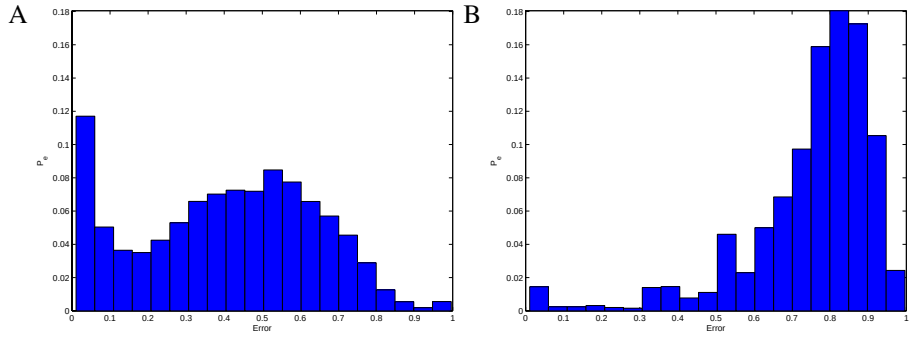

Figure 6: PDF estimate of the segmentation error for best-fit hand-segmentations of the same image (A) and randomly-selected segment pairs from hand-segmentations of the same image (B). The bin-sizes are selected using [32].

tree and the human segmented ground-truth (Figure 7A). Each tree contains many regions so to reduce computation to a reasonable level we extracted the first 52 regions ignoring the root from a breadth-first parsing of the scale-tree starting from the root. There are 50 images and 38 ground-truth segmentations of each image giving 1900 measurements. The primary difference between this distribution and the ground-truth segments compared

with themselves (Figure 6A) is the mode at zero in the ground-truth segments. Figure 7B shows the distribution of the distances of the top 52 regions from the tree to the ground-truth but with each segment in the ground-truth translated randomly within in the image. This is quite a demanding test since we constrain the segments to always lie within the image boundaries thus large segments cannot move every far. Again a χ^2 test allows us to confidently reject the hypothesis tat Figure 7A and Figure 7B are drawn from the same distribution. Figure 7C shows the scatter plot of this distance (tree to randomly perturbed ground-truth) to the ground-truth distance showing the common image introduces correlation ($\rho = 0.84$) between the distances. Measured over all the data, the probability that

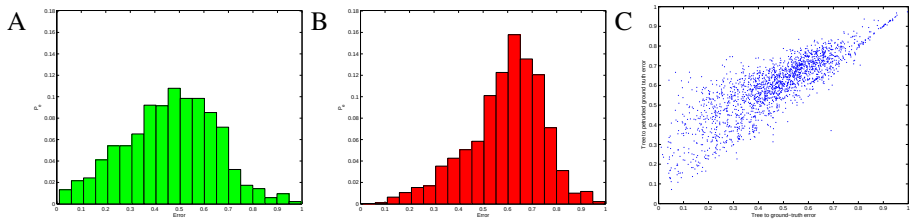

Figure 7: The distance distribution of the tree versus corresponding ground-truth segmentations (A), against perturbed ground-truth (B) and (C) the scattergram of tree to perturbed ground-truth (ordinate) versus tree to ground-truth (abscissa).

the tree to ground-truth data distance is smaller than the tree-to-perturbed ground-truth distance is 0.92. Since there are equal numbers of ground-truth and perturbed ground-truth one can model selection-at-random as binomial ($p = 0.5$) so that the probability of this figure arising by chance is negligibly different from zero. In a similar contest where we compare the distance of the tree to ground-truth with the distance between ground-truths on the same image with the segments chosen randomly the probability of the tree to ground-truth distance being the winner is 0.89 whereas where the contest is between the tree and ground-truth and the best-matching ground-truths the tree wins with a probability of 0.43 indicating that the tree does not model a ground-truth as well as another human segmenter with their segments put into minimum-error correspondence but that the tree beats human segmenters with their segments put into random correspondence.

6 Discussion

This paper has introduced a new colour scale-space processor that was devised by introducing colour equivalents of the extrema-identification step and the merging step. The new method of identifying extrema has considerable appeal because it is invariant to typical colour space transformations. The merge rule is less satisfactory for two reasons: it is not invariant to the colour space and it forces a recursive implementation in which we recheck for extrema. Nevertheless the algorithm executes in a few minutes for a 384x256 pixel image(24-bit pixels) on a domestic PC.

The tree provides agreement with the ground-truth segmentations but there is better agreement among the people that produced the ground-truth. There are several reasons for this. Firstly, we choose the significant regions from the colour tree in a very simple way – the top 52 regions ignoring the root. This means that we could have missed significant

regions further down the tree. Secondly, the tree uses only colour where it is reasonable to assume that the human segmenters used other information too. In at least one of the images in our collection the object shares a colour level-set with the background so segmentation on colour alone is impossible.

For the future we hope to make our dataset publically available as an adjunct to [29] (our dataset has more human segmentations per image) so that other researchers may compare their work with ours.

7 Acknowledgements

We would like to thank Luke Jefferson for his computer simulations.

References

[1] T. Iijima, "Basic theory of pattern normalization (for the case of a typical one-dimensional pattern," *Bulletin of the Electrotechnical Laboratory*, vol. 26, pp. 368–388, 1962.

[2] A. P. Witkin, "Scale-space filtering," in *8th Int. Joint Conf. Artificial Intelligence*, pp. 1019–1022, IEEE, 1983.

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

[4] B. M. ter Harr Romeny, ed., *Geometry-driven diffusion in Computer vision*. Dordrecht, Netherlands: Kluwer Academic, 1994.

[5] J. Koenderink, "The structure of images," *Biological Cybernetics*, vol. 50, pp. 363–370, August 1984.

[6] P. T. Jackway and M. Deriche, "Scale-space properties of the multiscale morphological dilation-erosion," *IEEE Trans. Patt. Anal. Mach. Intelli.*, vol. 18, pp. 38–51, January 1996.

[7] R. van den Boomgaard and A. Smeulders, "The morphological structure of images: the differential equations of morphological scale-space," *IEEE Trans. Patt. Anal. Mach. Intelli.*, vol. 16, pp. 1101–1113, November 1994.

[8] J. Bangham, R. Harvey, P. Ling, and R. Aldridge, "Morphological scale-space preserving transforms in many dimensions," *Journal of Electronic Imaging*, vol. 5, pp. 283–299, July 1996.

[9] M. Turner and N. Wiseman, "Efficient lossless image contour coding," *Computer Graphics Forum*, vol. 15, no. 2, pp. 107–118, 1996.

[10] J. A. Bangham, "Data-sieving hydrophobicity plots," *Anal. Biochem*, vol. 174, pp. 694–700, 1988.

[11] P. Salembier and J. Serra, "Flat zones filtering, connected operators, and filters by reconstruction," *IEEE Transactions on Image Processing*, vol. 4, pp. 1153–1160, August 1995.

[12] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide baseline stereo from maximally stable extremal regions," in *BMVC*, (Cardiff, UK), pp. 384–394, September 2002.

[13] S. Acton and D. Mukherjee, "Scale space classification using area morphology," *IEEE Transactions on Image Processing*, vol. 9, pp. 623–635, April 2000.

[14] R. Harvey, J. A. Bangham, and A. Bosson, "Scale-space filters and their robustness," in *Scale-Space Theory in Computer Vision, First International Conference, Scale-Space'97, Utrecht, The Netherlands* (B. M. ter Haar Romeny, L. Florack, J. J. Koenderink, and M. A. Viergever, eds.), vol. 1252 of *Lecture Notes in Computer Science*, pp. 341–344, Springer, July 1997.

[15] P. Salembier and L. Garrido, "Binary partition tree as an efficient representation for image processing, segmentation and information retrieval," *IEEE Transactions on Image Processing*, vol. 9, pp. 561–576, April 2000.

[16] J. A. Bangham, J. Hidalgo, R. Harvey, and G. Cawley, "The segmentation of images via scale-space trees," in *BMVC*, (Southampton, UK), pp. 33–43, September 1998.

[17] D. Dupplaw and P. Lewis, "Content based image retrieval with scale space object trees," in *SPIE Proceedings: Storage and Retrieval for Media Databases 2000* (M. Yeung, B.-L. Yeo, and C. Bouman, eds.), vol. 3972, pp. 253–261, 2000.

[18] K. Moravec, R. Harvey, and J. Bangham, "Scale trees for stereo vision," *IEE Proceedings: Vision, Image and Signal Processing*, vol. 147, pp. 363–370, August 2000.

[19] B. ter Haar Romeny, J.-M. Geusebroek, P. van Osta, R. van den Boomgaard, and J. Koenderink, "Color differential structure," in *Scale-Space 2001*, p. 353 ff, Springer, 2001.

[20] R. Kimmel, N. A. Sochen, and R. Malladi, "From high energy physics to low level vision," in *Proc. First Int. Conf. on Scale-space theory*, pp. 236–247, 1997.

[21] M. Comer and E. Delp, "Morphological operations for color image processing," *Journal of Electronic Imaging*, vol. 8, no. 3, pp. 279–289, 1999.

[22] P. Salembier and L. Garrido, "Binary partition tree as an efficient representation for filtering, segmentation and information retrieval," in *International Conference on Image Processing (ICIP'98)*, vol. II, (Chicago, USA), pp. 252–256, October 1998.

[23] S. Gibson, R. Harvey, and G. Finlayson, "Convex colour sieves," in *Proceedings 4th International Conference, Scale-Space 2003* (L. D. Griffin and M. Lillholm, eds.), vol. 2695 of *Lecture Notes in Computer Science*, pp. 537–549, Springer, June 2003.

[24] J. Astola, P. Haavisto, and Y. Neuvo, "Vector median filter," *Proceedings of the IEEE*, vol. 78, pp. 678–689, April 1990.

[25] "Common datasets and queries in MPEG-7 color core experiments," Tech. Rep. ISO/IEC JTC1/SC29/WG11/MPEG99/M5060, MPEG-7 Standards Body, October 1999.

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

[27] Y. Zhang, "A survey on evaluation methods for image segmentation," *Pattern Recognition*, vol. 29, no. 8, pp. 1335–1346, 1996.

[28] M. Everingham, H. Muller, and B. Thomas, "Evaluating image segmentation algorithms using monotonic hulls in fitness/cost space," in *BMVC*, (Manchester, UK), pp. 363–372, September 2001.

[29] D. Martin, C. Fowlkes, D. Tal, and J. Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *ICCV*, vol. 2, (Vancouver, Canada), pp. 461–423, July 2001.

[30] X. Yu, T. Bui, and A. Kryzak, "Robust estimation for range image segmentation and reconstruction," *IEEE Trans. Patt. Anal. Mach. Intelli.*, vol. 16, pp. 530–538, May 1994.

[31] A. Hoover, G. Jean-Baptiste, X. Jiang, P. J. Flynn, H. Bunke, D. B. Goldgof, K. Bowyer, D. W. Eggert, A. Fitzgibbon, and R. B. Fisher, "An experimental comparison of range image segmentation algorithms," *IEEE Trans. Patt. Anal. Mach. Intelli.*, vol. 18, pp. 673–698, July 1996.

[32] D. Scott, "On optimal and data-based histograms," *Biometrika*, vol. 66, pp. 605–610, December 1979.