Reducing the Dimensions of Texture Features for Image Retrieval Using Multi-layer Neural Networks

Article *in* Pattern Analysis and Applications · June 1999

DOI: 10.1007/s100440050028 · Source: DBLP

Jesse Sheng Jin
University of Newcastle

Tom Gedeon
Australian National University

Pattern Analysis & Applications (1999)2:196–203

Reducing the Dimensions of Texture Features for Image Retrieval Using Multi-layer Neural Networks

J. Antonio Catalan, J.S. Jin and T. Gedeon

School of Computer Science and Engineering, University of New South Wales, Sydney, Australia

Abstract: This paper presents neural network-based dimension reduction of texture features in content-based image retrieval. In particular, we highlight the usefulness of hetero-associative neural networks to this task, and also propose a scheme to combine the hetero-associative and auto-associative functions. A multichannel Gabor-filtering approach is used to derive 30-dimensional texture features from a set of homogeneous texture images. Multi-layer feedforward neural networks are then trained to reduce the number of feature dimensions. Our results show that the methods lead to a reduction of up to 30% while keeping or even improving the performance of similarity ranking. This has the benefit of alleviating the ill-effects of the high dimensionality of features in current image indexing methods and resulting in significant speeding up retrieval rates. Results using principal component analysis are also provided for comparison.

Keywords: Gabor filter; Image indexing; Image retrieval; Multi-channel filtering; Neural networks; Texture analysis

1. INTRODUCTION

The rapid development of digital imaging and storage technology has led to an explosive growth in the amount of available digital image data. An emergent issue in the management of such data is the ability to quickly access and retrieve images based on their content. The significance of content-based indexing lies in its ability to carry intrinsic image information, which is quite difficult with text-based annotation, and it is amenable to automation. Current image database management systems commonly employ texture, shape and colour descriptors to characterise image data. Of these three, texture is the more difficult one to describe using words, hence content-based indexing methods are particularly helpful. This paper focuses on content-based retrieval by texture. However, the technique developed can be easily applied to other high dimensional features.

Several methods to extract texture features for indexing and retrieving images are described in the literature [1–3]. One of these, the multi-channel filtering approach using 2D Gabor functions has been applied successfully to various problems in texture analysis including segmentation, classification and image retrieval [1,4–6]. The popularity of this method can be linked in part to findings that Gabor functions conform well to the receptive field profiles of the mammalian visual cortex [7]. Furthermore, these functions have the property of optimal joint resolution in the space and spatial-frequency domains which is important for texture discrimination tasks [8]. Since it is not uncommon for features obtained with this approach to have 20 or more dimensions, some issues arise due to high dimensionality.

Many of currently available multi-dimensional indexing techniques fail to handle high-dimensional data efficiently [9–11]. One of the more robust methods, the R*-tree works well up to 10 dimensions [12], but its performance degrades quickly with increasing dimension and data sizes. An obvious solution to the problem of exponential growth of multi-dimensional indexing structures with increasing dimension of data is to reduce the dimension.

Broadly speaking, dimension reduction is a process of selecting a low-dimensional subspace that would best approximate a phenomenon that is observed in a high-dimensional space. Applying it to the problem of image indexing and retrieval means selecting a subset or low-dimensional transformed set of image feature data with minimal or no loss of discrimination power in the query

Received: 6 July 1998
Received in revised form: 6 November 1998
Accepted: 15 December 1998

domain. If general applicability with the prevailing indexing structures is of concern, the dimension reduction should, if not improve, at least preserve the distance relationships between the data elements in the reduced space. Applying dimensionality reduction to a data set should ideally produce retrieval accuracy, which is at least at par with that of the original set.

A popular technique of dimension reduction that immediately comes to mind is Principal Component Analysis (PCA) [13] which is widely used because of its simplicity and the availability of efficient computational algorithms. It is a mapping technique that involves transforming data such that variables are decorrelated and a few variables account for most of the variances. Dimension reduction is achieved by dropping variables with insignificant variances. An improvement on linear PCA is provided by non-linear generalisations of PCA which extend the ability of PCA to incorporate non-linear relationships in the data [14].

A method of dimension reduction that does not involve transforming data is described in Liu and Picard [2]. Here, a subspace of the feature vector is selected depending on some salient properties of the query image [2]. Filter channels with high discriminatory power are identified, while weak ones are discarded leading to a reduced search space.

Neural networks, in particular, associative neural networks, form another class of solutions to the problem of dimensionality reduction. There are two kinds of associative neural networks: auto-associative and hetero-associative networks. The former takes the training patterns, memorises them and tries to reproduce them as the output. The latter takes input patterns distinct from the training patterns, and the output gives the best classification of the input in terms of the training categories. Note that both are supervised networks.

Some of the auto-associative networks, the so-called 'PCA-type' networks, produce results very similar but not superior to PCA [15] in terms of reconstruction error. Both linear and non-linear PCA transforms can be implemented by multi-layer feed-forward neural networks [14,16]. We will use multi-layer network throughout the rest of the paper because our discussion is restricted to networks which have only feed-forward connections from input to output. However, readers should keep in mind that there are other more general multi-layer networks, namely recurrent networks, with connections allowed both ways between a pair of units, and even from a unit to itself. The auto-associative network performs the identity mapping through a compression layer, and the goal is to keep as much information in this layer as possible in order to reconstruct the input data at the output layer. The use of one hidden layer in these networks allows dimensionality reduction akin to PCA while increasing the number of hidden layers to three allows non-linear PCA mapping. Auto-associative networks can achieve high dimension reduction rate. However, the similarity preservation is poor because the network is only sensitive to the training patterns. When there are more than one hidden layers, the network is difficult and slow to train.

Unlike auto-associative networks, hetero-associative neural networks develop arbitrary internal representations in the hidden layers to associate inputs to class identifiers usually in the context of pattern classification. The hetero-associative network preserves similarity well, though it has a complicated structure. The training can be easily trapped into local minima, especially when there are overlaps among classes and training starts from patterns in the overlapping area. It is difficult to avoid this by selecting the training patterns because of the high dimensionality. Besides, the dimension reduction rate is not as high as the auto-associative networks.

The adaptive property of auto-associative networks and the classifying property of hetero-associative networks remind us that a more robust scheme can be achieved by combining auto-associative and hetero-associative functions in one network. In this way, more complete information is embedded in the compression layer, which is adjusted to jointly satisfy both the requirements of dimensionality reduction and the similarity in the application domain. The objective of this paper is to investigate the usefulness of the hetero and the hybrid nets in image indexing and retrieval. We included results of one-hidden-layer auto-associators and the standard PCA to demonstrate how well the hetero and the hybrid nets fare against them. The use of better auto-associators (e.g. three-hidden layer nets) has been left to future investigation.

2. REDUCING DIMENSION OF TEXTURE FEATURES

Image features, such as texture, shape and colour, are often of high dimension. We will concentrate our discussion on texture features, in particular, multi-channel Gabor texture.

2.1. Gabor Texture

The multi-channel texture analysis involves decomposing an image into channels using frequency and/or orientation selective filters. This approach has been inspired by results of psychophysical experiments, which suggest that different cells in the visual cortex of the Macaque monkey are tuned to different combinations of frequencies and orientations [17]. The response of these cells have also been found to be modeled closely by 2D Gabor functions. It is, therefore, not surprising to see the use of multi-channel Gabor filtering in many researches on texture analysis.

The Gabor function, first introduced by Gabor in one dimension [18], is basically a Gaussian modulated by a complex sinusoid. Daugman [19] extended it to two dimensions. A general 2D Gabor function is given by

$$g(x,y) = \frac{1}{2\pi ab}\exp\left\{-\frac{1}{2}\left[\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2\right]\right\}\exp[j2\pi\,(Ux+Vy)] \tag{1}$$

where (x,y) define the spatial coordinates and (a,b) specify the width of the Gaussian in the spatial domain and equivalently the bandwidth in the frequency domain. U and V specify the 2D frequency of the complex sinusoid.

In our experiments, however, we used a simpler form:

$$g(x,y) = \exp[-4\ln 2(x^2 + y^2)/w^2]\cos[2\pi f(x\cos\theta + y\sin\theta) + \varphi] \qquad (2)$$

where w is the width, f is the frequency, θ is the orientation and φ is the phase. This defines an even-symmetric Gabor function which is described in Watson [20]. While some researchers use complex-valued filters [21] or filter-pairs with quadrature phase [22], Malik and Perona give justifications on the use of only even-symmetric filters [23]. The frequencies used are 48, 24, 12, 6, 3 and 1 cycles per image width. The orientations are 0, 36, 72, 108 and 144 degrees. The width is set as 1.324 cycles of the sinusoid.

The texture feature vector used to characterise each image in our experiments is derived with the setup shown in Fig. 1. First, the image goes through 30 channels by convolving it with Gabor functions defined by a combination of one of six frequencies and one of five orientations. A single value, which becomes an entry in a 30-dimensional texture feature vector, is then extracted from each channel. We take the average value of the magnitude of the filtered image in each channel.

2.2. Reducing Dimension Using Hybrid-Associative Neural Networks

Many neural network architectures are capable of universal function approximation and can perform arbitrary feature abstraction when mapping inputs to outputs. Multi-layer perceptron (MLP) networks, for example, are believed to form internal representations of the input data in the hidden layers. If the number of nodes in a hidden layer, called the representation layer, is less than the number of input nodes, the network finds a more compact representation complete enough to allow the desired hidden-layer-to-output-layer mapping. This merits consideration of these networks for the task of dimension reduction. Another neural network model that is commonly linked to dimensionality reduction is Kohonen's Self-Organising Feature Map (SOFM) [24]. The SOFM is an unsupervised neural network and performs non-linear topology-preserving projection from input spaces of arbitrary dimensions to one or two dimensional feature maps. Although possible, extending this to maps of greater than two is quite awkward, if not daunting. The following discussions focus mainly on supervised neural networks.

Dimension reduction can be achieved in auto-associative MLP networks when a hidden layer, called the representation layer, has a dimension l less than the input (and output) dimension d [25,26]. It has been shown [15,27] that if the hidden nodes are linear, the network converges to a solution (albeit not perfect) which is a projection onto the l-dimensional subspace spanned by l principal components. It is also shown that adding non-linearities in the hidden layer does not help. The optimum solution that can be found with only one hidden layer is that given by PCA. In this respect, the use of auto-associators with a single hidden layer does not offer any advantage considering that it takes a long time to train, whereas PCA can be done using efficient techniques for matrix computation.

In contrast with auto-associative networks, the output pattern in hetero-associative networks is not the same as the input pattern for each training pair. If each distinct output pattern corresponds to a class which one or more of the input patterns are associated to, the network is called a classifier. However, dimension reduction is not a classification process. Hetero-Associative MLP (HAMLP) networks perform the input-to-output mapping through one or more hidden layers. The activation outputs of the hidden layers are considered as arbitrary internal representations or features formed by the network in a multi-stage mapping process. If *any* of the hidden layers have a dimension less than the input pattern dimension, the network is forced to produce a more compact representation of the data sufficient enough to carry through with the desired output mapping in the subsequent layers, in which the dimension reduction is achieved.

An HAMLP network trained as a classifier would allow the transformation of image data in such a way that similar images are bunched near each other while the distances between dissimilar groups are increased depending on the criteria used for the classification. This should lead to an improvement in similarity retrieval efficacy. The difficulty of using HAMLPs is to choose the training samples. Improper samples will lead the system to a local minima.

Considering the completeness and the weakness of the auto-associative and hetero-associative networks, we propose a hybrid hetero-associative/auto-associative (hybrid-associative) neural network, which combines a hetero-associative network and an associative auto-associative network into one, as shown in Fig. 2. This network is constructed with a set of auto-associative outputs and a set of hetero-associative outputs. Both sets of output nodes are fully

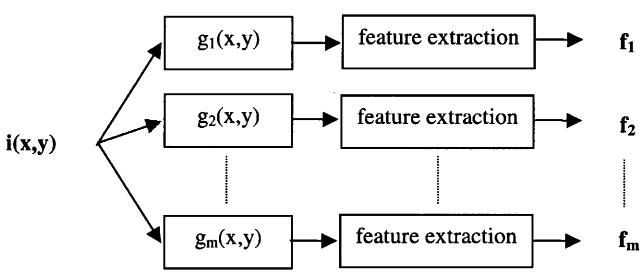

Fig. 1. Multi-channel filtering strategy.

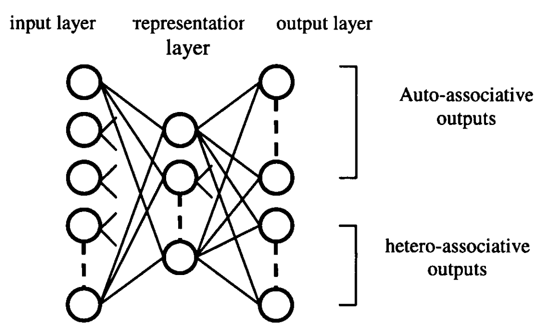

Fig. 2. Hybrid-associative networks.

connected to the same bottle-neck hidden layer. The network is trained to simultaneously implement an identity mapping from the input layer to the auto-associative subset of the output layer, and a hetero-associative mapping to the rest of the output layer. After training, the hybrid network should have learned both the hetero-associative and the auto-associative functions. Consequently, the advantages to similarity retrieval due to hetero-associative reorganisation should be attained and the representation should retain much of the input information, since it also allows demapping to a set of auto-associative outputs. The hybrid network may thus provide a more robust approach to dimensionality reduction than just a hetero-associative network.

Our hybrid-associative neural networks are trained using the backpropagation algorithm [28,29], a gradient descent learning rule. Starting with an initial set of small random weights, the algorithm searches through the weight space moving in the direction of the steepest descent, to minimise the cost function

$$E = \frac{1}{2}\sum_{k=1}^{n} (y_k - d_k)^2 \tag{3}$$

after presentation of each training pattern, where y_k is the response of output node k, d_k is the desired output response, and n is the total number of output nodes. The weights are adjusted according to the update rule

$$w_{ij}(t+1) = w_{ij}(t) + \eta\delta_j x_i \tag{4}$$

where w_{ij} is the weight from node i in layer L-1 to node j in layer L, η, is the learning parameter, δ_j is the error at node j and x_i is response of node i. If L is the output layer,

$$\delta_j = (d_j - x_j)\, x_j(1 - x_j) \tag{5}$$

or if L is a hidden layer,

$$\delta_j = x_j\,(1 - x_j) \sum_k \delta_k\, w_{jk} \tag{6}$$

where δ_k is the error at node k in layer $L+1$. The response of a node j is given by

$$x_j = \frac{1}{1 + e^{-\sum_m w_{mj} x_m}} \tag{7}$$

where m is a node at the layer preceding the node j layer. The training cycle continues until the cost function reaches a specified level.

3. EXPERIMENTS AND RESULTS

3.1. Empirical Methods on Dimension Reduction

We apply four methods: PCA, auto-associators, hetero-associators and hybrid-associators, to dimension reduction and compare their performance. We use the Karhunen-Loeve transformation in the PCA method to map initial features **x** to uncorrelated components **y**. The first component exhibits the largest variance and the last component, the least variance. This allows insignificant components to be dropped off with minimal resultant residual error and data reduction is achieved. Thus, it is possible to map the original d-dimensional data onto an l-dimensional subspace spanned by l most significant covariance eigenvectors where $l < d$.

The auto-associator used in the experiment is one hidden layer Associative MLP, the hetero-associator used is one hidden layer hetero-associative MLP, and the hybrid-associator used is one hidden layer hybrid hetero-associative/auto-associative MLP.

There are three databases used in the experiments, each consists of 530 homogeneous 128×128 pixel texture images with 53 distinct classes having 10 members each. A sample set of the image with one from each class is shown in the Appendix. Database A is a reference set, and used as the training set for the neural network methods and the base set for PCA computations. Database B is basically the same as A except that the images are shifted 10 pixels to right and 10 pixels down. Database C has the images of A rotated 90° clockwise. The feature data for each of the images are normalized to the range [0,1] separately per dimension. Because similarity is subjective, we choose these test data to eliminate the side effect of the similarity ranking problem.

Each type of neural network was trained with backpropagation using all the images in database A until the root-mean-square error is 0.01 or the number of epochs reach 5000. Both thresholds are sufficient enough to distinguish converging or non-converging situation as we are interested in the comparison of the reduction rate. For the hetero-associative network, 1-of-53 encoding is used, that is, there are 53 outputs, each corresponding to a class. Only one of the outputs should be high while the others remain low if the class corresponding to it is detected. To illustrate, if an input vector for an image belonging to class 1 is presented, output 1 is trained to go high (value 1), while all the other are trained to output low (value 0). On the other hand, if an input for an image belonging to class 2 is presented, output 2 is trained to go high, and so on. Training proceeds until all output nodes respond correctly (i.e. rms error reaches desired level) for each input in the training set, or until the maximum allowed number of epochs has elapsed.

The hybrid hetero-associative/auto-associative network has a total of 83 outputs of which 30 are auto-associative and 53 are hetero-associative. It simply combines the required outputs of the auto-associative net and the hetero-associative net into one larger network. It is trained so that both parts are jointly satisfied. Once trained, the hybrid network can function either as an auto-associator or hetero-associator, or both.

Once each network is trained, the feature data from each of the databases are used to activate the network and the transformed or reduced data is obtained from the hidden layer. The reduced data is used to form the texture index.

Similarity retrieval proceeds by common k-nearest neighbours search (i.e. Finding k nearest neighbours in terms of distances) over an index tree. Euclidean distances between the reduced texture data of the query image and the k

nearest neighbour images are calculated and sorted in ascending order. Retrieval results are listed in the same order with distance serves as similarity. The retrieval performance is measured in terms of percentage of correct retrieval (i.e. belonging to the appropriate class), when each image in the query set is presented. Ideally, if each member of a class is presented as a query, all the members of the class should be retrieved before members of the other classes are.

3.2. Experimental Results

3.2.1. Retrieval Scores. Table 1 shows a comparison of the retrieval scores for each of the dimensionality reduction methods used. Group I entries are the results when database A is used as both query and reference sets. Group II and Group III results were obtained with database A as reference and databases B and C are used as query sets, respectively. Retrieval performances for each database versus the number of top matches considered are given in Figs 3 to 5. In summary,

- Due to the linearity, PCA and auto-associative networks retains much of the original information. They can achieve very high reduction rate. For example, down to dimension 5, they perform better than hetero-associators and hybrid-associators.
- PCA performance stays close to those obtained with the raw data but maintained it quite consistently down to five dimensions.
- The auto-associative network registers the poorest results among the network-based methods but stays within 2–8% off the performance using the raw data at lower dimensions.
- The hetero-associative network is clearly superior to the other approaches as seen in its performance in all the databases used. It is able to maintain a retrieval performance better than with the normalised raw data using the training set as query set even down to 10 dimensions – a reduction of 33%. This performance, however, is not sustained when the query images do not belong to the training set.
- The hybrid network performs better than PCA overall, but does not really match the performance of the purely hetero-associative network as would have been expected. However, the hybrid network produces higher inter-class distance than the hetero network (Table 2). It is a very important property in developing index tree of image features.

3.2.2. Browsing Results. Browsing results, given in Figs 6 and 7, are obtained by one-to-one correlation and show the significance of our approach. In Fig. 6, where the raw feature data are used, only 8 of a possible of 10 members are retrieved in the top 13 matches in the class to which the query image (upper left image) belongs. Only 7 are retrieved in the top 10 matches. If the feature data is transformed and reduced to 15 dimensions by a hetero-associative network, the browsing result in Fig. 7 is obtained. It shows that 10 of 10 class members are retrieved within the top 13 and

Table 2. Average inter-class distances normalised by dimension

	Raw	Auto-asso.	Hetero-asso.	Hybrid-asso.
Average normalised inter-class distance	0.022	0.070	0.019	0.059

Table 1. Retrieval performance scores (% correct retrieval)

		No. of Hidden units					
		5	10	15	20	25	30
	Raw			65.64			
	Hetero-asso.	59.45	75.77	81.30	78.85	78.85	77.75
Group 1	Auto-asso.	60.30	62.89	58.30	62.42	60.23	63.23
	Hybrid-asso.	55.77	73.62	73.42	75.06	72.75	75.68
	PCA	63.25	65.43	65.83	65.62	65.60	65.64
	Raw			62.91			
	Hetero-asso.	53.53	66.28	68.36	72.06	72.38	72.43
Group II	Auto-asso.	55.72	60.45	56.00	56.66	56.38	60.68
	Hybrid-asso.	50.32	62.96	67.49	68.79	67.53	69.17
	PCA	62.42	64.26	64.58	64.68	64.72	64.70
	Raw			48.98			
	Hetero-asso.	41.94	60.09	57.06	57.32	60.36	57.68
Group III	Auto-asso.	43.75	46.94	41.25	41.15	40.87	43.58
	Hybrid-asso.	47.75	57.89	51.06	55.04	50.17	54.28
	PCA	51.38	52.85	52.02	52.25	51.57	51.53

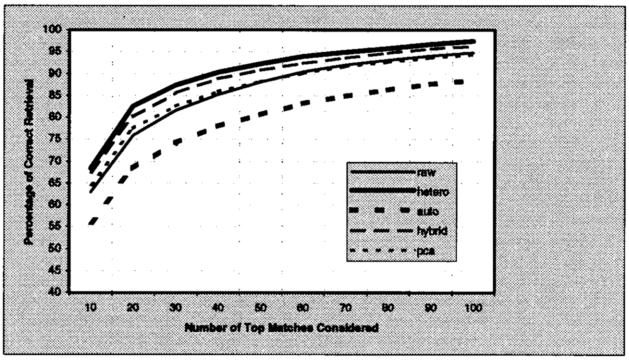

Fig. 3. Retrieval performance vs. number of top matches considered for database A using 15 dimensions.

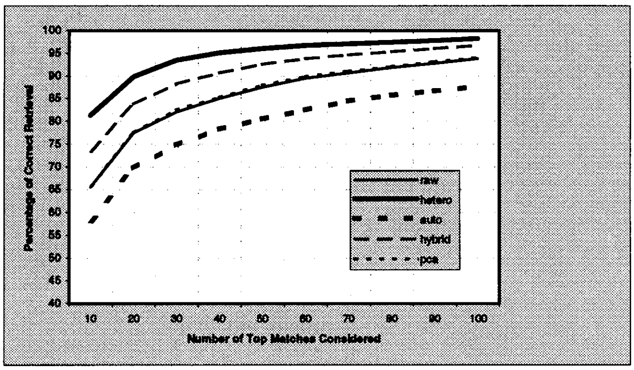

Fig. 4. Retrieval performance vs. number of top matches considered for database B using 15 dimensions.

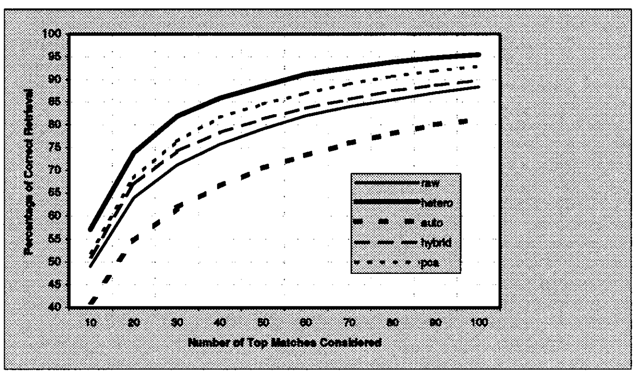

Fig. 5. Retrieval performance vs. number of top matches considered for database C using 15 dimensions.

9 members within the top 10. The former shows 15% improvement on recall rate and the latter shows 20% improvement on recall precision. Browsing performance depends upon the characteristics of the image classes and how well the set of texture features used can discriminate between the classes.

3.2.3. Discussions. The improvements seen using hetero-associative training can be attributed to the reorganisation

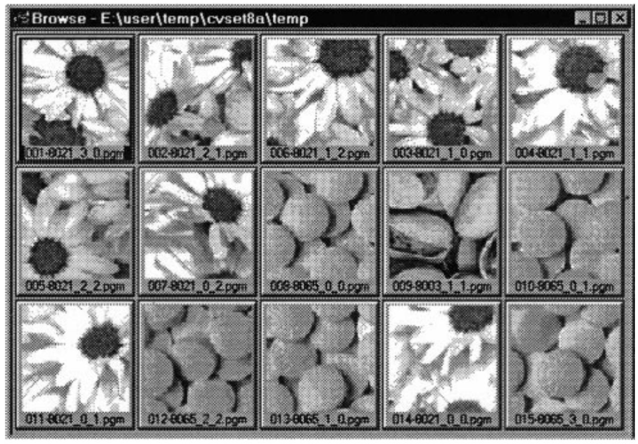

Fig. 6. Sample browsing result for the raw dataset A.

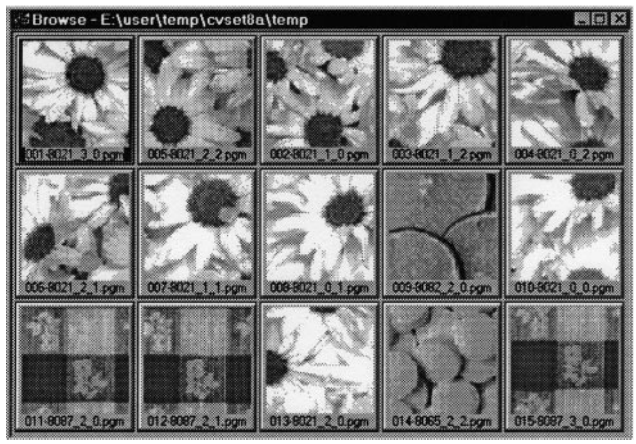

Fig. 7. Sample browsing result for dataset A reduced to 15 dimensions by hetero-association.

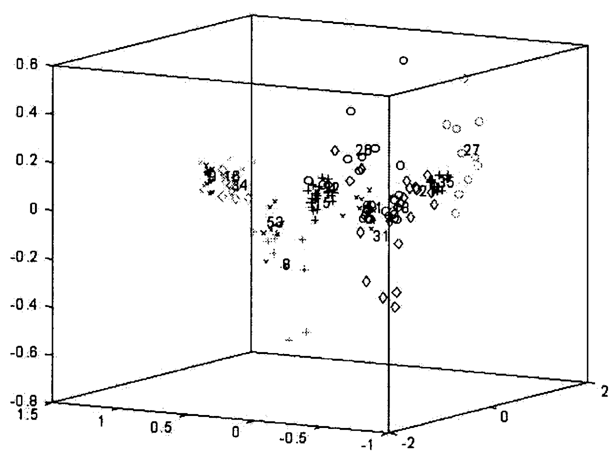

Fig. 8. Plot of the first three principal components of the raw feature vectors of selected classes in dataset A.

done by the network in the representation layer space. Figures 8 and 9 show plots of the first three principal components of the raw vector and 15-dimensional vector at the hidden layer of a trained hetero-associative network, respectively, for few of the classes in database A. The plot in Fig. 9 indicates that the hetero-associative network generally increases the inter-class distances, but tend to decrease the intra-class distances in heavily populated subspaces. These effects are echoed in Table 2, which lists the average inter-class distances normalized by the number of dimensions. The inter-class distances for the data transformed by hetero-association and the combination of auto-association and hetero-association are shown to be greater than that of the raw data. This translates to an increase in performance for similarity retrieval among the classes involved.

The poorer retrieval performance of hetero-associative dimensionality reduction with database B queries in comparison with those using database A indicates that the network does not completely learn the necessary generalisations. The hetero-associative network conformed more to the specific features of the training set, which in this case, is database A. Since the database B images are just shifted versions of those in A and are homogeneous images, most of the classes should generally exhibit the same patterns. In view of this, the ability to discriminate should more or less be the same for both databases. A possible sources of failure is a limitation in discrimination ability of the texture measure used, which just the simple average grey values of the Gabor filtered images. The use of more sophisticated measures may feed more meaningful information to the network for it to learn the required generalisation. Another possible cause of failure is the insufficiency of the network architecture used. Adding at least one more layer may improve the network's ability to discriminate between textures. Nonetheless, the three-layer hetero-associator has posted improvements over PCA even if the query images are 90° rotated versions of the training images.

Table 3 gives an insight into the relationship between the retrieval and classification performances of the networks with hetero-associative functions. It can be seen that the

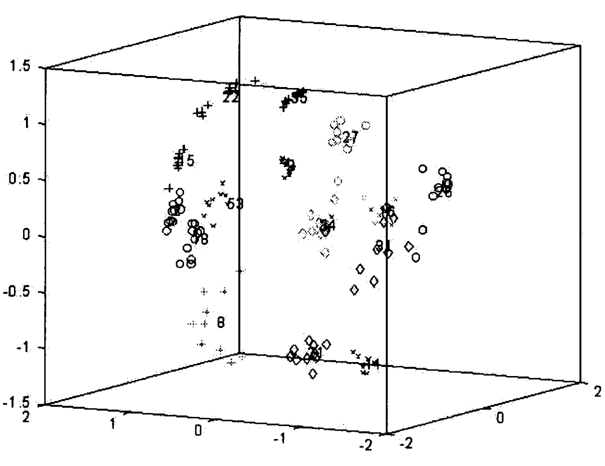

Fig. 9. Plot of first three principal components of selected classes in dataset A transformed by hetero-association.

Table 3. Hetero-associative classifier performance (% correct classification)

		Classifier thresholds (low/high)		
		0.4/0.6	0.3/0.7	0.2/0.8
Dataset A	Heter-asso.	98.49	98.49	98.49
	Hybrid-asso.	95.66	95.28	92.83
Dataset B	Hetero-asso.	60.38	56.23	50.19
	Hybrid-asso.	53.40	46.79	39.25
Dataset C	Hetero-asso.	29.25	25.85	23.02
	Hybrid-asso.	23.58	18.68	15.28

higher the classification scores, the higher the retrieval performance is. It does not mean, however, that if the classification scores are less than 50%, so are the retrieval scores.

As confirmed by our results, using auto-associators with a single hidden layer is not an attractive option. First, an optimally trained auto-associative network can only match the capability of PCA. Secondly, even if the performance of PCA is obtained, PCA is more computationally efficient. There is no real advantage in using these auto-associators.

4. CONCLUSION

We have shown that even with the use of simple texture measures, such as the average grey value of the filtered images, hetero-associative neural networks, and hybrid-associative neural networks, can provide significant reduction in dimensionality and provide a retrieval performance better than PCA. The ability to provide dimensionality reduction without sacrificing performance is most welcome. It would allow the use of high-dimensional feature vectors on current indexing structures such as the R*-tree.

References

1. Manjunath M, Ma W. Texture features for browsing and retrieval of image data. IEEE Trans Pattern Analysis and Machine Intelligence 1996; 18(8): 837–842
2. Liu F, Picard R. Periodicity, directionality, and randomness: Wold features for image modeling and retrieval. IEEE Trans Pattern Analysis and Machine Intelligence 1996; 18(7): 722–733
3. Swets D, Weng J. Using discriminant eigenfeatures for image retrieval. IEEE Trans Pattern Analysis and Machine Intelligence 1996; 18(8): 831–836
4. Jain A, Farrokhnia F. Unsupervised texture segmentation using Gabor filters. Pattern Recognition 1991; 23: 1167–1186
5. Augusteijn M, Clemens L, Shaw K. Performance evaluation of texture measures for ground cover identification in satellite images by means of a neural network Classifier. IEEE Trans Geoscience and Remote Sensing 1995; 33(3): 616–626
6. Ma LM, Jin JS. A generalized content-based image retrieval system. Proc of ACM Symposium on Applied Computing, Atlanta, GA, 1998
7. Jones J, Palmer L. Simple receptive fields in cat striate cortex:

a comparison with Gabor functions in two dimensions of space and two dimensions of spatial frequency. Abstr Soc Neurosci 10:800
8. Daugman J. Uncertainty relation for resolution in space, spatial-frequency, and orientation optimized by two-dimensional visual cortical filters. J Opt Soc America 1990; 2(7): 1160–1169
9. Lin KI, Jagadish HV, Faloutsos C. The TV-tree: an index structure for high-dimensional data. VLDB Journal 1994, 3(4): 517–549
10. Berchtold S, Keim DA, Kriegel HP. The X-tree: an index structure for high-dimensional data. Proc 22th Int Conf on Very Large Data Bases, India, 1996
11. White DA, Jain R. Similarity indexing with the SS-tree. Proc 12th IEEE Int Conf on Data Engineering, New Orleans, Louisiana, 1996
12. Niblack W, Barber R, Equitz W, Flickner M, Glasman E, Petrovic D, Yanker P, Faloutsos C, Taibin G. The QBIC Project: Querying images by content using color, texture and shape. Proc SPIE 1993; 1908: 173–187
13. Joliffe I. Principal Component Analysis. Springer-Verlag, New York, 1986
14. Kramer M. Nonlinear principal component analysis using autoassociative neural networks. AIChe Journal 1991; 37: 233–243
15. Baldi P, Hornik K. Neural networks and principal component analysis: Learning from examples without local minima. Neural Networks 1989; 2: 53–58
16. Kambhatla N, Leen T. Fast Non-linear dimension reduction. In: Cowan J, Tesauro G, Alspector J (eds), Advances In Neural Information Processing Systems 6. Morgan Kaufman, San Mateo, CA, 1994
17. De Valois R, Albrecht D, Thorell L. Spatial frequency selectivity of cells in macaque visual cortex. Vision Research 1982; 22: 545–559
18. Gabor D. Theory of communications. J Institute of Electronics Engineers 1946; 93: 429–457
19. Daugman J. Two-dimensional spectral analysis of cortical receptive field profiles. Vision Research 1980; 20: 847–856
20. Watson A. Detection and recognition of simple spatial forms. In: Brady J, Sleigh A (eds), Biological and Physical Processing of Images. Springer-Verlag, Heidelberg, New York, 1983
21. Bovik A, Clark M, Geisler W. Multichannel texture analysis using localized spatial filters. IEEE Trans Pattern Analysis and Machine Intelligence 1990; 12: 55–73
22. Turner M. Texture discrimination by Gabor functions. Biological Cybernetics 1986; 55: 71–82
23. Malik J, Perona P. Preattentive texture discrimination with early vision mechanisms. J Opt Soc America A 1990; 7: 923–932
24. Kohonen T. Self-organizing and Associative Memory. Springer Series in Information Sciences, vol. 8, 1994
25. Cottrel G, Munro W, Zipser D. Learning internal representations from gray-scale images: an example of extensional programming. Proc Ninth Annual cognitive Science Society Conference 1995, pp 461–473
26. Cottrell G, Metcalfe J. EMPATH: Face, emotion and gender recognition using Holons. In: Lippmann R, Moody, J and Touretzky (eds), Advances in Neural Information Processing Systems 3, San Mateo, CA, 1991, pp 564–571
27. Bourlard H, Kamp Y. Auto-association by multilayer perceptrons and singular value decomposition. Biological Cybernetics 1988; 59: 291–294
28. Hertz J, Krogh A, Palmer R. Introduction to the Theory of Neural Computation. Addison-Wesley, 1991
29. Pandya A, Macy R. Pattern Recognition with Neural Network in C++. CRC Press, 1996

Dr Jesse Jin graduated with a PhD from the University of Otago, New Zealand. He was a founder of computer vision research at the University of Otago. Currently, he is a senior lecturer and the Director of the Visual Information Processing Laboratory in the School of Computer Science and Engineering, University of New South Wales. He has more than 100 publications and has received a large number of research grants. His research interests include multimedia, image processing and medical imaging.

Mr. Jose A. Catalan is a PhD candidate in the School of Computer Science and Engineering, The University of New South Wales, Australia. His research interests include neural networks, multimedia and high dimensional analysis.

Tamás (Tom) D. Gedeon received the BSc (Hons.) and PhD degrees in Computer Science from The University of Western Australia (U.W.A.) in 1981 and 1989, and a Graduate Diploma in Management from the Australian Graduate School of Management, the University of New South Wales (U.N.S.W.) in 1994.

He has held academic positions at U.W.A., Flinders University in Adelaide, and at Brunel University in London. He has held visiting appointments at the University of Kent in Canterbury and the Technical University of Budapest. He is currently Head of the Department of Information Engineering in the School of Computer Science and Engineering at U.N.S.W.

His research is focussed on the development of automated systems for information extraction, and its synthesis into humanly useful information resources in multimodal information processing systems. The primary techniques have been neural and fuzzy techniques, and applied to legal information retrieval and petrophysical inferencing.

Correspondence and offprint requests to: Dr Jesse Jin, School of Computer Science and Engineering, University of New South Wales, Sydney 2052, Australia. Email: jesse@cse.unsw.edu.au

APPENDIX

TEXTURE PATTERNS IN DATABASE A (ONE FROM EACH OF 53 CLASSES)