

Transforming Pixel Signatures into an Improved Metric Space

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 to provide a rich description of local structure. However, when treated as vectors for statistical classification, the Euclidean space they define has unsatisfactory metric properties. In this paper we describe a scheme that makes use of a previously described signature similarity measure to define a non-linear transformation of pixel signatures into a space with improved metric properties. K-means clustering is used to select prototype signatures and multidimensional scaling uses inter-prototype distances to form a new improved metric space. Finally, support vector regression is used to learn the transformation between the original and new spaces.

1 Introduction

It has been shown that prompting in mammography can improve a radiologist's performance, even if the prompting system makes errors [5]. Prompting involves using computer-based image analysis to locate potential abnormalities so that the radiologist's attention can be drawn to them. Recently, non-linear scale-orientation signatures have been used for the detection of malignant tumours in mammographic images [9]. The approach can in principle be extended to detect other normal and abnormal structures. However, when the signatures are treated as vectors for statistical classification, the Euclidean vector space they define has unsatisfactory metric properties - a small change in underlying structure may produce a large change in the vector. This provides an unsatisfactory basis for statistical analysis. We have previously introduced a novel measure of signature similarity known as the BPM distance that is robust to small changes in scale and orientation of signature structure as well as intrinsic variability and the presence of other potentially confounding structures [4]. The distance measure is based upon the transportation ('earth-mover') algorithm [3]. In this paper, we present a scheme

that uses the BPM distance to define a non-linear transformation of pixel signatures into a space with improved metric properties. Firstly, k-means clustering is used to select a set of signature prototypes. A distance matrix is then constructed by measuring the BPM distance from every signature prototype to every other signature prototype. Multidimensional scaling is then used to construct a new space in which measurement of the inter-point *Euclidean* distances produces an approximation to the distance matrix. Finally, support vector regression is used to learn the mapping between the original and new spaces. The regression parameters may be applied to any signature to transform it into the new space.

2 Scale-orientation pixel 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 columns are values for the same orientation, 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 [9, 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.

3 Measuring signature similarity using BPM distance

In real applications, signatures may be obtained from one or more structures embedded in a variable mammographic background, resulting in partial similarity between signatures. This suggests that a fraction of a given signature should match to another obtained from the same structure. Therefore, a useful measure of similarity should be able to recognise when similar structures are represented

in two signatures despite intrinsic variability and regardless of the presence of other structures. We have addressed this problem using a *transportation problem* framework [4].

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 a number of warehouses to a number of markets whilst balancing supply and demand [3]. 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. Choosing a suitable set of costs to capture the two-dimensional nature of the signatures facilitates a meaningful solution to the problem. 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. We have previously described a method based on the transportation problem to handle partial similarity between signatures [4] but reproduce it here briefly for completeness.

We can implement partial signature similarity by supplying a dummy warehouse and market as in figure 2.




Figure 2: Partial matching using the transportation problem.

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 of the first signature, fS_1 , to describe the second signature. This dictates what the values in the dummy warehouse and market should be, as shown in figure 2. We can solve the problem of finding an appropriate value for the partial match fraction f by considering the transportation solution unit cost as a function of f . We observe a minimum in the transportation cost at the value of f that results in the best description of one signature in terms of the other. This is termed the best-partial-match (BPM) fraction. For example, when a structure is common to both signatures, the BPM fraction corresponds to a solution that finds the common structure but imports and/or exports the remaining material in the two signatures. We use a polynomial fit iterative search to find the BPM fraction at which the minimum unit cost occurs. This minimum unit transportation cost is termed the BPM distance.

4 Signature transformation scheme

Our transformation scheme consists of three separate components; signature prototyping, multidimensional scaling and support vector regression.

4.1 Selecting signature prototypes

Standard k-means clustering is a method of selecting a set of n prototypes from an underlying probability distribution of N samples. In our case, we choose to find a set of 256 prototype pixel signatures that are representative of pixel signatures formed from a complete range of mammographic structure. 256 initial prototypes are randomly selected and then a large number of signatures (2^{16}) are examined in turn. Each examined signature is assigned to its nearest prototype and the mean of that prototype is adjusted accordingly. This continues for all N signatures, resulting in a set of prototypes that are representative of the examined signatures. Figure 3 shows a mammogram containing a wide variety of structures. For each of the prototype signatures obtained by k-means clustering, the image pixel with the most similar signature is highlighted, given a visual representation of the prototype distribution. As expected, relatively few prototypes are required to represent background tissue. In contrast, the prototype density is higher in regions containing more structure. Now that we have generated a representative set of signature prototypes, we may use multidimensional scaling to construct an improved metric space.




Figure 3: A mammogram containing a variety of structure. The markers indicate the closest image signatures to the set of signature prototypes.

4.2 Constructing a new space

There are many types of multidimensional scaling but they all attempt to solve the following problem; given a set of distances, d_{rs} , find a configuration of points that will reproduce an approximation to these distances. The distances may or may not be *Euclidean* and they need *not* satisfy the triangle inequality

$$d_{rt} \leq d_{rs} + d_{st} \quad (1)$$

satisfied if they were produced by a metric. Our work uses an analytical technique which is briefly presented below. The reader is referred to Ripley [6] for a more detailed explanation.

A configuration X is an $n \times p$ matrix of data, considered as n points in p dimensions. We note that the distance matrix may be constructed from X as follows

$$\text{if } M = XX^T, \quad (2)$$

$$d_{rs}^2 = M_{rr} + M_{ss} - 2M_{rs}. \quad (3)$$

This observation leads to the following implementation, known as classical or metric multidimensional scaling. If L is the given distance matrix, we define

$$T = L^2. \quad (4)$$

Next, we remove row and column means and add back the overall mean. This follows from equation 3.

$$T' = -\frac{1}{2} \left[T - \frac{(T\mathbf{1})\mathbf{1}^T}{n} - \frac{\mathbf{1}(T\mathbf{1})^T}{n} + \frac{\mathbf{1}^T T \mathbf{1}}{n^2} \right]. \quad (5)$$

The rank of T' gives the number of dimensions that the constructed configuration will have. Next, we find the eigendecomposition of T' using schur decomposition

$$T' = CD^2C^T \quad (6)$$

and find D by taking the matrix square root. To finish, we must find the order of the eigenvalues i.e. the diagonal values of D with the largest first and take the corresponding columns of CD to be our configuration X . Given a Euclidean distance matrix, this technique produces a perfect configuration i.e. the given distance matrix is perfectly reproduced. If the distance matrix is non-Euclidean, the technique produces an approximation to the distance matrix.

Figure 4(a) is a randomly generated set of 2D points. The distance matrix in figure 4(b) is generated from the data points by measuring the inter-point Euclidean distances. The distance matrix is symmetric and has a zero diagonal. Multidimensional scaling then generates a new configuration from the distance matrix, as shown in figure 4(c). If the distance matrix is Euclidean, that is it contains Euclidean distances, then the new configuration will exactly reproduce the generating distance matrix when the inter-point distances are measured. However, if the distance matrix is not Euclidean, the new configuration can only reproduce

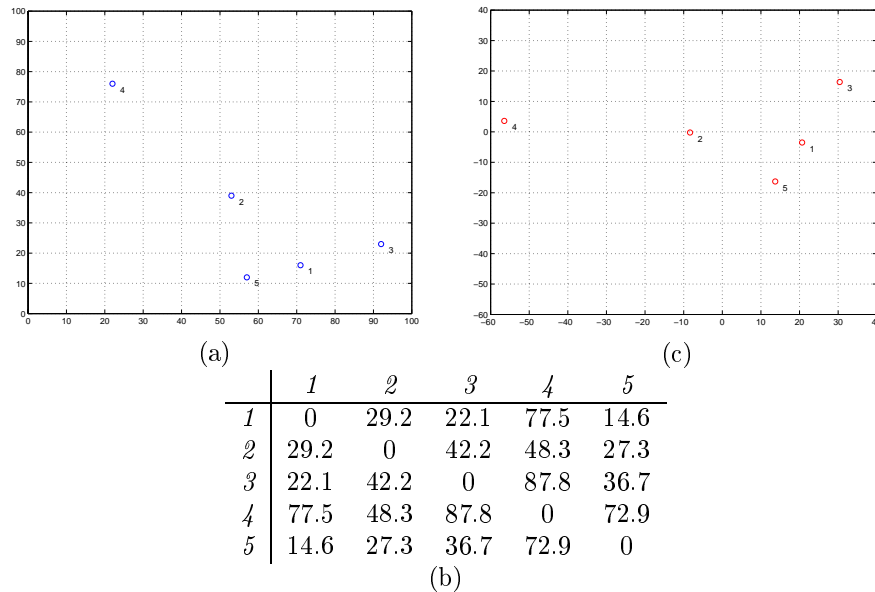


Figure 4: (a) a random configuration of points, (b) the distance matrix generated by the configuration and (c) a configuration reconstructed from the distance matrix.

an approximation to the generating distance matrix. It is interesting to note that the distance matrix is invariant to translation, rotation and reflection of the configuration of data points. In our case, a distance matrix is constructed by measuring the inter-prototype BPM distances. Multidimensional scaling then transforms each of these prototypes into a new space, from the distance matrix.

4.3 Learning the transformation

The use of support vector machines (SVM) in empirical data modelling is becoming widespread as they generalise well and address the curse of dimensionality. They provide promising empirical performance and have many attractive features. They can be used in both classification (SVC) and regression (SVR) problem domains. The reader is referred to Vapnik [8] for a more exhaustive treatment of SVMs and some recent applications. When the input space is not linearly separable, a mapping into a high dimensional feature space, where the data is linear, is required. However, a kernel function can be used to avoid explicit evaluation of this computationally intensive mapping. The kernel function used in our work is the popular Gaussian radial basis function. The regression input space is defined as the original prototype signatures, formed using k-means clustering and the output space is defined as the reconstructed prototypes in the new space. After performing SVR, the regression parameters can be applied to any signature to transform it into the new space.

5 Experimental evaluation

In the following experiment, pixel signatures were constructed by applying a morphological M filter, centred on every image pixel, at different scales and orientations. Ten scales were used that increased logarithmically and were chosen to encompass the range of structure sizes that occur in mammograms. As orientation is periodic, for each scale the filter was applied at twelve orientations, separated by fifteen degrees. A region of interest (ROI) was selected from a mammogram (mdb245ls) in the MIAS database [7] and a pixel was chosen from a salient linear structure. Figure 5(a) shows the ROI and the selected pixel. The ROI exhibits a variable background and contains a variety of structures that are commonly found in mammograms. Euclidean and BPM distances were then measured between the chosen pixel's signature and every pixel signature in the ROI using raw or transformed signatures. Plotting the similarity distance at each pixel forms a similarity image for each method. These images were inverted so that bright regions correspond to similar pixels and dark regions corresponded to pixels that were less similar. Three similarity images were generated. Firstly, Euclidean distance was measured between raw signatures. This allows us to see whether our scheme offers any improvement in metric properties compared to the original Euclidean vector space. Secondly, BPM distance was measured between raw signatures and finally, Euclidean distance was measured between transformed signatures. A regression (or Shepherd) plot was also generated by plotting the value of each pixel in the BPM distance image for raw signatures against the same pixel value in the Euclidean distance image for transformed signatures. This permits assessment of the efficiency of the transformation scheme.

6 Results and Discussion

Figures 5(b), (c) and (d) show the similarity images generated by the experiment described in section 5.

Figure 5(b) shows the similarity image generated by measuring Euclidean distance between raw signatures. This image clearly demonstrates the poor metric properties of the vector space defined by the raw signatures. The method has not been able to segment other linear structures in the image, although it has emphatically rejected blob-like structures. This image confirms the need for the transformation scheme described in this paper. Figure 5(c) shows the similarity image generated by explicitly measuring the BPM distance between raw signatures. This image took approximately a day to generate. We have previously discussed the suitability of the BPM measure, as indicated by this result [4] but for completeness we note that the method has successfully detected linear structures and suppressed other structures and background, thus measuring similarity in a practically useful way. This result encourages us to examine the performance of our transformation scheme as shown in figure 5(d). This figure shows the similarity image generated by transforming all image signatures using the scheme described in section 4 and measuring the Euclidean distance between the line pixel signature and every other image pixel. Having already constructed the new space and learnt the mapping from the original to the new space 'off-line', this image took a quarter of a second




Figure 5: Pixel signature similarity images are generated from (a) a pixel from a linear structure in the original image using (b) Euclidean distance between raw signatures, (c) BPM distance between raw signatures and (d) Euclidean distance between transformed signatures.

to generate, showing one of the advantages of our scheme. It is readily apparent that as well as offering significant improvement in computational efficiency, the scheme also offers a high degree of accuracy in approximating the BPM distance measurement.

We can further investigate the performance of the transformation scheme by examining the regression plot shown in figure 6. This shows the raw signatures BPM distance image pixel values (figure 5(c)) plotted against the transformed signatures Euclidean distance image pixel values (figure 5(d)). A perfect transformation scheme would produce an identical image and all of the points in figure 6 would lie on the dotted $y = x$ line. It is clear that the majority of the points lie on, or close to, the dotted line showing the efficiency of the scheme. A closer examination reveals that larger distances (from the chosen linear structure pixel)




Figure 6: Regression plot for similarity image.

are more efficiently modelled by the scheme than smaller distances. As the chosen pixel lies on a salient structure in the image, large distances from this pixel correspond to background mammographic structure. This agrees with our observation in section 4.1 that few prototypes are required to adequately model background structure. Smaller distances correspond to pixels that are more similar to the chosen linear structure pixel. The regression plot indicates that these are modelled less efficiently. A larger number of prototypes would improve the performance of the scheme but would increase computational requirements. At present, it takes approximately 24 hours to construct an improved metric space from 256 prototypes.

7 Conclusions

The results described in section 6 indicate that our scheme provides a comprehensive and efficient method of transforming pixel signatures into a space with improved metric properties. Future work will investigate ways to improve the scheme and provide quantitative results. We will then examine its use in classification performance when detecting focal abnormalities and normal structure in mammograms.

8 Acknowledgements

The principal author is funded by the EPSRC on a quota studentship.

References

- [1] J. A. Bangham, P. D. Ling, and R. Young. Multiscale recursive medians, scale-space and transforms with applications to image processing. *IEEE Transactions on Image Processing*, 5(6):1043–1048, 1996.
- [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] A. S. Holmes and C. J. Taylor. Developing a measure of similarity between pixel signatures. In *British Machine Vision Conference*, pages 614–622, 1999.
- [5] I. W. Hutt. *The Computer-aided Detection of Abnormalities in Digital Mammograms*. PhD thesis, University of Manchester, 1996.
- [6] B. D. Ripley. *Pattern recognition and neural networks*. Cambridge University Press, 1996.
- [7] 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.
- [8] V. Vapnik. *The nature of statistical learning theory*. Springer-Verlag, New York, 1995.
- [9] R. Zwigelaar, 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.