Classifying Textile Designs using Region Graphs

Wei Jia
weijia@computing.dundee.ac.uk
Stephen J. McKenna
stephen@computing.dundee.ac.uk
Annette A. Ward
award@computing.dundee.ac.uk
Keith Edwards
kedwards@computing.dundee.ac.uk

School of Computing
University of Dundee
Scotland, UK

Markov random field pixel labelling is often used to obtain image segmentations in which each segment or region is labelled according to its attributes such as colour or texture [4]. This paper explores the use of such a representation for image classification. In particular, the problem of classifying textile images according to design type is addressed.

Figure 1(a) shows an example of an image segmented into groups of regions by assigning each pixel a label; the label image is shown in the centre. Given such a labelling, the image can be represented as a *bag of shapes* by computing shape descriptors for each connected component [3]. However, a bag of shapes model ignores relationships between the groups of regions. In order to retain information about these relationships, we construct undirected weighted graphs as shown in Figure 1(b). Each vertex is associated with a group of regions (bag of shapes). Edges in the graph denote either the extent to which the groups' regions are spatially adjacent or the dissimilarity of their respective bags of shapes.

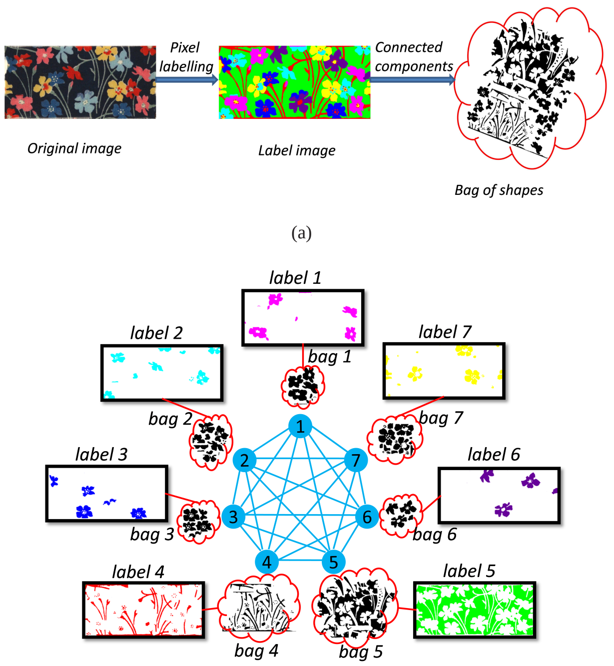

(a) (b)

Figure 1: (a) An image is segmented into labelled regions. It can then be represented using a *bag of shapes*. (b) Alternatively, a weighted graph can be constructed in which each vertex is associated with a group of regions that share the same label. Each group of regions is represented as a bag of shapes. Edge weights encode relationships between the groups.

In graph theory, the chromatic number of a graph is the smallest number of colours needed to colour the vertices without adjacent vertices sharing the same colour [7]. Consider toy examples with three groups of regions as vertices of a graph, fully connected to each other. The edge weights are assigned values proportional to the arc lengths of the common boundaries shared by the groups of regions. Figure 2(a) shows an example in which the groups of regions are equally adjacent to each other so that all edges are assigned the same weight. Deleting edges in order of weight generates the chromatic number sequence $3 \rightarrow 1$. This is an extreme case. Figure 2(b) shows an example in which edges are assigned different weights. Deleting edges in order of weight generates the chromatic number sequence $3 \rightarrow 2 \rightarrow 2 \rightarrow 1$.

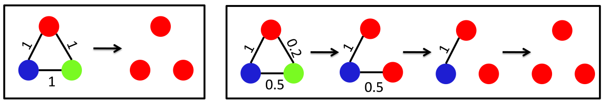

(a) (b)

Figure 2: Graph minimal colouring sequences obtained by deleting edges in the order of weight.

These examples illustrate that deleting edges by weight results in sequences of chromatic numbers that depend on the adjacency relationships of the region groups. Similarly, sequences of chromatic numbers can be computed from the complement graphs.

The sequence of (normalised) weights of those edges whose removal changes the chromatic number constitutes a feature vector. Feature vectors can also be computed based on sequences of graph domination numbers, another measure from graph theory [2].

The shape of each region (connected component) was described using generic Fourier descriptors (GFD) [8]. Given a collection of shapes from training images, a codebook was calculated by running k-means on the shape descriptors. Codewords were defined as the centres of the clusters [5]. A given shape can be assigned to the nearest codeword. A set of shapes can be represented as a histogram of the codewords. Feature vectors combined the shape descriptors and the graph-theoretic features.

In order to test the effectiveness of the algorithms, single text keywords assigned to the images were used to compare human labelling with machine classification. The method was applied to image classification using examples from a commercial textile archive owned by Liberty Fabric Ltd. Seven-fold cross validation was used to evaluate the accuracy of classification. We compared different feature sets based on a linear SVM classifier [1]. Since representing images as bags of local patch descriptors [5] such as SIFT [6] is popular, we also ran the experiment using SIFT features instead of GFD. In general, SIFT appears to be slightly more accurate than GFD. However, GFD has lower dimensionality. No matter what kinds of graph features were used, the results suggest that classification accuracy was better than using GFD or SIFT features alone.

[1] C. C. Chang and C-J. Lin. LIBSVM: A library for support vector machines.

[2] T. W. Haynes, S. T. Hedetniemi, and P. J. Slater. *Domination in Graphs-Advanced Topics*. New York: Dekker, 1998.

[3] W. Jia and S. J. McKenna. Classifying textile designs using bags of shapes. In *International Conference on Pattern Recognition (ICPR)*, 2010.

[4] W. Jia, S. J. McKenna, and A. A. Ward. Extracting printed designs and woven patterns from textile images. In *Int. Conf. Computer Vision Theory and Application (VISAPP)*, 2009.

[5] F. F. Li and P. Perona. A Bayesian hierarchical model for learning natural scene categories. In *CVPR*, 2005.

[6] D. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60 (2):91–110, 2004.

[7] S. Skiena. *Computational Discrete Mathematics: Combinatorics and Graph Theory with Mathematica*. Reading, MA: Addison-Wesley, 1990.

[8] D. S. Zhang and G. J. Lu. Shape-based image retrieval using generic Fourier descriptor. *Sig. Proc.: Image Communication*, 17:825–848, 2002.