

Single Color Extraction and Image Query

John R. Smith and Shih-Fu Chang

Columbia University, Center for Telecommunications Research, Image and Advanced
Television Laboratory,
New York, N.Y. 10027

jrsmith@ctr.columbia.edu and sfchang@ctr.columbia.edu

Abstract

In this paper we propose a method for automatic color extraction and indexing to support color queries of image and video databases. This approach identifies the regions within images that contain colors from predetermined color sets. By searching over a large number of color sets, a color index for the database is created in a fashion similar to that for file inversion. This allows very fast indexing of the image collection by color contents of the images. Furthermore, information about the identified regions, such as the color set, size, and location, enables a rich variety of queries that specify both color content and spatial relationships of regions. We present the single color extraction and indexing method and contrast it to other color approaches. We examine single and multiple color extraction and image query on a database of 3000 color images.

1 INTRODUCTION

There is an increasing need for ways to organize and filter the growing collections of image and video data. It is an extremely time consuming task to assign text descriptions to images and the inadequacy of textual annotations for visual data has been recognized. Recently, researchers have begun to investigate "content-based" techniques for indexing images using features such as color, texture and shape [1][2]. A successful content-based image database system requires the following components:

- identification and utilization of intuitive visual features
- effective feature representation and discrimination
- automatic extraction of spatially localized features
- techniques for efficient indexing

In this paper we investigate the use of color for organizing and retrieving images and videos from databases. We maintain that color is an intuitive feature for which it is possible to utilize an effective and compact representation. Our approach automatically extracts the color content of isolated regions within images and builds efficient indexes to retrieve the regions over a large collection of images. The spatial localization of the color regions also allows for queries to include spatial positions and relationships between color regions. This gives a great power of expression for database queries that include both specification of color sets and relative and absolute spatial locations.

Queries supported by the single color technique include the following examples: *give me all images containing...*

- a.) a large dark green area near top of image, i.e., trees
- b.) a yellowish-orange spot surrounded by blue, i.e., a sunset
- c.) a region composed of red, white and blue, i.e., a flag
- d.) an area with red and white in equal amounts, i.e., a checkered table cloth.

We will explain how these queries can be answered using color sets with one or more colors, and/or by specifying spatial relationships and composition of regions.

2 COLOR FEATURES

Color may be one of the most straight-forward features utilized by humans for visual recognition and discrimination. However, people show the natural ability of using different levels of color specificity in different contexts. For example, people would typically describe an apple as being 'red', probably implying some type of reddish hue. But in the context of describing the color of a car a person may choose to be more specific instead using the terms 'dark red' or 'maroon'. Color extraction by computer is performed without benefit of a context. Lack of knowledge also makes it difficult to cull the color information from the color distortion. The appearance of the color of real world objects is generally altered by surface texture, lighting and shading effects, and viewing conditions. Image database systems that use color retrieval must grapple with these problems of automated color image analysis.

2.1 Color histogram

One common method for characterizing image content is to use color histograms. The color histogram for an image is constructed by counting the number of pixels of each color. Retrieval from image databases using color histograms has been investigated in [1][2][3]. In these studies the formulations of the retrieval algorithms follow a similar progression: (1) selection of a color space, (2) quantization of the color space, (3) computation of histograms, (4) derivation of the histogram distance function, (5) identification of indexing shortcuts. Each of these steps may be crucial towards developing a successful algorithm. But there has been no consensus about what are the best choices for these parameters. In [1] we evaluated the retrieval performance when several of these parameters were varied on a database of 500 color images.




FIGURE 1. *Car color Image*: (1) conversion to HVS color space, (2) quantization of HVS space, (3) color median filtering, (4) conversion to indexed RGB space, (5) the processed color image has dominant color regions emphasized.

There are several difficulties with histogram based retrieval. The first of these is the high dimensionality of the color histograms. Even with drastic quantization of the color space, image histogram feature spaces can occupy over 100 dimensions in real valued space. This high dimensionality ensures that methods of feature reduction, pre-filtering and hierarchical indexing must be implemented. The large dimensionality also increases the complexity and computation of the distance function. It particularly complicates ‘cross’ distance functions that include the perceptual distance between histogram bins. Another challenge with the use of color histograms is to enable the extraction of localized features.

2.2 Color image segmentation

The extraction of spatially localized features is an extremely important aspect of image indexing. The isolated regions of interest within images should be identified and extracted independently from other regions in the image. For example, an image should be retrieved even when the user can describe only part of the image. If each image is represented by a single color histogram, this aspect of retrieval performance declines significantly. This is because extraneous information such as background colors may dominate the histogram.

Several attempts have been made to improve performance. In [1] images were segmented into fixed blocks and each block was indexed separately. In this way some blocks may still retain a reasonable characterization of objects of interest. On the other hand, the QBIC system [2] requires manual segmentation of images. In QBIC the color histograms are computed as attributes of the regions that have been cut out by hand. This reduces the potential contribution of background and other irrelevant colors but requires extensive human involvement in creation of the indexed data. Automated segmentation of images using color histograms may eventually provide useful results but has not yet been integrated into large image retrieval systems.

3 SINGLE COLOR EXTRACTION

The goal of the single color extraction method is to reduce the dimensionality of the color feature space while gaining the ability to localize color information spatially within images. Illustrated in Figure 1, we accomplish this through the following means: reduction of the full gamut of colors to a set of manageable size (~100 carefully selected colors). We avoid mapping unacceptably dissimilar colors into the same bins. We also allow higher tolerance in color lightness

and color saturation while reserving the most fine quantization for hue. We utilize a ‘colorizing’ algorithm to paint the color images using the reduced palette and a broad brush. This ensures that the most dominant colors and regions are emphasized. The processed images retain a visibly acceptable and compact representation of the color content. After this processing, we search over the set of colors remaining in the image, and map the regions that sufficiently contain the selected colors into a database index. The next section discusses the process in more detail.

3.1 Color space

The RGB color format is the most common color format for digital images. The primary reason for this is because it retains compatibility with computer displays. However, the RGB space has the major drawback in that it is not perceptually uniform. Because of this, uniform quantization of RGB space gives perceptually redundant bins and perceptual holes in the color space. Furthermore, ordinary distance functions defined in RGB space will be unsatisfactory because perceptual distance is a function of position in RGB space.

Other color spaces, such as CIE-LAB, CIE-LUV and Munsell offer improved perceptual uniformity [4]. In general they represent with equal emphasis the three color variants that characterize color: hue, lightness, and saturation. This separation is attractive because color image processing performed independently on the color channels does not introduce false colors [5]. Furthermore, it is easier to compensate for many artifacts and color distortions. For example, lighting and shading artifacts will typically be isolated to the lightness channel. In general, these color spaces are often inconvenient due to the basic non-linearity in forward and reverse transformations with RGB space. For color extraction we utilize the more tractable HSV color space because it has the above mentioned characteristics and the transformation from RGB space is non-linear but easily invertible.

The next issue after color space selection is quantization. The HSV color space can be visualized as a cone. The long axis represents value: blackness to whiteness. Distance from the axis represents saturation: amount of color present. The

TABLE 1. Statistics on 3000 indexed color images

	quantized HSV colors/ image	colors > threshold/ processed image	single color regions/ image	double color regions/ image	triple color regions/ image
$\bar{\mu}$	37.5	13.7	16.5	79.7	73.2
σ	18.1	5.3	5.5	38.3	40.5

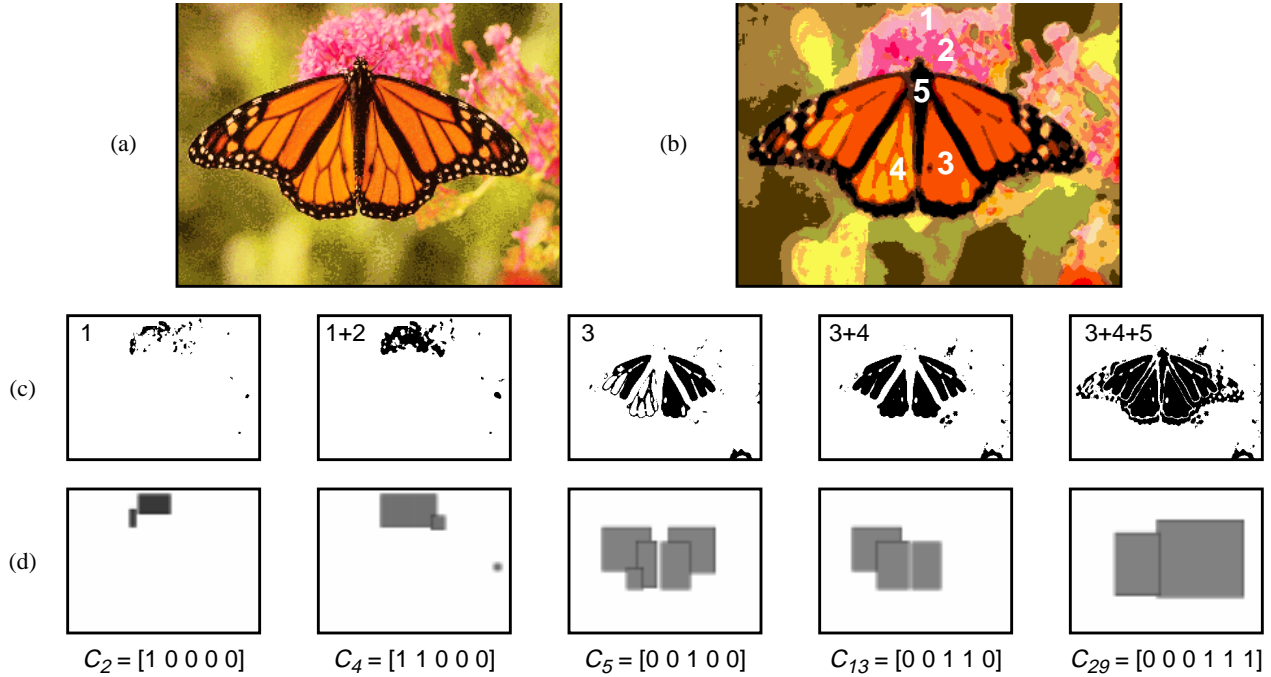


FIGURE 2. (a) *Butterfly* color image, (b) processed color image with 30 colors, (c) pixels from image (b) belonging to color set C_p , (d) minimum bounding rectangles (MBRs) for extracted regions used to index the image collection.

angle around the axis is the hue: tint or tone. Quantization of hue requires the most attention. The hue circle consists of the primaries red, green and blue separated by 120 degrees. A circular quantization at 20 degree steps sufficiently separates the hues such that the three primaries and yellow, magenta and cyan are represented each with three sub-divisions. Saturation and value are each quantized to three levels yielding greater perceptual tolerance along these dimensions. The quantized HSV space appears in Figure 3.



FIGURE 3. Quantized HSV color space, 18 hues, 3 saturations and 3 values + 4 grays = 166 colors.

3.2 Color processing

To identify color regions, the images are transformed to the quantized HSV space with 166 color bins and subsampled to approximately 196x196 such that correct aspect ratio is preserved. This generally reduces the image content to less than 50 colors. Even after the transformation it is still premature to isolate color regions because small details and spot noises interfere. We reduce most of this insignificant detail by using a colorizing algorithm. This processing is accomplished using a 5x5 median filter on each of the HSV channels. This non-linear filtering in HSV space does not introduce false hues. The color image is then converted back

to an indexed RGB space. Table 1 reports the statistics of color processing of 3000 color images.

3.3 Color region labelling

The next step involves the extraction of the color regions from the images. This is done by systematically selecting from the colors present in the image one at a time, and in multiples, each time generating a bi-level image. The levels correspond to the selected and un-selected pixels for the specified color set. Refer to Figure 2 and Table 2 for an illustration of region extraction and representation of the *Butterfly* color image. Next follows a sequential labelling algorithm that identifies the isolated regions within the image. The characteristics of each color region are evaluated in regards to several thresholds to determine whether the region will be added to the database. The first threshold is one for region size. In our system the region must contain more than 64 pels to be significant. This value still allows for sufficiently small regions to be indexed.

TABLE 2. Color Region Table -- *Butterfly* color image

<u>COLID</u>	<u>IMID</u>	<u>REGID</u>	<u>X</u>	<u>Y</u>	<u>SZX</u>	<u>SZY</u>	<u>AREA</u>
[0 0 1 0 0]	1	0	51	58	75	64	1753
[1 0 1 0 0]	1	1	153	58	72	67	2205
[1 0 1 1 0]	1	2	51	58	75	64	2307
[1 0 1 1 0]	1	3	142	80	44	69	1761
[1 0 1 1 1]	1	4	133	47	133	114	6943
[1 0 1 1 1]	1	5	66	137	68	93	2943

If more than one color is represented in the color set we utilize two additional thresholds. The first threshold is the absolute contribution of each color. If a color does not contribute at least 64 pels to the region, the region is not added.

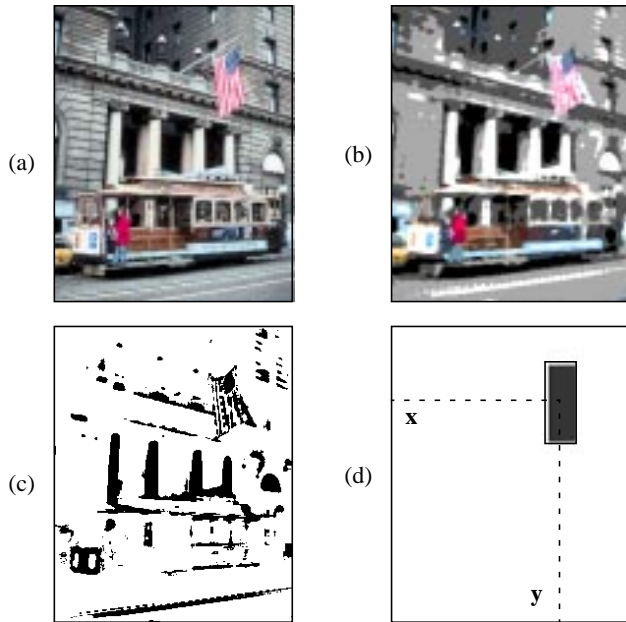


FIGURE 4. (a) *San Francisco* color image, (b) processed image with 73 colors, (c) pixels belonging to color set = {red, white, blue}, (d) extracted color region as present in index.

Furthermore, the relative contribution is also measured. All colors must contribute to at least 20% of the region area. Notice that this produces a firm limit of 5 colors per color region although, we use only up to 3 colors at a time. If a color region does not pass one of these thresholds then it will not be indexed by that color set. If a region is rejected because one of the colors from the color set is not sufficiently represented, the region still has a chance to be extracted using a reduced color set leaving out the under-represented color. Enforcing the color set thresholds prevents the unnecessary and redundant proliferation of indexed multiple-color regions.

3.4 Color image mining

Even with the reasonably small color gamut it is necessary to search systematically for multiple color regions. Otherwise, it will require 2^m passes over the image to test all combinations of m colors. We utilize a heuristic similar to that used for database mining [6]. The algorithm makes multiple passes over each image, expanding only the color sets that meet minimum support constraints. A color set C_i of binary colors is explored for an image only if for all colors k in C_i , where $C_i[k]=1$, there are at least τ_0 pixels in the image of color k such that τ_1 pixels of color k have not yet been allocated to a color region. We use τ_0 and $\tau_1 = 64$. If τ_0 is not met then C_i will have colors that cannot be represented sufficiently by any color regions. Exploring this color set and all supersets of it would be futile. If τ_0 is met while τ_1 is not, then a color region containing all of the colors in C_i can alternatively be reconstructed using subsets of C_i and spatial composition. Therefore, exploration of C_i and its supersets generate redundant information.

Figure 4 illustrates an example of the extraction of an American flag in the *San Francisco* color image. The region was extracted in whole while searching over color sets in the extraction process. The region and color set met the constraints to allow the region to be extracted. The users request for a {red, white, blue} region is answered with the minimum bounding rectangle in Figure 4(d) that represents the region.

3.5 Color specification and spatial positions

The color characteristics specified by the user are represented using the m -dimensional binary color vector. The values may be obtained by picking colors from a color chooser, by navigating visually through 3-D color space, or by textual specification. The binary color vector will be quickly matched to region data because we allow only up to three colors per color set for each indexed region. This sparse binary vector representation of the color sets makes it far easier to index the color distributions than that needed for the color histogram techniques.

After the color characteristics of the regions have been determined, the spatial positions and relationships between regions can be specified by the user. The spatial characteristics of the color region query can be handled using one of several techniques that have been devised for representing and querying spatial information [7][8].

4 CONCLUSIONS AND FUTURE WORK

Single color query is an extremely useful content-based query tool for users of image and video databases. We proposed a method for automatically extracting the single and multiple color regions within images. The color extraction approach allows interesting queries to be formulated based on size, shape and spatial relations of the color regions. The single color approach allows the user to specify the color content and spatial positions of region within images. Single color extraction and indexing is supported in the Content-Based Visual Query System being developed at Columbia University for a variety of image and video applications.

REFERENCES

- [1.] John R. Smith and Shih-Fu Chang, "Tools and Techniques for Color Image Retrieval," submitted to *ACM Multimedia '95*.
- [2.] C. Faloutsos, et. al., "Efficient and Effective Querying by Image Content," *IBM RJ 9453* (83074), August 3, 1993.
- [3.] M. Swain and D. Ballard, "Color Indexing," *International Journal of Computer Vision*, 7:1, 1991, p. 11 -- 32.
- [4.] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods*, John Wiley & Sons, 1982.
- [5.] John C. Russ, *The Image Processing Handbook*, IEEE Press, 1995.
- [6.] R. Agrawal, et. al., "Mining Association Rules between Sets of Items in Large Database," *ACM SIGMOD-93*, Washington, DC, May, 1993.
- [7.] T. Gevers and A.W.M. Smeulders, "An Approach to Image Retrieval for Image Databases," *Database and Expert System Applications (DEXA-93)*, 1993.
- [8.] S. K. Chang, et. al., "An Intelligent Image Database System," *I.E.E.E. Transactions on Software Engineering*, Vol. 14, No. 5, May 1988.