Pergamon

Pattern Recognition, Vol. 31, No. 9, pp. 1241–1253, 1998

0031-3203/98 $19.00 + 0.00

PII: S0031-3203(97)00164-7

VISUAL QUERYING BY COLOR PERCEPTIVE REGIONS

A. DEL BIMBO,* M. MUGNAINI, P. PALA and F. TURCO
Dipartimento di Sistemi e Informatica, Università di Firenze, 50139 Firenze, Italy

(*Received* 31 *March* 1997; *in revised form* 17 *November* 1997)

Abstract—A major research subject in image databases is to support efficient and effective access to images based on their visual content. In color image databases, this requires to support image retrieval by both global and local chromatic features. Image retrieval by color regions is complicated by the fact that regions which are requested in the query should correspond to relevant colored regions in the retrieved images. In the following we present the PICASSO system, which supports image indexing and retrieval based on colors. The system exploits a pyramidal representation of images. Multiple descriptions of image properties are created at different levels of resolution thus allowing effective retrieval through specific queries as well as imprecise ones. © 1998 Pattern Recognition Society. Published by Elsevier Science Ltd. All rights reserved

Image databases   Color image retrieval   Multi-resolution analysis

1. INTRODUCTION

The emerging of Multimedia as the new technology and the possibility of sharing and distributing image data through large-bandwidth computer networks have emphasized the importance of tools for retrieving visual information. Image databases are now currently employed in an eclectic range of different areas such as entertainment, art history, advertising, medicine and industry among others. In all these contexts, the main problem is related to the need of an efficient access to image visual content. Generally, the visual content of an image should be intended as what humans remember after looking at it for a couple of minutes. It may include shapes of relevant objects, color distribution or color patches, textured surfaces, or the arrangement of visual elements on the image plane.

Conventional attempts to cast visual features into textual keywords[1] have been far recognized to be inadequate in indexing pictures. In fact, the minor expressiveness of text with respect to visual features does not fully exploit capabilities of human memory, and items retrieved through a textual query could not be relevant at all for user's expectation. This is the reason why retrieval based on visual content has been identified as the means to overcome this modal clash.

The relevance of visual elements depends on the user's subjectivity (which is not known in advance, when the database is created) and on the context of application (which is known, instead). To cope with user's subjectivity, specific criteria need to be developed, which are able to manage imprecision and lack of knowledge about image content. The fact that the context of application is known in advance, instead, can drive the choice of pattern recognition algorithms and functions that must be included in the system. For example, spatial arrangement is relevant to a doctor examining a heart section; color arrangement is important when looking at paintings, but the presence of particular color tones or the disposition of color patches on the canvas will be better remembered by an expert than by an occasional user.

Generally speaking, an efficient system for retrieval by visual content should be able to

(1) provide a query paradigm that allows users to naturally specify both *selective* and *imprecise* queries;
(2) define retrieval facilities that are meaningful for the context of application;
(3) define similarity metrics which are satisfactory for user perception.

Querying by visual example is the interaction paradigm that exploits human natural capabilities in picture analysis and interpretation. It requires that visual features are extracted from images and used as indexes in the retrieval phase. Querying by visual example allows imprecision and incompleteness of expression, by letting users draw a sketch of their memory of the image (for example colored shapes arranged in a specified pattern). Interactivity with visual content is essential to visual information retrieval. In retrieving visual information according to its perceptual elements, one cannot expect that results obtained in response to a query are fully satisfactory. The traditional technique is to improve the quality of retrieval by keeping as low as possible the number of misses—at the expense of a larger number of false retrievals—allowing a form of interaction called *relevance feedback*. In this way, the original query of

*Author to whom correspondence should be addressed.

the user can be refined or changed on the basis of the retrieved pictures.

In recent years, prototype systems have been proposed which implement retrieval by content from image databases by addressing different facets of the informative content of pictorial data, such as object texture organization,[2] shape similarity with visual sketches,[3–5] semantic relationships between imaged objects.[6–8] With the increasing availability of devices supporting acquisition and visualization of color images, a growing attention is also being focused on chromatic features.

In the following, we present the PICASSO system, developed at the Visual Information Processing Lab of the University of Florence, Italy. The system is a framework including facilities for image indexing and retrieval based on shapes, colors, and spatial relationships. Only color retrieval facilities are discussed in this paper. For retrieval based on shape and spatial relationship the reader can refer to.[5,9] Concerning colors, the system supports both retrieval by global color similarity and retrieval by similarity of color regions. Shape, size, and position of color regions are considered as optional features that the user can select in the query.

2. COLOR-BASED IMAGE RETRIEVAL EXPERIENCES

Previous experiences in color-based retrieval essentially address two kinds of problems. The first considers the problem of finding database images whose color distribution is globally similar to that in a query image. In this problem, the user's interest lies on the chromatic content of the whole image. What is represented in a picture has no particular relevance. In painting databases, for example, this kind of query could help in finding paintings of the same author, or of a certain author's period, or perceptually similar paintings. The second problem is concerned with finding a certain object in a complex scene, using as clue its chromatic properties. The rest of the scene is not relevant for the user's purposes. Only the presence and location of the object are interesting.

Image retrieval based on global color distribution has been formerly proposed by Swain and Ballard.[10] Image color distribution is represented through color histograms, which are obtained by sampling the color space, and counting the number of pixels falling in each discrete color. Color histograms have the property of being invariant to translation and rotation. They change smoothly when objects are partially occluded, or when there are changes in scale and angle of view. Retrieval is performed by evaluating the intersection between the global color histogram of the user-provided example and stored images.

The QBIC database system[11] also evaluates similarity of global color properties using histograms. A weighted distance measure is used to evaluate the similarity of color histograms. The RGB color space is reduced to the best 256 colors, extracted by considering the mathematical transform to Munsell coordinates of each of 4096 initial color levels, and hence performing a clustering operation. The *distance measure* considers a weighted cross correlation between histogram bins. Weights represent the extent to which two histogram bins are perceptually similar to each other. A. K. Jain and A. Vailay in[12] used both color and shape features, analyzed over the entire image. In their system, the query is formulated through an example image and retrieval is accomplished by a similarity measure based on the image color histogram and edges.

A few authors have allowed retrieval of images based on similarity of color patches or objects. This is due to the inherent major complexity of this kind of retrieval. Searching for localized color regions, which may correspond to relevant objects, requires to use effective algorithms to locate uniform color regions. Moreover, it needs more complex data structures to represent color image properties, and more complex algorithms to evaluate perceptual similarity.

In[12] images are partitioned into blocks of equal size, each associated with its own local histogram. Similarity matching considers adjacency conditions among blocks with similar histograms. However, blocks are created according to a static partitioning of the image. This is generally inadequate to reflect the original arrangement of colors in a complex image.

An iterative technique is proposed in[14] to identify image regions which can potentially represent a given object, based on color features. Color Histogram Intersection is used to evaluate the match between image regions and query objects. Regions of potential interest are extracted considering a square window, and shifting it on the image by a fixed number of pixels in each direction.

A different and more reliable solution requires that the whole image is segmented into homogeneous color regions. Chromatic and geometric features of these regions are matched against corresponding regions in the query model. Segmentation is typically the hardest problem to be solved. Matching of segmented images with query color patches or segmented objects is also difficult. In a query by example, the color patches sketched by the user correspond to his approximated view of the image searched, rather than to the true image patches. Shapes of color regions of database images, as resulted from automatic segmentation, could not fit shapes of regions specified in the query.

The PICASSO system exploits a hierarchical multiresolution segmentation and a graph matching technique in order to support effective retrieval based on color regions. The system has been designed mainly for content based retrieval of paintings and art images.

3. HIERARCHICAL COLOR IMAGE SEGMENTATION

Image segmentation is the partition of an image into a set of non-overlapped regions—homogeneous with respect to some criteria—whose union covers the entire image. The number r of regions which are produced in the segmentation process determines the level of precision of the image partitionment. Traditionally, this level is set so as to find a trade-off between computational complexity and adherence of representation. The value of r can be determined adaptively on the basis of a statistical analysis of color distributions in the sampled color space.[15] However, in the absence of specific assumptions about images being stored, one such approach may lead to segmentations which do not meet user's expectations about perceptual groups. Moreover, at storage time it is impossible to forecast the level of precision which will be requested in the detection of image properties expressed by users in specific queries. If the segmentation process fails to detect a region in an image or performs a segmentation into regions which is different from that expressed by the user in the query, that image will not be retrieved in the searching process. Both these considerations evidence that, at storage time, the optimal level of precision cannot be defined.

In the PICASSO system, this hurdle is circumvented by creating multiple descriptions of each data, each one covering a different level of precision. Each database image is segmented into uniform color regions at different degrees of resolution, so as to obtain a pyramidal *multi-resolution representation*.[15]

Chromatic qualities are represented in the *CIE* $L^* u^* v^*$ space, through a set of 128 reference colors, obtained through competitive learning as discussed in the following section.

Regions are aggregated at a higher level if they are adjacent and their average colors differ less than a predefined threshold. Regions R_i are aggregated in a higher-level region by iteratively updating region clusters at each level separately—starting from the lowest level—so as to minimize the following energy associated with the image:

$$F = \sum_{R_i} \left\{ \alpha \frac{1}{A(R_i)} + \beta D(R_i) + \gamma \sum_{R_j} \frac{1}{D(R_i \cup R_j)} \right\} \quad (1)$$

where $A(R_i)$ is the area of region R_i, $D(R_i)$ is a measure of color uniformity of region R_i, $D(R_i \cup R_j)$ is a measure of color uniformity of R_i and its adjacent regions R_j and α, β, γ are control parameters.

The first term of the image energy (1) takes into account region area: the segmented image should not be composed of too many regions. According to this, small regions are penalized. The second term in equation (1) accounts for uniformity of color within the region: regions composed of pixels with distant colors are penalized. The third term in equation (1) is concerned with contrast of color between adjacent regions: since regions should be uniform in color, regions composed of pixels with distant colors are penalized.

With high α values, large regions are favored with respect to small ones. High values of β result into regions with very homogeneous colors. High values of γ avoid the presence of adjacent regions with similar colors.

The level of resolution of the segmentation is determined by the combination of values of α, β and γ. Low values of α and γ, and high values of β result into a fine resolution. The higher α and γ and the lower β are, the coarser the resolution is. Image segmentation starts at the finest level of resolution, where each region corresponds to a pixel. Minimization of the image energy (1) is achieved through an *Heuristic search* approach: every pair of adjacent regions is checked in order to verify if their merge decreases the image energy. The two regions that provide the maximum decrease of image energy are merged. This process is continued until a minimum of image energy is reached. At this step, parameters α, β and γ are changed, so that a new minimum is reached with a lower number of regions. Parameters α, β and γ are updated according to

$$\alpha(n+1) = c_1 * \alpha(n)$$

$$\beta(n+1) = \beta(n)/c_2$$

$$\gamma(n+1) = c_3 * \gamma(n) \quad c_i > 1.$$

This procedure is recursively applied until the coarsest resolution is reached. At this resolution the entire image is represented by a single region. In Fig. 1, the image energy and corresponding control parameters are shown in a segmentation process. Discontinuities of the image energy correspond to changes of the control parameters.

At the end of the segmentation process, for a color image I, N segmented images I_n, $n = 1, \ldots, N$, are obtained, such that I_n has a lower number of regions than I_{n-1}. Image I is described through a multi-layered graph as that sketched in Fig. 2. Each region R_k in image I_n is represented by a node $V_{n,k}$ of the graph $\mathcal{G}$. Node $V_{n,k}$ is connected through *intra-level links* to nodes which represent neighboring color regions. $V_{n,k}$ is connected through *inter-level links* to its son nodes at layer $n-1$.

The graph $\mathcal{G}$ is a multi-resolution index of the chromatic content of image I, which stores both the chromatic and the positional information extracted through the segmentation process.

4. COLOR REGION REPRESENTATION

In order to support effective retrieval by content of paintings, the perceptually uniform CIE $L^*u^*v^*$ color space has been chosen, such that close distances in the color space correspond to close distances for the user's perception. This property holds almost everywhere in this space, except for some hues (Red exhibits

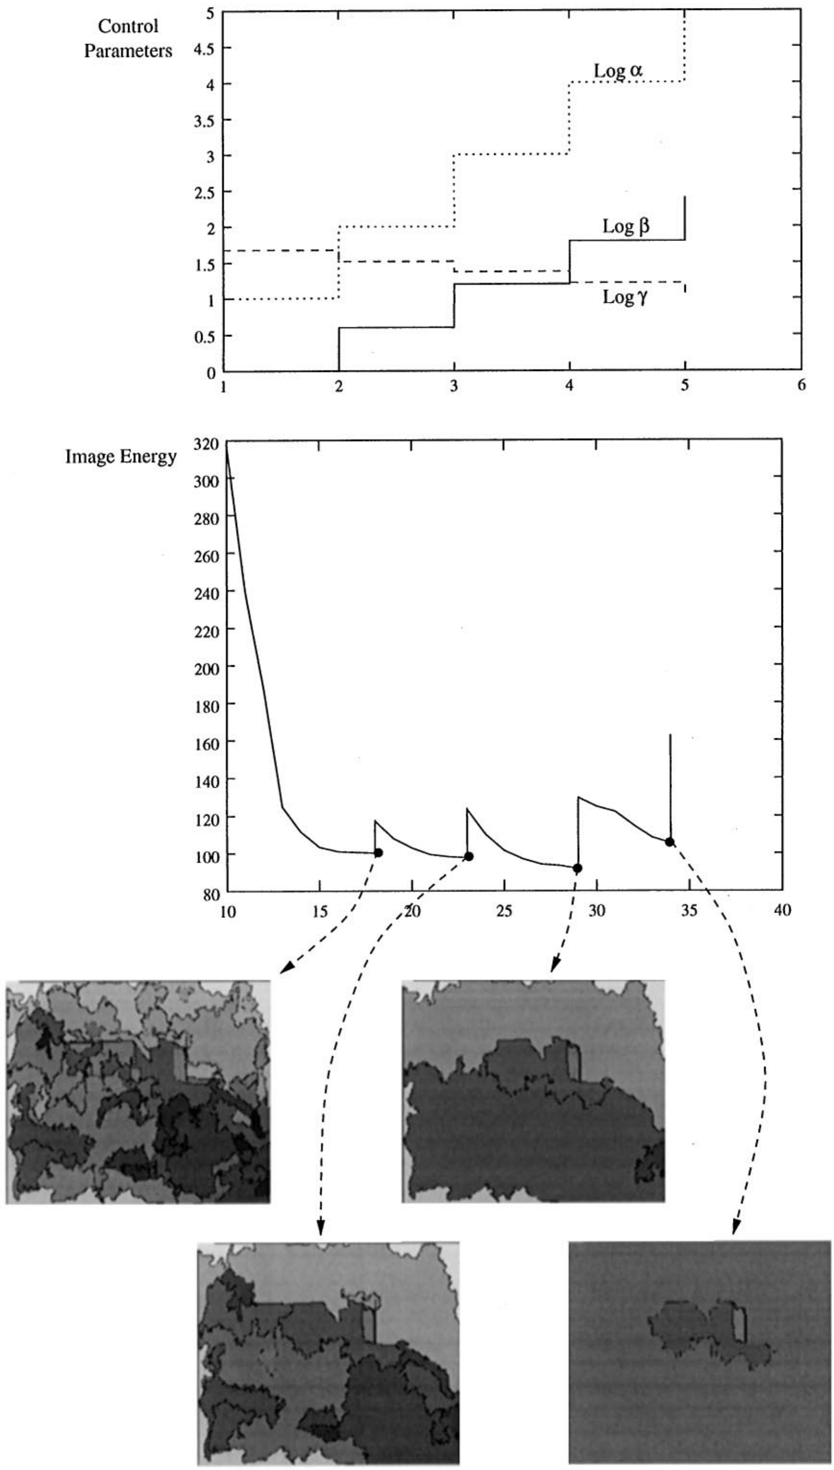

Fig. 1. Image energy behavior in a segmentation process.

a larger range than Green and Blue; colors close to grey are poorly discriminated). The computation of a perceptually meaningful color distance between two generic points in the $L^*u^*v^*$ space, requires to evaluate the length of the shortest path linking the two points.[16] Due to its complexity, such computation cannot be carried out in real time. Alternative ways must be foreseen to correct the error introduced by the approximation of the color distance with the Euclidean distance, in the case of distant colors.

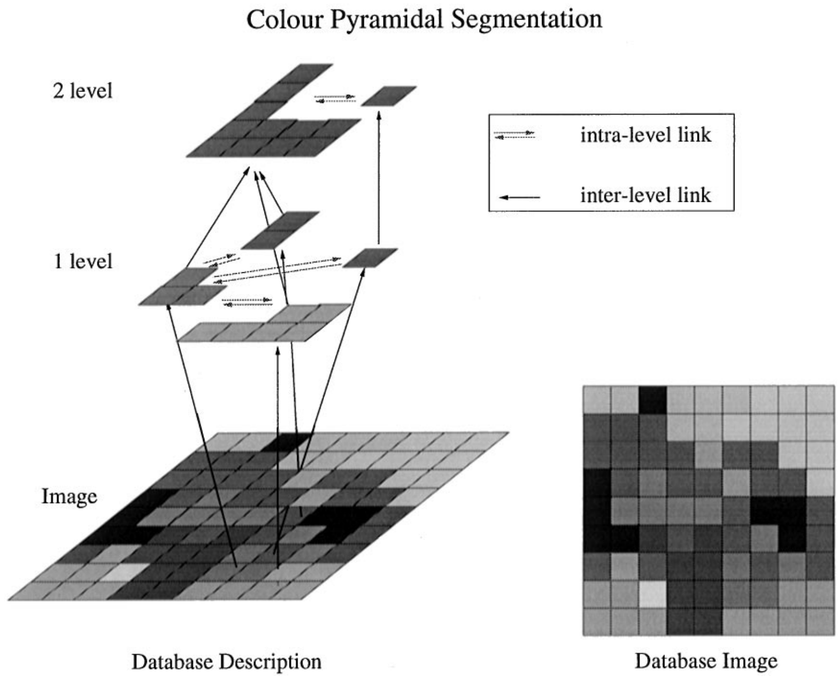

Fig. 2. Intra-level and inter-level links between color regions at different resolution levels.

In our approach, a uniform tessellation of the $L^*u^*v^*$ space has been performed and the number of colors has been reduced to a small set of *reference colors* $\mathbf{c}_i$. Clustering is obtained using an improved version of the standard K-means algorithm.[17] *Competitive learning* has been adopted for grouping points in the color space, as in Reference (18). The number of reference colors is chosen so that the distance between two generic colors belonging to the same tessel is well approximated by the Euclidean distance. From our experiments, we found that 128 colors suffice to achieve a reasonable compromise between accuracy and computational effort. Distances between reference colors are computed according to Reference (16) and stored in a matrix A.

4.1. Region description

Color regions are modeled through their spatial location, area, shape, average color, and a binary 128-dimensional *color vector*. Each entry in the color vector corresponds to a reference color. If a reference color is present in the region, its corresponding entry in the color vector is set to 1 (0 otherwise).

At the coarsest resolution, the image is represented by a single region and the color vector retains the global color characteristics for the entire image. As the resolution increases, regions correspond to smaller areas in the image and therefore have a smaller number of reference colors. For a generic region R^n (at level n of resolution), with k child regions $R_1^{n-1}, \ldots, R_k^{n-1}$ (at level $n-1$ of resolution) the color vector $C_v(R^n) = \{c_1^{R^n}, \ldots, c_{128}^{R^n}\}$ is computed as the union of the color vectors associated with child regions. Each entry $c_i^{R^n}$ is therefore modeled as

$$c_i^{R^n} = c_i^{R_1^{n-1}} \,|\, c_i^{R_2^{n-1}} \,|\, \ldots \,| c_i^{R_k^{n-1}}$$

Figure 3 shows different steps of segmentation and corresponding incremental descriptions for a sample painting.

Region *area* is evaluated as the ratio between the number of region pixels $\# R$ and image pixels $\# I$:

$$A(R) = \frac{\# R}{\# I}$$

Region *position* is evaluated as the absolute position of its centroid $(\bar{x}, \bar{y})$:

$$\bar{x} = \frac{1}{\# R} \sum_{(x,y) \in R} x, \qquad \bar{y} = \frac{1}{\# R} \sum_{(x,y) \in R} y$$

Region *shape* is modeled using the first 13 central moments $\mu_{i,j}$ defined as

$$\mu_{i,j} = \sum_{(x,y) \in R} (x - \bar{x})^i \, (y - \bar{y})^j$$

5. COLOR IMAGE RETRIEVAL

PICASSO system supports retrieval by visual example of images with one or more colored regions.

Queries are formulated through visual examples. Regions can be either sketched and then filled with appropriate colors (as in Fig. 4), or extracted from images, for example taken from a previous query

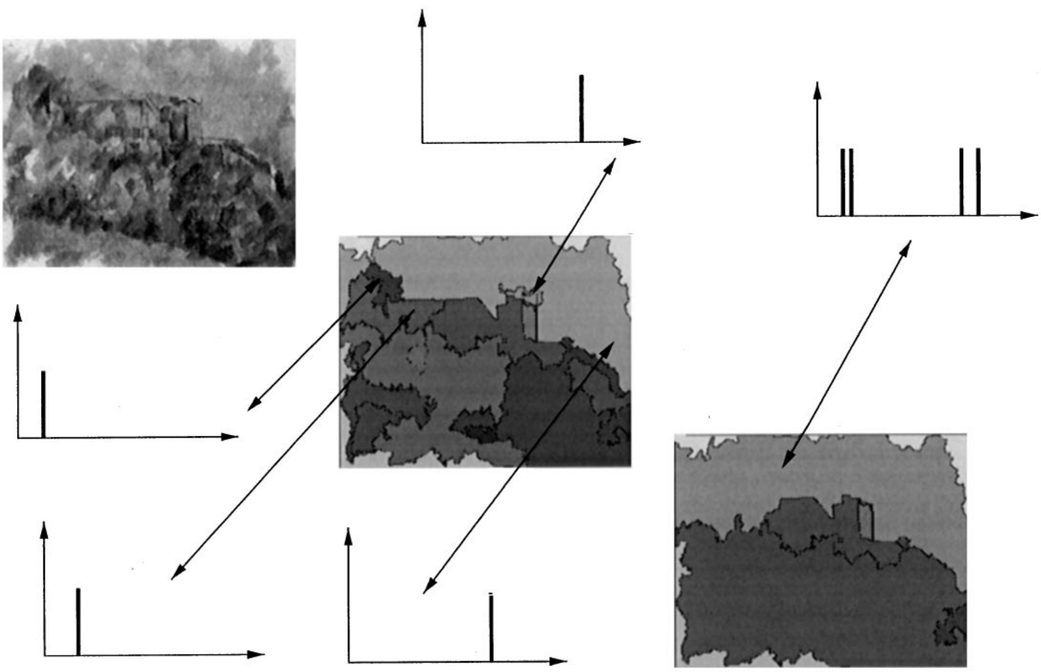

Fig. 3. Segmentation of a painting. The original image and segmented images are shown. The color vector of each region is computed as the union of the color vectors of its child regions.

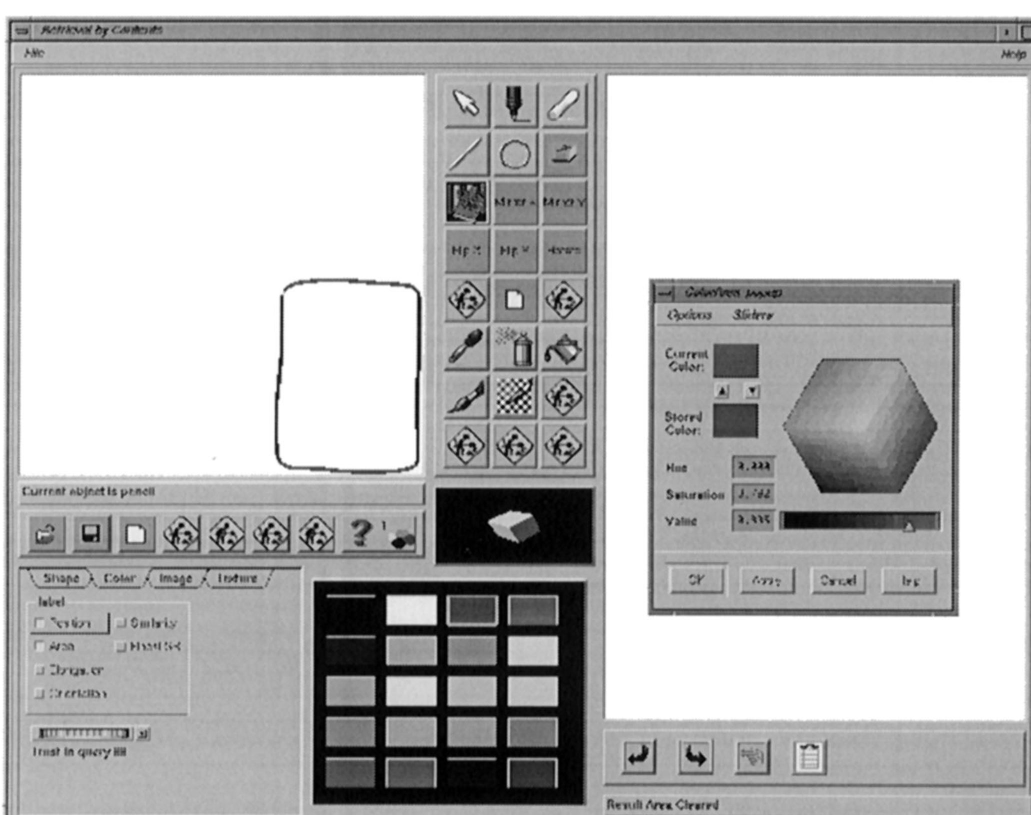

Fig. 4. The PICASSO interface. The user can sketch region contours and select the appropriate region color from a color picker. In this example region color, position and area have been selected for the retrieval process (lighter buttons in the left bottom part of the screen).

Fig. 5. Region contours can be sketched from database images.

output (as in Fig. 5). Similarity of color regions takes into account either chromatic qualities of sketched regions, or a combination of chromatic and spatial attributes (*region area*, *location*, *elongation*, *orientation*).

At database population time, images are automatically segmented and modeled through a pyramid structure as expounded in Section 3. The highest node of each pyramid includes both the binary color vector associated with the whole image, and the full image color histogram. Color histograms are used to compute the similarity between two images according to their global chromatic content. A *color index file* is built, storing the color vectors associated with the highest node of each image pyramid. It helps in selecting the set of candidate images containing regions with the same colors as the query. The color index file has 128 entries, one for each reference color. Each entry storing a list of database images where that reference color is present.

Given a query, the color index file is used, to select a set of candidate images that contain regions with the same colors as the query. Unrelevant images which do not contain some region with the same color as the query are filtered out.

The pyramid structure of each candidate image is analyzed in order to find the best matching region R_I for each query region R_Q.

Given a query region Q_R, the method used to find the best matching region of an image can be summarized as follows:

procedure *match* (R_Q, $\mathcal{G}$)
Let V be the vertex of the pyramid $\mathcal{G}$
Initialize the match $M=0$
analyze (V, R_Q)
return M
end

procedure *analyze* (V, R_Q)
Let C be the color of R_Q
if C is contained in the color vector of V
Compute the match T between R_Q and V, according to (2)
if $T > M$
update M with T
if V is a leaf
stop
else
let $V_1, \ldots, V_K$ be the children of the current node
for $k = 1$ **to** K
analyze (V_k, R_Q)
end

A matching score M between a query region R_Q and a database image region R_I is computed as the weighted sum of the distances (d_{feature}) between attributes considered:

$$M = w_c * d_c(R_Q, R_I) + w_l * d_l\,(R_Q, R_I) + w_a * d_a(R_Q, R_I) + w_s * d_s(R_Q, R_I) \qquad (2)$$

where $d_c(R_Q, R_I)$ is the distance between colors of query region R_Q and database image region R_I; $d_l(R_Q, R_I)$ the distance between positions of regions R_Q and R_I; $d_a(R_Q, R_I)$ the distance between areas of regions R_Q and R_I; $d_s(R_Q, R_I)$ the distance between shapes of regions R_Q and R_I; and w_i weights associated with color, position, area and shape distances.

A similarity coefficient for the whole image is evaluated as the sum of scores of the best matching image regions for each query region, $S = \sum_{i=1}^{N} M_i$.

Results of the query in Fig. 4 are shown in Fig. 6. All the retrieved images include a red-like color patch with position and area similar to the query. The system allows retrieval of images with non-uniform reddish color, with different degrees of brightness and saturation. Figures 7 and 8 show results of the same query when either the area or the position are selected as relevant features. Figure 9 shows the results obtained for the query in Fig. 5.

Figure 10 shows a query with multiple color regions. Retrieved images are shown in Fig. 11. The effects of the multi-resolution segmentation can be noticed by considering that in most retrieved images the green patch is not uniform, but is rather a mosaic of different green colors. However, these different colors are perceived as a uniform green patch by most users.

PICASSO system also supports retrieval by global color similarity. The user can switch back and forth from querying by color regions to querying by global color similarity. Querying by global similarity allows retrieval of images with similar global chromatic content, without taking into account the location of color regions in the image plane. This is useful, for example, to investigate the use of similar colors in different painters, or the use of similar chromatic content in different paintings of the same author. Retrieval by global color similarity is carried out by evaluating the correlation between the color histogram of the query image and that of database images. In Fig. 12, one of the output images of Fig. 11 has been picked up and used as a new query. It can be noticed from this example that the output includes many paintings of the same author (Cezanne).

Effectiveness of image retrieval based on color regions was tested on a set of users, to analyze to what extent the similarity, as it is estimated by the system, is close to the similarity perceived by humans. We

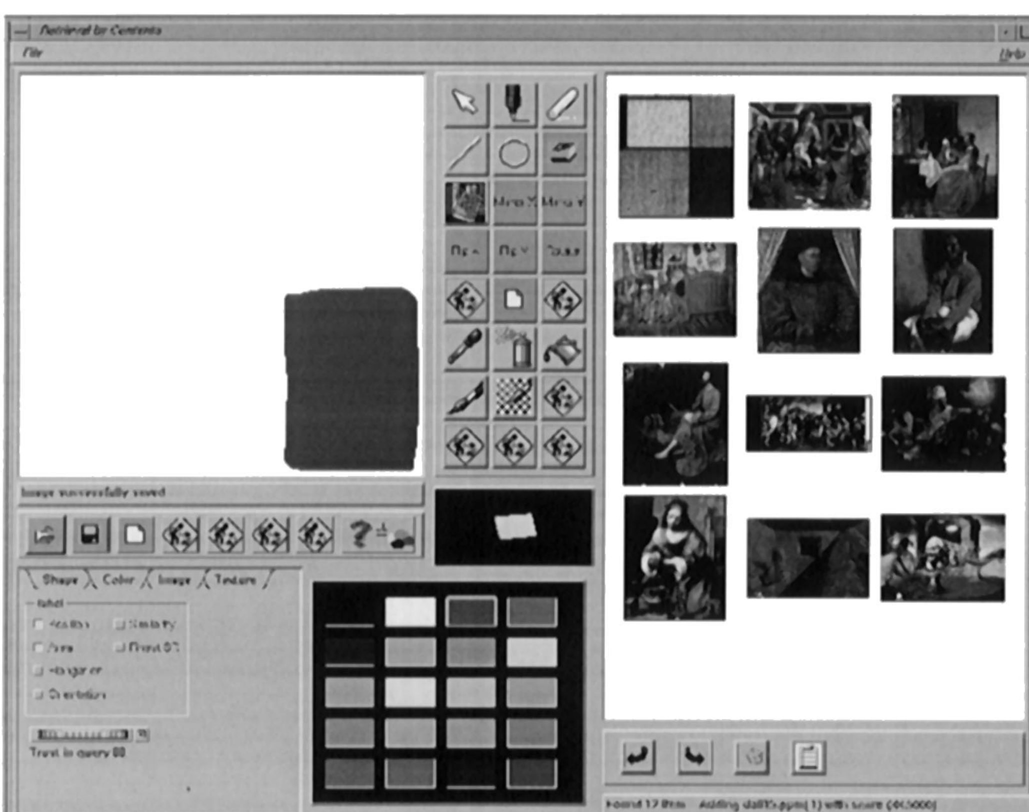

Fig. 6. Retrieved images of the query of Fig. 4.

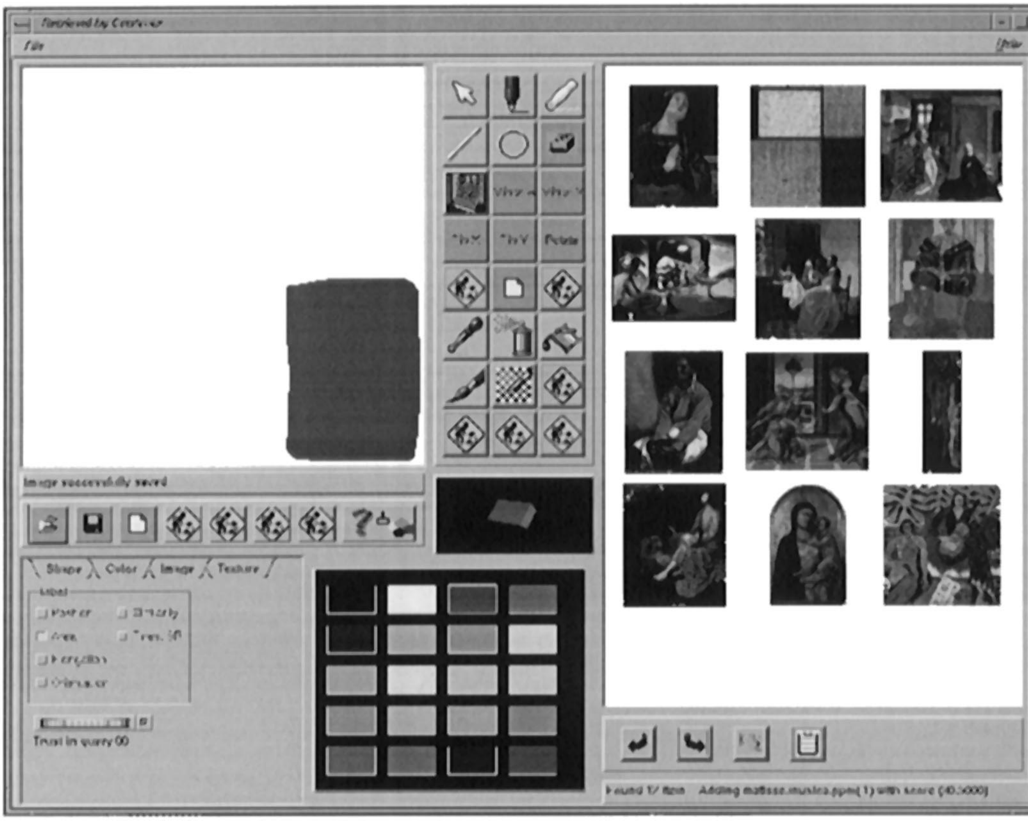

Fig. 7. Retrieved images of the same query as in Fig. 4 but considering only color and area features.

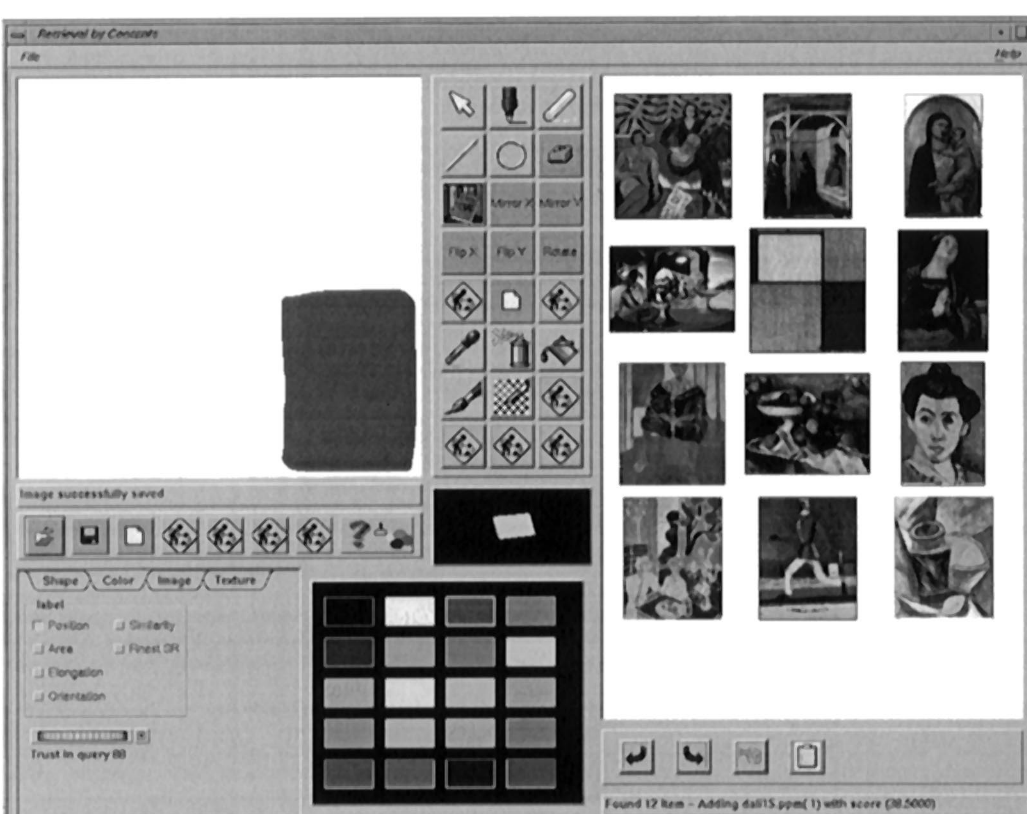

Fig. 8. Retrieved images of the same query as in Fig. 4 but considering only color and position features.

Fig. 9. Retrieved images of the query of Fig. 5.

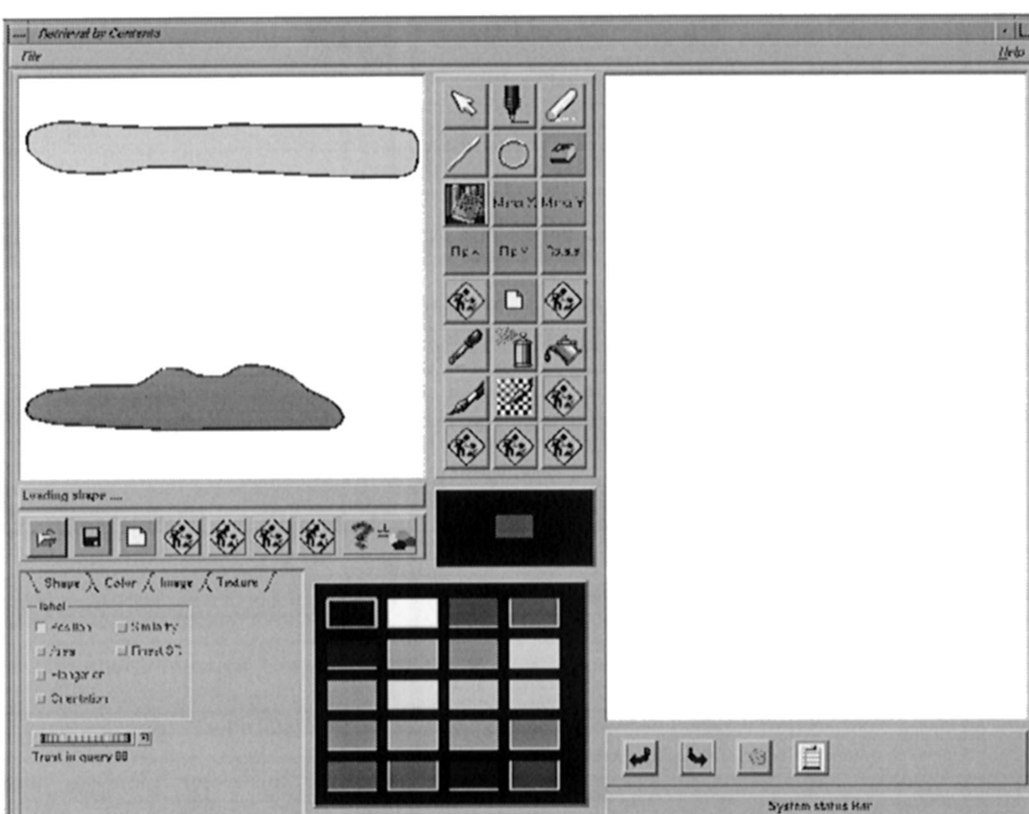

Fig. 10. User sketched query looking for paintings with a blue region in the upper part and a green region in the lower part.

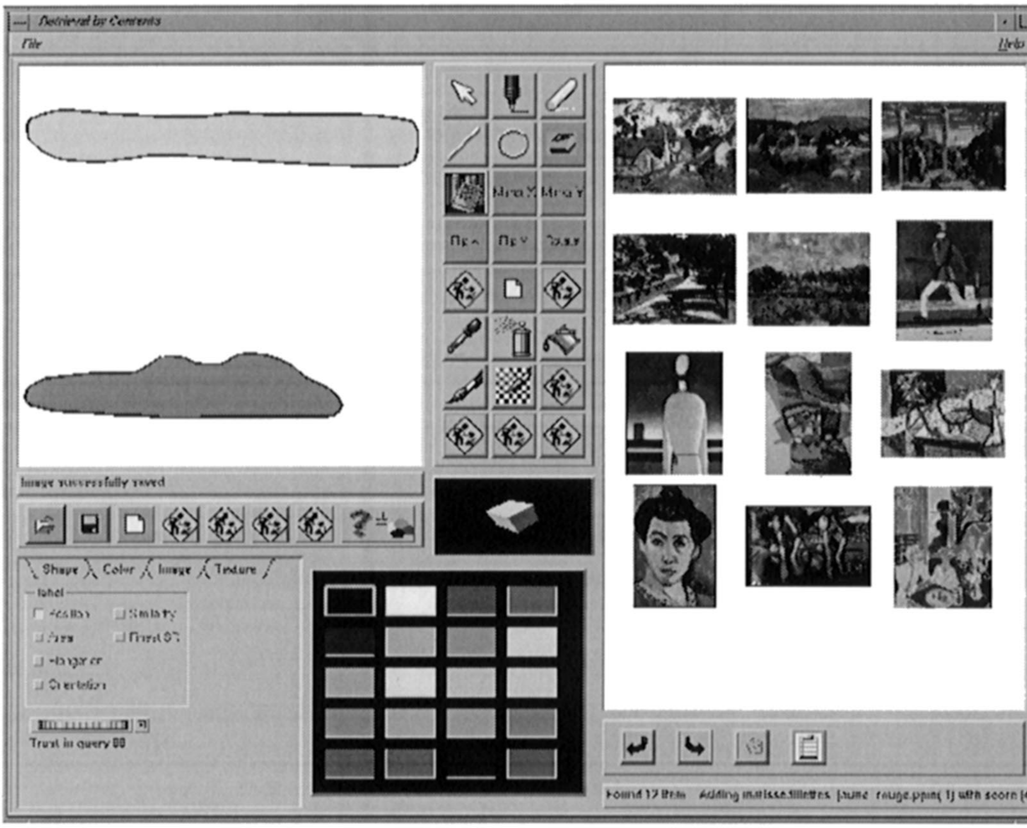

Fig. 11. Retrieved images of the query of Fig.10.

Fig. 12. Example of image retrieval by global color similarity.

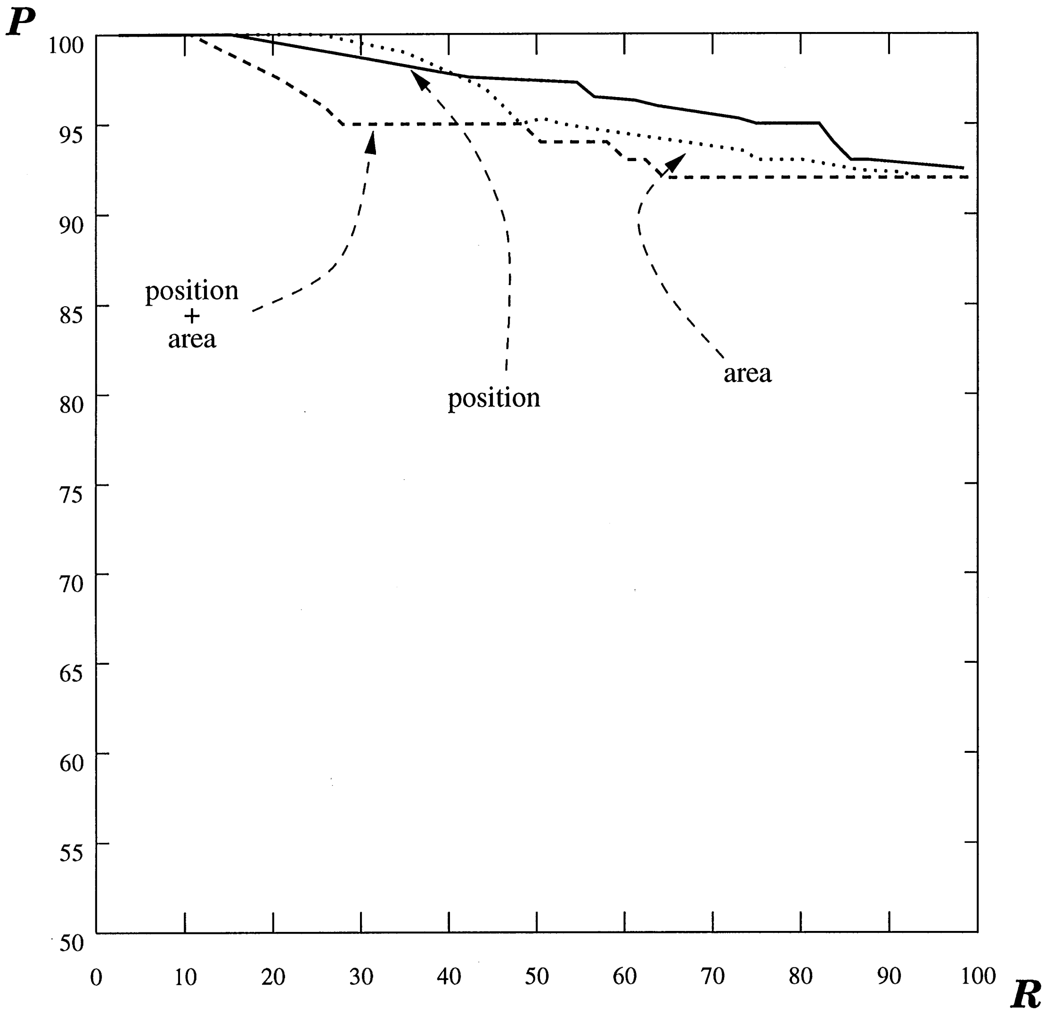

Fig. 13. System precision and recall.

collected answers from 30 people with university education. These were students in the fine art field (65%), and students in other fields like engineering and literature education (35%). In the test, we used a database of 400 color images (selected by experts on the basis of the objective of the test) and 10 sample queries. For each query, we asked the system to retrieve the most similar database images. Each query was processed three times, corresponding to the selection of the following features: color/area, color/position and color/position/area. Retrieved images were shown to the users which were asked to evaluate how many retrieved images were unrelevant and how many non-retrieved images were relevant. By averaging users answers, we derived a set of performance figures in terms of *precision* $\mathscr{P}$ and *recall* $\mathscr{R}$. In Fig. 13, values of $\mathscr{R}$ and $\mathscr{P}$ are reported. Values of these parameters are presented for the three different types of query: by color/position, by color/area, by color/area/position. The best values of $\mathscr{R}$ and $\mathscr{P}$ were obtained for color/position queries. In fact, position is an average measure of pixels coordinates. Therefore, the match is not too much sensitive to differences between query and image regions. If the area is taken into account, high matching scores are obtained only if the region sketched in the query closely corresponds to the region detected in the segmentation process.

6. CONCLUSIONS

In this paper, we presented a system that supports effective retrieval based on color regions through a hierarchical multi-resolution image segmentation. To cope with human perception of colors, image chromatic features are represented through a set of reference colors, obtained from a uniform tessellation of the $L^*u^*v^*$ color space.

The system allows retrieval of images based both on local and global chromatic features. Queries are performed either by sketching colored regions or picking up images from a previous output. Multi-resolution segmentation allows the system to cope with dissimilarities between regions expressed in the query and those in the database images.

An application of the PICASSO system to a sample database of paintings has been presented. Performance has been analyzed in terms of system precision and recall for different types of queries.

Presently, the PICASSO system is used by *Alinari Archives* in Florence.

REFERENCES

1. R. K. Srihari, Automatic indexing and content-based retrieval of captioned images, *IEEE Computer* **28**(9), 49–56 (1995).
2. F. Liu and R. W. Picard, Periodicity, directionality, and randomness: Wold features for image modeling and retrieval. M.I.T. Tech. Rep., no. 320 (1995).
3. K. Hirata and T. Kato, Query by visual example: content-based image retrieval, in *Advances in Database Technology – EDBT* '92, A. Pirotte, C. Delobel and G. Gottlob, Eds, Lecture Notes on Computer Science, Vol. 580. Springer, Berlin, 1992.
4. R. Mehrotra and J. E. Gary, Similar-shape retrieval in shape data management, *IEEE Computer* **28**(9), 57–62 (1995).
5. A. Del Bimbo and P. Pala, Visual image retrieval by elastic matching of user sketches, *IEEE Trans. Pattern Anal. Mach. Intell.* **19**(2), 121–132 (1997).
6. S. K. Chang and S. H. Liu, Picture indexing and abstraction techniques for pictorial databases, *IEEE Trans. Pattern Anal. Mach. Intell.* **6**(4), (1984).
7. S. K. Chang, C. W. Shi and Q. Y. Yan, Iconic indexing by 2-D strings. *IEEE Trans. Pattern Anal. Mach. Intell.* **9**(3), 413–427 (1987).
8. A. Del Bimbo, E. Vicario and D. Zingoni, Symbolic description of image sequences with spatio-temporal logic, *IEEE Trans. Knowledge Data Engng.* **4**(7), 609–621 (1995).
9. A. Del Bimbo and P. Pala, Image indexing using shape based visual features, *Int. Conf. on Pattern Recognition*, Wien, August (1996).
10. M. J. Swain and D. H. Ballard, Color indexing, *Int. J. Comput. Vision*, **7**(1), 11–32 (1991).
11. W. Niblack *et al.*, The QBIC project: querying images by content using color, texture and shape. Res. Report 9203, IBM Res. Div. Almaden Res. Center (1993).
12. A. K. Jain and A. Vailaya, Image retrieval using color and shape, *Pattern Recognition* **29**(8), 1233–1244 (1996).
13. A. Nagasaka and Y. Tanaka, Automatic video indexing and full video search for object appearances, in *IFIP Trans., Visual Database Systems II*, W. Knuth, eds, pp. 113–127. Elsevier, Amsterdam (1992).
14. V. Vinod and H. Murase, Focussed retrieval of color images. *Pattern Recognition* (to appear).
15. J. M. Corridoni and A. Del Bimbo, Multi-resolution color image segmentation. Tech. Rep. no. 7–96, Dip. Sistemi e Informatica, Università di Firenze, 1996.
16. A. K. Jain, Color distance and geodesics in color 3 space, *J. Opt. Soc. Am.* **62**(11), 1287–1291 (1972).
17. A. K. Jain, *Algorithms for Clustering Data.* Prentice Hall, Englewood Cliffs, NJ (1991).
18. T. Uchiyama and M. A. Arbib, Color image segmentation using competitive learning, *IEEE Trans. Pattern Anal. Mach. Intell.*, **12**(16), 1197–1206 (1994).
19. R. C. Carter and E. C. Carter, Object-oriented query languages: the notion and the issues, *IEEE Trans. Knowledge Data Engng.* **4**(4), 223–227 (1992).
20. J. Itten, *Kunst der Farbe* (*The Art of Color*). Otto Maier Verlag, Ravensburg, Germany (1961).
21. C. L. Huang, T. Y. Cheng and C. C. Chen, Color images' segmentation using scale space filter and Markov random field, *Pattern Recognition* **10**(25), 1217–1229 (1992).
22. Y. W. Lim and S. U. Lee, On the color image segmentation algorithm based on the thresholding and the fuzzy *c*-means techniques, *Pattern Recognition* **9**(23), 935–952 (1990).
23. J. Liu and Y. H. Yang, Multiresolution color image segmentation, *IEEE Trans. Pattern Anal. Mach. Intell.* **7**(16), 689–700 (1994).

About the Author—ALBERTO DEL BIMBO graduated in Electronic Engineering at the University of Florence, Italy in 1977. Since 1977 to 1988 he has been with IBM Italy SpA. Presently, he is Full Professor at the Faculty of Engineering of the University of Florence, Italy. His scientific interests cover image sequence analysis, shape and object recognition, image databases, visual languages, and advanced man–machine interaction. Professor A. Del Bimbo is presently an Associate Editor of Pattern Recognition and Journal of Visual Languages and Computing. Professor A. Del Bimbo is a member of IEEE, the Institute of Electrical and Electronics Engineers and IAPR, the International Association for Pattern Recognition. He presently serves as the Chairman of the Italian Chapter of IAPR. He has been in the Program Committee of many International Conferences in the field of Imaging Technology and Multimedia Computing. Professor A. Del Bimbo is the author of over 120 scientific publications appeared in International Journals and Conference Proceedings and has been Guest Editor of several special issues on Imaging Technologies.

About the Author—MAURO MUGNAINI received the Laurea degree in Electronic Engineering from the Università di Firenze, Italy, in 1996. He is presently a technician of Consorzio Nexus in Florence, Italy. His current research interests include image databases, internetworking and client/server applications.

About the Author—PIETRO PALA received the Laurea degree in Electronic Engineering from the Università di Firenze, Italy, in 1994. He is presently a Ph.D. student at the Dipartimento di Sistemi e Informatica of the Università di Firenze, Italy. His current research interests include pattern recognition, image databases, neural networks, and related applications.

About the Author—FRANCESCO TURCO received the Laurea degree in Electronic Engineering from the Università di Firenze Italy, in 1996. He is presently a external consultant of Agenzia per l'Alta Tecnologia CESVIT S.p.A. of Florence, Italy. His current research interests include image databases, multimedia applications, technology for persons with disability, and related activities.