

Developing a Measure of Similarity between Pixel Signatures

Anthony Holmes and Chris Taylor
Imaging Science and Biomedical Engineering
Stopford Building
University of Manchester
Oxford Road
Manchester, M13 9PT, UK
ash@sv1.smb.man.ac.uk
<http://www.isbe.man.ac.uk>

Abstract

Our previous work in computer-aided mammography has used scale-orientation pixel signatures that provide a rich description of local structure. We describe work using the transportation ('earth mover') algorithm to define a measure of similarity between two signatures that recognises similar structures whilst remaining robust to background variability and the presence of other structures. Specifically, we investigate an un-normalised, normalised and novel best-partial-matching (BPM) approach to measuring this *transportation distance* and compare it with Euclidean distance. Receiver operating characteristic (ROC) methodology is used to compare performance and an example of clinical application is given. The BPM approach outperforms the other methods.

1 Introduction

It has been shown that prompting in mammography can improve a radiologist's performance, even if the prompting system makes errors [4]. Prompting involves using computer-based image analysis to locate potential abnormalities and generate markers to attract the radiologist's attention to them. One approach is to compute some form of local feature signature for each pixel and then apply a statistical classifier [5]. Recently, scale-orientation signatures have been used for the detection of blob-like structures in mammographic images [11] and can also be used as a general technique for structure detection. However, when the signatures are treated as vectors for statistical classification, the Euclidean space they define has unsatisfactory metric properties - a small change in underlying structure may produce a large change in the vector. This produces an unsatisfactory basis for statistical analysis. We describe work using the transportation ('earth mover') algorithm to measure the extent to which values in two signatures need to be redistributed to make them identical. This measurement is called the

transportation distance and is a ‘natural’ measure of similarity between the two signatures. Specifically, we investigate an un-normalised, normalised and novel best-partial-matching approach to measuring the transportation distance between two signatures and compare them with Euclidean distance. Our aim is to find a measure that recognises similar structures in signatures whilst remaining robust to background variability and the presence of other structures.

2 Background

2.1 Scale-orientation signatures

Morphological M or N filters belong to a class of filters, known as sieves, that remove peaks or troughs smaller than a specified size [1]. By applying sieves of increasing size at a number of orientations, a scale-orientation signature can be constructed for each pixel in an image. The signature is a 2-D array in which the columns are values for the same orientation, the rows are values for the same scale and the values themselves are the grey-level change at the pixel, resulting from the application of the filter at a particular scale and orientation. Sieves have been shown to have desirable properties when compared to other methods of constructing scale-orientation signatures [11, 2]. In particular, the results at different positions on the same structure are similar (local stationarity) and the interaction between adjacent structures is minimised. Figure 1 shows examples of scale-orientation signatures for the centre pixel of some simple structures.




Figure 1: Scale-orientation signatures for (a) a Gaussian blob and (b) a Gaussian line.

2.2 Transportation problems

Transportation problems are a class of linear programming problems in which one attempts to minimise the cost of delivering integral quantities of goods produced at n warehouses to m markets whilst balancing supply and demand [3]. This generates a trans-shipment problem with no intermediate nodes, with each warehouse and market connected, and a set of unit costs, c_{ij} , of shipping from warehouse i to market j . There are efficient solutions to this class of optimisation problem [8]. We have previously used transportation methods to detect breast asymmetries [7]. More recently they have been applied to comparing histograms [9]. Signature elements can be thought of as warehouses, each containing goods proportional to the pixel intensity of the element. Alternatively, the elements can be thought of as

markets, each requiring goods proportional to the pixel intensity. By considering one signature as supplying goods and another signature as demanding goods, we create a transportation problem. A meaningful solution to the problem is facilitated by choosing a suitable set of costs, c_{ij} , to capture the two-dimensional nature of the signatures. Our work uses costs based on Euclidean distance, taking into account the periodicity of the orientation axis of the signatures. Thus, localised movement in both scale and orientation is favoured above larger scale movements. In each of the techniques described in section 3, the similarity measure is defined as the mean cost of moving a single unit of intensity; the total cost of transportation divided by the number of units (goods) transported.

3 Developing a suitable measure of similarity

Signatures may be obtained from one or more structures embedded in a variable mammographic background. A suitable measure of similarity should be able to recognise whether similar structures exist in both signatures, regardless of the presence of other structures or different backgrounds in either signature. The techniques introduced in this section attempt to address this issue using a transportation framework. Figure 2 shows two signatures, taken from (a) a binary line and (b) a binary blob. Visualisations of the solutions to the transport problem are shown for the (a) un-normalised, (b) normalised and (c) partial matching techniques discussed below.

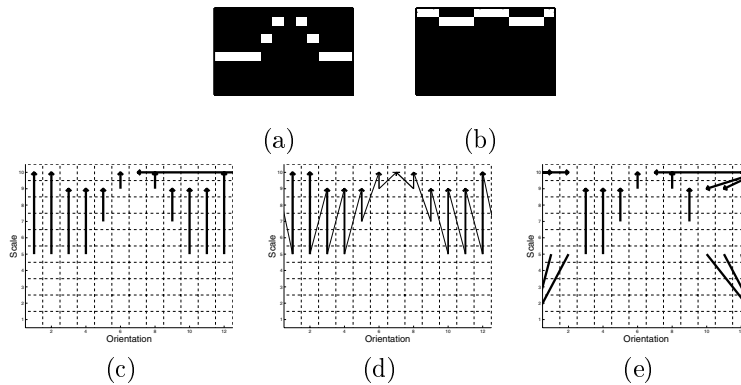


Figure 2: Transportation between signatures from (a) a binary line and (b) a binary blob using (c) un-normalised signatures, (d) normalised signatures and (e) partial matching at 50%. The arrows show the movement of ‘material’ from (a) to (b).

3.1 Un-normalised signatures

This method of obtaining a similarity measure involves no pre-processing of the signatures. That is, the ‘raw’ signatures define the transportation problem. If the

sum of each signature is different, we have an unbalanced problem. This is more commonly known as the ‘nonstandard’ transportation problem and is solved by introducing an extra warehouse (or market) to provide (or absorb) the difference. Movement to or from this extra location is free (zero cost). A visualisation of the solution is shown in 2(c). As the signatures are unbalanced, when transporting from signature (a) to signature (b), the extra goods required have been obtained from a *dummy* warehouse, indicated by an arrow entering the plot.

3.2 Normalised signatures

This method pre-processes the signatures to equalise their totals. The simplest way to do this is, for both signatures, to multiply each element in one signature by the total of the elements in the other signature. Thus a balanced problem is created and a solution exists without the need for an extra location to supply or absorb excess material. Figure 2(d) shows the solution for matching the two binary signatures. As the signatures are normalised, no external goods are required.

3.3 Partial matching

This new method is motivated by the observation that, in real applications, it is likely that similar structures will be embedded in different backgrounds, resulting in partial similarity between signatures. A specified fraction of one signature is matched to the other. Figure 3 shows schematically how partial matching works. The larger boxes are the signatures and the smaller boxes are a dummy warehouse and market. The first signature has a total supply of S_1 goods and the second signature requires S_2 goods. We choose a specified fraction fS_1 of the first signature to describe the second signature. This dictates what the values in the dummy warehouse and market should be to create a balanced problem. The inter-signature costs are defined as in section 2.2. The cost of moving between a dummy and signature location is chosen to be the same as moving to an adjacent 4-connected square within a signature. Varying this value determines the ratio of internal to external movement. Movement between the two dummy locations is prohibited. Figure 2(e) shows the transportation solution at a specified matching fraction of 50%. Six of the goods have been moved ‘internally’ and the rest of the goods have been ‘exported’ or ‘imported’.




Figure 3: Partial matching using the transportation problem.

3.4 Best-partial-matching

If we examine the transportation cost as a function of the partial match fraction, we find a minimum at the value that best describes one signature in terms of the other. For example, when a structure is common to both signatures, the best-partial-match (BPM) fraction corresponds to a solution that finds the common structure but imports and/or exports the remaining material in the signature. We currently use exhaustive search to find the BPM fraction at which the minimum cost occurs.

4 Experimental evaluation

4.1 Synthetic example

A synthetic image data set was constructed from combinations of a Gaussian line, Gaussian blob and fractal-generated background similar to those found in mammograms. Using synthetic data allows the exact definition of ground truth and so produces objective quantitative results. In addition, when comparing images containing the Gaussian line, there is no need to search over orientation as all occurrences of the line are identical. Thus, each signature comparison is faster and more signatures can be examined. Figure 4 shows the synthetic images used.




Figure 4: The synthetic image test set consists of a (a) blob, (b) line, (c) blob and line, (d) background, (e) blob and background, (f) line and background and (g) blob, line and background.

Two experiments were performed with the synthetic data. Both experiments involve comparing pairs of images from the test set. For each comparison, 200 pixels are randomly chosen from a structure of interest in one of the images. Similarity distances are then measured between the 200 chosen pixel signatures and their counterparts in the other image for the methods mentioned in sections 3.1, 3.2 and 3.4. By comparing similarity distances between image pairs that contain the same structure and image pairs that don't, ROC analysis may be employed to analyse the results.

The first experiment examines the ability of the different techniques to detect a structure in a signature that is *contained* in another signature. The Gaussian blob, figure 4(a), is compared with each of the other images that contain the blob, ie. 4(c), (e) and (g). It is then compared with each of the images that don't contain the blob, ie. 4(b), 4(d) and 4(f). This is repeated using the Gaussian line; comparing it with images that contain the line and with images that don't. The aim is to obtain lower similarity distances for image pairs containing the same structure. Figure 5(a) shows the results of this experiment.

The second experiment examines the ability of the different techniques to detect a structure that is *common* to both signatures when both signatures contain more data than just the structure in question. Firstly, the image consisting of a blob and a line, 4(c), is compared with the image containing a blob on a fractal background, 4(e). Thus, the blob is the common structure. It is then compared with the image containing a line on a background 4(f) where the line is the common structure. Then, it is compared with the background image, 4(d) where no common structure exists. This final comparison is repeated to ensure that an equal number of samples are taken from image pairs with common structure and image pairs with no common structure. Figure 5(b) shows the results of this experiment.

4.2 Clinical example

In order to investigate the performance of the methods in a clinical environment, a region of interest (ROI) was selected from a mammogram (mdb245ls) in the MIAS database [10]. The ROI contains several blob-like structures and several linear structures. Two suitable pixels are chosen, one from a blob-like calcification and the other from a linear structure. Figure 6(a) shows the ROI with an arrow indicating the chosen calcification pixel and figure 7(a) shows the ROI with an arrow indicating the chosen pixel from a linear structure. For each of the two chosen pixels, the ‘similarity distance’ between the chosen pixel’s signature and every pixel signature in the ROI is measured for each of the methods described in section 3. In contrast to the experiments using synthetic data, real structures may occur at any orientation and thus it is necessary to search over all orientations when comparing signatures. Plotting the similarity distance at each pixel forms a similarity image for each method. These images are inverted such that bright regions correspond to similar pixels and dark regions correspond to pixels that are less similar. The image intensity is scaled to the largest and smallest similarity distance present in the image.

Figure 6 shows the results for the pixel from a calcification and figure 7 shows the results for the pixel from a linear structure.

5 Results and discussion

5.1 Synthetic example

Figure 5 shows the results of the experiments described in section 4.1 using synthetic data. The ROC curves show the ability of the different measures to discriminate between signatures where a structure was or was not present [6]. The closer a curve is to the top-left corner, the better its discrimination ability.

The curves in figure 5 indicate that at high specificities, or low false positive fractions (FPF), the normalised and BPM methods have higher sensitivities (true positive fractions (TPF)) than either the un-normalised or Euclidean methods. This means that they will delineate chosen structures more successfully than the Euclidean method. It is clear that the normalised and BPM methods generally outperform the un-normalised and Euclidean methods in both experiments. More surprisingly, the un-normalised method produces a similarity distance that

BMVC99

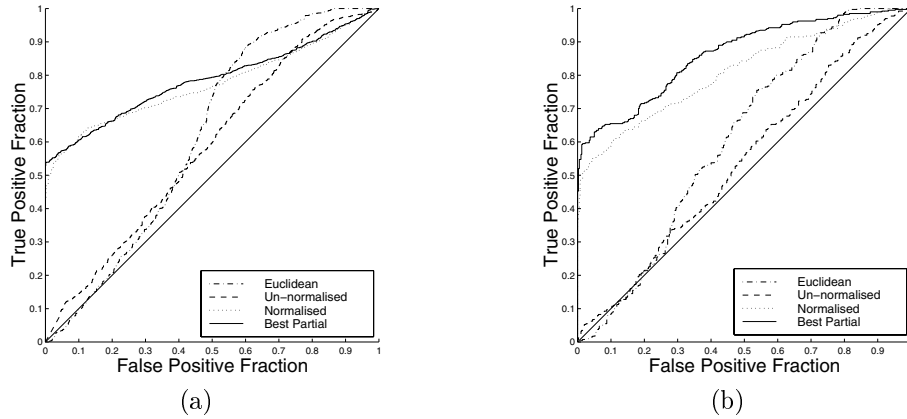


Figure 5: ROC curves for comparisons between pixel signatures (a) where one signature is *contained* in the other and (b) where both signatures contain a *common* structure.

performs *less* efficiently than the original Euclidean distance. This may be due to significant contrast differences between signatures from the synthetic images. In 5(a), there is little difference between the normalised and BPM methods but in 5(b), the BPM method is clearly superior to all the other methods. This confirms the expected superiority of BPM to detect common structure. The second experiment is a fairer representation of the clinical applications in which the methods would be used, so these results encourage us to examine the performance of the methods in the clinical environment.

5.2 Clinical example

Figure 6 and figure 7 show the results of the clinical experiment described in section 4.2. The two images generated using the un-normalised transportation method demonstrate markedly inferior performance when compared to the Euclidean, normalised and BPM methods and so they are not included in this paper. This result is not surprising, given the performance in figure 5.

5.2.1 Blob-like structure

Figures 6(b),(c) and (d) show the ‘similarity images’ generated from comparison with the pixel from the blob-like calcification indicated in figure 6(a).

The first observation from figure 6 is that the similarity images generated by the Euclidean and BPM methods are remarkably similar. They have both successfully delineated blob-like structures and rejected other material. A closer examination suggests that blob-like structures have been delineated marginally more successfully by the BPM method than the Euclidean method but that the BPM method has generated a similarity image with a marginally higher background intensity. This is in agreement with the results from the synthetic example in figure 5. It is




Figure 6: Similarity images for a pixel signature taken from (a) a calcification using (b) Euclidean, (c) normalised transportation and (d) BPM transportation distances.

probable that, as explained in section 5.1, the Euclidean method is primarily sensitive to intensity contrast in the ROI and as the blob-like structures are generally quite lucid and don't sit on very different backgrounds, the Euclidean method has found them quite successfully. The image generated by the normalised method is quite noisy but it has successfully delineated blob-like structures. Figure 6 also shows that the images generated by the normalised and BPM methods have higher background intensities than the Euclidean method. Again, this is in agreement with the results in figure 5.

5.2.2 Linear structure

Figures 7(b),(c) and (d) show the 'similarity images' generated from comparison with a pixel taken from the linear structure indicated in figure 7(a).

The difference between the images generated by the Euclidean and BPM methods in figure 7 is considerably greater than in figure 6. In fact, the most striking feature of the results is that the Euclidean method has emphatically rejected the blob-like structures whilst remaining reasonably indifferent to the presence of the linear structures. This supports the hypothesis that the Euclidean method is primarily sensitive to intensity contrast. Again, the normalised method contains a

BMVC99




Figure 7: Similarity images for a pixel signature taken from (a) a linear structure using (b) Euclidean, (c) normalised transportation and (d) BPM transportation distances.

significantly higher level of noise than the other methods (with the exception of the un-normalised method that is not reproduced here) but the linear structures have been detected with considerable clarity. Most importantly, the BPM method has successfully detected linear structures and suppressed other structures and background. These results indicate that our approach measures signature similarity in a practically useful way.

6 Future Work

Further investigation of the performance of our approach with clinical data is necessary before the next stage of research may be undertaken. This will use more clinical data where ‘structure’ truth is known, as it is with synthetic data. This permits pixel classification ROC curves to be generated for clinical data which give a quantitative measure of performance. Once completed, future work will measure the BPM distance between all signatures in a set, permitting the reconstruction of the signature positions in a new space with Euclidean separations equal to BPM distances. A neural network will make the non-linear transformation between the new and original spaces, providing an efficient run-time method of transforming

signatures into a space suitable for statistical analysis.

7 Acknowledgements

We would like to thank Reyer Zwiggelaar of the Division of Computer Science at the University of Portsmouth for helpful comments and the fractal background generation code.

References

- [1] J. A. Bangham, T. G. Campbell, and R. V. Alridge. Multiscale median and morphological filters for 2D pattern recognition. *Signal Processing*, 38:387–415, 1994.
- [2] R. Harvey, A. Bosson, and J. A. Bangham. The robustness of some scale-spaces. In *British Machine Vision Conference*, pages 11–20, 1997.
- [3] F. L. Hitchcock. The distribution of a product from several sources to numerous localities. *J. Math. Phys.*, 20:224–230, 1941.
- [4] I. W. Hutt, S. M. Astley, and C. R. M. Boggis. Prompting as an aid to diagnosis in mammography. *Digital Mammography*, pages 389–403, 1994.
- [5] D. J. Marchette, R. A. Lorey, and C. E. Priebe. An analysis of local feature extraction in digital mammography. *Pattern Recognition*, 30(9):1547–1554, 1997.
- [6] C. E. Metz. ROC methodology in radiologic imaging. *Investigative Radiology*, 21:720–733, 1986.
- [7] P. Miller and S. Astley. Automated detection of breast assymetries. In J. Illingworth, editor, *British Machine Vision Conference*, pages 519–528, 1993.
- [8] A. Ravindran, D. T. Phillips, and J. J. Solberg. *Operations research: principles and practice*. John Wiley & Sons, 2 edition, 1987.
- [9] Y. Rubner, C. Tomasi, and L. J. Guibas. A metric for distributions with applications to image databases. In *6th International Conference on Computer Vision*, pages 59–66, 1998.
- [10] J. Suckling, J. Parker, D. Dance, S. Astley, I. Hutt, C. Boggis, I. Ricketts, E. Stamatakis, N. Cerneaz, S. Kok, P. Taylor, D. Betal, and J. Savage. The mammographic images analysis society digital mammogram database. In *Excerpta Medica. International Congress Series*, volume 1069, pages 375–378, 1994. `mias@sv1.smb.man.ac.uk`.
- [11] R. Zwiggelaar, T. C. Parr, J. E. Schuum, I. W. Hutt, S. M. Astley, C. J. Taylor, and C. R. M. Boggis. Model-based detection of spiculated lesions in mammograms. *Medical Image Analysis*, 3(1):39–62, 1999.