UNIVERSITY OF MINNESOTA

This is to certify that I have examined this copy of a master's thesis by

Anand Sivaraman

and have found that it is complete and satisfactory in all respects,
and that any and all revisions required by the final
examining committee have been made.

Dr. Donald B Crouch

Name of Faculty Advisor

Signature of Faculty Advisor

Date

GRADUATE SCHOOL

Improving the Efficiency of a Color-Based Image Retrieval System

A thesis
submitted to the faculty of the graduate school
of the University of Minnesota
by

Anand Sivaraman

in partial fulfillment of the requirements
for the degree of
Master of Science

December 2003

Department of Computer Science
University of Minnesota Duluth
Duluth, Minnesota 55812
U.S.A

ABSTRACT

Content-based image retrieval systems retrieve images from a database that are determined to be similar to a query image based only on features extracted from the images. This thesis focuses on color-based image retrieval. We define methods to improve the efficiency and effectiveness of color-based retrieval. We test our hypotheses using a collection of color images and query images. Color histograms are used to extract and store the color content of the images. Our empirical results are very encouraging. We are able to reduce substantially the total color space without degrading retrieval performance. In addition, we are able to improve performance by conducting object retrieval based solely on color.

ACKNOWLEDGEMENTS

I would like to express my gratitude to my research advisor, Dr. Donald Crouch, for his guidance and constant support throughout my graduate work and my stay in the USA so far. His encouragement and cheerful spirit has helped motivate me to complete my master's degree.

I would also like to thank Dr. Doug Dunham and Dr. Carolyn Crouch for being my examining committee members and reviewing this thesis.

I would like to thank my parents and my sister for their continuing support and encouragement in all my endeavors.

TABLE OF CONTENTS

List of Tables

List of Figures

1. INTRODUCTION

Image Retrieval is in effect an extension of traditional information retrieval to include images [1]. Content-Based Image Retrieval (CBIR) is the process of retrieving images from a database on the basis of features that are extracted automatically from the images themselves [2]. In CBIR, a query is an image or portion of an image; relevant images are retrieved based on the similarity of the features of the query and the features of the individual images in the database. Possible features include color, texture, shape, orientation, or a combination thereof. This thesis focuses on the use of a single image feature for retrieval purposes, namely color.

Color is a visual feature that one immediately perceives when looking at an image [1]. Thus, one can oftentimes readily identify and distinguish images based solely on their color content. Images that appear similar to the human eye are typically very similar in their color content. Hence, color would appear to be a good choice as a computerized feature by which to identify, compare, and retrieve images automatically.

However, although so-called *color-based image retrieval* has the advantage that it is easy to identify and extract the color content of any image, the representation of color is storage intensive and retrieval based on color is computationally intensive. Color representation is a high dimensional feature, with each image representing a combination of approximately 16 million colors. In addition, calculating the distance between images based on color is computationally complex. In this thesis, we explore methods to improve the efficiency of color-based image retrieval without adversely affecting retrieval performance. We experiment with various methods to reduce the storage required of the color content and to improve retrieval efficiency.

The thesis is organized as follows. We introduce concepts underlying image retrieval and content-based image retrieval. We then discuss color-based retrieval per se, emphasizing techniques to extract and manipulate color content and methods for retrieving images based on color content. We describe our empirical approach to the study of color-based retrieval and its experimental results. A summary of our work and suggestions for future research in this area conclude the thesis.

2. IMAGE RETRIEVAL

The advent of the World Wide Web (WWW) and the development of highly economical devices to capture, store and transmit images have led to the creation of huge image libraries. Thus, we are faced with the inevitable problem of having to retrieve useful information from these collections, both effectively and efficiently. This has led to a renewed interest in image retrieval and its practical applications.

2.1 TEXT-BASED RETRIEVAL

Until rather recently, traditional image retrieval employed text as the primary means by which to represent and retrieve images from databases [1]. Text-based image retrieval can be traced back to the 1970's; images were represented by textual descriptions and subsequently retrieved using a text-based database management system [3]. Images were stored along with string attributes – keywords prepared by an annotator that reflected in a relatively broad manner the content of the image. A user would search the database by specifying a Boolean combination of keywords describing the characteristics of images to be retrieved. The system would search the database looking for images whose string attributes matched the image query; images similar in *content* to the query (based solely on the textual descriptions) would be retrieved.

Although text-based image retrieval took advantage of already well-established information retrieval algorithms and mechanisms, its disadvantages as an effective tool to retrieve images became readily apparent. Annotators could not adequately describe images rich in content. Textual descriptions cannot capture the perceptually salient visual features of an image, such as texture, shape, or visual effects [1]. In addition, the annotator's perception of an image may not necessarily reflect that of the user, since description and meaning of a visual image is open to individual interpretation [4].

Researchers concluded that some means of interaction with the visual content of the images themselves was needed to perform effective retrieval. This led to the development of content-based image retrieval, wherein images are retrieved based on their visual features [5].

2.2 CONTENT-BASED RETRIEVAL

Initial research in the retrieval of images based on their inherent features has been reported in [2]. Content-based image retrieval utilizes representations of features that are automatically extracted from the images themselves. Most of the current CBIR systems allow for querying-by-example, a technique wherein an image (or part of an image) is selected by the user as the query. The system extracts the features of the query image, searches through the database for images with similar features, and displays relevant images to the user in order of similarity to the query [6][7][8][9].

Content-based image retrieval systems attempt to exploit the visual information inherent in images, thus providing a more realistic perceptual representation of an image. In this context, content includes among other features, perceptual properties such as color, texture, shape, and spatial relationships [1]. Many CBIR systems have been developed that analyze, compare and retrieve images based on one or more of these features. Some systems have achieved various degrees of success by combining both content-based and text-based retrieval [10][11][12]. In all cases, however, there has been no definitive conclusion as to what features provide the best retrieval.

2.3 COLOR-BASED RETRIEVAL

Since color is a low-level image feature that does not appear to classify images distinctly, few CBIR systems exist that utilize only color as the image retrieval feature [13]. Yet color does have its advantages for image retrieval. It provides multiple measurements at a single pixel of the image, enabling classification to be done without the need for complex spatial decision-making [14] [15]. Color content is also independent of view and resolution and is easy to extract from an image and to manipulate [16].

In this thesis, we revisit the issue of using color as the dominant image feature in a CBIR system. We demonstrate that color can be used effectively for image retrieval and that an efficient CBIR color-based system can be implemented without adversely affecting retrieval performance.

3. COLOR-BASED IMAGE RETRIEVAL

Since color is the dominant visual feature of any image [1] [17], images that appear similar to the human eye oftentimes have similar color content. In this section, we provide the background necessary for understanding how color can be used to support image representation and retrieval.

3.1 COLOR REPRESENTATION

In CBIR systems, colors are usually represented as points in a 3-dimensional color space. Each point is a mixture of the three colors represented by the axes in amounts proportional to the distance of the point along the respective axes. Many different geometric quantitative color models exist. Such models may be classified as follows [1]:

- Colorimetric models represent the physical measurement of spectral reflectance using colorimeters.
- Physiologically-inspired models are based on the fact that the human eye recognizes color based on the three distinct types of cones present in the retina.
- Psychological models are based on how colors appear to a human observer.

Most color models in color-based retrieval are distinguished according to whether they are hardware-oriented or user-oriented. Hardware-oriented models are based on how optical devices such as color monitors and printers reproduce colors. These models do not necessarily correspond to the intuitive perceptions of colors. User-oriented models correspond to the human perception of colors based on wavelength, amount of white light (saturation), and brightness.

To retrieve images based on color, we need models of color representation that effectively capture the chromatic properties of an image and also model the distance between colors to correspond to that perceived by the human eye [1]. In this thesis, we use the RGB color space to recognize and represent colors. RGB is a hardware-based color representation that is physiologically inspired.

3.2 RGB COLOR REPRESENTATION

RGB is the most widely used representation of color in hardware. It is based on the physiology of the human eye, which recognizes colors based on the three color cones (Red, Green and Blue) present in the retina of the eye. RGB colors are obtained as a mixture of these three primary colors. The RGB space is represented as a solid cube with each of the three axes representing one of the colors, red, green, and blue. The colors black and white and the so-called *maximally saturated colors* red, green, blue, yellow, magenta, and cyan, represent the corners of the cube [1]. The gray scale is from black to white – along diagonally opposite corners. An example of the RGB color cube is shown in Figure 1.

Figure 1: RGB color cube

Most computerized graphic systems use a trichromatic RGB encoding. Every pixel is represented as a mixture of the three primaries red, green, and blue. The color to be displayed for each pixel is encoded using three bytes, one for each color with a scale ranging from 0-255. One can visualize the digitized RGB color cube as a cube containing 256 units along each axis, as illustrated in Figure 2. An encoding of (255,0,0) represents red, (0,255,0) represents green, and (0,0,255) represents blue. The maximum combination of all three colors (255,255,255) represents white, and the absence of all colors (0,0,0) represents black. This model can be used to display $(2^8)^3$ distinct colors, that is, approximately 16 million colors

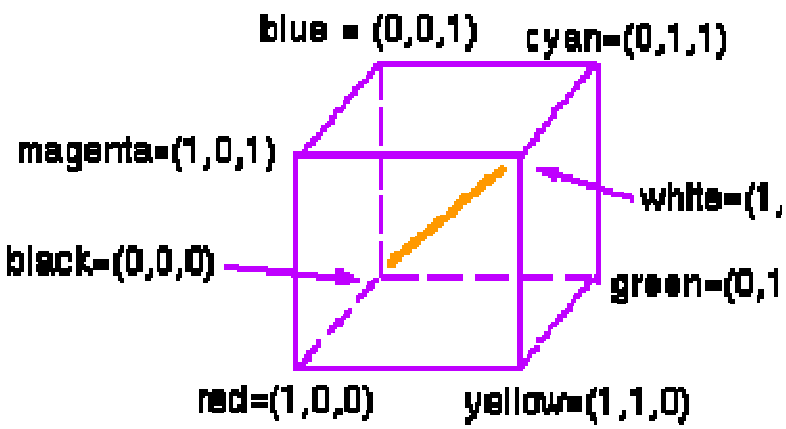

Figure 2: Digitized RGB color cube

.

3.3 COLOR HISTOGRAMS

Histograms may be used to represent the color content of an image. Given a discrete color space, like the RGB color space examined above, the color histogram of an image is obtained by counting the number of pixels in the image that corresponds to each discrete color [16]. Thus, a color histogram is a vector whose length is equivalent to the number of distinct colors in the model and whose value for a given vector component is the number of pixels in the image as a whole containing the color represented by the component. Vector normalization yields the percentage of occurrence of each color in the image.

Most computer-based image retrieval systems that employ color as an image feature utilize histograms for image representation and retrieval. Histograms are computationally simple to produce. Furthermore, techniques for measuring the distance between histograms are well defined [14]. In addition, histograms are fairly insensitive to translation and rotation about an axis perpendicular to the image plane. The angle of view of an image, scaling, and occlusion have minimum impact on the resulting histogram [16].

Histograms consist of long vectors with relatively small values for each vector component. To generate vectors of manageable size, the number of components of a vector is reduced by grouping components into distinct bins. This quantification must be such that the bins are not too large that highly distinct colors fall into the same bin. A high number of bins are usually necessary for a fair description of the images. Typically, quantification with around 100 bins is good enough for retrieval purposes [1]. Figure 3 illustrates a color image and a corresponding 64-bin color histogram of the image.

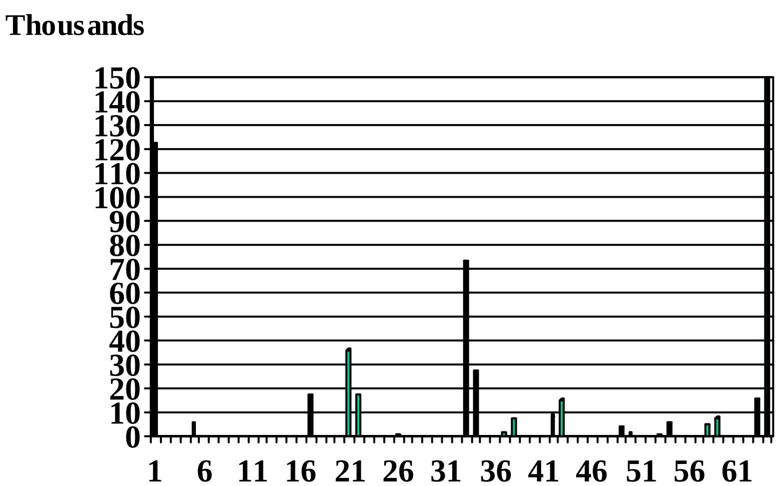

Figure 3: An image and its 64-bin histogram

A potential disadvantage of histograms for image representation is that they do not include spatial information. Images with very different connotations may have similar histograms because their color content is the same. We will illustrate a means to minimize this problem.

In this thesis, we use histograms as the primary means of representing the color properties of images. We attempt to refine the generation of histograms, to minimize the size of histograms without reducing retrieval performance, and to reduce the computational intensity of comparing histograms. Our primary objective is to improve the efficiency of image retrieval without affecting retrieval performance.

4. EXPERIMENTAL RETRIEVAL SYSTEM

We designed and conducted a set of experiments that resulted in the development of procedures for improving color-based image retrieval. In order to conduct our experiments, we needed a color-based retrieval system. Since none were available to us, we designed and built our own modest experimental retrieval system. The system consists of a test set of color images, sample image queries, a known set of relevant images, methods for generating and displaying color histograms, and search mechanisms. To conduct retrieval experiments using this system, a query image is chosen and the system returns all images in ranked order of perceived relevance to the query according to the search algorithms employed.

We chose as our database a set of images pertaining to Christmas, a collection of 72 images in total. All images are a depiction of the Christmas season with snow, decorative items, Santas, etc. Most images are very similar in color content, thus supporting a more rigorous testing of our system's ability to retrieve relevant images based solely on color.

We chose as our query image a Santa, an image that is dominantly red, white, and black and similar in color to most other Christmas holiday images. Fifteen images in the database contain a Santa in whole or as a part of the image. These Santa depictions are dominant in red and white, although other colors are present in significant amounts, depending on the primary subject of the image. 27 additional images in the database are very similar in color content to a Santa image even though they have no relevance to a Santa, and 30 images are totally irrelevant to the query with regards to color or theme.

Our retrieval system uses a color histogram to represent the color content of an image. To determine the distance between two images, we elected to use the *histogram intersection* similarity measure proposed by Swain and Ballard [Swain and Ballard, Color Indexing]. The intersection between the query image histogram $H(I_Q)$ and the histograms of the images in the database $\boldsymbol{H}(I_D)$ is defined as

$$D_H(I_Q, I_D) = \frac{\sum_{j=1}^{n} min(H(I_Q, j), H(I_D, j))}{\sum_{j=1}^{n} (H(I_D, j)} .$$

From the equation above, it is obvious that colors not present in the query image do not contribute to the distance measure obtained by this formula. Thus, the greater the similarity value the greater the similarity of the two images.

Although a digitized image is typically a combination of 16 million colors, comparing images using histograms containing this number of color bins is computationally intensive. Furthermore, since the human eye cannot distinguish 16 million colors [14], most retrieval systems elect to represent images with far fewer colors.

In many early retrieval systems, 256 colors were commonly used in the comparison of images [6]. These early systems verified that retrieval performance was generally unaffected by this scale of reduction in the number of color bins in the histogram. In fact, it is now commonly accepted that a reduction in the number of colors represented to just 64 distinct colors still yields good retrieval results. The basis for this conclusion may appear more obvious if one compares an original color image with an equivalent image transformed into one containing only 64 distinct colors. Figure 4 illustrates such a comparison. The two images appear to be very similar in composition. Most details and color content that would aid in the recognition or distinguishing of the images are maintained.

We exploit this fact in our research. To this effect, we divide the RGB color cube into 64 smaller sub-cubes. This is accomplished by dividing each axis of the RGB cube described earlier into just 4 units, instead of 256 units. Each of the resulting 64 color sub-cubes is represented by the color whose pixel is at the center of the sub-cube. A picture is transformed into a 64-color representation by replacing the color of each pixel in the image with the color of the sub-cube in which the pixel's color resides.

This image transformation allows us to represent images with color histograms containing 64 color bins. In our experimental retrieval system, we represent all images in the database with such histograms. Histogram intersection is a simple and effective measure of the distance between histograms, and it has been shown that its performance in assessing the relevance of images based on color content is good for histograms of length 64 [1] Thus, we will use this metric in our experiments and explore the impact of much shorter vector lengths on retrieval performance.

(a) Original Image

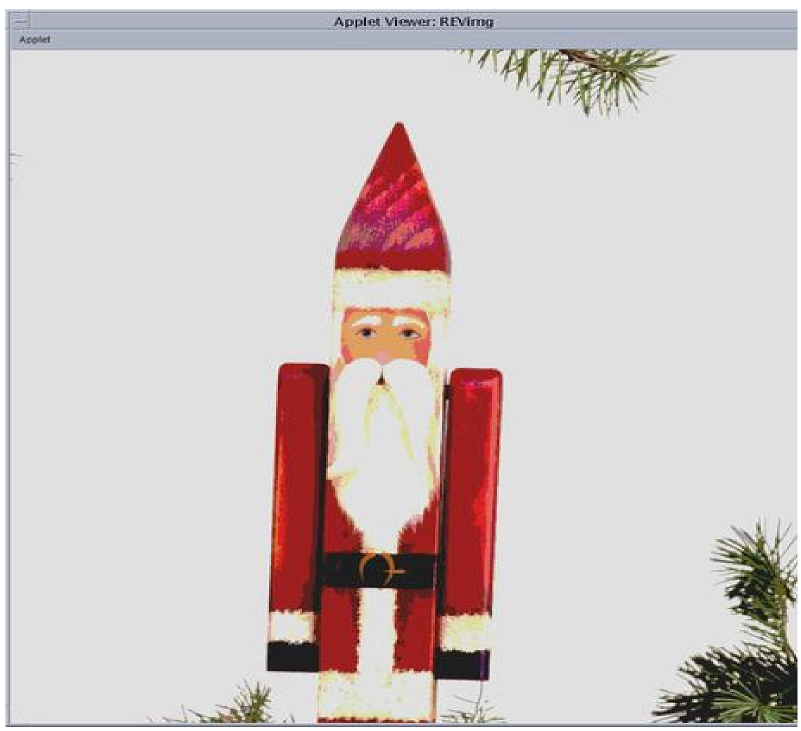

(b) Transformed Image

Figure 4: A color image and its transformation to 64 colors

To conduct a retrieval experiment, our system compares a histogram representation of the Santa image of Figure 4 with the histograms of each image in the database. The resulting images are displayed in the order of perceived relevance with the query, that is, in decreasing order of the histogram intersection similarity values. The ranked list of images represents the system's response to the query.

To generate a base case for retrieval, the case against which our experiments will be compared, we conducted a retrieval experiment using the 64-bin color histograms. The top 5 items returned by this search are displayed in Figure 5. Included in the results are the images themselves, the color histograms, and the histogram intersection similarity values. The query image is included as the first image in the list. As may be noted, performance results are as shown in Table 1.

These are considerably good results considering the similarity in color among all images in the database. Our experiments will assess the degree to which we can further reduce the length of the histograms or modify the search process to improve efficiency of the search process without sacrificing performance.

No. of Images	Relevant	Non-Relevant	Recall	Precision
5	3	2	20.0%	60.0%
10	5	5	33.3%	50.0%
15	6	9	40.0%	40.0%
20	7	13	46.7%	35.0%

Table 1: Results for a sample retrieval with 64-bin

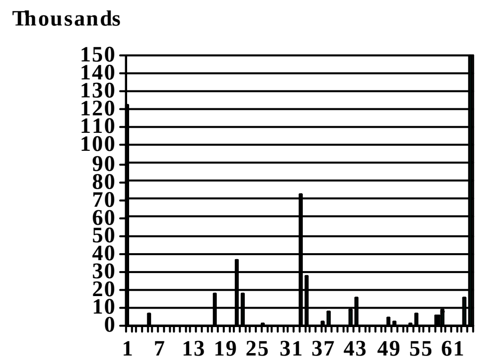

Query

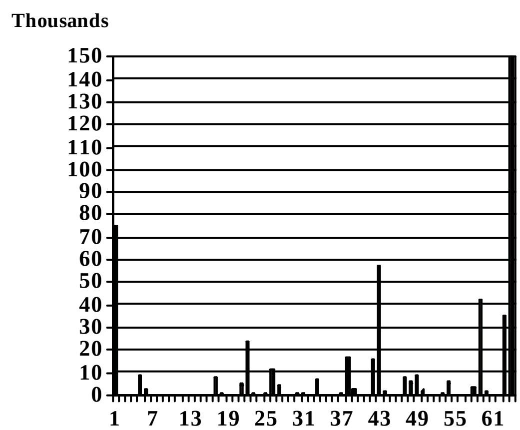

0.85780

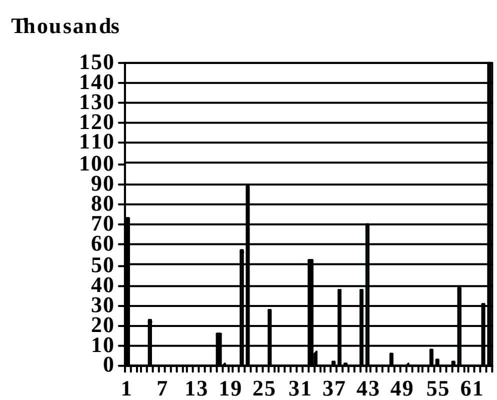

0.72835

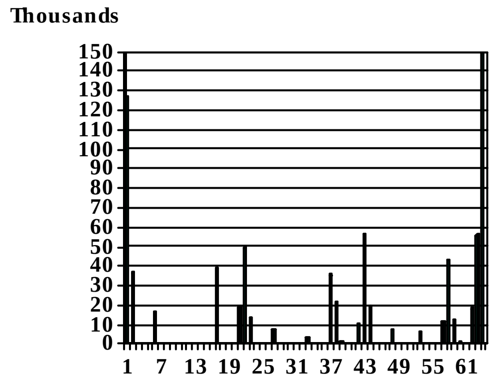

0.67550

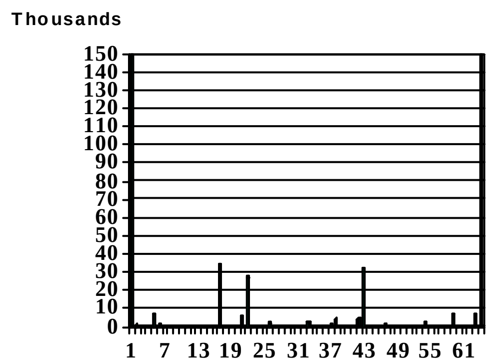

0.67512

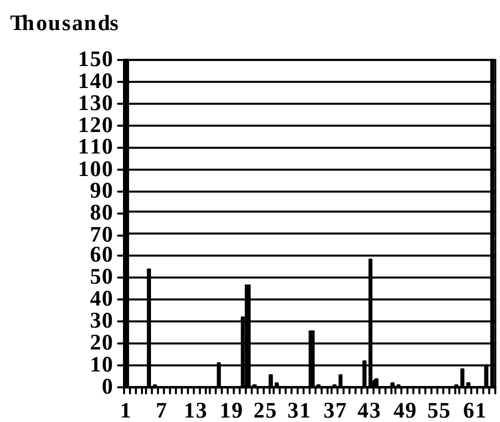

0.65640

Figure 5: Sample retrieval using 64 colors

5. EMPIRICAL STUDY

In this chapter we describe and discuss a set of experiments exploring methods for improving the efficiency of retrieval by reducing the color representation of images and improving the effectiveness of retrieval by performing a search based on image segments.

5.1 RETAINING HIGHLY SIGNIFICANT COLORS

A cursory analysis of the images and their 64-bin histograms in Figure 5 reveals that most colors represented by the histogram do not contribute to the predominant image features. In other words, relatively few distinct colors actually contribute to an image's color content and hence the uniqueness of an image. Most colors in the histogram bins do not play a significant role in the comparison of the image with the query via the intersection similarity measure. Therefore, it would appear that such colors could be removed from the image representations without affecting retrieval effectiveness.

To test this hypothesis, we analyzed the query vector and identified the dominant colors in its histogram. We identified the top n colors in the query, using values of n equal to 20, 15, 10, and 5. Figures 6 through 9 illustrate a comparison of a 64-color Santa image with each n-color image. The resulting images based only on the dominant colors are still very identifiable and lend support to our hypothesis that only a few colors affect image retrieval.

(a) Image with all colors

(b) Image with 20 colors

Figure 6: Original image and its 20-color representation

(a) Image with all colors

(b) Image with 15 colors

Figure 7: Original image and its 15-color representation

(a) Image with all colors

(b) Image with 10 colors

Figure 8: Original image and its 10-color representation

(a) Image with all colors

(c) Image with 5 colors

Figure 9: Original image and its 5-color representation

We conducted four retrievals for the four values of n. In each retrieval experiment, we reduced the 64-bin query histogram and the histograms of each image to a length n by removing the color bins that represented the non-dominant colors. We then conducted retrieval experiment applying the histogram intersection similarity measure to the reduced histograms of length n. The precision/recall results are shown in Table 2.

Retrieval performance did not degrade. The retrieval results obtained from this experiment for all values of n were as good as the retrieval results obtained using all the colors in the histograms of the images. This validated our proposal that only a few dominant colors in the images contribute to its uniqueness and thus aided in comparison and retrieval. In fact, reduction of the histogram by a factor of 68% to 92% did not degrade retrieval performance.

This experiment confirms that very few colors in an image dominate and contribute toward the content of an image in a way that would be responsible for the recognition and uniqueness of the images. We can exploit this fact and reduce the computational cost of comparing the histograms and hence improve the efficiency of a search without adversely affecting retrieval performance.

No. of Images	Relevant	Non-Relevant	Recall	Precision
5	1	4	6.7%	20.0%
10	5	5	33.3%	50.0%
15	5	10	33.3%	33.3%
20	7	13	46.7%	35.0%

Retrieval using 64 bin and top 20 colors

No. of Images	Relevant	Non-Relevant	Recall	Precision
5	1	4	6.7%	20.0%
10	5	5	33.3%	50.0%
15	5	10	33.3%	33.3%
20	7	13	46.7%	35.0%

Retrieval using 64 bin and top 15 colors

No. of Images	Relevant	Non-Relevant	Recall	Precision
5	1	4	6.7%	20.0%
10	4	6	26.7%	40.0%
15	5	10	33.3%	33.3%
20	7	13	46.7%	35.0%

Retrieval using 64 bin and top 10 colors

No. of Images	Relevant	Non-Relevant	Recall	Precision
5	1	4	6.7%	20.0%
10	5	5	33.3%	50.0%
15	7	8	46.7%	46.7%
20	8	12	53.3%	40.0%

Retrieval using 64 bin and top 5 colors

Table 2: Results for retrieval with dominant n colors of the query

5.2 REMOVING NON-SUBSTANTIVE COLORS

An analysis of Figure 5 reveals that some colors occur so frequently among all the images that they probably do not contribute to the uniqueness of an image. Colors, such as the backgrounds of the images, occur in relatively significant amounts in each image and in the database as a whole. Hence, it would appear that these colors are of little use in distinguishing images based on their color content alone. In order to test this hypothesis, we designed an experiment to eliminate such colors during the retrieval process.

To identify non-substantive colors, we generated the *centroid vector* of the histograms of all the images in the database. The centroid vector is a single histogram representing the entire collection; that is, a histogram generated by averaging the histograms of all the images in the collection. The color with the highest weight in the centroid vector, the one color that dominated all the images, was chosen for elimination.

We performed retrieval experiments using the 64-bin color histograms and the reduced histograms. For these retrievals, we eliminated the non-substantive color when calculating the intersection between histograms. Results of these experiments conducted are shown in Table 3.

Retrieval performance improved for the 64-bin color representations and was maintained for the 20-bin color representation. Slight degradation in performance occurred as the number of colors was further reduced. These experiments support our hypothesis that the dominant color did not contribute to retrieval effectiveness for larger histograms and could be eliminated from the decision process.

Further experimentation in which additional non-substantive colors were removed adversely affected retrieval performance. This was due to the fact that the other dominant colors were those that actually contributed to the content of the image. These colors did not appear as frequently in each of the images in the database, but rather dominated individual images.

Experiments	Relevant Images in top 20	Non-Relevant Images in top 20	Recall	Precision
Retrieval after eliminating top color	8	12	53.3%	40.0%
Retrieval based on top 20 colors after eliminating top color	7	13	46.7%	35.0%
Retrieval based on top 10 colors after eliminating top color	5	15	33.3%	25%
Retrieval based on top 5 colors after eliminating top color	6	14	40.0%	30.0%

Table 3: Results after removing the dominant non-substantive color

5.3 OBJECT RETRIEVAL

The first two experiments focused on color-based image retrieval rather than retrieval of a particular object. Even though images relevant to the given query contained Santas, the retrieval process itself involved retrieving those images whose color content as a whole was similar to the color content of the query image. The retrieval process did not consider the color content of a particular object of interest within the query, namely the Santa. It would appear that by focusing on the color content of a given object retrieval performance could be improved.

To test this hypothesis, we designed an experiment in which the query represented a subimage contained within a larger image in the database. For example, we selected a rectangle containing the Santa within the original query image used in the previous experiments. We created a histogram for this subimage and compared it with rectangles of similar size and position within each of the other images in the database as a whole. The results are shown in Table 4. As may be noted, performance increased significantly.

Experiments	Relevant Images in top 20	Non-Relevant images in top 20	Recall	Precision
Retrieval with query sub image compared to same sub-image of images in database	9	11	60.0%	45.0%

Table 4: Results for object retrieval

6. Conclusions and Future Work

The objective of this thesis was to investigate how to improve the efficiency and effectiveness of content-based image retrieval based solely on color representation of the images. Improvement in efficiency was considered successful only if the performance of the system did not degrade as a result of our modifications to the retrieval system. By reducing the number of non-substantive colors and retaining only the primary colors in a query vector, color-based image retrieval performance was shown not to be impaired. Furthermore, by focusing only on sub-images containing the query object, retrieval effectiveness was shown to improve. The experimental database used in this thesis was limited in size. Future work should consider applying the techniques we developed to larger databases with differing characteristics.

7. References

[1] A Del Bimbo, *Visual Information Retrieval*, Morgan Kaufmann Ed., 1999

[2] J Eakins, M E Graham, Content-based Image Retrieval, *A report to the JISC Technology Applications Programme*, 1999

[3] Thomas S. Huang, Yong Rui, and Shinh-Fu Chang, Image retrieval: Past, Present and Future, *International Symposium on Multimedia Information Processing*, 1997.

[4] P. G. B Enser, Query analysis in a visual information retrieval context, *Journal of Document and Text Management*, pg. 25 -52, 1993

[5] C. C. Venters and M. Cooper, Content-based image retrieval, Technical Report, JISC Technology Application Program, 2000

[6] M. Myron Flickner, H. Sawhney, W. Niblack, J. Ashley, Q. Huang, B. Dom, M. Gorkani, J. Hafner, D. Lee, D.Petkovic, D.Steele and P.Yanker, "*Query by image content: The QBIC system*", In IEEE Computer, pp. 23-31, Sept.1995

[7] Patrick M. Kelly, Michael Cannon and Donald R. Hush, Query by Image Example: The CANDID approach. In *Storage and Retrieval for Image and Video Databases III*, volume 2420, pp 238-248, SPIE, 1995

[8] Chad Carson, Serge Belongie, Hayit Greenspan, Jitendra Malik, *Blobworld: Image segmentation using Expectation-Maximization and its application to image querying*, Third International Conference on Visual Information Systems, 1999

[9] M. Das, E. M. Riseman, and B. A. Draper. Focus: Searching for multi-colored objects in a diverse image database, In *IEEE Conference on Computer Vision and*

*Pattern Recognition***,** pp 756-761, Jun 1997

[10] Ma W. Y, B. Manjunath, NeTra: A toolbox for navigating large image databases, In *Proceedings of IEEE Intl. Conference on Image Processing*, pp. 568-571, 1997

[11] V. E. Ogle and M. Stonebraker, *Chabot: retrieval from a relational database of images*, IEEE Computer, vol. 28, no. 9, pp. 40-8, Sept. 1995

[12] Michael Ortega, Yong Rui, Kaushik Chakrabarti, Sharad Mehrotra and Thomas S. Huang, *Supporting Similarity Queries in MARS*. In Proceeding of the ACM International Multimedia Conference, pp. 403-413, 1997

[13] Petteri Kerminen, Moncef Gabbouj, *The Visual Goodness Evaluation of Colors-Based Retrieval Processes*

[14] G. Stockman and L. Shapiro, *Computer Vision*, Prentice Hall, 2001

[15] Vishal Chitkara, Color-Based Image Retrieval Using Compact Binary Signatures, 2001

[16] M. J. Swain and D. H. Ballard, Indexing via Color Histograms, ICCV'90, pp. 390-393, 1990

[17] Tat-Seng Chua, Wai-Chee Low, and Chun-Xin Chu, Relevance feedback techniques for color-based image retrieval, *In Proceedings of Multi-Media Modeling'98*, IEEE Computer Society, pp 24-31, 1998

[18] M. J. Swain and D. H. Ballard, Color indexing, *International Journal on Computer Vision*, vol. 7, no. 1, pp. 11--32, 1991