A COMPARATIVE REVIEW OF COLOUR FEATURES FOR CONTENT-BASED IMAGE RETRIEVAL

Elena González $^{(1)}$**, Francesco Bianconi**$^{(2)}$ **and Antonio Fernández** $^{(1)}$

$^{(1)}$ *Universidad de Vigo*
Departamento de Diseño en la Ingeniería
Vigo / España
E-mail: {elena,antfdez}@uvigo.es

$^{(2)}$ *Università degli Studi di Perugia*
Dipartimento Ingegneria Industriale
Perugia / Italia
E-mail: bianco@unipg.it

Abstract

Content-based image retrieval (CBIR) consists in searching for digital images in large databases by analyzing the visual content of the image. In this paper we propose a categorization scheme for colour texture descriptors used in CBIR, focusing on colour-based features, i.e. those models that only take into account the colour of an image, disregarding the spatial distribution of the pixel values. A representative set of salient colour-based methods has been deeply investigated, implemented and tested. Comparative results are presented along with a discussion of the pros and cons of each technique.

Keywords: colour features, texture features, content-based image retrieval, visual appearance of industrial materials

1 Introduction

Computer vision can be defined[1] as "a branch of artificial intelligence and image processing concerned with computer processing of images from the real world". The aim of this processing or analysis of images may be quite different, depending on the particular application considered (for a general reference on the topic, the interested reader is referred to the work of Branch and Olague [7]). The wide variety of image analysis techniques proposed in literature can be grouped into three main categories, namely: *image classification* (IC), *image segmentation* (IS) and *content-based image retrieval* (CBIR).

Classification consists in grouping images into classes of similar visual properties. It is widely used for product grading (i.e. grouping items into lots of similar visual appearance). Segmentation means dividing a single image into meaningful areas. It finds widespread application in surface inspection (i.e. identification of stains, veins, cracks, etc.). Content-Based Image Retrieval, which is the topic we are concerned with in this paper, can be defined as "any technology that in principle helps to organize digital picture archives by their visual content" [8]. Most commonly CBIR is intended as the problem of searching for digital images in large databases by analyzing the visual contents of the image, instead of using meta-data such as annotations or keywords. CBIR finds a major industrial application in automated stock search engines. Such systems help in searching the most similar items in stock with respect to a query sample. This is very attractive for companies which produce/sell product with high aesthetic content [18]: granite, marble, parquet and fabric are just some examples.

It is commonly accepted that colour and texture are two key elements which determine the human perception of colour images. For this reason the above mentioned tasks are based, to a great extent, on the study of these elements. Texture analysis has been traditionally performed by extracting texture features from the gray-level images, and thus discarding colour information (for a comprehensive review of texture analysis methods readers are referred to reference [24]). Nevertheless, there is strong evidence that incorporating colour information into the texture model yields improved performance [10, 2]. Previous experiments

[1] `http://www.hyperdictionary.com/search.aspx?define=computer+vision`

conducted by the authors are in agreement with this assertion [4, 5]. Despite dozens of colour texture descriptors have been recently presented, to the best of our knowledge, a general framework to classify the wide variety of colour texture models has not been developed yet. In this paper we propose a categorization scheme for colour texture descriptors, based upon an extensive literature review. Our framework groups colour texture descriptors into *spatial*, *spectral* and *hybrid models*. Herein we focus on colour features, a class of spectral models for RGB images. The remaining of the paper is organized as follows: section 2 presents a framework for classification of colour features; sections 3 and 4 provide a thorough description of features based on colour statistics as well as histogram-based features; sections 5 and 6 describe the set-up and report the results of a comparative experiment. Final considerations are presented in section 7.

2 Proposed framework for colour texture descriptors

There is a substantial difference between colour and texture stimuli. Texture refers to the variation of the intensity in a neighbourhood of pixels, whereas colour focuses on the spectral content of an image. Traditionally texture and colour have been regarded as separate phenomena, but in the last years various approaches have emerged to take into account both elements jointly. Some attempts to establish a framework for colour texture analysis have been recently published [21, 27]. Unfortunately these works only cover a limited number of aspects, and fail to provide a general overview. We propose to classify the wide variety of existing colour texture descriptors for IC, IS and CBIR into three main categories: *spatial*, *spectral* and *hybrid*.

Spatial methods are based on the relative variation of pixel values in the spatial domain. Spectral methods rely solely on the colour content of the image, irrespective of the spatial distribution and, consequently, they are invariant to any spatial permutation of the pixels in the image. Hybrid methods share properties of both spatial and spectral methods, although they cannot be considered pure spatial nor pure spectral. There is a wide variety of ways to combine spatial and spectral features into an hybrid model, such as feature concatenation, fusion of classifiers and joint colour-spatial features.

In this work we focus on colour based features, a class of spectral methods devised for three channel images. Colour based features are advantageous in that they are invariant to translation and rotation, and only slightly dependent on the viewing angle. However, it is well-known that the ability of colour based features to characterize colour images drops drastically in case of varying illumination. It is also well-known that images with different spatial layout can originate similar colour features. Colour based features can be divided into two main groups, namely *colour statistics* and *colour histogram based features*, which are described in turn in the next two sections.

3 Colour statistics

With the term *colour statistics* we refer to global statistical parameters (such as mean value, standard deviation, median, percentiles, etc.) computed directly from the colour data.

Kukkonen et al. [13] used the mean values of the *R*, *G*, and *B* colour channels to classify ceramic tiles. In order to compensate for illumination variations the authors also use the rgb space, where the original *R*,*G* and *B* values are normalized by $1/(R+G+B)$. These features perform better, as one could reasonably expect.

A similar approach was presented by López et al. [16], who proposed a set of statistical features for surface appearance such as the mean, the standard deviation and the average deviation.

Based on the consideration that the hue component (*H*) represents an angle in the the IHLS colour space (an improved HLS colour space), Hanbury [11] proposed the use of circular statistics. The central idea is to consider colour content of each pixel as a two-dimensional vector $(cosH, sinH)$. Accordingly, to calculate the hue average of an image one has to determine the argument of the resultant vector of the sum of the unit vectors corresponding to each pixel, instead of averaging directly the hue values. To take into account the relationship between the chrominance coordinates, the foregoing unit vectors can be weighted by the saturation *S*, resulting in $\mathbf{C} = (ScosH, SsinH)$.

4 Colour histogram based features

The *colour histogram* is an estimation of the probability of occurrence of colours in an image. Hence, features computed from colour histogram can be used to model colour texture. The colour histogram depends both on the quantization of the colour space and on the colour space itself. Quantization plays an important role: a coarse quantization (few bins) leads to over-smoothed histograms, whereas a fine quantization (many bins) results in sparse histograms and high-dimensional feature spaces. A compromise between these two extreme situations has to be found.

4.1 Joint 3D colour histograms

The procedure to compute 3D colour histograms consists in establishing a set of representative colours (*palette*) and counting how many times each colour of the palette appears in the image. There are two main approaches to generate the colour palette. On the one hand, a data-independent palette can be generated by partitioning the entire colour space into regions, independently of the colour content of each image. This results in the same predefined palette for all the images to be processed. We refer to this approach as *image-independent quantization*. On the other hand, *image-dependent quantization* considers the colour content of each image. This can be done through suitable colour quantization procedures, such as clustering and related methods. In doing so, different input images give rise to different colour palettes.

4.1.1 Image-independent quantization

Histogram-based methods are based on the *colour histogram*, a first-order statistics of the distribution of colour inside an image which was originally proposed by Swain and Ballard [29]. In this work the authors used the rg-by-wb colour space, which resembles the opponent colour space of the human visual system.
After the pioneering work of Swain and Ballard, various authors proposed variations and improvements over the original colour histogram. Since one of the principal drawbacks of the colour histogram is the high dimensionality of the resulting feature space, many approaches aimed at reducing such dimensionality and/or permitting efficient storage.
Boukovalas et al. proposed a solution where the colour histogram is stored dynamically in the form of a binary tree [6], in order to make the technique more efficient in terms of memory requirements, without any performance penalty. To reduce histogram dimensionality Viet Trand and Lenz [30] proposed a modified version of the Karhunen-Loève transform which takes into account the characteristics of the CIELAB colour space. Paschos and Petrou [23] introduced the *histogram ratio features*, second-order statistics obtained by combining pairs of bins and computing the corresponding colour ratios. The rationale behind this method is that it gives an estimation of the relative importance of the basic colours of an image. A significant drawback of this approach is that the dimension of the feature vector is image-dependent, which makes it difficult to use standard classification methods.

The selection of an adequate colour space and an appropriate quantization scheme are obviously crucial points of each colour histogram method. In the implementation of Swain and Ballard [29] a coarse quantization (8 intervals) is used for the achromatic (wb) axis and a finer quantization (16 intervals) for the chromatic axes (rg) and (by). An extensive comparative experiments on the effects of colour space and quantization was carried out by Lee et al. [14]. They tested uniform quantization over six different colour spaces (RGB, CIEXYZ, CIELAB, CIEL*C*h, opponent colour space and HSV), and found that the CIELAB colour space in general works better, and that a higher number of bins gives better performance, even if it becomes saturated as the number of bins approaches 512.

4.1.2 Image-dependent quantization

The main advantage of the techniques described in the previous section is that, being image-independent, they are also general-purpose. Conversely they are not optimized for a specific image database. A different subdivision of the colour space can be obtained considering the actual colour content of each image. Image-dependent colour quantization methods can be classified into: *splitting algorithms* and *clustering-based algorithms*. The splitting algorithms recursively subdivide the colour space of the original image into separate subspaces according to predefined criteria; the clustering-based methods extract the colour palette through suitable clustering algorithms, such as the K-means. Even if image-dependent colour quantization may provide better retrieval results, this approach is usually more timeconsuming than image-independent colour quantization, since the colour palette has to be recalculated every time a new image is added to the database. The interested reader may refer to the work of Orchard and Bouman [20] and Scheunders [28] for further details about this class of methods.

4.2 Marginal colour histograms

Marginal colour histograms take into account the probability distribution of colours as a function of one or two channels, ignoring information about the other channels. Marginal histograms can be considered as projections of the joint 3D colour histogram onto lower dimensional subspaces (2D or 1D). The "raison d'être" of such methods is the assumption that colour coordinates are fairly uncorrelated. However this assumption may be sufficiently accurate for uncorrelated colour spaces such as the Ohta's (referred to as I1I2I3 in table 1), but not entirely acceptable for the RGB space [25], due to the correlation existing between the R, G and B channels of common imaging devices. Marginal colour histograms can be used as colour features directly, as described in section 4.2.1, or after some kind of processing, such as compression, as described in section 4.2.2. Approaches have also been proposed where the colour features are statistics (i.e. moments) extracted from the marginal histograms, as discussed in section 4.2.3.

4.2.1 Raw marginal histograms

Pietikäinen et al. [25] compared the performance of the joint 3D colour histogram with three 1D marginal histograms (one for each colour axis), in the classification of printed colour paper. Their results indicate that the marginal histograms perform almost as well as the joint 3D colour histogram, and that the Ohta's colour space performs slightly better than the RGB.
Lepistö et al. [15] used marginal colour histograms computed on the *H* and *V* channels of the HSV colour space for classification of natural rock images. Different number of bins for channel where tested: 4, 16 and 256. In our comparative experiment we used the settings that gave the best results: histograms of the *H* and *V* channels quantized with 256 levels.
Konstantinidis et al. [12] proposed the *fuzzy histogram linking* method to project the colour histogram onto one single dimension. This technique focuses on the differences between the regions that divide the colour space in a perceptual sense. To this end, on each component of the CIELAB space, different fuzzy sets are defined by using triangular-shaped built-in membership functions and a set of fuzzy rules. Through fuzzy inference each original colour triplet is assigned a colour membership value from 0 to 1, which is related to the common-sense idea of colour (from 0 to 1: *black*, *darkgrey*, *red*, *brown*, *yellow*, *green*, *blue*, *cyan*, *magenta*, *white*). The resulting feature space is a 10-bin histogram representing the probability distribution of the colour membership values.
Hanbury [11] introduced the *saturation-weighted hue histogram* (SWHH) and the *colour statistics histogram* (CSH). Both are marginal histograms computed over IHLS colour space. The former is a histogram of the hue value of each pixel weighted by the saturation value of the same pixel; the latter is a concatenation of the two histograms of the saturation-weighted mean hue ($|\mathbf{C}|$) and mean length ($arg(\mathbf{C})$) as a function of luminance, where $\mathbf{C}$ denotes the vector defined in section 3.

4.2.2 Compressed 2D histograms

Compressed chromaticity histograms have been proposed by Drew et al. [9]. In their approach the original images are previously submitted to a colour angle normalization procedure for illumination-invariance. Afterward the dimensionality is reduced by converting the original RGB images to the rgb space, and retaining only the *r* and *g* components. The resulting 2D chromaticity histograms, which can be considered as *images*, are compressed using a three-step procedure: a wavelet-based reduction step for low pass filtering followed by *discrete cosine transform* (DCT) and truncation of the DCT-transformed image. The experimental results show that the method is accurate and more resilient to noise than the joint 3D histogram.

In a similar way Berens et al. [3] proposed another approach to compress the colour histograms in order to reduce the computational cost in CBIR. In their procedure the original RGB images are first converted into the *red-green/blue-yellow* colour space, which better decorrelates colour information and is more perceptually uniform than the RGB. The resulting 2D chromaticity histograms are then compressed through standard image coding techniques, such as the *discrete cosine transform*, the *Hadamard transform*, the *Karhunen-Loève transform* and hybrid methods. The experimental results show that compression rates as high as 250 to 1 can be achieved without affecting retrieval performance.

4.2.3 Statistics from marginal histograms

Theses features are usually referred to as *soft colour texture descriptors* [26, 17]. The main advantage of these methods is that the dimension of the feature vector is usually low. This reduces the computational burden and makes these techniques particularly suited for real-time applications.

The chromaticity moments proposed by Paschos [22] are moments computed from 2D chromaticity histograms. The original RGB images are first converted to the xyY colour space, and then they are represented through the chromaticity histograms, which are the probability distribution of the *(x,y)* values. From the 2D chromaticity histograms a set of moments (up to 10) is computed and used as feature vector. In the original formulation the chromaticity moments are not invariant to image dimension. This makes the method inapplicable in CBIR. In order to cope with this problem we introduced, in our experimental comparison, a normalized version of the chromaticity moments.

A set of statistics, computed from the CIELUV colour space, for the classification of burn images, has been proposed by Acha et al. [1]. In this work a calibration procedure is described to move from the RGB space to the device-independent CIELUV. These statistics are: the mean and the standard deviation of each of lightness (L^*), hue ($atan(v^*/u^*)$), and chroma ($\sqrt{(u^*)^2 + (u^*)^2}$); the skewness and the kurtosis of each of L^*, u^* and v^*. With these features the authors report an accuracy higher than 80% in the classification of burn images.

López et al. [16] tested various combinations of statistical descriptors computed from the RGB and the CIELAB spaces. In addition to the features discussed in section 3 (mean, standard deviation and average deviation) they also considered two blocks of marginal histogram moments from the 2^{nd} to the 5^{th} degree and from the 6^{th} to the 10^{th} degree respectively. With the best combination the authors claim a classification accuracy higher than 98% with the CIELAB space and 93% with the RGB space in surface grading of decorated ceramic tiles.

In [26] Prats-Montalbán et al. used a similar set of statistical descriptors, namely: mean, standard deviation and moments from the 2^{nd} to the 5^{th} degree of each colour channel for inter-class surface grading of ceramic tiles. In this work three different colour spaces are used (CIELAB, CIELUV and RGB) and two compression/classification approaches *Soft Independent Modeling of Class Analogy* (SIMCA) and *Partial Least Squares Discriminant Analysis* (PLS-DA). An average classification accuracy of over 94% is reported in the inter-class classification of ceramic tiles.

5 Comparative study

An experimental campaign was carried out to compare a significant subset of the above described techniques. The experiment consisted in a content-based image retrieval task. As a first step we formed a database of 300 colour textures images of industrial materials (i.e. ceramic tiles, natural stone tiles, fabric, wood, leather, etc.). From this database we selected a subset of 6 query images representing fabric (2), granite (2), ceramic (1) and wood (1). The CBIR task consisted in retrieving, for each query image, the five "most similar" images in the database. In order to determine the "ground truth" of the experiment, we submitted the query images and the database to a group of 32 human observers. Each observer was asked to select the five most similar images in the database to each of the six query images, and to rank them in descending similarity order. The images were presented to the observer through a web application, available at `http://dismac.dii.unipg.it/∼CBIR`. Following the same approach proposed in [14], observers were given no guidance about whether to focus more on the spatial or chromatic differences between the images. The feedback from the observers was processed as a voting system: each time an image was selected by one user, we considered that image as receiving one weighted vote, where the weight is $(1/r)$, being r the rank given by the observer to the selected image. The ground truth was then obtained by sorting the images by weighted vote in descending order and retaining the first five images for each query image. The resulting ground truth is shown in fig. 1(a), where the images in the first column represent the query images, and the other images the most similar in order of descending weighted vote from left to right. The performance of the different methods was compared through a minimum distance classifier. For each query image $I_q, q \in \{1, \ldots, 6\}$ the classifier sorts the images of the database in order

(a) Ground truth

(b) Images retrieved with the method *chromaticity moments 3-3*, xyY colour space (ref. [22])

(c) Images retrieved with the method *mean value*, RGB colour space (ref. [13])

(d) Images retrieved with the method *joint 3D colour histogram*, rg-by-wb colour space (ref. [29])

Figure 1: Ground truth and retrieved images.The first column of each mosaic contains the query images. The other columns contain the retrieved images, in descending similarity order from left to right.

of ascending distance from the query image (herein we used the L_1 norm, also known as the *Manhattan* distance), and returns the first five images for each query image. To estimate the effectiveness of each method we used two parameters: 1) *precision* and 2) *average rank of relevant images* [19]. Precision is given by:

$$P = \frac{N_c}{N_{gt}} \tag{1}$$

where N_c is the number of relevant images (i.e. retrieved images which are in the ground truth), and N_{gt} is the number of images which form the ground truth (herein $N_{gt} = 30$). Average rank of relevant images is given by:

$$R = \frac{1}{N_c} \sum_{i=1}^{N_c} r_i \tag{2}$$

where r_i is the rank of the i-th relevant image.

6 Results and discussion

The results obtained with the different methods are summarized in table 1. The invariant version of the chromaticity moments 3-3 gave the best results (highest precision and lowest average rank), followed by various combinations of soft colour descriptors (mean value + other statistics), and by the joint 3D colour histogram. As one would expect, the original version of the chromaticity moments ([22]) does not provide good results. Surprisingly the techniques belonging to the marginal histograms group do not show good performance. Another interesting result is that high-dimensional feature spaces do not perform better than low-dimensional ones. It results, on average, that the precision of all the methods considered in the experiment is rather low. To explain this result it is worth considering the ground truth and the images retrieved by three of the methods which perform best (fig. 1). These images put it evident how different are the search criteria used by the humans and by the colour-based algorithms. The ground truth suggests that the search criteria used by the human observers are somewhat semantic: the observer classifies the

Group	Author	Method	Colour space	Dim.	*P*	*R*
Colour statistics	HANBURY-2003 [11]	Mean resultant chrominance vector	IHLS	2	9/30	1,9
	KUKKONEN-2001 [13]	Mean value	RGB	3	12/30	2,4
		Mean value	rgb	3	11/30	1,6
	LOPEZ-2005 [16]	Mean value + standard deviation	RGB	6	12/30	2,8
		Mean value + average deviation	RGB	6	12/30	2,8
Histogram-based features — Joint 3D	SWAIN-1991 [29]	Joint 3D colour histogram	rg-by-wb	2048	11/30	2,1
Histogram-based features — Marginal histograms	PIETIKAINEN-1996 [25]	Three concatenated marginal histograms	RGB	768	6/30	1,7
			rgb		5/30	2,0
			I1I2I3		5/30	2,0
	LEPISTO-2005 [15]	Marginal histogram (H)	HSV	256	7/30	3,1
		Marginal histogram (V)		256	7/30	1,6
		Two concatenated marginal histograms (H+V)		512	10/30	2,4
	KONSTANTINIDIS-2005 [12]	Histogram from fuzzy linking	Lab	10	6/30	1,7
	HANBURY-2003 [11]	Saturation-weighted hue histogram	IHLS	360	7/30	1,9
		Colour ststistics histogram		100	8/30	2,5
Histogram-based features — Statistics from marginal histograms	PASCHOS-2000 [22]	Chromaticity moments 3-3 (original implementation)	xyY	6	3/30	3,3
		Chromaticity moments 5-5 (original implementation)		10	6/30	2,8
		Chromaticity moments 3-3 (image size invariant)		6	12/30	1,9
		Chromaticity moments 5-5 (image size invariant)		10	11/30	2,0
	LOPEZ-2005 [16]	Mean value + standard dev. + marginal histogram moments	RGB	21	12/30	2,8

Table 1: Results of the comparative experiment.

image as belonging to a certain type of material (i.e. "granite" or "wood"), and searches for images representing materials of the same type. This is particularly evident for the search images 3, 4 (granite) and 6 (wood). We can argue that, in doing this semantic search, an important clue is represented by the spatial structure (texture) of the image, which is not considered in colour-based features. On the contrary the spatial structure is not taken into account by colour features. As a consequence, though the retrieved images can be very similar, in colour, to the search images, they may represent completely different materials and thus they do not match the ground truth.

7 Conclusions

In this paper we presented a review of colour features for content-based image retrieval. A classification criteria is also proposed to catalog the various colour-based techniques into meaningful categories. A comparative experiment for the retrieval of materials of industrial interest according to the criterion of visual appearance is presented. The experiment is based on a ground truth which was previously established through the average response of a group of human observers. The ground truth suggests that, in the process of finding the most similar images, the procedure followed by humans is somewhat "semantic". The experimental results show that retrieval algorithms based on colour features are rather inadequate to replicate this behaviour. Future work will involve the analysis of methods based on the analysis of colour and texture jointly.

References

[1] B. Acha, C. Serrano, J.I. Acha, and L.M. Roa. Segmentation and classification of burn images by color and texture information. *Journal of Biomedical Optics*, 10, 2005.

[2] M.A. Akhloufi, W. Ben Larbi, and X. Maldague. Framework for color-texture classification in machine vision inspection of industrial products. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*, pages 1067–1071, 2007.

[3] J. Berens, G.D. Finlayson, and G. Qiu. Image indexing using compressed colour histograms. *IEE Proceedings - Vision, Image and Signal Processing*, 147:57–71, 2000.

[4] F. Bianconi, A. Fernández, E. González, and J. Armesto. Robust color texture features based on ranklets and discrete Fourier transform. *Journal of Electronic Imaging*, 18:043012–1–8, 2009.

[5] F. Bianconi, A. Fernández, E. González, D. Caride, and A. Calviño. Rotation-invariant colour texture classification through multilayer CCR. *Pattern Recognition Letters*, 30:765–773, 2009.

[6] C. Boukouvalas, J. Kittler, R. Marik, and M. Petrou. Color grading of randomly textured ceramic tiles using color histograms. *IEEE Transactions on Industrial Electronics*, 46:219–226, 1999.

[7] J.W. Branch and G. Olague. La visión por computador. Una aproximación al estado del arte. *Dyna*, 133:1–16, 2001.

[8] R. Datta, D. Joshi, J. Li, and J.Z. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40:5:1–5:59, 2008.

[9] M.S. Drew, J. Wei, and Z. Li. Illumination-invariant color object recognition via compressed chromaticity histograms of color-channel-normalized images. In *Proceedings of the International Conference on Computer Vision*, pages 533–540, 1998.

[10] A. Drimbarean and P.F. Whelan. Experiments in colour texture analysis. *Pattern Recognition Letters*, 22:1161–1167, 2001.

[11] A. Hanbury. Circular statistics applied to colour images. In *Proceedings of the 8th Computer Vision Winter Workshop*, pages 27–32, Valtice (Czech Republic), February 2003.

[12] K. Konstantinidis, A. Gasteratos, and I. Andreadis. Image retrieval based on fuzzy color histogram processing. *Optics Communications*, 248:375–386, 2005.

[13] S. Kukkonen, H. Kälviäinen, and J. Parkkinen. Color features for quality control in ceramic tile industry. *Optical Engineering*, 40:170–177, 2001.

[14] S.M. Lee, J.H. Xin, and S. Westland. Evaluation of image similarity by histogram intersection. *Color Research and Application*, 30:265–274, 2005.

[15] L. Lepistö, I. Kunttu, and A. Visa. Color-based classification of natural rock images using classifier combinations. *Lecture Notes in Computer Science*, 3540:901–909, 2005.

[16] F. López, J.M. Valiente, R. Baldrich, and M. Vanrell. Fast surface grading using color statistics in the CIE Lab space. *Lecture Notes in Computer Science*, 3773:13–23, 2005.

[17] F. López, J.M. Valiente, J.M. Prats, and A. Ferrer. Performance evaluation of soft color texture descriptors for surface grading using experimental design and logistic regression. *Pattern Recognition*, 41:1744–1755, 2008.

[18] M.J. Álvarez, E. González, F. Bianconi, J. Armesto, and A. Fernández. Colour and texture features for image retrieval in granite industry. *Dyna*. To appear.

[19] H. Müller, W. Müller, D. M. Squire, S. Marchand-Maillet, and T. Pun. Performance evaluation in content-based image retrieval: overview and proposals. *Pattern Recognition Letters*, 22:593–601, 2001.

[20] M. Orchard and C. Bouman. Color quantization of images. *IEEE Transactions on Signal Processing*, 39:2677–2690, 1991.

[21] C. Palm. Color texture classification by integrative co-occurrence matrices. *Pattern Recognition*, 37:965–976, 2004.

[22] G. Paschos. Fast color texture recognition using chromaticity moments. *Pattern Recognition Letters*, 21:837–841, 2000.

[23] G. Paschos and M. Petrou. Histogram ratio features for color texture classification. *Pattern Recognition Letters*, 24:309–314, 2003.

[24] M. Petrou and P. García Sevilla. *Image Processing. Dealing with Texture*. Wiley Interscience, 2006.

[25] M. Pietikäinen, S. Nieminen, E. Marszalec, and T. Ojala. Accurate color discrimination with classification based on features distributions. In *Proceedings of the 13th International Conference on Pattern Recognition*, volume 3, pages 833–838, Vienna, Austria, August 1996.

[26] J.M. Prats-Montalbán, F. López, J.M. Valiente, and A. Ferrer. Multivariate statistical projection methods to perform robust feature extraction and classification in surface grading. *Journal of Electronic Imaging*, 17:031106–1–10, 2008.

[27] R. Schettini, G. Ciocca, and S. Zuffi. A survey of methods for colour image indexing and retrieval in image databases. In L. W. MacDonald and M. R. Luo, editors, *Color Imaging Science: Exploiting Digital Media*. John Wiley & Sons Ltd, 2001.

[28] P. Scheunders. A comparison of clustering algorithms applied to color image quantization. *Pattern Recognition Letters*, 18:1379–1384, 1997.

[29] M.J. Swain and D.H. Ballard. Color indexing. *International Journal of Computer Vision*, 7:11–32, 1991.

[30] L.V. Tran and R. Lenz. Compact colour descriptors for colour-based image retrieval. *Signal Processing*, 85:233–246, 2005.