Path based colour image segmentation

Clement Fredembach and Graham Finlayson
School of computing Sciences
University of East Anglia, Norwich, U.K.
{***cf,graham***}***@cmp.uea.ac.uk***

Abstract

The 2 pass raster segmenter is simple, fast and is often quoted in the literature. Unfortunately, it tends to oversegment images even in the presence of small amounts of noise. In this paper we present a generalization of this approach where we discover regions by taking multiple random paths through an image. This approach fares better but still over segments an image. Yet, an analysis of region density shows that the underlying image structure can be discovered from the path based segmentation. Indeed, the discovered edges are comparable to those discovered by the widely used mean shift algorithm.

1. Introduction

From Land's Retinex to scale-space processing [1], path-based methods have often been used with success in image processing and computer vision. Those paths can usually be divided in three categories: short random walks (as in [2] and [3]), partially complete (in the sense that they cover the entire image), such as Frankle MacCann spiral path for retinex [4] or complete as for raster paths for segmentation [5]. Most examples of complete paths are instances of the more general class of Hamiltonian paths, whose definition is "A path in a graph such that every vertex is visited once and once only" [6] (or, in terms of images we visit each pixel once and visit all pixels in the image). A framework to derive random Hamiltonian paths in images has recently been proposed in [7] and the authors have used them successfully to remove shadows from images. In this paper we present a framework for color image segmentation based on Hamiltonian paths, in order to obtain a simple yet accurate edge map of a color image.

The problem of image segmentation has been studied for a long time and has spawned a wide variety of approaches ([8],[9] and [10] among others). The best performing algorithms currently make use of a combination of color, texture and scale features and usually have many of parameters that can be adjusted for optimum segmentation. As a result, many of these algorithms are either difficult to implement and/or computationally expensive to use. One of the goals of this paper is to develop a segmentation algorithm that is equally powerful to antecedent methods but due to its simple path based implementation is simpler and easier to implement.

We first introduce the 2-pass raster segmentation commonly cited in the literature. We then review the Fredembach and Finlayson method to obtain random Hamiltonian paths [7]. We then show how to use simple first and second order statistics on color channels calculated along a path to group similar pixels. After a single Hamiltonian path through the image there are many line like segmented structures (as oppose to desired regions). We group these linear structures and discover arbitrarily shaped regions by repeating our segmentation along different paths where now we group together the linear structures. After a small number of path segmentations we can discover arbitrarily shaped regions. We provide a detailed discussion of the convergence of our method.

Section 2 presents the standard two pass path based raster segmentation. We also review how arbitrary random Hamiltonian paths can be created and discuss our path based segmentation idea. In section 3 we look at how path based segmentation works in experiments and this allows us to elaborate on the basic algorithm. Results on real images are presented in section 4 for our path based approach and for the widely used mean shift algorithm. For the images tested both algorithms provided broadly similar performance, with the former being delivered much more quickly. The paper concludes in section 5.

2. Background

2.1. Raster Segmentation (Sequential Labelling)

Sequential labelling is a technique used in computer vision for efficient segmentation of images [5]. Two orthogonal raster paths (such as the ones shown in figure 1) are used sequentially to connect pixel belonging to a same region. This method, first based on binary images, where determining the connectivity of pixels is straightforward [11] was then extended to encompass grayscale and color images [10].

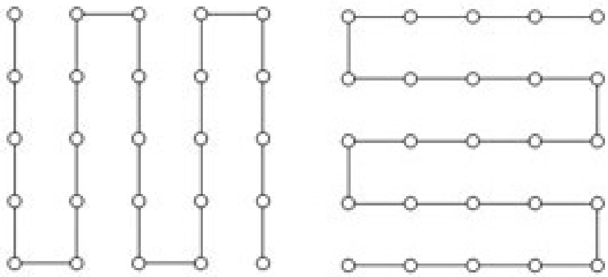

Figure 1. The 2 orthogonal raster paths used in the original sequential labelling method.

The sequential algorithm proceed as follows. The image is examined according to the paths shown in figure 1. If neighboring pixels are connected, they are then assigned the same label. When a pixel can be connected to more than one of its neighbors, the labels are considered to be equivalent (and are therefore merged).

To determine whether neighboring pixels are similar we will use Nayar and Bolle's reflectance ratio criterion (see eqn 1 below).

$$\frac{I_a - I_b}{I_a + I_b} \leq \theta \qquad (1)$$

This reflectance ratio, taken for image pixels a and b has the advantage that, for grey scale, it is independent of intensity. And, if computed on R, G, and B separately the triplet of ratios is independent of illumination [10]. And, so, supports segmentations which are independent of the lighting conditions.

2.2 Hamiltonian Paths

The problem of generating Hamiltonian paths in a general graph has been shown to be NP-complete [12]. Images, however, can be considered as a special class of graphs, namely grid graphs. In [7], a method to find Hamiltonian paths in such graphs in a linear time has been proposed. Briefly described, Hamiltonian paths are found in 4 steps, illustrated in figure 2. Downsampling (reducing the size of the graph by a factor 4 (the image is reduced by half in the x and y directions), finding a minimum spanning tree on this downsampled graph, upsampling (increasing the size of the graph by a factor 4) the tree and finally completing the graph. Randomness can be ensured by weighting all edges in the original graph with random weights prior to computing the minimum spanning tree. Refer ro [7] for a more complete description and proof that the method always generates a complete Hamiltonian path.

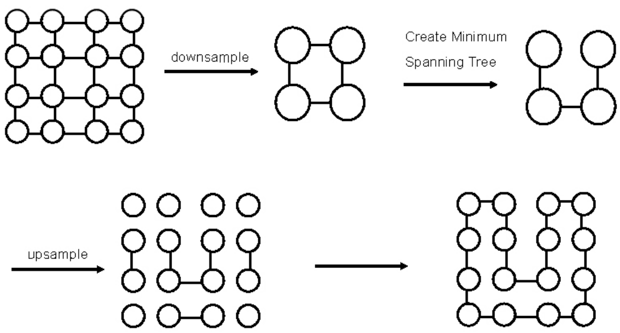

Figure 2. From the original graph to the final Hamiltonian cycle, all the steps used in creating such a path.

Since this method can generate a large number of random paths, we propose that can segment images with more accuracy than the 2-pass algorithm: we can use multiple paths to discover region connectivity. In the 2 pass approach to get large regions one needs to be "optimistic" about the underlying image structure and so use a fairly large threshold to determine pixel (and hence region) similarity. With multiple paths we can be "pessimistic" and use a smaller threshold since we are secure in the knowledge that we can joint pixels in multiple different ways. Using a large number of paths results in a area-like processing of the image, despite it not being explicitly defined in the segmentation algorithm.

Finally, we note that while the paths can be efficiently computed, they can also be pre-computed for a certain image size. Thus the algorithm cost is the number of pixels multiplied by the number of paths. Typically, the latter is small and so the algorithm is very fast.

3. Segmenting Images

Let us now consider how images are segmented. Before proceeding further we are interested in the plausibility of our approach. If we an image with 2 regions that are hard to segment can we automatically find the segmentation?

3.1. Convergence of the Algorithm

Let us create an image that consists of a double spiral. The two spirals are 1 pixel wide, while the image itself is of size 256x256, as illustrated in figure 3a. The first step in sequential labelling is to label all pixels in the image as belonging to a different region; here we have 256x256 pixels so we have 65536 different regions. We then recursively use the different pre-computed paths to process the image, using the color reflectance-ratio merging criterion [10], i.e. 2 labels a and b are equivalent if

$$\max\{\frac{R_a - R_b}{R_a + R_b}, \frac{G_a - G_b}{G_a + G_b}, \frac{B_a - B_b}{B_a + B_b}\} \leq \theta \qquad (2)$$

Where we defined θ to be 0.035

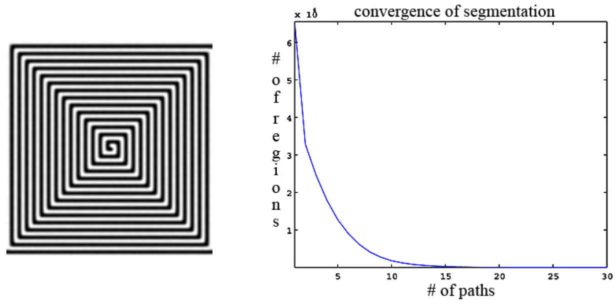

Figure 3. (a): The spiral figure used in the convergence experiment. (b): The curve showing the actual convergence.

If the structure of the different paths is random enough, and if the set of paths is complete with respect to the image size, then the segmentation should converge towards two distinct labels. Figure 3b displays the number of distinct labels (i.e. regions) after each path. After 29 paths, the algorithm has converged to 2 distinct regions, each of them containing one spiral. Due to the random nature of the paths, we repeated this experiment 50 times. The mean number of paths of convergence was 26 and the highest number was 31. From this example, it can be inferred that since "real" images generally have much larger regions, the algorithm should then converge in most cases with less paths. For equivalently sized images, we have used 15 different paths, since the improvement in quality beyond was not significant.

Moreover, it is simple matter to prove convergence in general. Consider we have an image with distinct regions where each region can be discriminated from one another using the ratio test. The segmentation fails if after n iterations we have two adjacent pixels that should belong to the same region but are labelled differently. By assumption, these adjacent pixels satisfy the ratio criterion and so if we considered a path that joined these pixels together then these pixels (and their associated regions) would be merged. Since our paths are generated randomly this must happen given enough paths.

3.2. Segmentation Experiment on a real image

We are now interested in the detail of our algorithm. How will it perform on a real image? Top left of Figure 4 shows a simple image with well defined colour regions. Let us now consider what happens when we recursively apply out path-based method using the ratio criterion. Figure 4 shows the evolution of the segmentation for an image (each white pixel is considered to belong to an edge between regions, the black areas are the regions). These different steps picture how the segmentation converges towards stability, usually after 15 steps or so. While the convergence is fast, as shown in Fig. 4 curve, it is really the steps between paths number 10 and 15 that effectively denoise the segmentation.

The method chosen to represent the segmentation is one based on regions density. The underlying assumption of this method is that a segmented image is composed of several regions

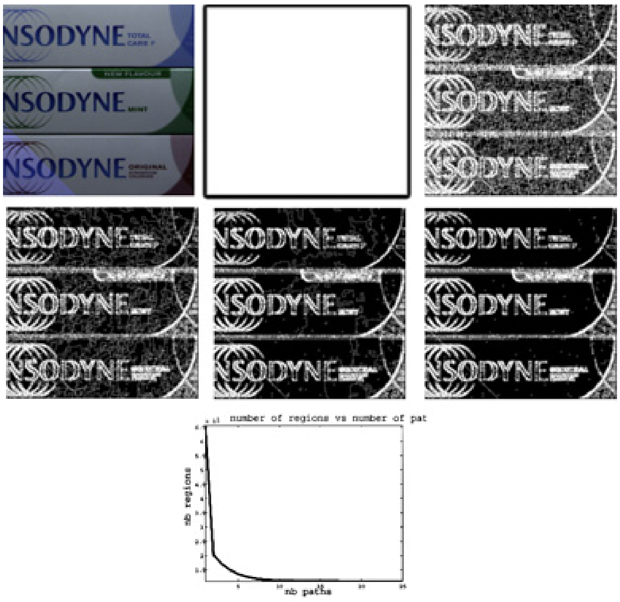

Figure 4. From left to right and top to bottom: The original image and the segmentations after 1 (all white since the first step is to label all pixels differently), 2, 5, 10 and 15 paths respectively. The curve on the 3rd row shows the speed of convergence for this particular image.

within which the pixels have the same "label". The region density is obtained by sliding a small $n \times n$ window over the image (in all our experiments, $n = 3$). The number of different regions (or labels) within this window expresses the region density for the center pixel. If we look in a small window and there is a single underlying region then we say this window has density 1. If there are two regions then we have density 2 and so a small edge, up to a region density of 9 (the maximal value) where all pixels within the window belong to a different region.

By definition, all pixels within a region have the same value (labels). A region density higher than one is therefore indicative of the presence of an edge. Additionally, since the segmentation is based on color ratios, we can encounter very high region densities in case of fast-changing reflectances, such as in grass or vegetation regions. However, most of the pixels belonging to such regions will appear solid white on the density map and edges can also therefore be extrapolated. This explains the fairly large regions among tree leaves and grass in figure 8.

We can use this approach because, in effect noise is not a significant factor in our edge maps. An illustration can be seen in figure 5, where the region density of the original image is shown on the left and the right image is the edges obtained with our method. The original image contains significant noise, but the use of several random paths in effect denoised the image while preserving edge information.

4. Results

We first compare our results with the ones obtained using the 2-pass raster scan method. From the convergence curves previously shown, we see that the main reduction in the number of regions occurs within the first step. We might therefore expect both methods to deliver similar number of regions and, using the density approach described above, the segmented images would have similar edge representations. The results, displayed in figure 6 are mitigated. While the strong edges of the image are present in both results, we also see that the edge density map for the raster segmentation is much noisier. And, this shows that it

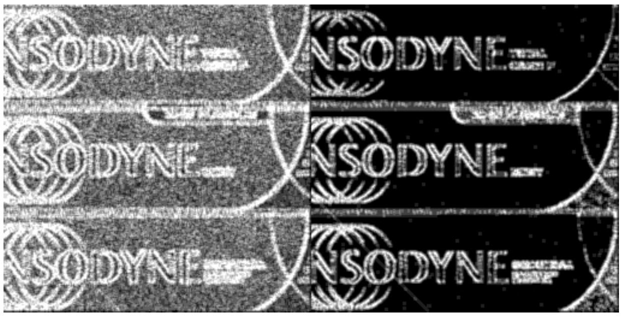

Figure 5. Left: region density of the original image. Right: region density of the segmented image.

has not merged regions as effectively as our multiple path approach.

Figure 6. Left: region density of the raster segmented image. Right: region density of the segmented image.

Figure 7 and 8 show results obtained with a variety of images. In figure 7, the first column contains the original images and the second edges obtained with the 2-pass raster approach. In figure 8, the first column are the edges obtained with our method while the second column are edges obtained with the widely used meanshift [9] algorithm (which we used with the default parameters).

From these results, two main aspects can be observed. The first one is that, as previously thought, the results from our algorithm are an improvement over the original sequential labelling formulation problem. The second one, when comparing our results to meanshift is that while our algorithm is intrinsically much simpler, the results are broadly comparable. Both the 2-pass approach and our method also contain large vegetation regions compared to the meanshift algorithm. Since those regions are rapidly changing reflectance-wise, their underlying region density will be high. Filtering the region density map with a simple point-based high pass filter allow us to extract edges for both black (low density) and white (high density) regions. The resulting edges therefore oversegment some parts of the image, while undersegmenting others. The only drawback however, is the presence of noisy regions that can be explained by the fact that we only use local color information to merge different labels/regions. We are currently developing second order metrics (rate of change within an area) in order to "clean up" the segmentations and obtain finer segmentations for highly textured regions.

5. Conclusion and Future Work

Up to this point, only first order statistics have been used in our segmentation framework. The obtained edges are, while accurate, sometimes either too thick or too noisy compared to the

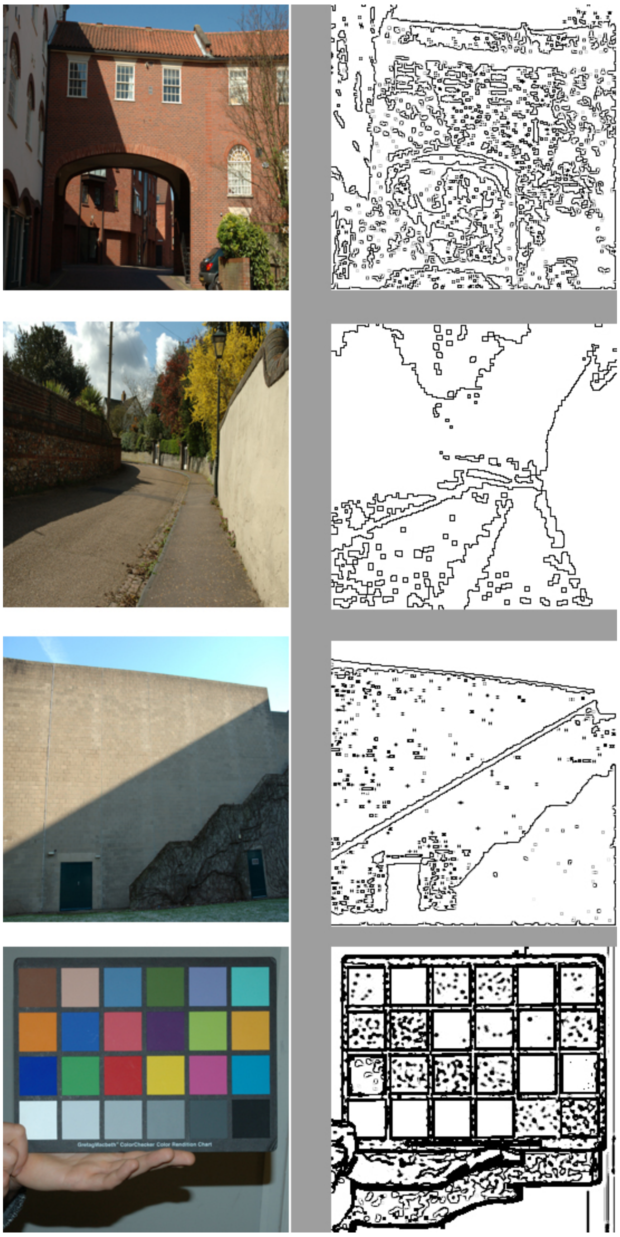

Figure 7. 1st Column: Original images, 2nd column: edges obtained with the raster method

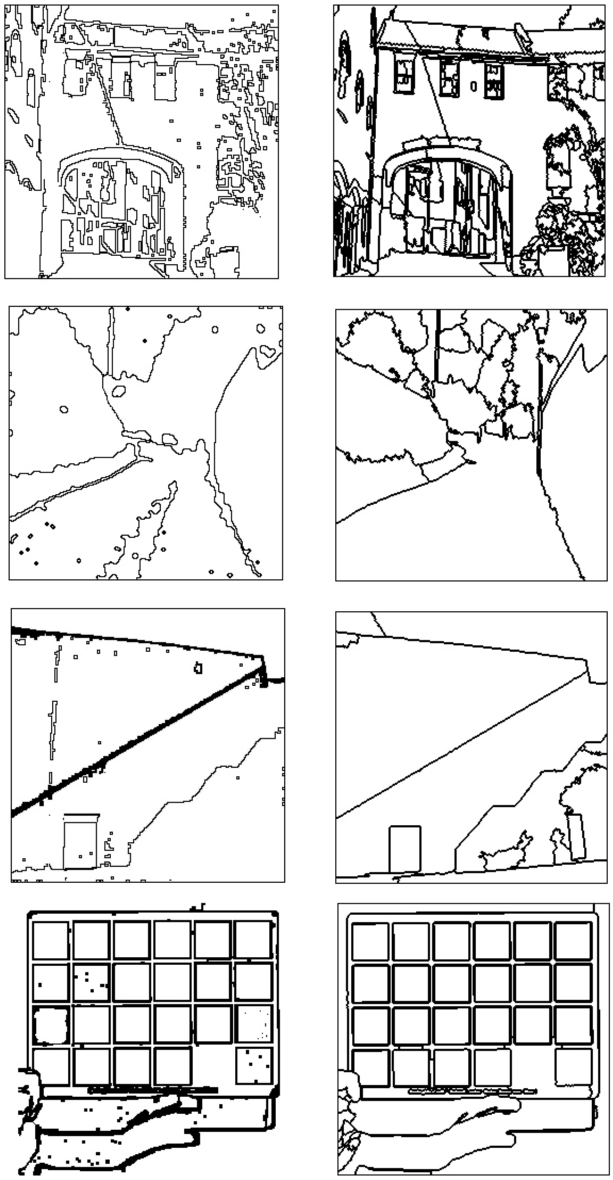

Figure 8. 1st Column: Edges obtained with our method, 2nd column: edges obtained with the meanshift algorithm

size of segmented regions. In [10], Nayar and Bolle discarded noisy or small regions in order to focus only on "valid" regions. Here, we however would like to obtain a full segmentation of the image. To improve current segmentations, one will have to look at higher order statistics, such as the rate of changes, in order to accurately detect and segment textures without adding too much complexity.

The 2 pass raster segmenter is simple and fast and is often quoted in the literature. Unfortunately, it tends to oversegment images even in the presence of small amounts of noise. In this paper we present a generalization of this approach where we discover regions by taking multiple random paths through an image. This approach fares better but still over segments an image. Yet, an analysis of region density shows that the underlying image structure can be discovered from the path based segmentation. Indeed, the discovered edges are comparable to those discovered by the widely used mean shift algorithm

Acknowledgments

Graham D. Finlayson gratefully acknowledges the support of the Leverhulme trust.

References

[1] J.A. Bangham, R. Harvey, P.D. Ling and R.V. Aldridge, Morphological scale-space preserving transforms in many dimensions, Journal of Electronic Imaging, 5, 283 (1996).
[2] E.H. Land, The retinex theory of color vision, Scientific American, pp. 108–129 (1977).
[3] M. Meila and J. Shi, Learning segmentation by random walks., in Advances in Neural Information Processing Systems (NIPS), pp. 873–879 (2000).
[4] J. Frankle and j. McCann, Method and apparatus for lightness imaging, US Patent No. 4,384,336 (1983).
[5] D.H. Ballard and C.M. Brown, Computer Vision, Prentice Hall (1982).
[6] B. Bollobas, Graph Theory, Springer Verlag (1979).
[7] C. Fredembach and G.D. Finlayson, Hamiltonian path

based shadow removal, in Proc. of the 16th British Machine Vision Conference (BMVC), pp. 970–980 (2005).
[8] W.Y. Ma and B.S. Manjunath, Edgeflow: a framework for boundary detection and image segmentation,, IEEE Trans. on Image Processing, 9, 1375 (2000).
[9] D. Comanicu and P. Meer, Mean shift: A robust approach towards feature space analysis, IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI), 24 (2002).
[10] S.K. Nayar and R.M. Bolle, Computing reflectance ratios from an image, Pattern Recognition, 26, 1529 (1993).
[11] S.B. Gray, Local properties of binary images in two dimension,, IEEE Trans. on Computers, 20, 551 (1971).
[12] Michael R. Garey and David S. Johnson, Computers and Intractability; A Guide to the Theory of NP-Completeness, W. H. Freeman & Co., New York, NY, USA (1990).

Author Biography

Clément Fredembach received the MS degree in Communication Systems from the Swiss Federal Institute of Technology in Lausanne (EPFL) in 2003. His MSc thesis, done at Gretag Imaging AG focused on image classification. He is currently a final year Ph.D. student at the University of East Anglia (UEA), Norwich, UK, working on shadow removal and path-based computation.