Content-based Retrieval of Segmented Images

Tat-Seng Chua, Swee-Kiew Lim and Hung-Keng Pung

Department of Information Systems and Computer Science
National University of Singapore
Kent Ridge, Singapore 0511
E-Mail: {chuats, limsweek, punghk} @ iscs.nus.sg

Abstract

Most general content-based image retrieval techniques use colour and texture as main retrieval indices. A recent technique use colour pairs to model distinct object boundaries for retrieval. These techniques have been applied to overall image contents without taking into account the characteristics of individual objects. While the techniques work well for the retrieval of images with similar overall contents (including backgrounds), their accuracies are limited because they are unable to take advantage of individual object's visual characteristics, and to perform object-level retrieval. This paper looks specifically at the use of colour-pair technique for fuzzy object-level image retrieval. Three extensions are applied to the basic colour-pair technique: (a) the development of a similarity-based ranking formula for colour-pairs matching; (b) the use of segmented objects for object-level retrieval; and (c) the inclusion of perceptually similar colours for fuzzy retrieval. A computer-aided segmentation technique is developed to segment the images' contents. Experimental results indicate that the extensions have led to substantial improvements in the retrieval performance. These extensions are sufficiently general and can be applied to other content-based image retrieval techniques.

Keywords: Content-based Retrieval, Partial-Match, Colour-Pairs, Image Segmentation.

Multimedia 94- 10/94 San Francisco, CA, USA

1. Introduction

With the advent of image technology, there has been an increasing interest in the development of video and image management and retrieval systems. Most early image retrieval systems are text-based (Bertino et al 1988, Al-Hawamdeh et al 1991 and Chua et al 1994a), where relevant text annotations are attached to each image and use as basis for retrieval. The users of such system are required to supply a textual search query to locate a list of images in the image database. Main problems with such system are that the text descriptions need to be entered manually, and they are usually incomplete and inconsistent. Also, the users are required to know the range of image descriptions beforehand in order to formulate query for effective retrieval. Attempts to improve the retrieval performance of such systems by using relevance feedback techniques to extend the query and/or image descriptions have not produced consistently satisfactory results (Price et al 1992).

To overcome the problems of text-based image retrieval system, the use of image content features as the basis for retrieval has been considered. The main image content features used are the colours, texture and shapes (Niblack et al 1993). While colours and textures have been used successfully in general and fully automated retrieval systems (Binaghi et al 1992, Chua et al 1994b), the use of shapes has largely been limited to specialised retrieval systems (Stanchev et al 1992) as domain-specific shape database needs to be setup manually beforehand for effective retrieval. In order to capture the local features of objects in images in a general retrieval system, the use of colour-pairs to model distinct object boundaries has been considered (Nagasaka & Tanaka 1992).

The general techniques based on colours and textures, including the colour-pair technique, have been applied to overall image contents without taking into account the characteristics of individual objects within the image. While these techniques work well for the retrieval of images with similar overall contents (including backgrounds), their retrieval accuracies are limited. This is because for most images, the background accounts for a large portion of their contents, and it tends to obscure the existence of objects that are of interests to the users. Thus, techniques must be developed to take advantage of individual object's visual characteristics, and to perform object-level retrieval. Also, in photographic images, the same object may exhibit a slightly different set of colours under different illumination and other environmental conditions, thus a robust colour-based retrieval scheme must be tolerant of such colour variations.

In this paper, we aim to extend the general image retrieval technique based on colours to perform fuzzy object-level retrieval. The basic retrieval technique chosen in our study is the colour-pair matching technique developed by Nagasaka & Tanaka (1992). This technique is chosen because it is sufficiently general and it is able to model the local features of individual objects. Three extensions to the colour-pair technique will be considered: (a) the incorporation of a ranking formula to compute the similarity between the query and the images; (b) the ability to perform colour matching within the objects' boundaries; and (c) the inclusion of perceptually similar colours in similarity computation. To facilitate the extraction of object segments from the images, a computer-aided segmentation approach will be developed. The system will be tested on an image database of over 100 images.

Briefly, the contents of the paper are as follows. Section 2 reviews the related works on image content retrieval and segmentation. Section 3 describes the overall design of our system. Section 4 summarises the colour-pair technique to be used as the main technique for retrieval in this investigation; while section 5 discusses our extensions to this technique to perform fuzzy, object-level retrieval with similarity ranking. Section 6 describes the implementation and evaluation of the system. Lastly, section 7 contains our concluding remarks.

2. Background and Review

2.1. Content-based Retrieval

Recently, multimedia information is increasing both its production and significance. This trend is most typical for image and video information. In most cases, however, large volumes of images and video clips are just kept there for possible future use without any good indices. Effective indexing and retrieval of these information require the identification of specified objects within the images and video sequences. Over the years, several techniques have been developed for object identification and retrieval.

One of the most fundamental technique for detecting an object within an image is by template matching (Pratt 1991). However, a template match is rarely exact because of the image noise, spatial and amplitude quantisation effects, and the a priori uncertainty as to the exact shape and structure of an object to be detected. Also, a large number of templates must be made available in order to cater for all objects to be matched with possible size and orientation variations. For this reason, template matching is usually limited to small local features, which are more invariant to size and shape variations. Several variants of template matching technique have been developed. Kato et al (1991) propose a cognitive approach to visual interaction that permits the users to make a rough sketch of the object that they want to retrieve and match it against those stored in the database. In HyperBook (Tabuchi et al 1991), templates are used for silhouette pattern matching. In Stanchev et al (1992), a fuzzy template matching approach is adopted to overcome the problems of image noise and uncertain image extent. These methods are likely to fail if the search objects change its orientation and/or direction.

Because of the inherent problems of using the topology of objects for retrieval, strategies based on the global features of the entire image, such as colours and textures, have been developed. In a prototype system developed by Niblack et al (1993), colours, texture and/or shapes are used as separate query types to retrieve images. Binaghi et al (1992) use solely the colour contents of an image as the basis for retrieval and employ the fuzzy set theory to express the intrinsic uncertainty of human in evaluating colour similarities. Chua et al (1994b) study the use of combined query involving both colour and texture in a similarity-based retrieval that also takes into account of perceptually similar colours.

The use of global image features in the above manner is unable to take advantage of the local features between and within the objects in the image. In order to overcome this problem while making the technique independent of object topology, Nagasaka & Tanaka (1992) use the colour pairs to model the distinct boundaries and relationships between objects in an image. It selects a set of distinct colour pairs from the query image to match against the images in the database. The method is robust against changes of object's size, location, orientation and direction. However, as the method is based on exact colour matching, the query and the target images must have the same set of colour pairs for the retrieval to be successful.

The general image retrieval techniques based on colour, texture and colour-pairs have been applied to full images. In most cases, the image background tends to dominate the characteristics of the whole image, thus making these techniques ineffective for retrieving images containing specific objects. This problem can be solved by segmenting the objects within the images so that similarity matching can be performed within the object boundaries.

2.2. Image Segmentation

Segmentation attempts to identify homogeneous regions within an image and classify them as segments. Numerous intra-frame image segmentation techniques have been developed over the past two decades to segment an image's content. Such techniques can be classified as: histogram mode seeking, region growing scheme, spatial clustering scheme and split and merge scheme (Haralick & Shapiro 1984). These techniques are concerned only with the intensity cue within an individual image. Although the techniques are effective in identifying homogeneous segments, they are unable to group segments belonging to same objects.

Recent progress in image processing technology has shown that visual motions provide a useful cue for grouping related segments belong to a moving object. Even when an object is camouflaged by its similarity in appearance to other object in the environment, any independent movement of the object immediately gives it away. Based on this premise, a number of inter-frame segmentation techniques have been developed. Diehl (1991) studies the segmentation of a sequence of pictures into a hierarchy of moving objects and sub-objects. His approach, however, has the problem that the resulting segmentation do not necessarily correspond to real objects especially when there are occlusions among the objects. Morikawa & Harashima (1992) proposed a 2-dimensional analysis-by-synthesis approach to successively estimate the segmentation and the relative depth of each object by combining the information from multiple frames.

These inter-frame segmentation techniques, however, are not very reliable. They may fail under unfavourable illumination conditions such as intensity changes due to motion, shadows, illumination, and noise. Furthermore when a group of occluding objects are moving with approximately the same displacements over a limited period of time, these techniques would not be able to segment these objects correctly. Worse still, the techniques are unable to handle stationary objects. In order to overcome these problems, semi-automated approaches are normally adopted in practical systems. In Daneels et al (1993), an interactive outlining approach called "snake" is employed. Alternatively, visual interface may be provided for users to group manually segments created from automated means to form meaningful object segments.

3. System Design

The overall design of the system is shown in Figure 1. The system consists of 4 main modules: (a) visual interface module; (b) image segmentation module; (c) image retrieval module; and (d) query result evaluator. The user interacts with the system through the Visual Interface Module to specify the query and to examine the retrieval results. The Segmentation Module is used to segment the images into the required object segments. The user may choose any of the segmented objects as query. The Retrieval Module then performs the retrieval using an extension of the colour-pair matching technique, and the results is channelled to the Query Result Evaluator to perform the ranking. The ranked list of images is presented to the users as a list of icons. An example interface of a retrieval session can be found in Figure 3.

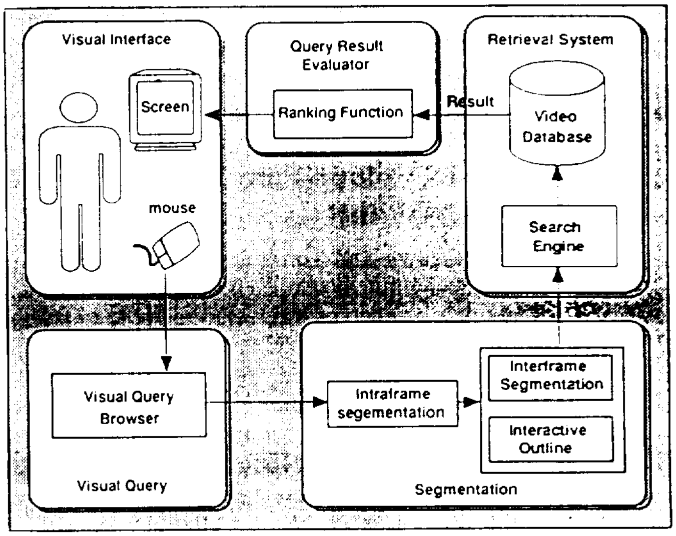

Figure 1: System Overview

3.1. The Image Segmentation Module

In the Image Segmentation Module, we adopt a computer-aided segmentation approach. The images are initially segmented using the automatic segmentation techniques. For static images, we use the spatial clustering infra-frame technique (Haralick & Shapiro 1984) to obtain an initial segmentation of homogeneous regions. If the image is part of a video sequence, then the inter-frame segmentation technique developed in Morikawa & Harashima (1992) is used to improve the segments obtained and to identify moving object segments.

Because of the limitations of automated segmentation techniques as discussed in Section 2.2, a visual interface is provided for the users to interactively improve the object segments created via automated means. The user may create new object segments by entering an approximate object outline via a mouse. A "shrink warp" approach is used to group those segments that are enclosed within the outline. Image segments that are partially wrapped up are removed from the group. The process can be repeated until the user is satisfied with the object segments created. In the case where the required object boundary is merged with the background, the user is allowed to specify an area to be included into the specified object irrespective of the original segmentation. The final set of segments is stored in the image database for subsequent retrieval.

3.2. Colour Space

Cognitive science shows that humans do not perceive colours in the form of Red, Green and Blue components commonly used in colour display devices. Hence, a device-independent colour space is required that represent colour differences along the human perceptual dimensions in a uniform scale. This colour space must have a metric significance, i.e. equal physical distance between two colours in the space corresponds to the equal perceptual difference sensed by the humans. This requirement can be satisfied by specifying colours in the CIE L*u*v colour space (Hall 1989), which has been reported to be closely approximating the colour differences sensed by the human (Carter & Carter 1983). The details of mapping an RGB colour into a CIE L*u*v colour can be found in Hall (1989).

4. The Basic Colour-Pair Image Retrieval Technique

This section reviews and summarises the colour-pair image retrieval technique developed in Nagasaka & Tanaka (1992). The main essence of the technique is to extract a set of distinct colour pairs from an image to model the adjacency relationships between regions and objects in that image. If two images contain similar sets of colour-pairs, then the contents of the images are likely to be similar. The technique is designed to be robust against changes in object size, orientation, and position. However, it is only an approximate method.

The technique can be broken down into three stage. Stage 1 pre-processes the images and sets up the logical database to facilitate retrieval. Each image is divided into a smaller number of cells of x by y pixels each (set x=y=30). For each cell, compute its colour histogram, and select a set of cell colours whose histogram values are larger than a specified threshold value (equal to 10). The threshold is used to eliminate those cell colours that arise from noise which are usually random and small in quantity. Store the resulting cell colour information for each image in the database.

The second stage selects the set of colour-pairs representing the query image. It processes the query image in the similar way as in stage 1 to obtain a set of cells with their own cell colours. For each set of adjacent cells, extract the set of all possible cell colour-pairs, one from each cell. Form a list of all possible cell colour-pairs for the entire image. Scan through the list of colour pairs and eliminate those colour pairs whose number of occurrences are less than a threshold (equal to 10). This is necessary to eliminate very small objects that are not of significance. Select Y most distinct colour pairs from the list (*Y* is set to 16). A colour pair is defined to be more distinct than the other if the Euclidean distance of its two colours in the CIE L*u*v colour space is larger. The *Y* colour pairs is then used to represent the query during retrieval. The process of selecting the query colour-pairs can be summarised in Figure 2.

The last stage searches for images in the database for the occurrences of the *Y* query colour-pairs. For each colour pair, search through the image's cell colour database for its appearances as colours in adjacent cells. The colour pair is marked to be present in the image if it appears in sufficiently large connected regions of more than k cells (k=8). Again, k is used to eliminate image noise. If more than 85% of the colour pairs are marked as present in a particular image, then the image is retrieved. The unranked set of images retrieved are presented to the users for selection.

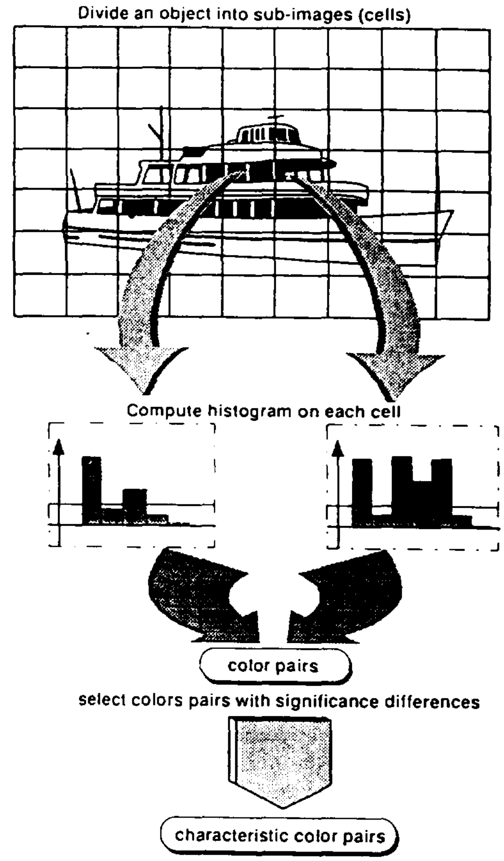

Figure 2: The selection of distinct colour pairs

5. Extensions to the Colour-Pair Retrieval Technique

There are several limitations to the basic colour-pair algorithm. Firstly, the algorithm does not provide any ranking of results. It therefore suffers from the same problems as exact-match retrieval (Willet 1988). Secondly, it is likely that the colours of the background may be used as colour pairs which may be totally irrelevant to the query object. This problem can be solved by performing segmentation on the images and considering only the colours within the objects' segments. Thirdly, the same objects may have different colours under different illumination conditions. This can partially be handled by allowing similarity-based retrieval which relaxed the matching of colours. Finally, different instances of the same objects may have different colours such as a red car vs a blue car. In such a case, we need to allow for colour substitution during colour matching for a particular object. Domain knowledge of the object is needed for this to be accomplished.

In an attempt to overcome these limitations, three extensions are added to the basic colour-pair retrieval algorithm: (a) the development of a similarity-based formula for ranking the results; (b) the use of segmented objects for object-level retrieval; and (c) the inclusion of perceptually similar colours for fuzzy retrieval. These extensions will address the first three limitations discussed above. They will not tackle the last limitation of objects featuring different colours in different images as domain knowledge of possible object colour distributions are needed. This is left as a topic for future investigation.

5.1. Similarity-based Ranking Formula

In the original colour-pair method, an image is considered retrieved if more than 85% of the query colour-pairs is "present" in that image. Given that we know the occurrences of each query colour-pairs in both the query and the images, one suitable measure of similarity between the query, Q, the image, G, is:

$$SIM(Q,G) = \frac{1}{N}\sum_{j\in 1}^{N}\left(1 - \frac{|q_j - g_j|}{MAX\{q_j, g_j\}}\right) \quad (1)$$

where N is the number of query colour pairs used; and, q_j and g_j are the number of occurrences of colour pair j in the query and image respectively. In arriving at this similarity measure, we first compute the absolute difference between q_j and g_j, and normalise the difference by dividing it by the larger of the two values. This results in a normalised difference of between 0 and 1. In order to reflect the *similarity* rather than the *difference*, we use 1 to minus the normalised difference such that the output is directly proportional to the similarity between them. Next, we average the similarity values for all colour-pairs by dividing the summation by N.

The similarity measure (1), however, is dependent on the size of the object as a larger object has a larger number of occurrences of each colour pair. To make it size-independent, we divide q_j (or g_j) by the total occurrences of all colour pairs found in query Q (or image G), giving rise to the normalised colour-pair count qn_j (or gn_j). The resulting similarity measure is:

$$SIM_N(Q,G) = \frac{1}{N}\sum_{j\in 1}^{N}\left(1 - \frac{|qn_j - gn_j|}{MAX\{qn_j, gn_j\}}\right) \quad (2)$$

5.2. Incorporation of Perceptually Similar Colours

The variations in object's colours due to illumination and other environment effects can be modelled accurately using the colour illumination models developed in physics and used successfully in computer graphics (Hall 1989). However, there are problems in employing these models to predict real-life illumination phenomena because many unknowns, such as the 3D configuration of the scene, the colours of the lights and the reflectivity coefficients of the objects etc, are needed. For this reason, we use only a heuristic method in an attempt to capture as much of the colour variations as possible. This heuristic method is similar to the one used in Chua et al (1994b) to model uncertainty in human colour perception.

Our method aims to capture those colour variations that are perceptually similar in the CIE L*u*v sense. For each query colour pair (C_1, C_2), instead of looking for the exact occurrences of (C_1, C_2) as adjacent colour pair in the image, we look for the occurrences of the relaxed colour pair (C_1^*, C_2^*). Here C_i^* (i=1,2) is defined as the set of colours, including C_i, that are within the distance of t from C_i. A suitable value for t found through experimentation is 0.02 (i.e. 2% variations), or such that the number of elements of C_i^* is less than 5. The threshold t can be changed during retrieval to ensure that the set C_i^* does not include any of the colours in other query colour-pairs. The colour pair containing one or two nonexact colours is known as the nonexact colour pair.

In enumerating g_j, which is the number of occurrences of colour pair j in image G, we add in the nonexact colour-pairs with a lower contribution factor of 0.7. In addition, we would like to differentiate whether the final g_j count includes contributions from nonexact colour pairs. If so, a lower weighting factor w_j will be used to sum its contribution in the similarity measure. In our experiment, we set w_j to 0.9 if nonexact colour-pairs are used; otherwise it is set to 1. By normalising g_j as before to obtain gn_j, the similarity measure (2) becomes:

$$SIM_{FN}(Q,G) = \frac{1}{N}\sum_{j\in 1}^{N}\left\{w_j * \left(1 - \frac{|qn_j - gn_j|}{MAX\{qn_j, gn_j\}}\right)\right\} \quad (3)$$

The final $SIM_{FN}(Q,G)$ value gives the normalised similarity between the query Q and image G, ie. it will return a value close to 1 (or 0) if Q and G are very similar (or dissimilar).

5.3. Object-based Image Retrieval

In order to eliminate unwanted details from the background, all colour-pair computation are performed within the object segments generated from the computer-aided segmentation module. Given a query image, the user can select any one of the segments in the image and use it as query object. The colour-pair information is extracted from within the query object and use as the set of query colour-pairs. Equation (3) is used to compute the similarities between the query object and all individual object segments in each image. For each image, a number of similarity values (equal to the number of segmented objects in that image) may be obtained. The retrieval of an image is determined on the basis of similarity between one of its segmented object and the query object.

In extracting the colour-pair information from an segmented object, we adopt the same approach as in the original algorithm by performing cell-subdivision. For the case of segmented object, the subdivision is carried out on its upright rectangular bounding box. Only those cell areas that intersect or fall within the object's boundary are considered. If possible, the same number of cells is used irrespective of the size of the bounding box.

6. Implementation and Testing

The system has been implemented on a UNIX platform using C++ language on the X-windows environment. The X-based user interface is developed using InterViews toolkits (Linton et al 1989). The actual machine used for the implementation is a SiliconGraphics Indigo System.

As the emphasis of this study is to evaluate the effectiveness of the technique on object-level retrieval, over 100 images and short video clips (to be called images throughout for simplicity) obtained mostly from the domain of animals are used for testing. Each of these images contains one or more clearly defined objects which facilitate query specification and performance evaluation. The images are divided into 10 categories as shown in Table 1, which also lists the number of relevant images in each category. The images are captured from the laser disc and stored in Iris RGB image format. The images are segmented into meaningful object segments using the computer-aided segmentation module. Over 90% of images require manual refinement through the visual interface in order to obtain the desired object segmentation.

Expt.	Description	Number of relevant images
1	Trains	8
2	Human head	14
3	Wolf	6
4	Car	10
5	Parrot	12
6	Elephant	11
7	Head	14
8	Sea lion	8
9	Dolphin	12
10	Apes	10

Table 1: Categories of image database

To test the effectiveness of the retrieval technique developed, 10 test queries, one from each image category, are chosen for evaluation purposes. The queries are tested on two separate retrieval techniques - the original colour pair technique as summarised in section 4, and the extended colour-pair technique. In order to facilitate meaningful comparisons, the ranking formula given in Equation (1) is used to rank the results in the original colour-pair technique. In both cases, segmented objects are used as queries rather than the full query image. Thus, the sets of query colour-pairs used in both techniques are identical.

The results of retrieval are presented in a browser as shown in Figure 3. The browser lists the retrieved images in decreasing order of similarity from left to right, and from top to bottom. The browser also shows the similarity value of each retrieved image to give the user an indication of the strength of confidence that the query object is found within the retrieved image.

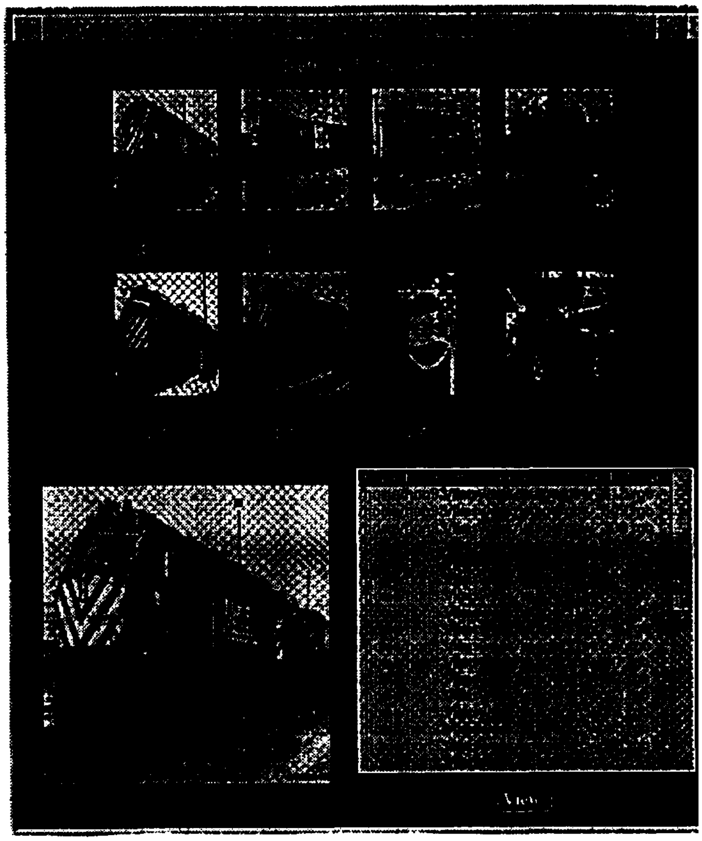

Figure 3: Screen layout of a retrieval session
(The bottom left corner shows the query image)

Table 2 summarises the results of retrieval in terms of recall and precision (see Salton & McGill 1983). For each query and for each method, we compute a pair of values - Recall(0.4) and Precision(avg). The Recall(0.4) values are computed by using a cut-off value of 0.4 for the similarity value. An image is regarded as being "retrieved" if its similarity value is at least 0.4; otherwise it is considered as not being retrieved at all. The Precision(avg) values listed are computed using the definition used in Salton & Buckley (1990). They are found by averaging the precision values at recall levels of 0.25 (low recall level), 0.5 (average recall level) and 0.75 (high recall level).

From Table 2, it can be seen that the original colour-pair retrieval technique has an average Recall(0.4) value of 0.44 and Precision(avg) value of 0.47. The average Recall(0.4) value of the extended technique has been increased to 0.78. This represents an improvement of over 77%. The average Precision(avg) value of the extended technique is 0.67, representing a 42% improvement in performance over the original technique. These results demonstrate that our extensions have led to substantial improvements in the performance of the system.

	Original Colour-Pair Technique		Extended Colour-Pair Technique	
Expt	Recall(0.4)	Precision(avg)	Recall(0.4)	Precision(avg)
1	0.67	0.77	0.83	0.80
2	0.54	0.40	0.77	0.55
3	0.15	0.30	0.78	0.50
4	0.15	0.20	1	0.73
5	0.67	0.67	0.67	0.76
6	0.56	0.64	0.78	0.67
7	0.24	0.50	0.79	0.66
8	0.16	0.46	0.80	0.69
9	0.60	0.44	0.63	0.71
10	0.69	0.31	0.79	0.60
Overall	0.443	0.469	0.784	0.667

Note: see the text for the meanings of Recall (0.4) and Precision (avg)

Table 2: Recall/Precision table of retrieval experiment

Analysis of results indicate that the improvements are achieved in the following two areas. Firstly, the use segmentation helps to remove unwanted colour-pairs formed by the background objects. The exclusion of these colour-pairs reduces the chances of retrieving irrelevant images thus improving the precision of the retrieval. Secondly, the inclusion of perceptually similar colour in the similarity computation increases the chances of the same objects with slightly different colours being retrieved, thus improving the recall. Furthermore, the use of partial-match similarity formula permits the results to be weighted and ranked, thus enhancing the interface and usability of the system. These extensions are sufficiently general and should be applicable to other content-based image retrieval techniques based on colours or other visual attributes.

7. Conclusions

We have described the extension of a content-based image retrieval technique to perform fuzzy object-level retrieval. The technique that we have chosen in our study is the colour-pair retrieval technique developed by Nagasaka & Tanaka (1992). Three extensions are added to the original technique: (a) the incorporation of similarity-based ranking formula to rank the results; (b) the use of segmentation to perform similarity computations only within the segmented objects' boundaries; and (c) the inclusion of perceptually similar colour in similarity computation. The performance of the extended system has been tested against the original system on an image database of over 100 images. The results of the experiment demonstrate that the extended system outperforms the original in both precision and recall by a substantial margin. The average precision obtained for the extended technique is 0.67 which is extremely encouraging. The extensions developed are sufficiently general and should be applicable to other content-based image retrieval techniques.

Several improvements to the system are evident. Firstly, the whole indexing and retrieval process of the extended method is almost fully automated. The only exception is the object segmentation process, which still requires a large amount of manual refinement effort. Thus better automated object segmentation method should be developed in order for this method to be scalable to handle large image databases. Secondly, the fact that colour is used as the only indexing feature limits its application and performance. The immediate task is therefore to extend the system to include other image attribute such as the texture. Thirdly, our approach to handle variations in colours resulting from illumination effects are limited, and could only capture a subset of colour variations due to these effects. A more detailed modelling of these effects using the techniques developed in the area of computer graphics (Hall 1989) should be considered. Fourthly, in order to handle instances of same object featuring different colours, colour substitution approach should be considered. However, domain knowledge on possible range and distribution of object colours should be studied and developed. Finally, if enough visual signatures, such as the colour pairs or texture, of known objects are set up with links to their (textual) semantic descriptions, then we could generate most of the semantic descriptions of an image by comparing its contents with the stored object signatures. This will lead to a semi-automatic image description and indexing system.

Acknowledgments:

The authors would like to thank the National University of Singapore for the provision of a research grant # RP920637 under which this research is carried out.

References

Al-Hawamdeh S, Ooi BC, Price R, Tng TH, Ang YH & Hui L (1991). Nearest Neighbour Searching in Picture Archival System. International Conference on Multimedia Information System '91, Singapore, McGraw-Hill, Jan, 17-34.

Bertino A, Rabitti F & Gibbs S (1988). Query Processing in a Multimedia Document System. ACM Trans. on Office Information System, 6(1), pp 1-41.

Binaghi E, Gagliardi I & Schettini R (1992). Indexing and Fuzzy Logic-Based Retrieval of Colour Images. IFIP Visual Database Systems, II, Elsevier Science Publishers B.V., 79-92.

Carter RC & Carter EC (1983). CIELUV Colour Difference Equations for Self-luminous Displays. Colour Research and Applications, 8, 252-553.

Chua TS, Pung HK, Lu GJ & Jong HS (1994a). A Concept-based Image Retrieval System. Proceedings of 27th Annual Hawaii Int'l Conference on Systems Sciences, 4-7 Jan, 590-598.

Chua TS, Chan SK, Pung HK & Lu GJ (1994b). A Content-based Image Retrieval System. Internal Report, Dept of Information Systems & Computer Science, National University of Singapore. Submitted for publication.

Daneels D, van Campenhout D, Niblack W, Equitz W, Barber R, Bellon E & Fierens F (1993). Interactive Outlining: An Improved Approach using Interactive Contours. Image and Video Storage and Retrieval. SPIE.

Diehl N (1991). Object-oriented Motion Estimation and Segmentation in Image Sequences. Signal Processing: Image Communication, Elsevier Science Publishers BV, 23-56.

Hall R (1989). Illumination and Colour in Computer Generated Imagery. Springer-Verlag, New York.

Haralick RM & Shapiro LG (1984). Survey : Image Segmentation Techniques. Computer Vision, Graphics, and Image Processing, 29, 100-132.

Kato T, Kurita T, Shimogaki H, Mizutori T & Fujimura K (1991). A cognitive Approach to Visual Interaction. International Conference on Multimedia Information System 91, McGraw-Hill, 109-120.

Linton MA, Calder PR & Vlissides JM (1989). Composing User Interfaces with InterViews. IEEE Computer, 22 (2), Feb, 8-22.

Morikawa H, Harashima H (1992). 2D Structural Descriptions for Video Handling. in 'Visual Computing', TL Kunii (Ed), Proceedings of Computer Graphics Int'l '92, Springer-Verlag, Jun, 25-37.

Nagasaka A & Tanaka Y (1992). Automatic Video Indexing and Full-Video Search for Object Appearances. Visual Database Systems, II, IFIP, Elsevier Science Publishers B.V., Oct, 113-127.

Niblack W, Barber R, Equitz W, Flickner M, Glasman E, Petkovic D, Yanker P, Faloutsos C & Taubin G (1993). The QBIC Project: Querying Images by Content Using Colour, Texture, and Shape. Research report, Report # RJ 9203 (81511), IBM Almaden Research Center, San Jose, CA, Feb.

Pratt WK (1991). Digital Image Processing. (2nd Ed). John Wiley.

Price R, Chua TS & Al-Hawamdeh S (1992). Applying Relevance Feedback to a Photo Archival System. Journal of Information Science, 18, 203-215.

Salton G & Buckley C (1990). Improving Retrieval Performance by Relevance Feedback". Journal of the American Society for Information Science, 41(4), 288-297.

Salton G & McGill MJ (1983). Introduction to modern Information Retrieval. McGraw-Hill.

Stanchev PL, Smeulders AWM and Groen FCA (1992). An Approach to Image Indexing of Documents. Visual Database Systems, II, IFIP, Elsevier Science Publishers B.V., 63-77.

Tabuchi M, Yagawa Y, Fujisawa M, Negishi A, Kojima K & Muraoka (1991). Hyperbook: A Multimedia Information System That Permits Incomplete Queries. International Conference on Multimedia Information System 91. McGraw-Hill, 3-16.

Willet P (Ed) (1988). Document Retrieval Systems. Taylor Graham.