

Making colors worth more than a thousand words

Jurandy Almeida

Anderson Rocha

Ricardo Torres

Siome Goldenstein

Instituto de Computação
Universidade Estadual de Campinas
CEP 13084-851, Campinas, SP – Brasil

{jurandy.almeida, anderson.rocha, rtorres, siome}@ic.unicamp.br

ABSTRACT

Content-based image retrieval (CBIR) is a challenging task. Common techniques use only low-level features. However, these solutions can lead to the so-called ‘semantic gap’ problem: images with high feature similarities may be different in terms of user perception. In this paper, our objective is to retrieve images based on color cues which may present some affine transformations. For that, we present CSIR: a new method for comparing images based on discrete distributions of distinctive color and scale image regions. We validate the technique using images with a large range of viewpoints, partial occlusion, changes in illumination, and various domains.

Categories and Subject Descriptors

H.3 [Information Storage and Retrieval]: H.3.3 Information Search and Retrieval

General Terms

Information Retrieval.

Keywords

Content-Based Image Retrieval (CBIR), color and scale representative features.

1. INTRODUCTION

Content-based image retrieval (CBIR) is a challenging task. Common techniques use low-level features and explore local shape and intensity information for viewpoint and occlusion [30]; wavelets and autoregressive models [35]; surface reflection [3]; and Gabor filters [9]. Even fractal transformations can hold interesting results [6].

Some CBIR techniques use segmentation as a pre-processing stage. However, experience has demonstrated that segmentation is suited only for narrow domains due to its

own difficulty [31,36]. The most preferred descriptors for retrieval in broad-image domains use color and texture information [8,25,31]. Some approaches rely on color histograms and color correlograms [11,34], color coherence vectors [26], and border/interior pixel classification [33]. Sometimes, features such as shape and silhouette [2,37], and moment invariants [14,28] can reduce the ‘semantic gap’ problem: images with high feature similarities may be different in terms of user perception.

Recent developments have used middle- and high-level information to improve the low-level features. Li et al. [17] have performed architectonics building recognition using color, orientation, and spatial features of line segments. Raghavan et al. [29] have designed a similarity-preserving space transformation method of low-level image space into a high-level vector space to improve retrieval. Some researchers have used bag of features for image categorization [10,13,21]. Others have used Bayesian approaches to unsupervised one-shot learning of object categories [7]. However, these approaches often require complex learning stages and can not be directly used for image retrieval tasks.

To address these problems, we present a new method to compare images using discrete distributions of distinctive Color and Scale representative Image Regions (CSIR). The key advantages of this method are: (1) it is robust to viewpoint, occlusion, and illumination changes; (2) it is invariant to image transformations such as rotation and translation; (3) it does not require any learning stage; and (4) it uses an effective metric to compare images with different number of features. Hence, it does not require a fixed number of features for an image.

To support such statements, and to show that the method can be used in CBIR tasks, we validate the technique using images with a large range of viewpoints, partial occlusion, changes in illumination, and various domains.

2. IMAGE DESCRIPTORS

In this section, we present some low-level feature descriptors widely used in the literature. In Section 3, we compare these descriptors with our technique.

In general, we classify color image descriptors into three categories: (1) global-based; (2) partition-based; and (3) region-based.

1. **Global-based.** It comprises methods that globally describe the color distribution of images. Such methods do not take into account the color spacial distribution. These methods are both time and space

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC’08 March 16-20, 2008, Fortaleza, Ceará, Brazil

Copyright 2008 ACM 978-1-59593-753-7/08/0003 ...\$5.00.

computational efficient. GCH (c.f., Sec. 2.1), and BIC (c.f., Sec. 2.2), are examples of such techniques.

2. **Partition-based.** It comprises methods that spatially decompose the image into a fixed number of regions. Regions are individually analyzed in order to capture the color spatial distribution. Such methods do not take into account image’s visual cues. CCV (c.f., Sec. 2.3), and LCH (c.f., Sec. 2.4) are examples of such approaches.
3. **Region-based.** It comprises methods that use segmentation to decompose images according to visual cues. The number of obtained regions as well as shape, size, and location vary from image to image. The objective is not to find and separate objects in the image but to find similar group of pixels. CBC (c.f., Sec. 2.5) is an example of such techniques.

2.1 Global Color Histogram (GCH)

The simplest approach to encode the information present in an image is the Global Color Histogram (GCH) [34]. A GCH is a set of ordered values, one for each distinct color, representing the probability of a pixel being of that color. Uniform quantization and normalization are used to reduce the number of distinct colors and to avoid scaling bias [34]. The L_1 (City-block) or L_2 (Euclidean) are the most used metrics for histogram comparison.

Histograms are effective for retrieval if there is uniqueness in the color pattern present in the images we want to compare. However, GCH can be sensitive to changes in viewpoint, occlusion, and illumination [36].

2.2 Border/Interior Classification (BIC)

Stehling et al. [33] have presented the border/interior pixel classification (BIC), a compact approach to describe images. BIC relies on the RGB color-space uniformly quantized in $4 \times 4 \times 4 = 64$ colors. After the quantization, the image pixels are classified as *border* or *interior*. A pixel is classified as *interior* if its 4-neighbors (top, bottom, left, and right) have the same quantized color. Otherwise, it is classified as *border*.

After the image pixels are classified, two color histograms are computed: one for border pixels and another for interior pixels. The two histograms are stored as single histogram with 128 bins. BIC compares the histograms using the $dLog$ distance function [33]

$$dLog(q, d) = \sum_{i=0}^{i < M} \|f(q[i]) - f(d[i])\| \quad (1)$$

$$f(x) = \begin{cases} 0, & \text{if } x = 0 \\ 1, & \text{if } 0 < x < 1 \\ \lceil \log_2 x \rceil + 1, & \text{otherwise} \end{cases} \quad (2)$$

where q and d are two histograms with M bins each. The value $q[i]$ represents the i^{th} bin of histogram q , and $d[i]$ represents the i^{th} bin of histogram d .

2.3 Color Coherence Vectors (CCVs)

Zabih et al. [26] have presented an approach to compare images based on color coherence vectors. They define color’s coherence as the degree to which pixels of that color are members of large similarly-colored regions. They refer to

these significant regions as coherent regions. Coherent pixels are part of some sizable contiguous region, while incoherent pixels are not.

In order to compute the CCVs, first the method blurs and discretizes the image’s color-space to eliminate small variations between neighboring pixels. Next, it finds the connected components in the image aiming to classify the pixels within a given color bucket as either coherent or incoherent.

CCV binary classification is based on a non-binary visual property of the images (the size of the connected components) and an empirical size threshold is needed. Hence the most of the useful information about the size of the connected components is lost in this reduction.

2.4 Local Color Histogram (LCH)

Tan et al. [20] have presented an approach based on local color histograms (LCH). This technique decomposes the image into equally-sized cells and individually describes each cell using a local color histogram.

The image contents are represented using a local color histogram matrix, one for each cell

$$h_{i,j,k} = \frac{a_{i,j,k}}{n} \quad (3)$$

where n is the number of image’s pixels, $a_{i,j}$ is a cell, and k is a quantized color. The LCH distance between two images is the difference of corresponding cell histograms using L_1 .

2.5 Color-based clustering (CBC)

Stehling et al. [32] have presented a region based approach to retrieve images named color-based clustering (CBC). This method decomposes the image into disjoint connected components. Each region presents a minimum size s_{min} and a maximum color dissimilarity d_{max} . Each region is defined in terms of the average color in the Lab space (L, a, b), its normalized horizontal and vertical center (h, v), and its size in pixels normalized with respect to the image’s size (s).

The L_2 Euclidean distance between two regions a_i of image A and b_j of image B , is

$$\mathcal{D}(a_i, b_j) = \alpha \times L_2^{color}(a_i, b_j) + (1 - \alpha) \times L_2^{center}(a_i, b_j) \quad (4)$$

where $L_2^{color}(a_i, b_j)$ considers the L, a, b color differences of images A and B . $L_2^{center}(a_i, b_j)$ considers the center differences. The parameter α measures which sum component needs to be most valued. The distance $d(A, B)$ of two images is the weighted distance $\mathcal{D}(a_i, b_j) \forall a_i, b_j \in A, B$. We have used IRM [16] for such computation.

3. THE CSIR FRAMEWORK

In presence of different viewpoint, occlusion, and illumination in broad-image domains, the direct use of color and texture descriptors can fail.

To address this problem, we present a new method for CBIR based on images’ discrete distributions of distinctive local representative features and color properties: CSIR. It is a region-based technique to retrieve images using color visual cues which are robust to pose, orientation, and scale changes. Our framework is based on three key steps: (1) feature region detection; (2) description; and (3) comparison metric; as we illustrate in Algorithm 1.

Our CSIR approach is different from previous literature [15,19], where the authors describe the images based on

Algorithm 1 The CSIR framework.

Require: Input image I ;1: **Feature region detection:** search for local scale and rotation invariant feature regions \mathcal{R} . \triangleright Sec. 3.12: **Description:** \triangleright Sec. 3.2i **Construct** a separate Gaussian pyramid \mathcal{G}_R , \mathcal{G}_G , and \mathcal{G}_B for each color channel (R,G,B) of the image I .ii **For each** feature region $r \in \mathcal{R}$

- **Extract** local scaled and oriented patches \mathcal{P} from the Gaussian color pyramids \mathcal{G}_R , \mathcal{G}_G , and \mathcal{G}_B .

- **For each** patch $p \in \mathcal{P}$
 - **Calculate** a local low-level color descriptor (e.g., BIC, GCH).

3: **Comparison metric:** use an appropriate metric to compare the images. \triangleright Sec. 3.3

histograms of gradient orientation and do not codify color information of the images. In other related work, the authors transform the image to an invariant color-space [1] while here we merge low-level information and local representative features using the image color-space.

3.1 Feature region detection

In this stage, we are interested in patterns that can be repeatedly found amongst similar images independent of some affine transformations (e.g., pose, orientation, and scale).

To find such regions, we use a feature region detector or operator [23]. The detectors provide regions which later we use as support regions to compute color descriptors.

Many different techniques for describing local image regions are available: the rotation invariant Harris points [23]; the rotation and scale invariant Harris-Laplace, Laplace-of-Gaussian [18], Difference-of-Gaussian [5, 19], and Hessian-Laplace regions [19, 22]. Furthermore, there are the affine transformations invariant Harris-Affine [24], and Hessian-Affine regions [24], among others.

Here, we use the Difference-of-Gaussian (DoG) operator. This idea was first proposed by Crowley and Parker [5]. We search for local extrema in the 3D scale-space representation of an image $I(x, y, \sigma)$ where σ is the scale. In this approach, we create a pyramid representation of an image using difference-of-Gaussian filters. We detect a feature point if a local 3D extremum is present and if its absolute value is higher than a threshold.

To build the scale-space pyramid, we successively smooth and sample the input image with a Gaussian kernel. We obtain the DoG representation by subtracting two successive smoothed images. The local 3D extrema in the pyramid representation determine the localization and the scale of the feature regions.

We have used the DoG operator for a number of reasons. First, it is very efficient: we build all DoG levels by us-

ing only smoothing and sub-sampling operations. Second, DoG operator provides a close approximation to the scale-normalized Laplace of Gaussian (LoG) regions [18]. This approximation is interesting because Mikolajczyk [22] have showed that the maxima and minima of LoG operator produces the most stable image features compared to a range of other possible operators, such as gradient, Hessian, or Harris corner. However, LoG is more computational intensive than DoG.

We perform the feature region detection in the luminance channel and do not codify color information. Our experiments have shown that two good choices are the V channel for HSV color representation or Y for YCbCr. In Figure 1, we present the result of this stage for an input image.




Figure 1: Feature regions detection.

Formally, let the scale space of an image be a function $L(x, y, \sigma)$, produced from the convolution of a Gaussian random variable $G(x, y, \sigma)$ with an input image $I(x, y)$

$$L(x, y, \sigma) = G(x, y, \sigma) * I(x, y), \quad (5)$$

where $*$ is the convolution operator in x and y , and

$$G(x, y, \sigma) = \frac{1}{2\pi\sigma^2} e^{-\frac{x^2+y^2}{2\sigma^2}}. \quad (6)$$

As proposed by Lowe [19], to find good representative and invariant regions in scale-space, we can search for local extrema in the DoG function convolved with the image, $D(x, y, \sigma)$. We can compute this step using two successive scales separated by a constant multiplicative k . The constant factor k is required for true scale invariance [18, 19].

$$\begin{aligned} D(x, y, \sigma) &= (G(x, y, k\sigma) - G(x, y, \sigma)) * I(x, y) \\ &= L(x, y, k\sigma) - L(x, y, \sigma). \end{aligned} \quad (7)$$

3.2 Description

In this stage, our objective is to describe each persistent region found in Stage 1. It is important that each region codifies color information that present repeatable patterns amongst similar images.

Stage 1 provides important information that enables us to retrieve similar images even when such images are slightly modified by some affine transformations. However, this approach still does not consider color information or even capture the color spatial distribution. In Stage 2, we use the Stage 1 resulting regions to find similar image color cues capturing their distribution and location in the images.

When we use color descriptors, we represent images' color patterns. Given a query image, we are interested in finding

similar images and we use repeatable color patterns to find the answers. CSIR approach introduces a new concept. Instead of using the color pattern analysis in the whole image (like previous approaches), we analyze color patterns in representative image’s regions. We show, experimentally, that analyzing only these regions, instead of the whole image, it is possible to improve effectiveness in CBIR tasks.

In order to describe the representative image’s regions, we construct a separate Gaussian pyramid $\mathcal{G}_R, \mathcal{G}_G, \mathcal{G}_B$ for each color channel (R,G,B) of the input image I . Figure 2 shows the RGB-composed resulting pyramid of an input image.



Figure 2: Six-octave RGB-composed resulting pyramid of the image in Figure 1(a). Each octave has 6 scales.

For each region of Stage 1, we extract a local scaled and oriented patch feature region from the Gaussian color pyramid. The patches capture the different illumination, viewpoint, and orientations of the image. Figure 3 shows the resulting patches for the input image in Figure 1(a). There are three patches for each scale, eight patches for each octave, from left to right, top to bottom.




Figure 3: Some resulting patches of the image in Figure 1(a).

Next, for each extracted patch, we calculate a local low-level image descriptor (e.g., BIC, GCH, CCV) that represents patches’ local color information. Figure 4 shows the resulting features for three patches for each scale, eight patches for each octave, from left to right, top to bottom. In this

case, we have used BIC to encode the low-level information on each patch. According to the BIC classification, white color represents *border* and black represents *interior*.




Figure 4: Some resulting low-level information patches of the image in Figure 1(a).

3.3 Comparison metric

CSIR method provides color and scale information regions that describe an image. The number of feature for each image is different. The more complex an image the more feature regions CSIR provides for it. Hence, we need to compare images with different number of features. For that, we model the image features as hyper points under an unknown distribution. Further, we use the Earth Mover’s Distance (EMD) metric to evaluate dissimilarity between two multi-dimensional distributions (image features). The advantage is that EMD “lifts” this distance from individual features to full distributions [27].

Intuitively, given two distributions \mathbf{B}_p and \mathbf{B}_q , we can view \mathbf{B}_p as a mass of earth properly spread in space, and \mathbf{B}_q as a collection of holes in that same space. Then, the EMD measures the least amount of work needed to fill the holes with earth. Here, a unit of work corresponds to transporting a unit of earth by a unit of ground distance.

EMD provides a way to compare images based on their discrete distributions of local features. Let $(\mathcal{X}, \mathcal{D})$ be a metric space, $\mathbf{B}_p, \mathbf{B}_q \subset \mathcal{X}$ be two equal-mass sets, and π be a matching between \mathbf{B}_p and \mathbf{B}_q . The EMD is the minimum possible cost of π and is defined as

$$EMD(\mathbf{B}_p, \mathbf{B}_q) = \min_{\pi: \mathbf{B}_p \rightarrow \mathbf{B}_q} \sum_{s \in \mathbf{B}_p} \mathcal{D}(s, \pi(s)). \quad (8)$$

The computation of \mathcal{D} is based on establishing the correspondence between two images’ unordered features. However, the complexity of finding the optimal exact correspondence between two equal-mass sets is cubic in the number of features per set. Hence, we have used a low-distortion EMD embedding [12] to reduce the problem of correspondence between sets of local features to an L_1 distance.

4. EXPERIMENTS

In this section, we compare our CSIR approach to the set of image descriptors described in Section 2. For each image, we compute the feature vector using the selected descriptors. We sort the resulting feature vectors using a proper comparison metric.

We show that our approach is more resilient to some affine transformations than previous approaches that use color information on the whole image. Further, we provide results that point out that CSIR is indeed suitable for CBIR tasks.

4.1 Methodology

In this work, we have used the *query-by-example* (QBE) paradigm [36]. In QBE, we give an image as a visual example to the system and we query for images that are similar to the given example. Clearly, the effectiveness of these systems is dependent on the properties of the example image.

In order to assess the system effectiveness, we have a database with reference models, a set of images that represent the queries, and a common metric to be used in the effectiveness retrieval assessment. Here, our reference models are equal to the query models, and we test all images in the database against the remaining images, one at a time.

To evaluate the descriptors we present in this paper, we have used two image databases described in the literature. To create a more realistic scenario, we have merged these two databases.

The first database is a selection of the Corel Photo Gallery and is the same as the reported in [33]. This database is highly heterogeneous and comprises images with different domains.

The second database is freely available¹ and comprises images with a common background and different viewpoint, occlusion, and illumination.

As our objective in this paper is to retrieve objects in a database, we have excluded images in both databases that do not represent an explicit object. The resulting combined database we use in the experiments comprises 1,320 broad-image domains spanned into 72 different object classes. Each class contains at least 5 images. We present some examples of the resulting database in Figure 5.



Figure 5: Resulting database. *Boats, Rodeo and Car* classes.

We use the $Precision \times Recall$ [31,36] metric to assess the retrieval effectiveness. *Precision* is the ratio of the number of

¹<http://www.mis.informatik.tu-darmstadt.de/Research/Projects/categorization/eth80-db.html>

relevant images retrieved to the total number of irrelevant and relevant images retrieved. *Recall* is the ratio of the number of relevant images retrieved to the total number of relevant images in the database.

Also, we have used some unique value measurements in the validation. The first one corresponds to the resulting precision when the number of retrieved images is enough to include all relevant images for a given query. This measurement is named *R-value* [4], hence p_R stands for the precision in this point.

We also evaluate the measurements p_{30} , r_{30} , p_{100} , and r_{100} . These values are estimatives of the number of retrieved images which a common user would assess in a practical retrieval system [4].

Finally, we considered the average value for three ($3P$) and eleven ($11P$) points in the precision/recall curve [4]. We obtain the $3P$ value averaging the precision through three predefined recall points (usually 20%, 50% e 80%). We obtain the $11P$ value averaging 11 predefined recall points (usually 0%, 10%, ..., 90%, 100%).

4.2 Overall results

Figure 6 shows the results for seven image descriptors [15, 20, 26, 32–34]. Here, CSIR is represented by two curves: CSIR_{BIC} and CSIR_{GCH} that uses BIC and GCH respectively as the local low-level image descriptor.




Figure 6: CSIR_{BIC}, GCH vs. existing approaches.

Table 1 presents the average result for seven unique value measures we use in this paper: $3P$, $11P$, p_{30} , r_{30} , p_{100} , r_{100} e p_R .

Table 1: Unique values measurements results.

	$3P$	$11P$	p_{30}	r_{30}	p_{100}	r_{100}	p_R
CSIR _{BIC}	.67	.58	.53	.66	.32	.90	.28
BIC	.49	.42	.39	.52	.23	.76	.17
CSIR _{GCH}	.44	.38	.37	.49	.21	.70	.13
CBC	.31	.27	.27	.38	.16	.58	.09
GCH	.27	.24	.23	.34	.14	.50	.10
LCH	.26	.23	.23	.34	.14	.51	.10
CCV	.26	.23	.23	.34	.14	.50	.10

In fact, the use of representative regions to better represent color cues in the image does improve the retrieval effectiveness for broad-domain images under different illumination, occlusion, and focus conditions as we see in Figure 6

and Table 1. The greater the value the better the descriptor. For instance, the CSIR_{BIC} is $\approx 37\%$ better than traditional BIC and CSIR_{GCH} is $\approx 63\%$ better than GCH.

4.3 Visual examples

In this section, we show two resulting queries Q_1 and Q_2 for CSIR_{BIC} and BIC. We show the query on top left and the resulting retrieved images sorted from left to right, top to bottom.

We show Q_1 and its top-11 results in Figures 7(a) and 7(b). The use of discrete distributions of distinctive color and scale image regions of CSIR_{BIC} yields better results than the BIC global analysis. For instance, BIC retrieves the non-relevant image R_9 .

We show Q_2 and its top-11 results in Figures 8(a) and 8(b). CSIR_{BIC} captures the variations in viewpoint, partial occlusion, and illumination. Note that BIC retrieves the non-relevant images R_2, R_4, R_5, R_7 and R_9 .

5. CONCLUSIONS

In this paper, we have presented CSIR: a new method for comparing images based on their discrete distributions of distinctive color and scale representative regions.

Our method is robust to viewpoint, occlusion, and illumination changes; it is invariant to image transformations such as rotation and translation; and it does not require any learning stage.

Our key contribution is that instead of using the color pattern analysis in the whole image (as previous approaches) we use distinctive color and scale representative patterns that can be repeatedly found amongst similar images independent of some affine transformations.

We have provided experiments showing that CSIR is suitable for CBIR tasks and that it provides good retrieval effectiveness.

Future work includes the evaluation of other feature region operators and low-level image descriptors to improve the image representation.

6. ACKNOWLEDGMENTS

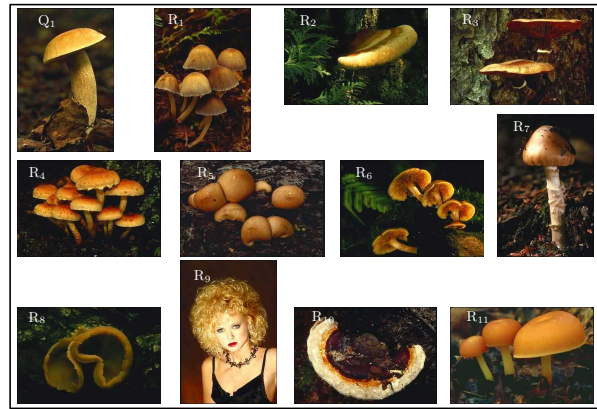
We thank the support of Fapesp (Grants 05/52959-3 and 05/58103-3), Capes (Grant 01P-05866/2007), CNPq (Grants 301278/2004, 311309/2006-2, and 477039/2006-5), and Microsoft EScience Project.

7. REFERENCES

- [1] A. E. Abdel-Hakim and A. A. Farag. CSIFT: A SIFT descriptor with color invariant characteristics. *CVPR*, 2:1978–1983, 2006.
- [2] N. Alajlan, M. S. Kamela, and G. Freeman. Multi-object image retrieval based on shape and topology. *Signal Processing: Image Communication*, 21(10):904–918, 2006.
- [3] E. Angelopoulou and L. B. Wolff. Sign of gaussian curvature from curve orientation in photometric space. *TPAMI*, 20(10):1056–1066, 1998.
- [4] R. A. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley, 1999.
- [5] J. Crowley and A. Parker. A Representation for Shape Based on Peaks and Ridges in the Difference of Low-pass Transform. *TPAMI*, 6(2):156–170, 1984.
- [6] L. M. K. et al. Fast texture database retrieval using extended fractal features. In *VI Storage and Retrieval for image and video databases*, volume 3312, pages 162–173, 1998.
- [7] L. F. Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *TPAMI*, 28(4):594–611, 2006.
- [8] D. A. Forsyth and M. M. Fleck. Automatic detection of human nudes. *IJCV*, 32(1):63–77, 1999.
- [9] X. Fu, Y. Li, R. Harrison, and S. Belkasim. Content-based image retrieval using gabor-zernike features. In *ICPR*, pages 417–420, 2006.
- [10] K. Grauman and T. Darrell. Efficient Image Matching with Distributions of Local Invariant Features. In *CVPR*, pages 627–634, 2005.
- [11] J. Huang, S. R. Kumar, M. Mitra, W.-J. Zhu, and R. Zabih. Spatial color indexing and applications. *IJCV*, 35(3):245–268, 1999.
- [12] P. Indyk and N. Thaper. Fast image retrieval via embeddings. In *3rd Intl. Workshop on Statistical and Computational Theories of Vision*, 2003.
- [13] J. Sivic and B. Russell and A. Efros and A. Zisserman and W. Freeman. Discovering objects and their location in images. In *ICCV*, pages 370–377, 2005.
- [14] K. Jarrah, M. Kyan, I. Lee, and L. Guan. Application of image visual characterization and soft feature selection in cbir. In *Multimedia Content Analysis, Management, and Retrieval*, volume 6073, 2006.
- [15] Y. Ke and R. Sukthankar. Pca-sift: A more distinctive representation for local image descriptors. In *CVPR*, volume 02, pages 506–513, 2004.
- [16] J. Li, J. Z. Wang, and G. Wiederhold. IRM: Integrated region matching for image retrieval. In *ACM Intl. Conf. on Multimedia*, pages 147–156, 2000.
- [17] Y. Li and L. G. Shapiro. Consistent line clusters for building recognition in cbir. In *16th ICPR*, volume 3, pages 30952–30957, 2002.
- [18] T. Lindeberg. Scale-space theory: a basic tool for analyzing structures at different scales. *J. of Applied Statistics*, 21(1):225–270, 1994.
- [19] D. G. Lowe. Distinctive image features from scale-invariant keypoints. *IJCV*, 60(2):91–110, 2004.
- [20] H. Lu, B. Ooi, and K. Tan. Efficient image retrieval by color contents. In *Intl. Conf. on Applications of Databases*, volume 819, pages 95–108, June 21–23 1994.
- [21] M. Marszałek and C. Schmid. Spatial Weighting for Bag-of-Features. In *CVPR*, pages 2118–2125, 2006.
- [22] K. Mikolajczyk and C. Schmid. Scale & Affine Invariant Interest Point Detectors. *IJCV*, 60(1):63–86, 2004.
- [23] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. *TPAMI*, 27(10):1615–1630, October 2005.
- [24] K. Mikolajczyk, T. Tuytelaars, C. Schmid, A. Zisserman, J. Matas, F. Schaffalitzky, T. Kadir, and L. V. Gool. A comparison of affine region detectors. *IJCV*, 65(1/2):43–72, 2005.
- [25] M. Mirmehdi and M. Petrou. Segmentation of color texture. *TPAMI*, 22(2):142–159, 2000.
- [26] G. Pass, R. Zabih, and J. Miller. Comparing images using color coherence vectors. In *ACM Multimedia*,




(a) CSIR_{BIC} retrieves all images correctly.




(b) BIC retrieves the non-relevant image R_9 .

Figure 7: Q_1 top-11 results.



(a) CSIR_{BIC} retrieves all images correctly.



(b) BIC retrieves the non-relevant images R_2 , R_4 , R_5 , R_7 , and R_9 .

Figure 8: Q_2 top-11 results.

pages 65–73, 1996.

- [27] Y. Rubner, C. Tomasi, and L. J. Guibas. The earth mover’s distance as a metric for image retrieval. *IJCV*, 40(2):99–121, 2000.
- [28] C. S. Sastry, A. K. Pujari, and B. L. Deekshatulu. A fourier-radial descriptor for invariant feature extraction. *IJWMIIP*, 4(1):197–212, 2006.
- [29] B. Shah, V. Raghavan, P. Dhatric, and X. Zhao. A cluster-based approach for efficient cbir using a similarity-preserving space transformation method. *JASIST*, 57(12):1694–1707, 2006.
- [30] C. Shimid and R. Mohr. Local grayvalue invariants for image retrieval. *TPAMI*, 19(5):530–535, 1997.
- [31] A. W. M. Smeulders, M. Worring, S. Santine, A. Gupta, and R. Jain. Cbir at the end of early years. *TPAMI*, 22(12):1349–1380, 2000.
- [32] R. O. Stehling, M. A. Nascimento, and A. X. Falcão. An adaptive and efficient clustering-based approach for content-based image retrieval in image databases.

In *IEEE Intl. Database Engineering and Applications Symposium*, pages 356–365, July 2001.

- [33] R. O. Stehling, M. A. Nascimento, and A. X. Falcão. A compact and efficient image retrieval approach based on border/interior pixel classification. In *11th CIKM*, pages 102–109, 2002.
- [34] M. J. Swain and B. H. Ballard. Color indexing. *IJCV*, 7(1):11–32, 1991.
- [35] P. M. Tardiff and A. Zaccarin. Multiscale autoregressive image representation for texture segmentation. In *VIII Image Processing*, volume 3026, pages 327–337, 1997.
- [36] R. S. Torres and A. X. Falcão. Cbir: Theory and applications. *RITA*, 13(2):161–185, 2006.
- [37] B. Wang and J. A. Bangham. Shape retrieval using matching pursuit decomposition. In *IEEE AVSS*, pages 98–104, 2006.