PROCEEDINGS OF SPIE

Face recognition by detection of matching cliques of points

Fred Stentiford

SPIE.

Face Recognition by Detection of Matching Cliques of Points

Fred Stentiford

University College London, Electronic & Electrical Engineering Dept, Gower St, London, UK

ABSTRACT

This paper addresses the problem of face recognition using a graphical representation to identify structure that is common to pairs of images. Matching graphs are constructed where nodes correspond to image locations and edges are dependent on the relative orientation of the nodes. Similarity is determined from the size of maximal matching cliques in pattern pairs. The method uses a single reference face image to obtain recognition without a training stage. The Yale Face Database A is used to compare performance with earlier work on faces containing variations in expression, illumination, occlusion and pose and for the first time obtains a 100% correct recognition result.

Keywords: Face recognition, pattern recognition, similarity, human vision, graph matching

1. INTRODUCTION

The use of intuitively plausible features to recognise faces is a powerful approach that yields good results on certain datasets. Where it is possible to obtain a truly representative set of data for training and adjusting recognition parameters, optimal performance can be attained. However, when facial images are distorted by illumination, pose, occlusion, expression and other factors, some features become inappropriate and contribute noise to the discrimination on unseen data. Indeed it can never be known in advance what distortions will be present in unseen and unrestricted data and so features that are applied universally are likely to reduce performance at some point.

Many approaches to face recognition are reported in the literature [1,2]. Graph matching approaches provide attractive alternatives to the feature space solutions in computer vision. Identifying correspondences between patterns can potentially cope with non-rigid distortions such as expression changes, pose angle and occlusions. However, graph matching is an NP-complete problem and much of current research is aimed at solving the associated computational difficulties.

SIFT feature descriptors are used by Leordeanu et al [3] to construct spectral representations of adjacency matrices whose nodes are feature pair correspondences and entries are dependent on feature separations. Objects in low resolution images are recognised by matching correspondences against a set of pre-trained models. Felzenszwalb et al [4] also match a graphical model of specific objects to images in which parts are matched according to an energy function dependent on colour difference and relative orientation, size and separation. Fergus et al [5] avoid the computational complexity of a fully connected shape model by adopting a "star" model that uses "landmark" parts. The model is trained using specific feature types and recognition is obtained by matching appearance densities of model parts. Kim et al [6] reduces the computational demands by first segmenting one of the images. Each region is mapped using SIFT descriptors and a function dependent on distortion, ordering, appearance and displacement is minimised to obtain appropriate candidate points and region correspondence.

A more general approach by Duchenne et al [7] uses graph matching to encode the spatial information of sparse codes for pairs of images. An energy function is maximised using a graph cuts strategy that is dependent on node feature correlation, reduced node displacement and discouraging node crossing. Duchenne et al [8] also uses a tensor based algorithm to match hypergraphs in which correspondences are identified between groups of nodes and hyperedges linking them. The method is illustrated by matching two similar faces using triples of SIFT descriptors. Celiktutan et al [9] also match hypergraphs connecting node triples in the spatial-temporal domain by minimizing an energy function. Computation is reduced by considering a single salient point in each video frame and limiting connections along the time dimension.

Kolmogorov et al [10] present a graph-cut algorithm for determining disparities that ensures that single pixels in one image are assigned single pixels in the second image and occlusions are handled correctly. An energy function is

employed that is minimised by reducing the intensity difference between pixels, by penalizing pixel occlusions, and requiring neighbouring pixels to have similar disparities.

Berg et al [11] sets up correspondences by identifying edge feature locations and measuring their similarity by using the correlation between feature descriptions and the distortion arising from local changes in length and relative orientation. An approximate version of Integer Quadratic Programming is used to detect faces. Cho et al [12] proposes a method for growing matching graphs where nodes represent features and edges the geometric relationships. The Symmetric Transfer Error is used to measure the similarity of node pairs and the reweighted random walk algorithm to match nodes.

Shape driven graphical approaches [13-15] including active appearance models assign fiducial points to nodes and maximise a similarity function to obtain recognition of candidate images.

This paper makes use of a fully connected graph matching representation in order to measure the similarity of pairs of patterns. Nodes are raw pixels and edges take the value of the relative orientation of the two pixels. Patterns represented by graphs being compared match if all node features and edge values match. This simple framework has the advantage that the graph matching process is much faster because it can be reasonably assumed that if local node edges match, the relative orientation of more distant nodes will not vary significantly within that locality and will therefore also match.

2.   PROPOSED APPROACH

The approach taken in this paper detects structure that is common between pairs of images and uses the extent of such structure to measure similarity. In this case the size of the largest structure found to match both patterns is the number of nodes in the corresponding fully connected maximal graph or clique.

A pictorial structure is represented as a collection of parts and by a graph $G = (V, E)$ where the vertices $V = \{v_1, ..., v_n\}$ correspond to the parts and there is an edge $(v_i, v_j) \in E$ for each pair of connected parts v_i and v_j. An image part v_i is specified by a location x_i. In this paper parts v_i correspond to individual pixels. Given a set of vertices $V^1 = \{v_1^1, ..., v_n^1\}$ in image 1 that correspond to a set of vertices $V^2 = \{v_1^2, ..., v_n^2\}$ in image 2 the following conditions are met by all parts to form a clique

$$d_g(\boldsymbol{x}_i^1) = d_g(\boldsymbol{x}_i^2) \qquad (1)$$

$$\left| d_b(\boldsymbol{x}_i^1) - d_b(\boldsymbol{x}_i^2) \right| \leq \varepsilon_1 \qquad (2)$$

$$\left| d_a(x_i^1, x_j^1) - d_a(x_i^2, x_j^2) \right| \leq \varepsilon_2 \quad \forall i,j \quad i \neq j \qquad (3)$$

where $d_g(x_i)$ is the grey level gradient direction at x_i, $d_b(x_i)$ is the intensity at x_i, and $d_a(x_i, x_j)$ is the angle subtended by the point pair (x_i, x_j). Clique generation begins with the selection of a random pair of pixels $(\boldsymbol{x}_i^1, \boldsymbol{x}_j^1)$ from reference image 1 and a pair $(\boldsymbol{x}_i^2, \boldsymbol{x}_j^2)$ from candidate image 2 that satisfy (1,2,3). A new pair of points $(\boldsymbol{x}_k^1, \boldsymbol{x}_k^2)$ is added where

$$d_g(\boldsymbol{x}_k^1) = d_g(\boldsymbol{x}_k^2) \qquad (4)$$

$$\left| d_b(\boldsymbol{x}_k^1) - d_b(\boldsymbol{x}_k^2) \right| \leq \varepsilon_1 \qquad (5)$$

$$\left| d_a(x_k^1, x_m^1) - d_a(x_k^2, x_m^2) \right| \leq \varepsilon_2 \qquad (6)$$

where $\boldsymbol{x}_k^1$ has not already been selected and $\boldsymbol{x}_m^1$ is the closest point to $\boldsymbol{x}_k^1$ from those already selected from reference image 1

$$m = \arg\min_p \left| \boldsymbol{x}_p^1 - \boldsymbol{x}_k^1 \right|$$

It is noted that all points further away from x_k^l than x_m^l in image 1 are very likely to satisfy the condition (6) and therefore do not need to be tested according to the same condition. It is assumed therefore that x_k^l is a member of the clique by satisfying conditions (4,5,6).

New candidate points (x_k^l, x_k^2) are selected randomly and added to the clique if conditions (4,5,6) are satisfied. Up to N attempts are made to find a new point after which the current clique is completed and the construction of a new clique started. The search proceeds on a trial and error basis and the selection is not guided by additional heuristics as these have always been found to damage performance. After the generation of P cliques the largest is retained. Let the number of nodes in the maximal clique extracted between the reference image for class c and candidate image i be n_c^i. The classification of image i is given by C_i where

$$C_i = \arg\max_c n_c^i$$

The process allows more than one point in the first image to be mapped into the same point in the second image, but not the reverse. This gives the search more freedom to navigate around occlusions in both images without introducing node crossing. The relationship between points is not dependent upon their separation or absolute position and therefore the similarity measure is translation and scale invariant. It also means that there is no special constraint placed on the disparity of points that is dependent on their separation. The measure is partially invariant to the rotation of the images to within the angle ε_2. It should also be noted that although the cliques are maximal in terms of the algorithm, there is no guarantee that the cliques extracted are the largest theoretically possible; the solution of an NP-complete problem would be necessary to confirm this.

3. YALE FACE DATABASE

In order to make a assessment of performance, the Yale Face Database A [16] is used in this paper. This database consist of 11 categories of expression and lighting from 15 individuals. The categories are Normal, Happy, Glasses, Surprised Wink, Sad, Sleepy, No Glasses, Left Light, Centre Light, and Right Light (Fig. 1).

Figure 1. Subject 1 expressions plus left, centre and right illuminations

Figure 2. Reference faces – Normal category

A great variety of face recognition techniques have been applied to the Yale database ranging from Fisher Linear Discriminant to PCA and SIFT features and published performance figures are given in Table 1. The error rates shown are for the performance using the smallest training set used in the method. Within some approaches performance improves with the size of the training set but with the potential of introducing the effects of overtraining that can affect generalization to unseen data.

Images from the Yale Faces database were reduced in size to 100x76 pixels. The category Normal was used as a reference set (Fig. 2) when measuring the similarity of the 15 faces in the expression, illumination and new occlusion categories. The background in the 15 reference images were was manually erased and set to white. The remaining candidate images were not changed and the background was left intact.

Table 1. Performances reported on the Yale Face Database A

Reference	Test set error rate	Size of training set
Tjahyadi [17]	11.6%	1
Sellahewa [18]	18.0%	1
Pozo-Banos [19]	4.34%	2
Ruiz-del-Solar [20]	7.7%	2
Aly [21]	9.9%	2
Liu [22]	29.0%	3
Cheng [23]	16.84%	3
Du [24]	33.55%	3
Li [25]	24.5%	4
Quintiliano [26]	17.0%	4
Rziza[27]	6.67%	5
Gudivada [28]	5.16%	5
Lu[29]	5.0%	5
Aroussi [30]	2.22%	5
Qi [31]	7.78%	5
Xia [32]	11.3%	6
Hua [33]	34.2%	8
Tang [34]	13.9%	10

Grey levels $d_g(x_i)$ are in the range of 0-255 and match if values differ by no more than $\varepsilon_1 = 160$. The threshold on the angular difference between matching pairs of points in each image is $\varepsilon_2 = 19°$. The grey level gradient is quantized into the four directions 0°, 90°, 180° and 270°. Up to N=100 attempts are made to add new points to a clique and P=100 cliques are generated for each image i, the maximal clique identified, and the classification C_i determined. This defined a fixed framework for clique extraction but with three very broad thresholds thereby enabling more points to become candidates for inclusion in a clique. There is therefore less emphasis placed on the information possessed by individual pixel properties than that contained in the structural relationships between the points forming the clique.

4. FACE RECOGNITION RESULTS

4.1 Expressions

The seven expressions Happy, Glasses, Surprised Wink, Sad, Sleepy, and Noglasses for the 15 subjects were all compared with the reference faces (Fig. 2) and maximal cliques extracted. Fig 3a shows the sizes of maximal cliques for each subject that ranged from 1265 to 2860 nodes. There were no errors, that is, the largest clique was always formed with the correct reference. Fig. 3b shows the totals of clique sizes for each category and reflects the overall distortion of the expression from the references. The Noglasses category contains several faces that are almost identical to those in the Normal category as indicated by high scores; dips in the Noglasses category for subjects 8 and 13 are due to their wearing glasses in the references. The two peaks with subject 4 in the Sad category and subject 8 in the Glasses category are due to identical copies of reference (Normal) images for those subjects being present in the respective Yale Database A categories.

Fig. 4 shows a 1985 node maximal clique extracted from the subject 1 reference and the same subject from the Surprise category. Graph edges are only shown between the four closest nodes for clarity. Fig. 5 shows an enlarged portion around the mouth region in which the grey level gradient direction is indicated by red radial lines.
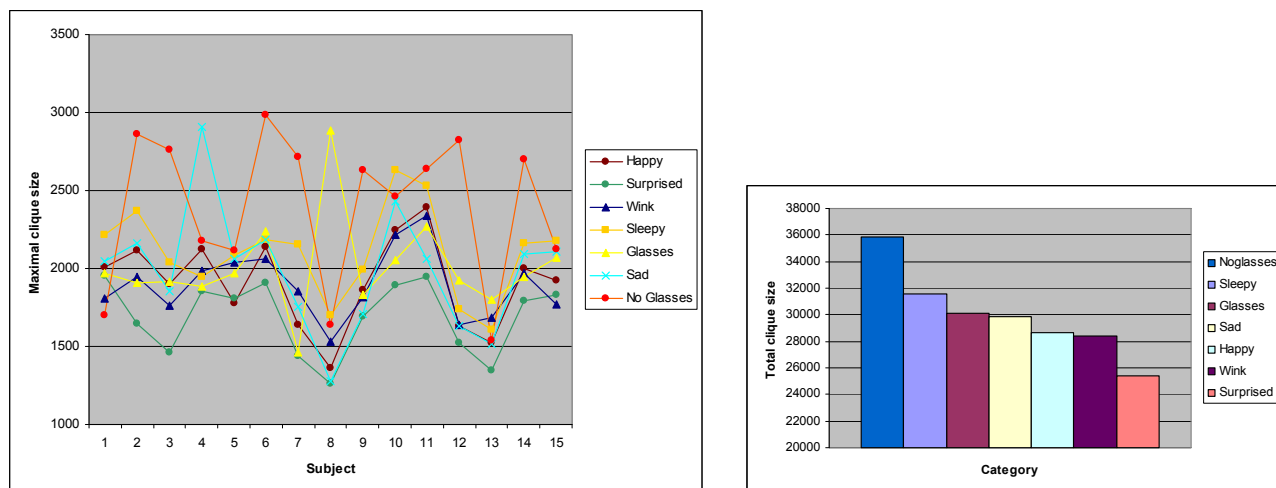
Figure 3. a) Sizes of maximal matching cliques for each subject and expression. b) Clique size totals for each expression

Figure 4. Reference face and Surprised version showing matching maximal clique with 1764 nodes
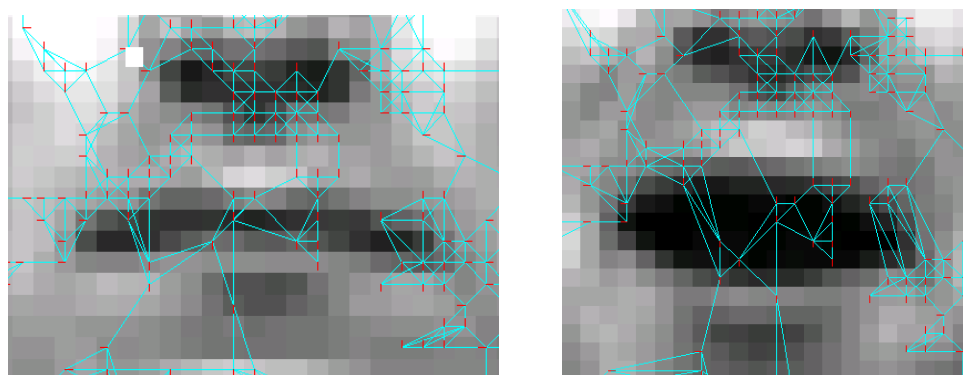

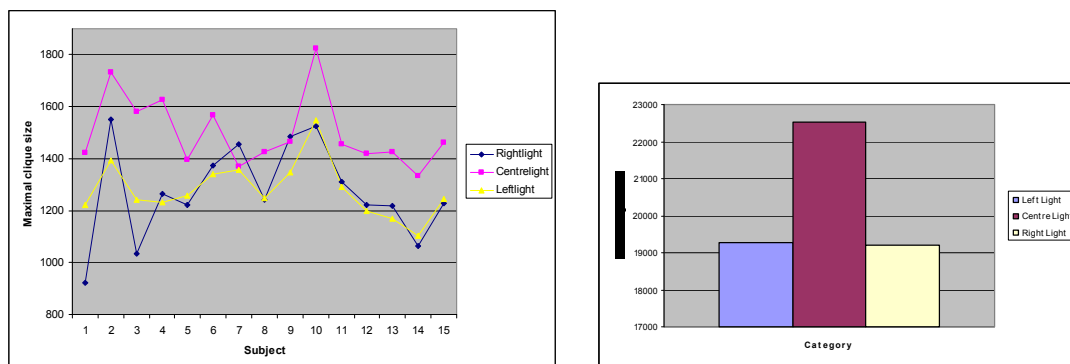Figure 5. Close-up of mouth region in Fig. 4

Figure 6. a) Sizes of maximal matching cliques for each subject and illumination. b) Clique size totals for each category

4.2 Illumination Changes

Illumination changes significantly distort the information available for subsequent processing as is reflected in Fig. 1. The direction of illumination affects the gradient directions in certain regions of the face and prevents points in these regions from being candidates for comparison with the reference face.   The maximal cliques sizes were generally lower than those obtained with the 7 expressions (Fig. 3).   Faces illuminated by Centrelight obtain the highest degree of similarity whilst Left and Right illuminations yield similar lower performances (Fig. 6). Fig. 9b shows a maximal clique matching the Left Light subject 1 face and ignoring the background. There were no errors in the illumination categories and it was noted that despite subject 14 wearing glasses but not in the reference, this face was classified correctly.
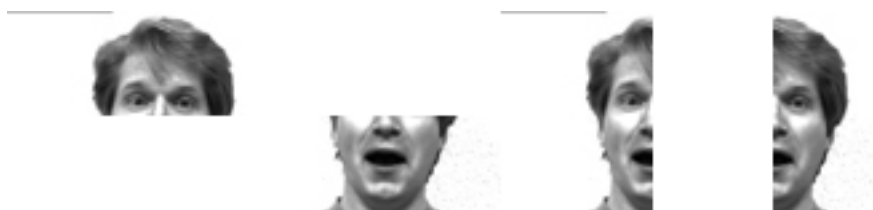
4.3 Occlusions

Figure 7. Subject 1 Surprised Top, Bottom, Left and Right sections

Occlusion is an important distortion that can considerably degrade recognition performance. Some further similarity measurements have been carried out using just the top, bottom, left and right halves of the Surprised set of faces (Fig. 7). The areas of the images were smaller and the relative effect of noise was therefore greater and so the relative angular threshold ε_3 was increased to 20°.
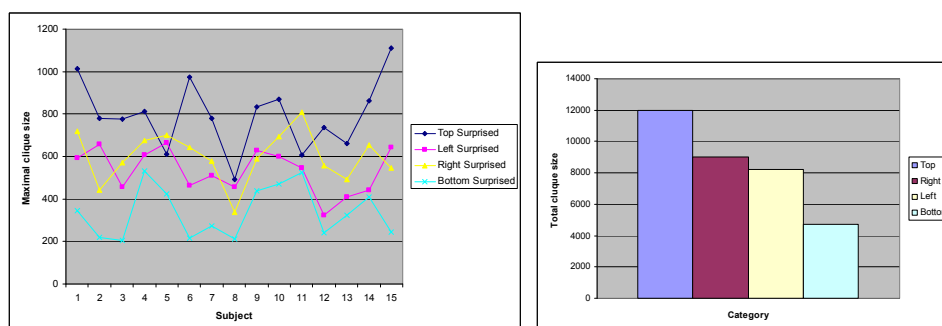

Figure 8: a) Sizes of maximal matching cliques for each subject and occlusion. b) Clique size totals for each category

Table 2. Error rates on occluded faces

	Top	Bottom	Left	Right
% error	0.0	40.0	0.0	13.3

The top sections of the Surprised set of faces yield the largest cliques in common with the reference set and the Bottom sections the lowest (Fig. 8a). Left and Right sections obtain intermediate scores and follow a similar pattern perhaps reflecting the symmetrical nature of facial images (Fig. 8b). Subjects 2 and 8 amongst the Right Surprised gave rise to errors and there were 6 errors in the Bottom surprised group (Table 2). This result is consistent with human performance which finds the recognition of facial identity easier if the eyes and top half of the face are visible, but not if they are obscured. Fig. 9a illustrates how the top half of the face was located in the reference with no points matched in the lower half of the face.

Figure 9. Matching maximal cliques for a) top half of Surprised face. b) Left Light face

4.4 Pose

Pose is another important aspect of facial recognition. A set of poses from MIT [35] (Fig. 10) were analysed for a single subject taking right, forward and left facing poses as references (Fig. 11). Again the noise levels were higher in this data and the relative angular threshold was increased to 23°.

The highest similarities were obtained near the angular locations of the reference poses with similarities dropping off as distances increase (Fig. 12). An analysis of the relative shifts of corresponding points in matching cliques would provide information on the pose position.

Figure 10. 20 poses taken from MIT database

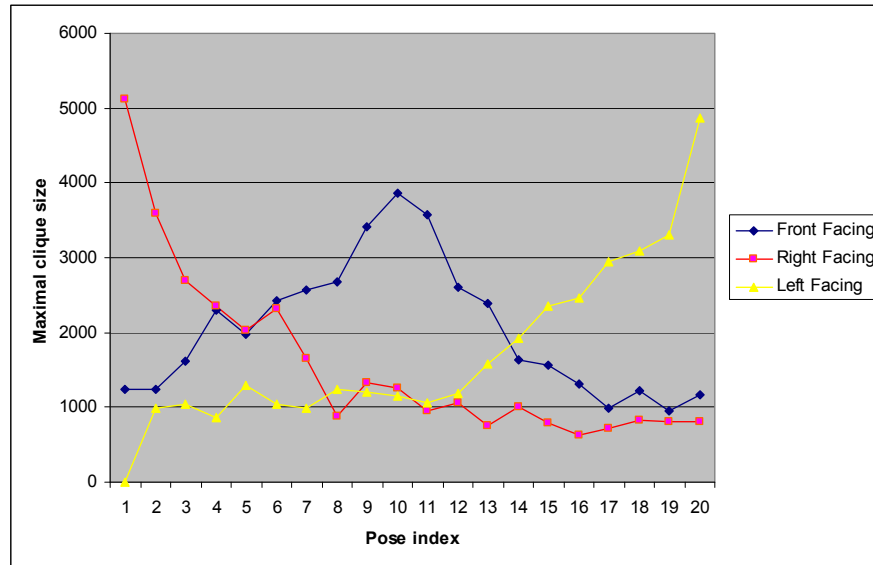Figure 11. Right, Forward and Left facing reference poses

Figure 12.   Sizes of maximal matching cliques for each facial pose

4.5 Clique Structure

The generation of matching maximal cliques in this paper does not guarantee that the cliques are theoretically maximal or that the internode relationships are all satisfied between patterns. The suboptimal size is acceptable if performance is also acceptable, but the reliability of the structural correspondence could be uncertain. As described in section 3 it is assumed that more distant angular relationships between nodes will be satisfied if the closest pairs of nodes satisfy the matching threshold.

Fig. 13 shows the distribution of the maximum relative angle differences between each node and all others in a typical pair of maximal matching cliques with 1985 points. Although not visible in the chart, there is just one instance of a relative angle of 27° that exceeds the threshold of 19°.
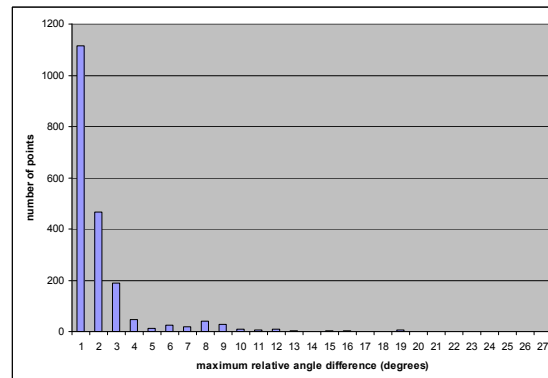

Figure 13. Distribution of maximum relative angular relationships

5.   DISCUSSION

The similarity measure used in this paper identifies structure that is common to pairs of patterns. This means that pattern classes can be accommodated that contain patterns that only possess features in common with just some of the other class members. Some patterns can therefore be members of the same class but have no features in common, a situation that is not permitted by several feature based approaches that are dependent on a metric.

The framework employed here requires a minimal number of representative pattern samples to obtain recognition. This reduces the likelihood of overtraining effects on generalisation. In addition the two parameters ($\varepsilon_1, \varepsilon_2$) that define the operation of the clique matching process are independent of the pattern content and were unchanged when processing Yale faces. The values of ε_1 and ε_2 were chosen by looking at a single pair of images and found not to be critical to performance. Increasing the number of cliques generated (*P*) and the number of attempts to add new nodes to a clique (*N*) increases the likelihood of discovering larger cliques. However, currently the search for maximal cliques automatically obtains registration and *P* and *N* therefore could be significantly reduced (with the associated reduction in processing) given prior information on the registration of the pattern pairs.

The clique extraction process is scale independent as is illustrated by the flexible matching in Fig. 5. However, at the same time the magnitude of the similarity score is dependent on the size of the area of interest in the image and the structural information available and therefore a lower definition image will yield a lower similarity score despite scale independent matching. As with changes in the pose angle, smaller clique sizes reflect less structure in common and indicate the lower recognition reliability arising from smaller images. The recognition is also largely independent of brightness; a very wide threshold (160/255) is used during pixel matching. Illumination direction, however, changes not only the brightness but also the gradients in various parts of the face. Nevertheless, unaffected regions and low level indirect lighting of other parts of the face still allowed useful matching to take place in those regions and obtained correct recognition.

It has been shown that common structure can still be identified if only part of the candidate image is available. This has benefits when patterns are occluded, but again the image size is reduced and the reliability of the recognition is reduced. This also applies to changes in pose angle when larger parts of the frontal view are obscured by the rotation. It is possible in future work that the pose angle could be determined from the information contained in the node correspondences between the maximal cliques in the frontal reference image and the posed image.

Earlier research on modelling visual attention in the human visual system has employed maximal cliques to determine the similarity of regions within the *same* pattern [36]. This has enabled the background in an image to be recognised thereby isolating and identifying salient objects. The operation of the model corresponds with human behaviour in many examples. In addition the maximal clique similarity measure has been shown to fail in the same manner as the human visual system on the Poggendorff illusion [37].

The computation takes approximately a minute in Matlab to compare two grey level images (100x76). This time is much less for face pairs that differ significantly because searches are abandoned while cliques are still small, but is larger if the images are very similar. This may be reduced as discussed above with prior registration. However, the extraction of each clique is an independent operation and may therefore be conducted in parallel both for each reference *c* and the *P* attempts at maximal clique construction. Furthermore parallel operations could also be introduced during graph matching itself by allowing additional nodes to be added simultaneously at different locations. This is possible because the conditions for addition are dependent only on local node properties. This means that the overall potential for a speedup of many orders of magnitude is possible in an appropriate implementation.

6. CONCLUSIONS

This paper has demonstrated the existence of matching clique-like structures that obtain for the first time no errors when applied to the classification of images in the Yale Face Database A and represents an improvement on error rates in earlier work. The approach requires no training stage and shows immunity to changes in translation, expression, illumination and occlusion. The effect of pose angle was also explored. Although the serial implementation is slow there is potential for very fast parallel operation.

Support for the approach is also obtained from its use in modelling aspects of human vision. Further work is necessary on larger and more challenging datasets. There is scope for including colour in the clique node properties as well as increasing the image resolution to improve discrimination on more difficult images.

REFERENCES

[1] Patel, R., Rathod, N. and Shah, A., "Comparative analysis of face recognition approaches a survey," Int. J. of Computer Applications, 57(17), 50-61 (2012).

[2] Naruniec, J., "A survey on facial features detection," Int. J. of Electronics and Telecommunications, 56(3), 267-272 (2010).

[3] Leordeanu, M. and Hebert, M., "A spectral technique for correspondence problems using pairwise constraints," ICCV, (2005)

[4] Felzenszwalb, P. F. and Huttenlocher, D. P., "Efficient matching of pictorial structures," CVPR, (2000).

[5] Fergus, R., Perona, P. and Zisserman, A., "A sparse object category model for efficient learning and exhaustive recognition," CVPR, (2005)

[6] Kim, J. and Grauman, K., "Asymmetric region-to-image matching for comparing images with generic object categories," CVPR, (2010)

[7] Duchenne, O., Joulin, A. and Ponce, J., "A graph-matching kernel for object categorization," ICCV, (2011).

[8] Duchenne, O., Bach, F., Kweon, I. and Ponce, J., "A tensor-based algorithm for high-order graph matching," IEEE Trans PAMI, 33(12), 2383-2395 (2011).

[9] Celiktutan, O., Wolf, C. and Sankur, B., "Fast exact matching and correspondence with hyper-graphs on spatio-temporal data," Technical Report LIRIS-RR-2012-002, INSA-Lyon, (2012).

[10] Kolmogorov, V. and Zabih, R., "Computing visual correspondence with occlusions using graph cuts," ICCV, (2001).

[11] Berg, A. C., Berg, T. L. and Malik, J., "Shape matching and object recognition using low distortion correspondences," CVPR, (2005).

[12] Cho, M. and Lee, K. M., "Progressive graph matching: making a move of graphs via probabilistic voting," CVPR, (2012).

[13] Wiskott, L., Fellous, J-M., Kruger, N. and von der Malsburg, C., "Face recognition by elastic bunch graph matching," IEEE Trans. Pattern Anal. Machine Intell., 16, 775-779 (1997).

[14] Cootes, T. F., Edwards, G. J. and Taylor, C. J., "Active Appearance Models," IEEE Trans. Pattern Anal. Machine Intell., 23, 681-685 (2001).

[15] Matthews, I. and Baker, S., "Active appearance models revisited," Int. J. Computer Vision, 60(2), 135-164 (2004).

[16] Yale Face Database, http://cvc.yale.edu/projects/yalefaces/yalefaces.html.

[17] Tjahyadi, R., Liu, W., An, S. and Venkatesh, S., "Face recognition based on ordinal correlation," Proc. Intelligent Sensors, Sensor Networks and Information Processing Conference, (2005).

[18] Sellahewa, H. and Jassim, S. A., "Illumination and expression invariant face recognition: toward sample quality-based adaptive fusion," IEEE International Conference on Biometrics: Theory, Applications and Systems, (2008).

[19] Pozo-Banos, M., Travieso, C. M., Alonso, J. B. and Ferrer, M. A., "Discriminative multi-projection vectors: modifying the discriminative common vectors approach for face verification," IEEE International Carnahan Conference on Security Technology, (2010).

[20] Ruiz-del-Solar, J. and Navarrete, P., "Eigenspace-based face recognition: a comparative study of different approaches," IEEE Trans. On Systems Man and Cybernetics – part C: Application and Reviews, 35(3), 315-325 (2005).

[21] Aly, M., "Face recognition using SIFT features," Caltech Technical Report, (2006).

[22] Liu, X., Yin, J., Feng, Z., Dong, J. and Wang, L., "Orthogonal neighbourhood preserving embedding for face recognition," ICIP, (2007).

[23] Cheng, M., Tang, Y. Y. and Pun, C., "Nonparametric feature extraction via direct maximum margin alignment," International Conference on Machine Learning and Applications, (2011).

[24] Du, H., Guo, X., Zhang, F. and Zhao, J., "Feature extraction based on fuzzy bidirectional maximum margin criterion for face recognition," IEEE International Conference on Intelligent Computing and Intelligent Systems, (2010).

[25] Li, W., Ruan, Q. and Wan, J., "Graph-preserving shortest feature line segment for dimensionality reduction," Neurocomputing, 110, 80-91 (2013).

[26] Quintiliano, P., Guadagnin, R. and Santa-Rosa A., "Practical procedures to improve face recognition based on eigenfaces and principal component analysis," Pattern Recognition and Image Analysis, 11(2), 372-375 (2001).

[27] Rziza, M., Aroussi, M. E., Hassouni, M. E., Ghouzali, S. and Aboutajdine, D., "Local curvelet based classification using linear discriminant analysis for face recognition," World Academy of Science, Engineering and Technology, (2009).

[28] Gudivada, S. and Bors, A. G., "Face recognition using ortho-diffusion bases," 20th European Signal Processing Conference (EUSIPCO), (2012).

[29] Lu, C., Zhang, C., Zhang, T. and Zhang, W., "Kernel based symmetrical principal component analysis for face classification," Neurocomputing, 70, 904-911 (2007).

[30] Aroussi, M. E., Ghouzali, S., Rziza, M. and Aboutajdine, D., "Face recognition using enhanced Fisher linear discriminant," 5th International Conference on Signal Image Technology and Internet Based Systems, (2009).

[31] Qi, Y. and Zhang, J., "Two-dimensional multiple maximum scatter difference method for face recognition," Applied Mathematics and Computation, 216, 3550-3557 (2010).

[32] Xia, Q., "A unified framework for PCA, LDA and LPP," Proc. SPIE 4th International Conference on Digital Image Processing, (2012).

[33] Hua, W. and He, X., "Discriminative concept factorization for data representation," Neurocomputing, 74, 3800-3807 (2011).

[34] Tang, H., Lyu, M. R. and King, I., "Face recognition committee machine," IEEE International Conference on Acoustics, Speech and Signal Processing, (2003).

[35] Weyrauch, B., Huang, J., Heisele, B. and Blanz, V., "Component-based Face Recognition with 3D Morphable Models," First IEEE Workshop on Face Processing in Video, Washington, D.C., (2004). (see also http://cbcl.mit.edu/software-datasets/heisele/facerecognition-database.html).

[36] Stentiford, F. W. M., "Saliency identified by absence of background structure," Proc. Human Vision and Electronic Imaging XVIII Conf., San Francisco, (2013).

[37] Stentiford, F. W. M., "Interest point analysis as a model for the Poggendorff illusion," Proc. Human Vision and Electronic Imaging XVII, SPIE Conf., San Francisco, (2012).