ELSEVIER

Pattern Recognition Letters 22 (2001) 323–337

Pattern Recognition
Letters

www.elsevier.nl/locate/patrec

A scheme of colour image retrieval from databases¹

J.M. Fuertes^{a,*}, M. Lucena^a, N. Pérez de la Blanca^b, J. Chamorro-Martínez^b

^a Departamento de Informática, Escuela Politécnica Superior, Universidad de Jaén, 23071 Jaén, Spain

^b Departamento de Ciencias de la Computación e I.A., E.T.S. de Ingeniería Informática, Universidad de Granada, 18071 Granada, Spain

Received 20 October 1999; received in revised form 25 September 2000

Abstract

This paper presents a scheme of image retrieval from a database using queries prompted by the colour and the shape of the objects present in different scenes. Of the whole scheme of image retrieval, we will focus attention on the modules that allow feature extraction of the component objects from the scenes and the matching of the objects among the different images. The defined scheme enables the indexing of images by measuring the similarity between the integral objects. © 2001 Elsevier Science B.V. All rights reserved.

Keywords: Colour image segmentation; Colour spaces; Deformable model; Image retrieving; Content-based retrieval

1. Introduction

Over the last decade we have seen an increasing need to save and share multimedia elements of high graphic content. This need has led to the development of mechanisms that allow for an efficient storage, access and retrieval of graphics and images organised around large databases. The interaction between these databases takes place through queries made up of key images (Swain, 1993; Jacobs et al., 1995; Kankanhalli et al., 1996; Gevers and Smeulders, 1999b, 2000; Del Bimbo and Pala, 1997; Del Bimbo et al., 1998; Chang, 1997).

Fig. 1 shows the general scheme of image retrieval. In this paper, we will show the development of the modules of feature extraction and image matching.

In (Chang and Fu, 1980), the concept of query by pictorial example was introduced. This shows a retrieval scheme in which the query to the database is made up of one or more images extracted by the system from the characteristics used in the comparison established between the query image and the input data. This is an adaptation of the query by example (QBE) mechanism for alphanumeric data, which laid down the bases of the interfaces of the image retrieval schemes developed later on.

In the selection process of the images that fulfil the conditions expressed in the query, it is necessary to use some kind of similarity ranking. The segmentation of scenes made up of different objects is in the majority of cases ill-defined. This is due to the phenomena of occlusion, translation,

¹ Electronic Annexes of Figs. 3, 4, 7–13 available. See <http://www.elsevier.nl/locate/patrec>.

* Corresponding author. Fax: +34-953002400.

E-mail address: jmf@ujaen.es (J.M. Fuertes).

Fig. 1. General scheme of image retrieval from a database.

rotation, scaling or deformations due to perspective, shading or changes in the illumination of the scenes (Nayar and Bolle, 1996; Baker et al., 1998; Funt and Finlayson, 1995; Finlayson et al., 1998). In this paper, we introduce a technique of automatic segmentation of scenes into regions of homogeneous colour. These regions should not be associated to the concept of object (this word has semantic connotations not contemplated in our scheme). The mechanism for joining regions defined in Section 2.3 allows for a perfect segmentation of the object if the colours are similar in the region it occupies. Therefore, the relationship *homogeneous region of coverage = segmented object* does not necessarily have to hold true.

The colour in a scene is a sufficiently discriminating attribute to allow the identification of the boundary points of interest by the human eye. Perceptual experiments show that humans can easily distinguish differences in colour, even when the intensity of the image is uniform. It therefore seems necessary to include the purely chromatic data of a scene and integrate it with the intensity so as to improve the results obtained in the detection of boundaries. However, traditionally, colour data have been avoided on many occasions either to simplify the computations or to reduce the storage space required to contain the studied image. Recent works reveal increasing interest in the use of colour as an attribute for image segmentation (Chapron, 1992; Healey, 1992; Schettini, 1993; Lambert and Carron, 1999; Borsotti et al., 1998; Zugaj and Lattuat, 1998; Yung and Lai, 1998) as well as for image retrieval from large databases (Swain and Ballard, 1991; Swain, 1993; Mehtre et al., 1995, 1998; Kankanhalli et al., 1996; Gevers and Smeulders, 1999a,b,c, 2000; Del Bimbo et al., 1998; Chang, 1997; Chang et al., 1997;

Ma and Manjunath, 1997; Pentland et al., 1996; Smith and Chang, 1996; Flickner et al., 1995).

The description of shape has been developed with the theory of non-rigid deformation. This approach considers two shapes to be *similar* according to the amount of deformation that would have to be applied to one of the objects in order to transform it into the other one. There are many references related to the idea of describing objects with a mode of non-rigid deformation (Terzopoulos et al., 1988; Bookstein, 1989; Pentland and Sclaroff, 1991; Jain and Vailaya, 1998). We use here the main ideas from the Nastar deformable model (Nastar and Ayache, 1996).

One of the main features to include in a content-based image retrieval scheme is the recovering of scenes containing some specific object defined in the query. In this sense, we propose a preliminary segmentation of images for obtaining homogeneous regions supported in a hierarchical structure of clustering, and then, we use the colour and shape attributes for establishing similarity between scenes. Our scheme, as contrasted with other works (Gevers and Smeulders, 1999a; Jacobs et al., 1995; Jain and Vailaya, 1998; Kankanhalli et al., 1999; Mehtre et al., 1995), allows us to recover images based on queries formulated in terms of complete scenes or homogeneous colour regions contained in them.

The layout of this paper is as follows. In Section 2, we will show the feature extraction module, which defines the mechanism for segmentation of scenes and the characteristics of colour and shape used. In Section 3, we will show the image matching module between query and database. In Section 4, we will detail our experiments, and finally, in Section 5, the conclusions of this paper will be presented.

2. Feature extraction module

Feature extraction plays an important role in the proposed scheme of image retrieval. The objective is twofold: firstly we will segment the scenes into their component homogeneous regions, and then we will determine the features which will allow for their identification.

The images that we are interested in are colour images, and in order to obtain the first part of our objective, we will need to define a representative colour space and establish a mechanism of segmentation. For this, we have used a two-step method of detecting discontinuities and growth regions. To reach the second part of our objective, we have used the average colour of the region and a modal description of its shape.

2.1. Perceptual colour spaces

In order to adapt as closely as possible to the way in which the human eye perceives colours, *perceptual colour spaces* or *uniform scale chromatic spaces* (UCS) were developed. They define colour data in the most similar way possible as perceived by a human being. In this sense, the intensity data are completely separated from those of the purely chromatic, so that for a constant value of the former we can group all the colours of the spectrum of the same intensity in only one plane. These spaces are, among others, $(L^*a^*b^*)$, $(L^*u^*v^*)$ and (HSV) .

In this paper, we use *HSI* as the representation space. Although the mechanisms of segmentation have been developed using this representation model, the general scheme of image retrieval presented would remain valid with any other form of representation.

The (*HSI*) colour space is a tricolour system that formalises the colour system developed by Munsell (Munsell, 1913, 1939). Its design shows the way in which humans see colour, and therefore it offers numerous advantages in image processing (Gonzalez and Woods, 1992; Taylor and Lewis, 1992; Toet, 1992; Kehtarnavaz et al., 1993; Liu and Yang, 1994).

In (Fuertes et al., 1999) we can find a more detailed development of the above space. Geometrically, we have a space made up of two cones joined at the base by a cylinder. This shape owes itself fundamentally to movement along the intensity axis towards either black or white (situated at the apexes of the cones). It becomes meaningless to consider the data concerning saturation or shape. Black has zero intensity and white has maximum intensity. Black or white with higher/lower intensity or different shades do not exist. Fig. 2 shows the *HSI* space used.

2.2. Edge detection in colour images

In order to determine local variations of colour in the image, it will be necessary to define a distance and an order inside the chosen colour space.

Fig. 2. (a) Three-dimensional representation of the defined space; (b) longitudinal projection of the *HSI* space (intensity (*I*) and saturation (*S*) axes); (c) side view of the same (saturation (*S*) and hue (*H*) axes).

2.2.1. Distance in the representation space

The colour space used is inscribed inside a cylinder for which the distance between the two points (colours) $P_1 = (H_1, S_1, I_1)$, $P_2 = (H_2, S_2, I_2)$ is given by

$$d = ((d_I)^2 + (d_C)^2)^{1/2}, \quad (1)$$

where

$$d_I = |I_1 - I_2| \quad (2)$$

and

$$d_C = \begin{cases} ((S_1)^2 + (S_2)^2 - 2S_1S_2 \cos \phi)^{1/2}, & \text{if } P_1 \text{ and } P_2 \text{ are chromatic,} \\ 0, & \text{if } P_1 \text{ and } P_2 \text{ are achromatic,} \\ S_1 - \Omega_{I_2}, & \text{if } P_1 \text{ is chromatic and } P_2 \text{ is achromatic,} \end{cases} \quad (3)$$

given

$$\Omega_{I_2} = \text{Chromaticity threshold for } I_2, \quad (4)$$

$$\phi = \begin{cases} |H_1 - H_2|, & \text{if } |H_1 - H_2| < \pi, \\ 2\pi - |H_1 - H_2|, & \text{if } |H_1 - H_2| > \pi. \end{cases} \quad (5)$$

The value d_I is the difference of intensity between the two points, d_C is the difference between the projection of the points in a chromatic plane, and ϕ is the smallest angle that separates H_1 and H_2 . The chromaticity threshold Ω establishes the separation of the colour representation space into two regions: the *chromatic region* and the *achromatic*. This separation responds to the perceptual necessity of using intensity, saturation or tone when determining borders using the colour points analysed. In (Fuertes et al., 1997) we can find an empirical approach to its calculation.

2.2.2. Order in the colour representation space

As there is no universally accepted natural way of ordering colours, we will determine the order in a simple practical way: the lexicographic order. In our case, given the two pixels (H_1, S_1, I_1) and (H_2, S_2, I_2) we will first compare the values of I , then those of S and finally those of H . To compare the latter component, the criterion used is the following: if the smallest angle to pass from H_1 to H_2 is positive, we choose (H_2, S_2, I_2) as the

larger element and in other cases, we choose (H_1, S_1, I_1) .

2.2.3. Gradient vector calculation

One of the classic methods for the detection of discontinuities is the use of the gradient vector (magnitude and direction). According to the colour representation space discussed in this paper, we have used first-order masks based on the Gaussian, extending the Canny filter (Canny, 1986), from monochrome images to colour images.

In the discrete case, the convolution of a grey level image with a Gaussian first partial derivative mask can be carried out through the successive convolution of the image with two perpendicular unidimensional masks in the following way:

$$\sum_k m_k x_{i+k} - m_k x_{i-k}, \quad (6)$$

where the values m_k represent the mask coefficients and (x_{i+k}, x_{i-k}) represent the image coefficients.

Grouping the terms and pairing the corresponding pixels equidistant from where the mask is centred, we obtain the following equation:

$$\sum_k m_k (x_{i+k} - x_{i-k}) = \sum_k m_k c(x_{i+k}, x_{i-k}) d(x_{i+k}, x_{i-k}), \quad (7)$$

$d(x_i, x_j)$ being the distance between the colours of the pixels x_i and x_j and

$$c(x_i, x_j) = \begin{cases} 1, & \text{if } x_i > x_j, \\ -1, & \text{if } x_i \leq x_j, \end{cases} \quad (8)$$

the ordering function described above.

2.3. Image segmentation

An edge can be described as a transition between adjacent regions. One of the main disadvantages that edge-detection-based segmentation algorithms present lies in the necessity of an a priori knowledge of the kind of objects to be segmented. When this knowledge is not available, the growth regions and *split and merge* mechanisms can solve the problem.

The segmentation algorithm proposed will use on one hand the information obtained from the gradients, associated to their corresponding boundaries, and on the other hand the lack of a priori knowledge of the object model to be segmented, using a growth region. The main difficulty in using any growth technique is the determination of the placement of the seeds around which the regions will grow.

The proposed iterative algorithm, detailed in (Fuertes et al., 1997), will calculate two basic steps:

- Placement of the seeds along the detected boundaries.
- Use of the gradient data to carry out the propagation of the seeds.

In Figs. 3(c)–(e) we show the iterative region-growing process from the edge map (Fig. 3(b)) in which we place the seeds.

When the entire image has been automatically segmented into regions of interest, we need to establish a way to organise the data within a structure which makes it possible to not only support the specific regions (with all their characteristics established) but to add one to another in order to form elements containing more information. This addition mechanism will segment the scenes in regions that are homogeneous in colour.

The structure of data chosen to represent the descriptive information of each image is that of a tree. On its leaves the different regions established in the growth process of the seeds are placed on the detected edges. These leaves are grouped together in such a way that as we move up the hierarchy of the tree we find that at each level both of the nodes of the lower level which are topologically adjacent and placed at a minimum distance in the colour representation space have been joined. Each of the nodes (tree leaf nodes or nodes formed by grouping together others) saves, among other descriptors, the average colour of the pixels grouped in the region defined in the node over the original image. By doing this we obtain a simple criterion for joining nodes. In the root node the whole image will appear as a composition of different regions (see Fig. 4).

Using the proposed grouping method, Fig. 3(f) shows the grouping of patches which contain the central element of the original image given in

Fig. 3(a). In black we can see node 103 of the tree with all its descendants grouped together to form one region. The remaining ungrouped nodes appear in different colours. With regard to this image, we should point out that, after detecting discontinuities, 635 different seeds were placed which, once grown, brought about a segmentation in which 53 regions appeared (the leaf nodes). A total of 105 nodes are generated when the tree structure is formed.

Section 3 shows how this simple structure enables us to carry out efficient search mechanisms which make adjustment possible between structures of data representing images contained in the database and structures of data representing the image or element which is the same as that which is used in the query.

2.4. Colour features

The average colour of the pixels contained in the same region of the original image will be saved in each node of the structure.

The determined regions will completely cover the image and hence the proposed method will allow us to

- Simplify the scenes. From each region, only the shape features (Section 2.5) and the average colour of the pixels contained in the same region of the original image will be saved.
- Organise data. The simplification of the data will allow for the design of structures of data capable of collecting the topological relations of the different regions, which in turn will allow for their later addition according to similarity criteria to form regions that are capable of representing the objects that make up the image.

2.5. Shape features

We have chosen a model of shape deformation proposed in Nastar (Nastar and Ayache, 1996) to establish the deformation we would have to apply to an object in order to transform it into another one, and in this way to be able to measure the energy used as a measure of object similarity. This

Fig. 3. (a) Original image; (b) map of edges obtained when applying the proposed algorithm for detecting discontinuities (mask 9×9 , $\sigma = 1.0$, upper threshold = 0.9, lower threshold = 0.5); (c) results of the region-growing process after the first step; (d) after step 8; (e) after step 10; (d) after step 10; (e) final image obtained after the regions growing and fusion processes; (f) visualisation of node 103 of the hierarchic structure of regions for the original image given in (a).

model gives us an analytical expression of the vibration modes of the nodes situated on a closed curve (we are interested in contours of closed

forms). We will define a mechanism to represent the value of the modal range as a function of this mode, as follows: $\tilde{u}_i(t) = f(i)$.

Fig. 4. Hierarchic structure for the addition of detected regions (leaf nodes) over the original image placed in the root node.

We will call this representation the *modal spectrum* (Nastar and Ayache, 1996). In summary, it describes which modes and what quantity of them need to be excited in order to deform a shape enough to make it appear like another. In this way, we can obtain the following indicator of the amount of *deformation energy* required for the process:

$$E_{\text{deformation}} = \frac{1}{2} \mathbf{U}^T \mathbf{K} \mathbf{U} = \frac{1}{2} \sum_{i=1}^N \omega_i^2 \tilde{u}_i^2, \quad (9)$$

where $\mathbf{K}$ is the stiffness matrix of the shape, $\mathbf{U}$ is the displacement matrix, ω_i is the frequency of vibration and $\tilde{u}$ is the range of the i th mode in consideration.

In this context, rigid movements have zero deformation energy and two deformations are deemed to be similar when their corresponding fields of displacement $\mathbf{U}_1$ and $\mathbf{U}_2$, regardless of rigid transformations, are similar. Expressed in another way, *two deformations are similar when their corresponding deformation spectra are similar*. The sought objective is to provide our scheme with a mechanism which will allow us to recognise the grade of similarity between two shapes from the deformation imposed on one to emulate the other.

As expected, the deformation energy that would have to be applied to a shape for it to emulate the other will be less, the greater the likeness between them.

As our aim is to define a vector of characteristics that will allow us to recognise the deformation of a given shape, we will have to define a prototype shape, common to all the shapes under consideration, against which to calculate the deformation. As a prototype, the circular shape with a known radius, r , has been selected.

In Fig. 5, we show two different shapes on which we apply a deformation from the common prototype of a circular shape and in Fig. 6 we show the deformation spectra (the first modes in each deformation).

Once (a) the deformation spectra of converting the given shape to a circular one and (b) the number of contour points for all the shapes under consideration have been calculated, we will obtain the deformation spectra of all shapes. All that remains is to establish a distance between the deformation spectra of the given shape as introduced and those of the shapes contained in the database. To establish an index of similarity, we use a Euclidean distance

Fig. 5. (a) Relation of deformation between a rectangle and the circular prototype; (b) relation of deformation between a pentagon and the circular prototype.

$$d(D_1, D_2) = \frac{1}{p} \sqrt{\sum_{i=1}^p (\tilde{u}_i(D_1) - \tilde{u}_i(D_2))^2}, \quad (10)$$

where D_1 and D_2 are the deformation labels for each of the p considered modes.

3. Image-matching module

The objective of the image-matching module within the general scheme of image retrieval from a database is to establish a mechanism of indexing the shapes of component objects of different scenes

according to the similarity that exists between them.

Similarity has been established with two relevant characteristics: colour and shape. Each of the nodes of the hierarchic structure established represents an image. Characteristics associated to the supported region are included in the same.

By proposing a scheme of image retrieval from databases containing characteristics which are not defined a priori we give the model, in the user's interface, the possibility of awarding greater importance (understood as a value of relevance) to colour and shape. This simple mechanism allows the user to decide, upon establishing the query, not only the characteristic (colour or shape) used but the relevance of one with regard to the other. This type of supervised query allows the user to surf the database tuning the retrieval.

We will have to generate hierarchical structures which correspond to all of the images contained in the database. The query is, in itself, a sole node with characteristics of colour and shape. The adjustment method is simple: we measure the similarity between the query node and all of the hierarchical structures of the database. The similarity used is the minimum distance in terms of colour (Eq. 1) and modal deformation (Eq. 10). Those images containing within their associated structure a node which is similar to the query node will be kept.

Fig. 6. (a) Deformation spectrum for the relation shown in Fig. 5(a); (b) deformation spectrum for the relation shown in Fig. 5(b).

Fig. 7. Some examples of images contained in the database.

The similarity measure presented in this paper supports an invariance for three different transformations: scaling, translation and rotation. Invariance to scaling transformation is assured by

the normalisation mechanism described in the algorithm shown in this section. Invariance to translation transformation is also assured due to the fact that, upon considering isolated objects

that are scaled from the original size of shape to the circular prototype, the shape's position within the image has no influence. Invariance to rotation transformation is also assured since the method used only responds to deformations that must be applied to pairs of points (prototype–object), so the possible rotation transformations will not affect the scheme.

4. Experimental results

The experiments on the general scheme of image retrieval proposed have been defined with a database composed of 160 images containing multiple objects of different shapes and colours.

More than 40 different objects were used to make up the different scenes. In each of the images we established all of the conditions including illumination in order to introduce phenomena of shading or changes in hue and the positioning of

the camera and objects and in order to introduce phenomena of partial concealment transformations to rotation, translation, scaling and weak perspective. In Fig. 7, we show 30 of the images contained in the database.

Three experiments are presented in this paper according to the three different methods of consulting that the defined scheme supports with the designated user interface:

- *Colour-based query.* This query will extract or retrieve the images that contain objects whose colour is the closest to the target object's. In Figs. 8 and 9, we can see the queries and the obtained answers.
- *Shape-based query.* Through this type of query, those images that most resemble the shape of the target object can be extracted. The query and answer can be seen in Fig. 10.
- *Colour and shape combination-based query.* Here, the images that most resemble the colour and shape of the target object can be retrieved.

Fig. 8. Example of colour-based query.

Fig. 9. Example of colour-based query (cont.).

Fig. 10. Example of shape-based query.

Figs. 11 and 12 show the queries and the answers obtained from this kind of query.

The results given in Figs. 8–12 show, for each query, the images retrieved and indexed in rows according to the similarity criteria that have been established in each case. At the same time, by the side of each image we can observe the segmented image that prompted the answer.

In addition, we proposed an experiment to determine the number of images retrieved in the established scheme, using the retrieval method based

on the combination of the characteristics of colour and shape. We establish 10 queries designed by different users. In Fig. 13, we show the different query images used in the retrieval scheme.

For the colour and shape combination-based query, we calculate the precision and recall rates for measuring the accuracy of the retrieval scheme. The precision rate can be defined as the percent of retrieved images similar to the query among the total number of retrieved images. The recall rate is defined as the percent of retrieved images which

Fig. 11. Example of colour and shape combination-based query.

Fig. 12. Example of colour and shape combination-based query (cont.).

are similar to the query among the total number of images similar to the query in the database. It can be easily seen that both precision and recall rates are a function of the total number of retrieved images. In order to have high accuracy, the system needs to have both high precision and high recall rate.

In Tables 1 and 2, we show the precision rate and the recall rate for the carried out queries. Rows show the queries and columns show the number of recovered images.

5. Conclusions

We have presented the modules of segmentation and matching of a scheme of image retrieval from a database through the use of queries about the *component* elements of the different scenes.

The defined scheme allows for both the automatic establishment of the regions which are homogeneous in colour in the images based on a colour representation space (*HSI* in our case) and an edge detector based on gradients established

Fig. 13. Some images used as query in the retrieval scheme.

Table 1

Precision rates obtained for the different queries when we recover the 5, 10, 15 and 20 *most similar* images

	5	10	15	20
Query 1	100.0	90.0	86.7	75.0
Query 2	100.0	80.0	60.0	45.0
Query 3	100.0	50.0	33.3	25.0
Query 4	100.0	70.0	60.0	45.0
Query 5	40.0	50.0	46.7	35.0
Query 6	100.0	90.0	66.7	50.0
Query 7	100.0	70.0	46.7	35.0
Query 8	100.0	100.0	100.0	80.0
Query 9	80.0	50.0	33.3	25.0
Query 10	100.0	100.0	66.7	50.0

Table 2

Recall rates obtained for the different queries when we recover the 5, 10, 15 and 20 *most similar* images

	5	10	15	20
Query 1	20.8	37.5	54.2	62.5
Query 2	55.6	88.9	100.0	100.0
Query 3	100.0	100.0	100.0	100.0
Query 4	55.6	77.8	100.0	100.0
Query 5	25.0	62.5	87.5	87.5
Query 6	50.0	90.0	100.0	100.0
Query 7	71.4	100.0	100.0	100.0
Query 8	29.4	58.8	88.2	94.1
Query 9	80.0	100.0	100.0	100.0
Query 10	50.0	100.0	100.0	100.0

over itself (extension of Canny's algorithm to colour spaces).

The formulation of the query can be established on the basis of colour, shape or colour and shape

combined as considered by the user. Regardless of the kind of query chosen by the user, the result obtained is a set of images indexed via a criterion of similarity specified in the query. Including weights on characteristics of colour and shape, through the query interface, in the retrieval scheme does not mean the user must have previous knowledge of the type of images contained in the databases. By including these, we believe the scheme provides a simple tuning mechanism over the extraction, thus allowing for a better iteration of the retrieved images.

The proposed scheme has been tried out on a database of 160 complex images. In our experiments, when we consider the colour and shape combination-based query, the recall rate is higher than 87.5% when we recover the 15 most similar images in nine of ten queries proposed.

This paper leaves the problem of grouping nodes unresolved, although mechanisms not based on similarity or colour could be used. The objective is to define a scheme of grouping the regions obtained so that the objects of interest *appear* in some level of the structure after applying the joining method of the same.

The mechanism for joining nodes, based on a criterion of neighbourhood and of minimum distance in *HSI* colour space, characterises each of the nodes by the average colour of the region established in the original image. This criterion means that a complex object made up of different elements of different colours would never be

correctly grouped together. Queries must be made up of only one object of uniform colour.

The use of both colour as well as a deformable model for shape characterisation allows the scheme to retrieve scenes in which the sought object presents slight deformations due to perspective and partial concealment, in addition to translation, rotation and scaling which have already been mentioned.

Although the size of the database is small – 160 images broken down into hierarchic structures containing some 6000 nodes – retrieval times for queries were never greater than 0.9 s on a SUN Enterprise workstation with a 200 MHz processor.

References

- Baker, S., Nayar, S.K., Murase, H., 1998. Parametric feature detection. *Internat. J. Comput. Vision* 27 (1), 27–50.
- Bookstein, F., 1989. Principal warps: thin-plate splines and the decomposition of deformations. *IEEE Trans. Pattern Anal. Machine Intell.* 11 (6), 567–585.
- Borsotti, M., Campadelli, P., Schettini, R., 1998. Quantitative evaluation of color image segmentation results. *Pattern Recognition Letters* 19 (8), 741–747.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Machine Intell.* 8, 679–698.
- Chang, N.S., Fu, K.S., 1980. Query-by-pictorial example. *IEEE Trans. Software Eng.* 6 (6), 519–524.
- Chang, S.F., 1997. Content-based indexing and retrieval of visual information. *SPMag* 14 (4), 45–48.
- Chang, S.F., Chen, W., Meng, H., Sundaram, H., Zhong, D., 1997. VideoQ: an automatic content-based video search system using visual cues. In: *ACM Multimedia Conf.*, Seattle.
- Chapron, M., 1992. A new chromatic edge detector used for color image segmentation. In: *Proc. 11th IAPR Internat. Conf. Pattern Recognition, The Hague, Vol. 3*. pp. 311–314.
- Del Bimbo, A., De Marsico, M., Levialdi, S., Peritore, G., 1998. Query by dialog: an interactive approach to pictorial querying. *Image and Vision Comput.* 16 (8), 557–569.
- Del Bimbo, A., Pala, P., 1997. Visual image retrieval by elastic matching of user sketches. *IEEE Trans. Pattern Anal. Machine Intell.* 19 (2), 121–132.
- Finlayson, G.D., Schiele, B., Crowley, J.L., 1998. Comprehensive colour image normalization. In: Burkhardt, H., Neumann, B. (Eds.), *Computer Vision, ECCV'98*. In: *Lecture Notes in Computer Vision*, Vol. 1406. Springer, Berlin.
- Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steel, D., Yanker, P., 1995. Query by image and video content: the QBIC system. *IEEE Comput.* 28 (9), 23–32.
- Fuertes, J.M., Lucena, M., Perez de la Blanca, N., Fdez-Valdivia, J., 1997. New methods to segmenting color images. In: *VII Symposium Nacional de Reconocimiento de Formas y Analisis de Imagenes, Barcelona*. pp. 233–238.
- Fuertes, J.M., Lucena, M., Perez de la Blanca, N., Fdez-Valdivia, J., 1999. Feature extraction for an image retrieving scheme. *J. Comput. Sci. Technol.* 1 (1), 40–58.
- Funt, B.V., Finlayson, G.D., 1995. Color constant color indexing. *IEEE Trans. Pattern Anal. Machine Intell.* 17 (5), 522–529.
- Gevers, T., Smeulders, A.W.N., 1999a. Color-based object recognition. *Pattern Recognition* 32, 453–464.
- Gevers, T., Smeulders, A.W.N., 1999b. Content-based image retrieval by viewpoint-invariant image indexing. *Image and Vision Comput.* 17 (7).
- Gevers, T., Smeulders, A.W.N., 1999c. Color based object recognition. *Pattern Recognition* 32, 453–464.
- Gevers, T., Smeulders, A.W.N., 2000. PicToSeek: combining color and shape invariant features for image retrieval. *IEEE Trans. Image Process.*, January.
- Gonzalez, R.C., Woods, R.E., 1992. *Digital Image Processing*. Addison-Wesley, Reading, MA.
- Healey, G., 1992. Segmenting images using normalized color. *IEEE Trans. Systems Man Cybernet.* 22, 64–73.
- Jacobs, C.E., Finkelstein, A., Salesin, D.H., 1995. Fast multi-resolution image querying. In: *Proc. SIGGRAPH'95*. pp. 277–286.
- Jain, A.K., Vailaya, A., 1998. Shape-based retrieval: a case study with trademark image databases. *Pattern Recognition* 31 (9), 1369–1390.
- Kankanhalli, M.S., Mehtre, B.M., Kankg Wu, J., 1996. Cluster-based color matching for image retrieval. *Pattern Recognition* 29 (4), 701–708.
- Kankanhalli, M.S., Mehtre, B.M., Huang, H.Y., 1999. Color and spatial feature for content-based image retrieval. *Pattern Recognition Letters* 20, 109–118.
- Kehtarnavaz, N., Griswold, N.C., Kang, D.S., 1993. Stop-sign recognition based on color/shape processing. *Machine Vision and Applications* 6, 206–208.
- Lambert, P., Carron, T., 1999. Symbolic fusion of luminance–hue–chroma features for region segmentation. *Pattern Recognition* 32 (11), 1857–1872.
- Liu, J.Q., Yang, Y.H., 1994. Multiresolution color image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.* 16 (7), 689–700.
- Ma, W.Y., Manjunath, B.S., 1997. NeTra: a toolbox for navigating large image databases. In: *Proc. ICIP'97*. pp. 568–575.
- Mehtre, B.M., Kankanhalli, M.S., Narasimhalu, A.D., Man, G.C., 1995. Color matching for image retrieval. *Pattern Recognition Letters* 16, 325–331.
- Mehtre, B.M., et al., 1998. Content based image retrieval using color shape approach. *Information Process. Management* 35 (1), 109–120.
- Munsell, A.H., 1913. *Munsell Book of Color*. Munsell Color Co., 2441 North Calvert St., Baltimore, MD.

- Munsell, A.H., 1939. A Color Notation, eighth ed.. Munsell Color Co., Boston, MA.
- Nastar, C., Ayache, N., 1996. Frequency-based nonrigid motion analysis: application to four dimensional medical images. *IEEE Trans. Pattern Anal. Machine Intell.* 18 (11), 1067–1079.
- Nayar, S.K., Bolle, R.M., 1996. Reflectance based object recognition. *Internat. J. Comput. Vision* 17 (3), 219–240.
- Pentland, A., Sclaroff, S., 1991. Closed-form solutions for physically-based shape modeling and recognition. *IEEE Trans. Pattern Anal. Machine Intell.* 13 (7), 715–729.
- Pentland, A.P., Picard, R.W., Sclaroff, S., 1996. Photobook: content-based manipulation of image databases. *Internat. J. Comput. Vision* 18 (3), 233–254.
- Schettini, R., 1993. A segmentation algorithm for color images. *Pattern Recognition Letters* 14, 499–506.
- Smith, J.R., Chang, S.-F., 1996. VisualSEEK: a fully automated content-based image query system. In: *ACM Multimedia Conf.*, Boston.
- Swain, M.J., Ballard, D.H., 1991. Color indexing. *Internat. J. Comput. Vision* 7 (1), 11–32.
- Swain, M.J., 1993. Interactive indexing into image database. In: *Storage and Retrieval for Image and Video Databases*. In: *Proc. SPIE*, Vol. 1908. pp. 95–103.
- Taylor, R.I., Lewis, P.H., 1992. Colour image segmentation using boundary relaxation. In: *Proc. 11th IAPR Internat. Conf. Pattern Recognition*, The Hague, Netherlands, Vol. 3. pp. 721–724.
- Terzopoulos, D., Witkin, A., Kass, M., 1988. Constraints on deformable models: recovering 3D shape and nonrigid motion. *Artificial Intell.* 36, 91–123.
- Toet, A., 1992. Multiscale colour image enhancement. *Pattern Recognition Letters* 13 (3), 167–174.
- Yung, H.C., Lai, H.S., 1998. Segmentation of color images based on the gravitational clustering concept. *OptEng* 37 (3), 989–1000.
- Zugaj, D., Lattuati, V., 1998. A new approach of color images segmentation based on fusing region and edge segmentations outputs. *Pattern Recognition* 31 (2), 105–113.