# Features for image retrieval: an experimental comparison

Thomas Deselaers · Daniel Keysers · Hermann Ney

Received: 19 May 2007/Accepted: 28 November 2007/Published online: 11 December 2007 © Springer Science+Business Media, LLC 2007

**Abstract** An experimental comparison of a large number of different image descriptors for content-based image retrieval is presented. Many of the papers describing new techniques and descriptors for content-based image retrieval describe their newly proposed methods as most appropriate without giving an in-depth comparison with all methods that were proposed earlier. In this paper, we first give an overview of a large variety of features for content-based image retrieval and compare them quantitatively on four different tasks: stock photo retrieval, personal photo collection retrieval, building retrieval, and medical image retrieval. For the experiments, five different, publicly available image databases are used and the retrieval performance of the features is analyzed in detail. This allows for a direct comparison of all features considered in this work and furthermore will allow a comparison of newly proposed features to these in the future. Additionally, the correlation of the features is analyzed, which opens the way for a simple and intuitive method to find an initial set of suitable features for a new task. The article concludes with recommendations which features perform well for what type of data. Interestingly, the often used, but very simple, color histogram performs well in the comparison and thus can be recommended as a simple baseline for many applications.

Keywords Image retrieval · Features · Image classification · Quantitative comparison

D. Keysers

T. Deselaers (⊠) · H. Ney

Human Language Technology and Pattern Recognition, Computer Science Department, RWTH Aachen University, Aachen, Germany e-mail: deselaers@cs.rwth-aachen.de

H. Ney e-mail: ney@cs.rwth-aachen.de

Image Understanding and Pattern Recognition, German Research Center for Artificial Intelligence (DFKI), Kaiserslautern, Germany e-mail: daniel.keysers@dfki.de

## 1 Introduction

Image retrieval in general and content-based image retrieval (CBIR) in particular are wellknown fields of research in information management in which a large number of methods have been proposed and investigated but in which still no satisfying general solutions exist. The need for adequate solutions is growing due to the increasing amount of digitally produced images in areas like journalism, medicine, and private life, requiring new ways of accessing images. For example, medical doctors have to access large amounts of images daily (Müller et al. 2004), home-users often have image databases of thousands of images (Sun et al. 2002), and journalists also need to search for images by various criteria (Markkula and Sormunen 1998; Armitage and Enser 1997). In the past, several CBIR systems have been proposed and all these systems have one thing in common: images are represented by numeric values, called features or descriptors, that are meant to represent the properties of the images to allow meaningful retrieval for the user.

Only recently have some standard benchmark databases and evaluation campaigns been created which allow for a quantitative comparison of CBIR systems. These benchmarks allow for the comparison of image retrieval systems under different aspects: usability and user interfaces, combination with text retrieval, or overall performance of a system. However, to our knowledge, *no quantitative comparison of the building blocks of the systems, the features that are used to compare images, has been presented so far.* In (Shirahatti and Barnard 2005) a method for comparing image retrieval systems was proposed relying on the Corel database, which has restricted copyrights, is no longer commercially available today, and can therefore not be used for experiments that are meant to be a basis for other comparisons.

Another aspect of evaluating CBIR systems are the requirements of the users. In (Markkula and Sormunen 1998) and (Armitage and Enser 1997) studies of user needs in searching image archives are presented and the outcome in both studies is that CBIR alone is very unlikely to fulfill the needs but that semantic information obtained from meta data and textual information is an important additional knowledge source. Although today the semantic analysis and understanding of images is much further developed due to the recent achievements in object detection and recognition, still most of the requirements specified are not satisfiable fully automatically. Therefore, in this paper *we compare the performance of a large variety of visual descriptors*. These can then later be combined with the outcome of textual information retrieval as described e.g., in (Deselaers et al. 2006).

The main question we address in this paper is: Which features are suitable for which task in image retrieval? This question is thoroughly investigated by examining the performance of a wide variety of different visual descriptors for four different types of CBIR tasks.

The question of which features perform how well is closely related to the question which features can be combined to obtain good results in a particular task. Although we do not directly address this question here, the results from this paper lead to a new and intuitive method to choose an appropriate combination of features based on the correlation of the individual features.

For the evaluation of the features we use five different publicly available databases which are a good starting point to evaluate the performance of new image descriptors.

Although today various initiatives for evaluation of CBIR systems have evolved, only few of them resulted in evaluation campaigns with participants and results: *Benchathlon*<sup>1</sup>

<sup>&</sup>lt;sup>1</sup> http://www.benchathlon.net/

was started in 2001 and located at the SPIE Electronic Imaging conference but has become smaller over time.  $TRECVID^2$  is an initiative by the TREC (Text Retrieval Conference) on video retrieval in which video retrieval systems are compared.  $ImageCLEF^3$  is part of the Cross-Language Evaluation Framework (CLEF) and started in 2003 with only one task aiming at a combination of multi-lingual information retrieval with CBIR. In 2004, it comprised three tasks, one of them focused on visual queries and in 2005 and 2006 there were four tasks, one and two of them purely visual, respectively. We can observe that evaluation in the field of CBIR is at a far earlier stage than it is in textual information retrieval (e.g., Text REtrieval Conference, TREC) or in speech recognition (e.g., Hub4-DARPA evaluation). One reason for this is likely to be the smaller commercial impact that (content-based) image retrieval has had in the past. However, with the increasing amount of visual data available in various form, this is likely to change in the future.

The main contributions of this paper are answers to the questions above, namely

- an extensive overview of features proposed for CBIR, including features that were
  proposed in the early days of CBIR and techniques that were proposed only recently in
  the object recognition and image understanding literature as well as a subset of features
  from the MPEG7 standard.
- a quantitative analysis of the performance of these features for various CBIR tasks (in particular: stock photo retrieval, personal photo retrieval, building/touristic image retrieval, and medical image retrieval).
- pointing out a set of five databases from four different domains that can be used for benchmarking CBIR systems.

Note that we do not focus on the combination of features nor on the use of user feedback for content-based image retrieval in this paper; several other authors propose and evaluate approaches to these important issues (Yavlinski et al. 2004; Heesch and Rüger 2003; Müller et al. 2000; Müller et al. 2000; MacArthur et al. 2000). Instead, we mainly investigate the performance of single features for different tasks.

#### 1.1 State of the art in content-based image retrieval

This section gives an overview on literature on CBIR. We mainly focus on different descriptors and image representations. More general overviews on CBIR are given in (Smeulders et al. 2000; Forsyth and Ponce 2002; Rui et al. 1999). Two recent reviews of CBIR techniques are given in (Datta et al. 2005; Lew et al. 2006).

In CBIR, there are, roughly speaking, two different main approaches: a *discrete approach* and a *continuous approach* (de Vries and Westerveld 2004). (1) The discrete approach is inspired by textual information retrieval and uses techniques like inverted files and text retrieval metrics. This approach requires all features to be mapped to binary features; the presence of a certain image feature is treated like the presence of a word in a text document. (2) The continuous approach is similar to nearest neighbor classification. Each image is represented by a feature vector and these features are compared using various distance measures. The images with lowest distances are ranked highest in the

<sup>&</sup>lt;sup>2</sup> http://www.nlpir.nist.gov/projects/trecvid/

<sup>&</sup>lt;sup>3</sup> http://www.imageclef.org

retrieval process. A first, though not exhaustive, comparison of these two models is presented in (de Vries and Westerveld 2004).

Among the first systems that were available were the QBIC system from IBM (Faloutsos et al. 1994) and the Photobook system from MIT (Pentland et al. 1996). OBIC uses color histograms, a moment based shape feature, and a texture descriptor. Photobook uses appearance features, texture features, and 2D shape features. Another well known system is Blobworld (Carson et al. 2002), developed at UC Berkeley. In Blobworld, images are represented by regions that are found in an Expectation-Maximization-like (EM) segmentation process. In these systems, images are retrieved in a nearest-neighborlike manner, following the continuous approach to CBIR. Other systems following this approach include SIMBA (Siggelkow et al. 2001), CIRES (Iqbal and Aggarwal 2002), SIMPLIcity (Wang et al. 2001), IRMA (Lehmann et al. 2005), and our own system FIRE (Deselaers et al. 2005; Deselaers et al. 2004). The Moving Picture Experts Group (MPEG) defines a standard for content-based access to multimedia data in their MPEG-7 standard. In this standard, a set of descriptors for images is defined. A reference implementation for these descriptors is given in the XM Software.<sup>4</sup> A system that uses MPEG-7 features in combination with semantic web ontologies is presented in Bloehdorn et al. (2005). In Di et al. (2002) a method starting from low-level features and creating a semantic representation of the images is presented and in Meghini et al. (2001) an approach to consistently fuse the efforts in various fields of multimedia information retrieval is presented.

In (Squire et al. 1999), the VIPER system is presented which follows the discrete approach. VIPER is now publicly available as the GNU Image Finding Tool (GIFT) and several enhancements have been implemented during the last years. An advantage of the discrete approach is that methods from textual information retrieval can easily be transferred as e.g., user interaction and storage handling. Nonetheless, most image retrieval systems follow the continuous approach often using some optimization, for example pre-filtering and pre-classification (Smeulders et al. 2000; Wang et al. 2001; Park et al. 2002), to achieve better runtime performance, e.g., (Faloutsos et al. 1994; Pentland et al. 1996; Carson et al. 2002; Siggelkow et al. 2001).

We can clearly observe that many different image description features have been developed. However, only few works have quantitatively compared different features. Interesting insights can also be gained from the outcomes of the ImageCLEF image retrieval evaluations (Clough et al. 2004; Clough et al. 2006) in which different systems are compared on the same task. The comparison is not easy because all groups use different retrieval systems and text-based information retrieval is an important part of these evaluations. Due to the lack of standard tasks, in many papers on image retrieval, new benchmark sets are defined to allow for quantitative comparison of the proposed methods to a baseline system. A problem with this approach is that it is simple to create a benchmark for which you can show improved results (Müller et al. 2002).

Recently, local image descriptors are getting more attention within the computer vision community. The underlying idea is that objects in images consist of parts that can be modelled with varying degrees of independence. These approaches are successfully used for object recognition and detection (Dorkó 2006; Fei-Fei and Perona 2005; Fergus et al. 2003; Opelt et al. 2006; Marée et al. 2005; Deselaers et al. 2005) and CBIR (Deselaers et al. 2004; Jain 2004; Schmid and Mohr 1997; van Gool et al. 2001). For the representation of local image parts, SIFT features (Lowe 2004) and raw image patches are commonly used and a bag-of-features approach, similar to the bag-of-words approach in natural language

<sup>&</sup>lt;sup>4</sup> http://www.lis.ei.tum.de/research/bv/topics/mmdb/e\_mpeg7.html

processing, is commonly taken. The features described in Section 3.7 also follow this approach and are strongly related to the modern approaches in object recognition. In contrast to the methods described above, the image is not modelled as a whole but rather image parts are modelled individually. Most approaches found in the literature on part-based object recognition learn (often complicated) models from a large set of training data. This approach is impractical for CBIR applications since it would require an enormous amount of training data on the one hand and would lead to tremendous computing times to create these models on the other hand. However, some of these approaches are applicable for limited domain retrieval, e.g., on the IRMA database (cf. Section 5.3) (Deselaers et al. 2006).

*Overview.* The remainder of this paper is structured as follows. The next section describes the retrieval metric used to rank images given a feature and a distance measure and the performance measures used to compare different settings. Section 3 gives an overview of 19 different image descriptors and distance measures which are used for the experiments. Section 4 presents a method to analyze the correlation of different image descriptor/distance combinations. In Section 5, five different benchmark databases are described that are used for the experiments presented in Section 6. The experimental section is subdivided into three parts: Section 6.1 directly compares the performance of the different methods for the different tasks, Section 6.2 describes the results of the correlation analysis, and Section 6.3 analyzes the connection between the error rate and the mean average precision. The paper concludes with answers to the questions posed above.

#### 2 Retrieval metric

The CBIR framework used to conduct the experiments described here follows the continuous approach: images are represented by vectors that are compared using distance measures. For the experiments we use our CBIR system FIRE.<sup>5</sup> FIRE was designed as a research system with extensibility and flexibility in mind. For the evaluation of features, only one feature and one query image is used at a time, as described in the following.

*Retrieval Metric.* Let the database  $\{x_1, \ldots, x_n, \ldots, x_N\}$  be a set of images represented by features. To retrieve images similar to a query image q, each database image  $x_n$  is compared with the query image using an appropriate distance function  $d(q, x_n)$ . Then, the database images are sorted according to the distances such that  $d(q, x_{n_i}) \le d(q, x_{n_{i+1}})$  holds for each pair of images  $x_{n_i}$  and  $x_{n_{i+1}}$  in the sequence  $(x_{n_1} \ldots, x_{n_i}, \ldots, x_{n_N})$ . If a combination of different features is used, the distances are normalized to be in the same value range and then a linear combination of the distances is used to create the ranking.

To evaluate CBIR, several performance evaluation measures have been proposed (Müller et al. 2001) based on the precision P and the recall R:

$$P = \frac{\text{Number of relevant images retrieved}}{\text{Total number of images retrieved}},$$
$$R = \frac{\text{Number of relevant images retrieved}}{\text{Total number of relevant images}},$$

Precision and recall values are usually represented in a precision-recall-graph  $R \rightarrow P(R)$  summarizing (R, P(R)) pairs for varying numbers of retrieved images. The most common

<sup>&</sup>lt;sup>5</sup> freely available under the terms of the GNU General Public Licencse at http://www-i6.informatik. rwth-aachen.de/~deselaers/fire.htm

way to summarize this graph into one value is the mean average precision that is also used e.g., in the TREC and CLEF evaluations. The average precision AP for a single query q is the mean over the precision scores after each retrieved relevant item:

$$AP(q) = \frac{1}{N_R} \sum_{n=1}^{N_R} P_q(R_n),$$

where  $R_n$  is the recall after the *n*th relevant image was retrieved.  $N_R$  is the total number of relevant documents for the query. The mean average precision *MAP* is the mean of the average precision scores over all queries:

$$MAP = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} AP(q),$$

where Q is the set of queries q.

An advantage of the mean average precision is that it contains both precision and recall oriented aspects and is sensitive to the entire ranking.

We also indicate the classification error rate ER for all experiments. To do so we consider only the most similar image according to the applied distance function. We consider a query image to be classified correctly, if the first retrieved image is relevant. Otherwise the query is misclassified:

$$ER = \frac{1}{|Q|} \sum_{a \in Q} \begin{cases} 0 & \text{if the most similar image is relevant/from the correct class} \\ 1 & \text{otherwise.} \end{cases}$$

This is in particular interesting if the database for retrieval consists of images labelled with classes, which is the case for some of the databases considered in this paper. For databases without defined classes but with selected query images and corresponding relevant images, the classes to be distinguished are "*relevant*" and "*irrelevant*" only.

This is in accordance with precision at document X being used as an additional performance measure in many information retrieval evaluations. The ER used here is equal to 1 - P(1), where P(1) is the precision after one document retrieved. In (Deselaers et al. 2004) it was experimentally shown that the error rate and P(50), the precision after 50 documents, are correlated with a coefficient of 0.96 and thus they essentially describe the same property. The precision oriented evaluation is interesting, because most search engines, both for images and text, return between 10 and 50 results, given a query.

Using the ER, the image retrieval system can be viewed as a nearest neighbor classifier using the same features and the same distance function as the image retrieval system. The decision rule of this classifier can be written in the form

$$q \rightarrow r(q) = \arg \min_{k=1,\dots,K} \{\min_{n=1,\dots,N_k} d(q, x_{nk})\}.$$

The query image q is predicted to be from the same class as the database image that has the smallest distance to it. Here,  $x_{nk}$  denotes the *n*-th image of class k.

#### **3** Features for CBIR

In this section we give an overview of the features tested, with the intention to include as many features as possible. Obviously we cannot cover all features that have been proposed in the literature. For example, we have left out the Blobworld features (Carson et al. 2002)

because for comparing images based on these features, user interaction to select the relevant regions in the query image is required. Furthermore, a variety of texture representations have not been included and we have not investigated different color spaces.

However, we have tried to make the selection of features as representative and at the state-of-the-art as possible. Roughly speaking, the features can be grouped into the following types: (a) color representation, (b) texture representation, (c) local features, and (d) shape representation.<sup>6</sup> The features that are presented in the following are grouped according to these four categories in Table 1. Table 1 also gives the timing information on feature extraction and retrieval time for a database consisting of 10 images.<sup>7</sup>

The distance function used to compare the features representing an image obviously also has a big influence on the performance of the system. Therefore, we refer to the used distance functions for each feature in the particular sections. We have chosen distance functions that are known to work well for the features used as the discussion of their influence is beyond the scope of this paper. Different comparison measures for histograms are presented e.g., in (Puzicha et al. 1999; Nölle 2003) and dissimilarity metrics for direct image comparison are presented in Keysers et al. (2007).

#### 3.1 Appearance-based image features

The most straight-forward approach is to directly use the pixel values of the images as features: the images are scaled to a common size and compared using the Euclidean distance. In this work, we have used a  $32 \times 32$  down-sampled representation of the images and these have been compared using the Euclidean distance. It has been observed that for classification and retrieval of medical radiographs, this method serves as a reasonable baseline (Keysers et al. 2007).

In Keysers et al. (2007) different methods were proposed to directly compare images accounting for local deformations. The proposed *image distortion model* (IDM) is shown to be a very effective means of comparing images with reasonable computing time. IDM clearly outperforms the Euclidean distance for optical character recognition and medical radiographs. The IDM is a non-linear deformation model, it was also successfully used to compare general photographs (Deselaers 2003) and for sign language and gesture recognition (Zahedi et al. 2005). In this work it is used as a second comparison measure to compare images directly. Therefore the images are scaled to have a common width of 32 pixels while keeping the aspect ratio constant, i.e., the images may be of different heights.

<sup>&</sup>lt;sup>6</sup> Note that no features that fully cover the shapes in the images are included since therefore an algorithm segmenting the images into meaningful regions is required, but since fully-automatic segmentation for general images is an unsolved problem, it is not covered here. The features that we mark to represent shape only represent shape in a local (for the SIFT features) and very rough global context (for appearance-based image features). There are however, overview papers on the shape features defined in MPEG7 which use databases consisting of segmented images for benchmarks (Bober 2001).

<sup>&</sup>lt;sup>7</sup> These experiments have been carried out on a 1.8 GHz machine with our standard C++ implementation of the software. The SIFT feature extraction was done with the software from Gyuri Dorko (http://www.lear.inrialpes.fr/people/dorko/downloads.html), the MPEG7 experiments were performed with the MPEG7 XM reference implementation (http://www.lis.ei.tum.de/research/bv/topics/mmdb/mpeg7. html), and the downscaling of images was performed using the ImageMagick library (http://www.imagemagick.org/). The timings include the time to load all data and initialize the system.

Section	Comp. measure	Туре	Extr.[s]	Retr.[s]
ures				
3.1	Euclidean	abcd	0.25	0.19
3.1	IDM	abcd	0.25	9.72
3.2	JSD	а	0.77	0.16
3.3	JSD	b	14.24	0.13
3.4	Euclidean	b	3.51	0.16
3.5	JSD	b	8.01	0.12
3.5	Euclidean	b	8.68	0.17
3.6	JSD	ab	28.93	0.16
3.6	JSD	ab	18.23	0.14
3.7	-	ac	4.69	7.13
3.7	JSD	ac	4.69 + 5.17	0.27
3.7	EMD	ac	4.69 + 3.37	0.55
3.7	-	cd	11.91	9.23
3.7	JSD	cd	11.91 + 6.23	0.27
3.7	EMD	cd	11.91 + 4.50	1.03
3.8.1	MPEG7-internal	а	0.48	0.42
3.8.2	MPEG7-internal	ad	0.20	0.33
3.8.3	MPEG7-internal	b	0.16	0.43
	Section           ures           3.1           3.2           3.3           3.4           3.5           3.6           3.6           3.7           3.7           3.7           3.7           3.7           3.7           3.7           3.7           3.7           3.7           3.8.1           3.8.2           3.8.3	Section         Comp. measure           ures         3.1         Euclidean           3.1         IDM         3.2           3.2         JSD         3.3           3.4         Euclidean         3.5           3.5         JSD         3.5           3.6         JSD         3.6           3.7         -         3.7           3.7         EMD         3.7           3.7         JSD         3.7           3.7         EMD         3.7           3.7         JSD         3.7           3.7         MPEG7-internal         3.8.1           3.8.2         MPEG7-internal         3.8.3	SectionComp. measureTypeures $3.1$ Euclideanabcd $3.1$ IDMabcd $3.2$ JSDa $3.3$ JSDb $3.4$ Euclideanb $3.5$ JSDb $3.6$ JSDab $3.6$ JSDab $3.6$ JSDab $3.7$ -ac $3.7$ JSDac $3.7$ JSDac $3.7$ EMDac $3.7$ JSDcd $3.7$ JSDcd $3.7$ JSDcd $3.7$ JSDcd $3.7$ JSDcd $3.8.1$ MPEG7-internala $3.8.2$ MPEG7-internalad $3.8.3$ MPEG7-internalb	SectionComp. measureTypeExtr.[s]urres $3.1$ Euclideanabcd $0.25$ $3.1$ IDMabcd $0.25$ $3.2$ JSDa $0.77$ $3.3$ JSDb $14.24$ $3.4$ Euclideanb $3.51$ $3.5$ JSDb $8.01$ $3.5$ JSDb $8.68$ $3.6$ JSDab $28.93$ $3.6$ JSDab $18.23$ $3.7$ -ac $4.69$ $3.7$ JSDac $4.69 + 5.17$ $3.7$ EMDac $4.69 + 3.37$ $3.7$ EMDcd $11.91 + 6.23$ $3.7$ EMDcd $11.91 + 4.50$ $3.8.1$ MPEG7-internala $0.48$ $3.8.2$ MPEG7-internalad $0.20$ $3.8.3$ MPEG7-internalb $0.16$

Table 1 Grouping of the features into different types

(a) color representation, (b) texture representation, (c) local features, (d) shape representation. The table also gives the time to extract the features from 10 images and to query 10 images in a 10 image database to give an impression of the computational costs of the different features (experiments were performed on a 1.8 GHz machine)

## 3.2 Color histograms

Color histograms are among the most basic approaches and widely used in image retrieval (Smeulders et al. 2000; Faloutsos et al. 1994; Deselaers 2003; Puzicha et al. 1999; Swain and Ballard 1991). To show performance improvements in image retrieval systems, systems using only color histograms are often used as a baseline. The color space is partitioned and for each partition the pixels with a color within its range are counted, resulting in a representation of the relative frequencies of the occurring colors. We use the RGB color space for the histograms. We observed only minor differences with other color spaces which was also observed in (Smith and Chang 1996). In accordance with (Puzicha et al. 1999), we use the Jeffrey divergence or Jensen-Shannon divergence (JSD) to compare histograms:

$$d_{JSD}(H,H') = \sum_{m=1}^{M} H_m \log \frac{2H_m}{H_m + H'_m} + H'_m \log \frac{2H'_m}{H'_m + H_m},$$

where H and H' are the histograms to be compared and  $H_m$  is the *m*th bin of H.

#### 3.3 Tamura features

In Tamura et al. (1978) the authors propose six texture features corresponding to human visual perception: *coarseness, contrast, directionality, line-likeness, regularity,* and *roughness.* From experiments testing the significance of these features with respect to human perception, it was concluded that the first three features are very important. Thus, in our experiments we use coarseness, contrast, and directionality to create a histogram describing the texture (Deselaers 2003) and compare these histograms using the Jeffrey divergence (Puzicha et al. 1999). In the QBIC system (Faloutsos et al. 1994) histograms of these features are used as well.

### 3.4 Global texture descriptor

In Deselaers (2003) a texture feature consisting of several parts is described: *Fractal dimension* measures the roughness of a surface. The fractal dimension is calculated using the reticular cell counting method (Haberäcker 1995). *Coarseness* characterizes the grain size of an image. It is calculated depending on the variance of the image. *Entropy* of pixel values is used as a measure of disorderedness in an image. The *spatial gray-level difference statistics* describe the brightness relationship of pixels within neighborhoods. It is also known as co-occurrence matrix analysis (Haralick et al. 1973). The *circular Moran autocorrelation function* measures the roughness of the texture. For the calculation a set of autocorrelation functions is used (Gu et al. 1989). From these, we obtain a 43 dimensional vector consisting of one value for the fractal dimension, one value for the coarseness, one value for the entropy and 32 values for the difference statistics, and 8 values for the circular Moran autocorrelation function. This descriptor has been successfully used for medical images in Lehmann et al. (2005).

## 3.5 Gabor features

Gabor features have been widely used for Texture analysis (Park et al. 2002; Squire et al. 1999). Here we use two different descriptors derived from Gabor features:

- Mean and standard deviation: Gabor features are extracted at different scales and directions from the images and the mean and standard deviation of the filter responses is calculated. We extract Gabor features in five different orientations and five different scales leading to a 50 dimensional vector.
- A bank of 12 different circularly symmetric Gabor filters is applied to the image, the energy for each filter on the bank is quantized into 10 bands and a histogram of the mean filter outputs over image regions is computed to give a global measure of the texture characteristics of the image (Squire et al. 1999). These histograms are compared using the JSD.

#### 3.6 Invariant feature histograms

A feature is called invariant with respect to certain transformations if it does not change when these transformations are applied to the image. The transformations considered here are translation, rotation, and scaling. In this work, invariant feature histograms as presented in (Siggelkow 2002) are used. These features are based on the idea of constructing invariant features by integration, i.e., a certain feature function is integrated over the set of all considered transformations. The feature functions we have considered are monomial and relational functions (Siggelkow et al. 2001) over the pixel intensities. Instead of summing over translation and rotation, we only sum over rotation and create a histogram over translation. This histogram is still invariant with respect to rotation and translation. The resulting histograms are compared using the JSD. Previous experiments have shown that the characteristics of invariant feature histograms and color histograms are very similar and that invariant feature histograms can sometimes outperform color histograms (Deselaers et al. 2004).

## 3.7 Local image descriptors

Image patches, i.e., small subimages of images, or features derived thereof currently are a very promising approach for object recognition, e.g., (Deselaers et al. 2005; Fergus et al. 2005; Paredes et al. 2001). Obviously, object recognition and CBIR are closely related fields (Vailaya et al. 2001; Antani et al. 2002) and for some clearly defined retrieval tasks, object recognition methods might actually be the only possible solution: e.g., looking for all images showing a certain person, clearly a face detection and recognition system would deliver the best results (Pentland et al. 1996; Deselaers et al. 2005).

We consider two different types of local image descriptors or local features (LF): (a) patches that are extracted from the images at salient points and dimensionality reduced using PCA transformation (Deselaers et al. 2005) and (b) SIFT descriptors (Lowe 2004) extracted at Harris interest points (Dorkó 2006, chapters 3, 4).

We employ three methods to incorporate local features into our image retrieval system. The methods are evaluated for both types of local features described above:

*LF histograms.* The first method follows (Deselaers et al. 2005): local features are extracted from all database images and jointly clustered to form 2,048 clusters. Then for each of the local features all information except the identifier of the most similar cluster center is discarded and for each image a histogram of the occurring patch-cluster identifiers is created, resulting in a 2,048 dimensional histogram per image. These histograms are then used as features in the retrieval process and are compared using the Jeffrey divergence. This method was shown to produce good performance in object recognition and detection tasks (Deselaers et al. 2005). Note that the timing information in Table 1 does not give the time to create the cluster model, since this is only done once for a database and can be computed offline.

*LF signatures.* The second method is derived from the method proposed in (Mikolajczyk et al. 2005). Local features are extracted from each database image and clustered for each image separately to form 32 clusters per image. Then for each image, the parameters of the clusters, i.e., the mean and the variance, are saved and the according cluster-identifier histogram of the extracted features is created. These "local feature signatures" are then used as features in the retrieval process and are compared using the Earth Mover's Distance (EMD) (Rubner et al. 1998). This method was shown to produce good performance in object recognition and detection tasks (Mikolajczyk et al. 2005).

*LF global search.* The third method is based on global patch search and derived from the method presented in (Paredes et al. 2001). Here, local features are extracted from all database images and stored in a KD tree to allow for efficient nearest neighbor searching.

Given a query image, we extract local features from the image in the same way as we did for the database images and search for the k nearest neighbors for each of the query-patches in the set of database-patches. Then, we count how many patches from each of the database image were found for the query patches and the database images with the highest number of patch-hits are returned. Note that the timing information in Table 1 does not include the time to create the KD tree, since this is only done once for a database and can be computed offline.

# 3.8 MPEG-7 features

The Moving Picture Experts Group (MPEG) has defined several visual descriptors in their standard referred to as *MPEG-7 standard*.<sup>8</sup> An overview of these features can be found in (Eidenberger 2003; Manjunath et al. 2001; Ohm 2001; Yang and Kuo 1999). The MPEG initiative focuses strongly on features that are computationally inexpensive to obtain and to compare and also strongly optimizes the features with respect to the required memory for storage.

Coordinated by the MPEG, a reference implementation of this standard has been developed.<sup>9</sup> This reference implementation was used in our framework for experiments with these features. Unfortunately, the software is not yet in a fully functional state and thus only three MPEG7 features could be used in the experiments. For each of these features, we use the comparison measures proposed by the MPEG standard and implemented in the reference implementation. The feature types are briefly described in the following:

## 3.8.1 MPEG 7: scalable color descriptor

The *scalable color descriptor* is a color histogram in the HSV color space that is encoded by a Haar transform. Its binary representation is scalable in terms of bin numbers and bit representation accuracy over a broad range of data rates. Retrieval accuracy increases with the number of bits used in the representation. We use the default setting of 64 coefficients.

# 3.8.2 MPEG 7: color layout descriptor

This descriptor effectively represents the spatial distribution of the color of visual signals in a very compact form. This compactness allows visual signal matching functionality with high retrieval efficiency at very small computational costs. It allows for query-by-sketch queries because the descriptor captures the layout information of color features. This is a clear advantage over other color descriptors. This approach closely resembles the use of very small thumbnails of the images with a quantization of the colors used.

## 3.8.3 MPEG 7: edge histogram

The *edge histogram descriptor* represents the spatial distribution of five types of edges, namely four directional edges and one non-directional edge. According to the MPEG-7

<sup>&</sup>lt;sup>8</sup> http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.html

<sup>&</sup>lt;sup>9</sup> http://www.lis.e-technik.tu-muenchen.de/research/bv/topics/mmdb/e\_mpeg7.html

standard, the image retrieval performance can be significantly improved if the edge histogram descriptor is combined with other descriptors such as the color histogram descriptor. The descriptor is scale invariant and supports rotation invariant and rotation sensitive matching operations.

#### 4 Correlation analysis of features for CBIR

After discussing various features, now let us assume that a set of features is given, some of which account for color, others accounting for texture, and maybe others accounting for shape. A very interesting question then is, how features that can be used in combination can be chosen. Automatic methods for feature selection have e.g., been proposed in (Vasconcelos and Vasconcelos 2004; Najjar et al. 2003). These automatic methods, however do not directly explain why features are chosen, are difficult to manipulate from a user's perspective, and normally require labelled training data.

The method proposed here does not require training data but only analyzes the correlations between the features themselves, and instead of automatically selecting a set of features it provides the user with information helping to select an appropriate set of features.

To analyze the correlation between different features, we analyze the correlation between the distances d(q, X) obtained for each feature of each of the images X from the database given a query q. For each pair of query image q and database image X we create a vector  $(d_1(q, X), d_2(q, X), ..., d_m(q, X), ..., d_M(q, X))$  where  $d_m(q, X)$  is the distance of the query image q to the database image X for the mth feature. Then we calculate the correlation between the  $d_m$  over all  $q \in \{q_1, ..., q_l, ..., q_L\}$  and all  $X \in \{X_1, ..., X_n, ..., X_N\}$ .

The M × M covariance matrix  $\Sigma$  of the  $d_m$  is calculated over all N database images and all L query images as:

$$\Sigma_{ij} = \frac{1}{NL} \sum_{n=1}^{N} \sum_{l=1}^{L} (d_i(q_l, X_n) - \mu_i) \cdot (d_j(q_l, X_n) - \mu_j)$$
(1)

with  $\mu_i = \frac{1}{NL} \sum_{n=1}^{N} \sum_{l=1}^{L} d_i(q_l, X_n).$ 

Given the covariance matrix  $\Sigma$ , we calculate the correlation matrix  $\mathcal{R}$  as  $\mathcal{R}_{ij} = \Sigma_{ij}/\sqrt{\Sigma_{ii}\Sigma_{jj}}$ . The entries of this correlation matrix can be interpreted as similarities of different features. A high value  $\mathcal{R}_{ij}$  means a high similarity between features *i* and *j*. This similarity matrix can then be analyzed to find out which features have similar properties and which do not. One way to do this is to visualize it using multi-dimensional scaling (Hand et al. 2001, p. 84ff). *Multi-dimensional scaling (MDS)* seeks a representation of data points in a lower dimensional space while preserving the distances between data points as well as possible. To visualize this data by multi-dimensional scaling, we convert the similarity matrix  $\mathcal{R}$  into a dissimilarity matrix  $\mathcal{D}$  by setting  $\mathcal{D}_{ij} = 1 - |\mathcal{R}_{ij}|$ . For visualization purposes, we choose a two-dimensional space for MDS.

#### 5 Benchmark databases for CBIR

To cover a wide range of different applications in which CBIR is used, we propose benchmark databases from different domains. In the ImageCLEF evaluations large image retrieval benchmark databases have been collected. However, these are not suitable for the comparison of image features as for most of the tasks textual information is supplied and necessary for an appropriate solution of the task. Table 2 gives an overview of the

Database	Images	Queries	Avg. rel	Query mode
WANG	1,000	1,000	99.0	Leaving-one-out
UW	1,109	1,109	59.3	Leaving-one-out
IRMA 10000	10,000	1,000	520.2	Test & database images are disjoint
ZuBuD	1,005	105	5.0	Test & database images are disjoint
UCID	1,338	262	3.5	Leaving-one-out

**Table 2** Summary of the databases used for the evaluation with database name, number of images in the database, number of query images, average number of relevant images per query, and a description how the queries are evaluated

databases used in the evaluations. Although the databases presented here are small in comparison to other CBIR tasks, they represent a wide variety of tasks and allow for a meaningful comparison of feature performances.

The WANG database (Section 5.1), as a subset from the Corel stock photo collection, can be considered similar to stock photo searches. The UW database (Section 5.2) and the UCID database (Section 5.5) mainly consist of personal images and represent the home user domain. The ZuBuD database (Section 5.4) and the IRMA database (Section 5.3) are limited domain tasks for touristic/building retrieval and medical applications, respectively.

# 5.1 WANG database

The WANG database is a subset of 1,000 images of the Corel stock photo database which have been manually selected and which form 10 classes of 100 images each. One example of each class is shown in Fig. 1. The WANG database can be considered similar to common stock photo retrieval tasks with several images from each category and a potential user having an image from a particular category and looking for similar images which have e.g. cheaper royalties or which have not been used by other media. The 10 classes are used for relevance estimation: given a query image, it is assumed that the user is searching for images from the same class, and therefore the remaining 99 images from the same class are considered relevant.

# 5.2 UW database

The database created at the University of Washington consists of a roughly categorized collection of 1,109 images. These images are partly annotated using keywords. The remaining images were annotated by our group to allow the annotation to be used for relevance estimation; our annotations are publicly available.<sup>10</sup>

The images are of various sizes and mainly include vacation pictures from various locations. There are 18 categories, for example "spring flowers", "Barcelona", and "Iran". Some example images with annotations are shown in Fig. 2. The complete annotation consists of 6,383 words with a vocabulary of 352 unique words. On the average, each image has about 6 words of annotation. The maximum number of keywords per image is 22 and the minimum is 1. The database is freely available.<sup>11</sup> The relevance assessment for

 $<sup>^{10}\</sup> http://www-i6.informatik.rwth-aachen.de/~deselaers/uwdb/index.html$ 

<sup>&</sup>lt;sup>11</sup> http://www.cs.washington.edu/research/imagedatabase/groundtruth/



africa

monuments

buses



flowers

horses

mountains

Fig. 1 One example image from each of the 10 classes of the WANG database together with their class labels

the experiments with this database were performed using the annotation: an image is considered to be relevant w.r.t. a given query image if the two images have a common keyword in the annotation. On the average, 59.3 relevant images correspond to each image. The keywords are rather general; thus for example images showing sky are relevant w.r.t. each other, which makes it quite easy to find relevant images (high precision is likely easy) but it can be extremely difficult to obtain a high recall since some images showing sky might have hardly any visual similarity with a given query.

This task can be considered a personal photo retrieval task, e.g., a user with a collection of personal vacation pictures is looking for images from the same vacation, or showing the same type of building.

# 5.3 IRMA-10000 database

The IRMA database consists of 10,000 fully annotated radiographs taken randomly from medical routine at the RWTH Aachen University Hospital. The images are split into 9,000 training and 1,000 test images. The images are subdivided into 57 classes. The IRMA



buildings clouds mountain people sand sky



bench, car, house, lantern, trees, window



trees, bushes, overcast sky, building, post



skv



struct, bushes, flowers, people



partially cloudy sky, hills, trees, grasses, ground, houses

Husky Stadium, north stands, people, football, field,...

sailboats, ice, water, buildings

mosque, tiles, people, sky, car

Fig. 2 Examples from the UW database with annotation

database was used in the ImageCLEF 2005 image retrieval evaluation for the automatic annotation task. For CBIR, the relevances are defined by the classes, given a query image from a certain class, all database images from the same class are considered relevant. Example images along with their class numbers and textual descriptions of the classes are given in Fig. 3. This task is a medical image retrieval task and is in practical use at the Department for Diagnostic Radiology of the RWTH Aachen University Hospital.

As all images from this database are gray value images, we evaluate neither the color histograms nor the MPEG7 scalable color descriptor since they only account for color information.

## 5.4 ZuBuD database

The "Zurich Buildings Database for Image Based Recognition" (ZuBuD) is a database which has been created by the Swiss Federal Institute of Technology in Zurich and is described in more detail in (Shao et al. 2003a, 2003b).

The database consists of two parts, a training part of 1,005 images of 201 buildings, 5 of each building and a query part of 115 images. Each of the query images contains one of the buildings from the main part of the database. The pictures of each building are taken from



Fig. 3 Example images of the IRMA 10000 database along with their class and annotation

different viewpoints and some of them are also taken under different weather conditions and with two different cameras. Given a query image, only images showing exactly the same building are considered relevant. To give a more precise idea of this database, some example images are shown in Fig. 4.

This database can be considered as an example for a mobile travel guide task, which attempts to identify buildings in pictures taken with a mobile phone camera and then



**Fig. 4** (a) A query image and the 5 images from the same building in the ZuBuD-database (b) 6 images of different buildings in the ZuBuD-database

obtains certain information about the building (Shao et al. 2003). The ZuBuD database is freely available.<sup>12</sup>

### 5.5 UCID database

The UCID database<sup>13</sup> was created as a benchmark database for CBIR and image compression applications (Schaefer and Stich 2004). In Schaefer (2004) this database was used to measure the performance of a CBIR system using compressed domain features. This database is similar to the UW database as it consists of vacation images and thus poses a similar task.

For 264 images, manual relevance assessments among all database images were created, allowing for performance evaluation. The images that are judged to be relevant are images which are very clearly relevant, e.g., for an image showing a particular person, images showing the same person are searched and for an image showing a football game, images showing football games are considered to be relevant. The used relevance assumption makes the task easy on one hand, because relevant images are very likely quite similar, but on the other hand, it makes the task difficult, because there are likely images in the database which have a high visual similarity but which are not considered relevant. Thus, it can be difficult to have high precision results using the given relevance assessment, but since only few images are considered relevant, high recall values might be rather easy to obtain. Example images are given in Fig. 5.

# 6 Evaluation of the features considered

In this section we report the results of the experimental evaluation of the features. To evaluate all features on the given databases, we extracted the features from the images and executed experiments to test the particular features. For all experiments, we report the mean average precision and the classification error rate. The connection between the classification error rate and mean average precision shows the strong relation between CBIR and classification. Both performance measures have advantages. The error rate is very precision oriented and thus it is best if relevant images are retrieved early. On the contrary, the mean average precision accounts for the average performance over the complete PR graph. Furthermore, we calculated the distance vectors mentioned in Section 4 for each of the queries performed to obtain a global correlation analysis of all features.

#### 6.1 Performance evaluation of features

The results from the single feature experiments are given in Figs. 6 and 7 and in Tables 3 and 4. The results are sorted by the average of the classification error rates. The results from the correlation analysis are given in Fig. 9. Note that the features 'color histogram' and 'MPEG7 scalable color' were not evaluated for the IRMA database because pure color descriptors are not suitable for this gray-scale database.

<sup>12</sup> http://www.vision.ee.ethz.ch/ZuBuD

<sup>13</sup> http://www.vision.doc.ntu.ac.uk/datasets/UCID/ucid.html



Fig. 5 Example images from the UCID database



**Fig. 6** Classification error rate [%] for each of the features for each of the databases (sorted by average error rate over the databases). The different shades of gray denote different databases and the blocks of bars denote different features

It can clearly be seen that different features perform differently on the databases. Grouping the features by performance results in three groups, one group of five features clearly outperforms the other features (average error rate < 30%, average mean average precision  $\approx 50\%$ ). A second group has average error rates of approximately 40%





**Fig. 7** Mean average precision for each of the features for each of the databases (sorted in the same order as Fig. 6 to allow for easy comparison)

(respectively average mean average precision 40%) and a last group performs clearly worse.

The top group is led by the color histogram which performs very well for *all* color tasks and has not been evaluated on the IRMA data. When all databases are considered, the global feature search (cf. Section 3.7) of SIFT features extracted at Harris points (Dorkó 2006, chapters 3, 4) performs best on the average. This good performance is probably partly due to the big success on the ZuBuD database, where features of similar type were observed to perform exceedingly well (Obdrzalek and Matas 2003). They also perform well on the UCID database, where relevant images, in contrast to the UW task, are very close neighbors. The possible high dissimilarity between relevant images in the UW database, thus explains the bad performance there. However, the patch histograms outperform the SIFT features on all other tasks as they include color information which obviously is very important for most of the tasks. They also obtain a good performance for the IRMA data. It can be observed that the error rates for the UCID database are very high in comparison to the other databases, so the UCID task can be considered to be harder than e.g., the UW task.

A similar result to the one obtained using color histogram is obtained by the invariant feature histogram with monomial kernel. This is not surprising, as it is very similar to a color histogram, except that it also partly accounts for local texture. It can be observed that the performance for the color databases is nearly identical to the color histogram. The relatively bad ranking of these features in the tables is due to the bad performance on the IRMA task.

Feature	WANG	UW	IRMA	UCID	ZuBuD	Average
Color histogram	16.9	12.3	_	51.5	7.8	22.1
LF SIFT global search	37.2	31.4	27.7	31.7	7.0	27.0
LF patches histogram	17.9	14.6	24.9	58.0	24.4	28.0
LF SIFT histogram	25.6	21.4	30.8	50.4	18.3	29.3
Inv. feature histogram (monomial)	19.2	12.9	55.80	53.8	7.8	29.9
MPEG7: scalable color	25.1	13.0	-	60.7	32.2	32.7
LF patches signature	24.3	17.6	42.7	68.7	36.5	38.0
Gabor histogram	30.5	20.5	44.9	74.1	24.4	38.9
$32 \times 32$ Image	47.2	26.4	22.8	82.8	27.0	41.2
MPEG7: color layout	35.4	21.2	47.7	75.2	27.0	41.2
$X \times 32$ image	55.9	26.7	21.4	83.2	20.9	41.6
Tamura texture histogram	28.4	16.8	33.0	63.7	84.4	45.2
LF SIFT signature	35.1	20.9	99.3	58.4	20.0	46.7
Gray value histogram	45.3	23.0	42.6	86.64	47.0	48.9
LF patches global	42.9	42.7	48.2	63.4	47.8	49.0
MPEG7: edge histogram	32.8	22.9	99.3	69.9	23.5	49.7
Inv. feature histogram (relational)	38.3	23.6	39.2	83.2	93.9	55.6
Gabor vector	65.5	37.9	42.5	95.8	73.0	62.9
Global texture feature	51.4	32.4	67.7	95.4	98.3	69.0

 Table 3
 Error rate [%] for each of the features for each of the databases (sorted by average error rate over the databases)

Leaving out the IRMA task for this feature, it would be ranked second in the entire ranking. The high similarity of color histograms and invariant feature histograms with monomial kernel can also directly be observed in Fig. 9 where it can be seen that color histograms (point 1) and invariant feature histograms with monomial kernel (point 11) have very similar properties.

The second group of features consists of four features: signatures of SIFT features, appearance-based image features, and the MPEG 7 color layout descriptor.

Although the image thumbnails compared with the image distortion model perform quite poorly for the WANG, the UW, and the UCID tasks, they perform extremely well for the IRMA task and reasonably well for the ZuBuD task. A major difference between these tasks is that the first three databases contain general color photographs of completely unconstrained scenes, whereas the latter ones contain images from limited domains only.

The simpler appearance-based feature of  $32 \times 32$  thumbnails of the images, compared using Euclidean distance, is the next best feature, and again it can be observed that it performs well for the ZuBuD and IRMA tasks only.

As expected, the MPEG7 color layout descriptor and  $32 \times 32$  image thumbnails obtain similar results because they both encode the spatial distribution of colors or gray values in the images.

Among the texture features (Tamura texture histogram, Gabor features, global texture descriptor, relational invariant feature histogram, and MPEG-7 edge histogram), the Tamura texture histogram and the Gabor histogram outperform the others.

order as rable 5 to allow for easy comparison)						
Feature	WANG	UW	IRMA	UCID	ZuBuD	Average
color histogram	16.9	12.3	-	51.5	7.8	22.1
LF SIFT global search	37.2	31.4	27.7	31.7	7.0	27.0
LF patches histogram	17.9	14.6	24.9	58.0	24.4	28.0
LF SIFT histogram	25.6	21.4	30.8	50.4	18.3	29.3
Inv. feature histogram (monomial)	19.2	12.9	55.80	53.8	7.8	29.9
MPEG7: scalable color	25.1	13.0	-	60.7	32.2	32.7
LF patches signature	24.3	17.6	42.7	68.7	36.5	38.0
Gabor histogram	30.5	20.5	44.9	74.1	24.4	38.9
$32 \times 32$ image	47.2	26.4	22.8	82.8	27.0	41.2
MPEG7: color layout	35.4	21.2	47.7	75.2	27.0	41.2
$X \times 32$ image	55.9	26.7	21.4	83.2	20.9	41.6
Tamura texture histogram	28.4	16.8	33.0	63.7	84.4	45.2
LF SIFT signature	35.1	20.9	99.3	58.4	20.0	46.7
Gray value histogram	45.3	23.0	42.6	86.64	47.0	48.9
LF patches global	42.9	42.7	48.2	63.4	47.8	49.0
MPEG7: edge histogram	32.8	22.9	99.3	69.9	23.5	49.7
Inv. feature histogram (relational)	38.3	23.6	39.2	83.2	93.9	55.6
Gabor vector	65.5	37.9	42.5	95.8	73.0	62.9
Global texture feature	51.4	32.4	67.7	95.4	98.3	69.0

 Table 4
 Mean average precision [%] for each of the features for each of the databases (sorted in the same order as Table 3 to allow for easy comparison)

#### 6.2 Correlation analysis of features

Figure 8 shows the average correlation of different features over all databases. The darker a field in this image is, the lower the correlation between the corresponding features, bright fields denote high correlations. Figure 9 shows the visualizations of the outcomes of multidimensional scaling of the correlation analysis. We applied the correlation analysis for the different tasks individually (4 top plots) and for all tasks jointly (bottom plot). Multidimensional scaling was used to translate the similarities of the different features into distances in a two-dimensional space. The further away two points are in the graph, the less similar the corresponding features are for CBIR, and conversely the closer together they appear, the higher the similarity between these features.

For each of these plots the according distance vectors obtained from all queries with all database images have been used (WANG database: 1,000,000 distance vectors, UW&UCID database: 194,482+350,557 distance vectors, IRMA database: 9,000,000 distance vectors, ZuBuD database: 115,575 distance vectors, all databases: 10,660,614 distance vectors).

The figures show a very strong correlation between color histograms (point 1) and invariant feature histograms with monomial kernel (point 11). In fact, they lead to hardly any differences in the experiments. For the databases consisting of color photographs they outperform most other features. A high similarity is also observed between the patch signatures (point 14) and the MPEG7 color layout (point 2) for all tasks.

Two other features that are highly correlated are the two methods that use local feature search for the two different types of local features (points 5 and 12). The different



Fig. 8 Correlation of the different features. Bright fields denote high and dark fields denote low correlation. Another representation of this information is given in Fig. 9

comparison methods for local feature histograms/signature have similar performances (3, 4 and 13, 14, respectively).

Another strong correlation can be observed between  $32 \times 32$  image thumbnails (point 18) and the MPEG7 color layout representation (point 2), which was to be expected as both of these have a rough representation of the spatial distribution of colors (resp. gray values) of the images.

Interestingly, the correlation between  $32 \times 32$  images compared using Euclidean distance (point 18) and the X × 32 images compared using the image distortion model (point 18) is low, with only some similarity for the IRMA and the ZuBuD task. This is partly due to the exceedingly good performance of the image distortion model for the IRMA task and partly due to the missing invariance with respect to slight deformations in the images for the Euclidean distance. For example in the ZuBuD task, the image distortion model can partly compensate for the changes in the viewpoints which leads to a much better performance.

Another interesting aspect is that the various texture features (MPEG7 edge histogram (6), global texture feature (10), Gabor features (8, 7), relational invariant feature histogram (15), and Tamura texture histogram (17)) are not correlated strongly. We conclude that none of the texture features is sufficient to completely describe the textural properties of an

Fig. 9 Correlation of the different features visualized using multi-dimensional scaling. Features that lie close together have similar properties. Top 4 plots: database-wise visualization, bottom plot: all databases jointly. The numbers in the plots denote the individual features: 1: color histogram, 2: MPEG7: color layout, 3: LF SIFT histogram, 4: LF SIFT signature, 5: LF SIFT global search, 6: MPEG7: edge histogram, 7: Gabor vector, 8: Gabor histograms, 9: gray value histogram, 10: global texture feature, 11: inv. feature histogram (monomial), 12: LF patches global, 13: LF patches histogram, 14: LF patches signature, 15: inv. feature histogram (relational), 16: MPEG7: scalable color, 17: Tamura texture histogram, 18:  $32 \times 32$  image, 19: X × 32 image



image. The Tamura texture histogram and the Gabor histogram outperform the other texture features, Tamura features being better in three and Gabor histograms being clearly better in two of the five tasks, both of them are a good choice for texture representation.

To give a little insight into how these plots can be used to select sets of features for a given task, we discuss how features for the WANG database could be chosen in the following paragraph. Combined features are linearly combined as described in Section 2. Here, all features are weighted equally, but some improvement of the retrieval results can be achieved by choosing different weights for the individual features. In Deselaers et al. (2007) we present an approach to automatically learning a feature combination from a set of queries with known relevant images using a discriminative maximum entropy model.

Table 5         Combining features           using the results from the	Features	ER [%]	MAP [%]
for the WANG database	Color histograms	16.9	50.5
	+ Global texture	15.7	49.5
	+ Tamura histograms	13.7	51.2
	+ Thumbnails	13.7	53.9
	+ Patch histograms	11.6	55.7

*Finding a suitable set of features.* Assume we are about to create a CBIR system for a new database consisting of general photographs. We extract features from the data and create the according MDS plot (Fig. 9, top left). Since we know that we are dealing with general photographs, we start with a simple color histogram (point 1). The plot now tells us that invariant feature histograms with monomial kernel (11) would not give us much additional information. Next, we consider the various texture descriptors (points 6, 10, 15, 17, 7, 8) and choose one of these, say global texture features (10) and maybe another: Tamura texture histograms (17). Now we have covered color and texture and can consider a global descriptor such as the image thumbnails (18) or a local descriptor such as one of (12, 13, or 14) or (3, 4, or 5). After adding a feature, the performance of the CBIR system can be evaluated by the user. In Table 5 we quantitatively show the influence of adding these features for the WANG database. It can be seen that the performance is incrementally improved by adding more and more features.

## 6.3 Connection between mean average precision and error rate

In Figs. 10 and 11 the correlation between mean average precision and error rate is visualized database-wise and feature-wise, respectively. The correlation of error rate and mean average precision over all experiments presented in this paper is 0.87. In the keys of the figures, the correlations per database and per feature are given, respectively.





Fig. 11 Analysis of the correlation between classification error rate and mean average precision for the features. The numbers in the legend give the correlation for the experiments performed using the individual features

From Fig. 10 it can be seen that this correlation varies between the tasks between 0.99 and 0.67. For the UCID task, this correlation is markedly strong with 0.99. The correlation is lowest for the UW task which has a correlation of 0.67 and which is the only task with a correlation below 0.8.

In Fig. 11, the same correlation is analyzed feature-wise. Here, the correlation values vary strongly between 0.4 and 1.0. The LF SIFT signature descriptor has the lowest correlation and the LF patches histogram descriptor also has a low correlation of only 0.6. The two different image thumbnail descriptors have a correlation of 0.7. All other features have correlation values greater than 0.8, thus it can be said that an image representation that works well for classification will generally work well for CBIR as well and vice versa. Exemplarily, this effect can be observed when looking at the results for the WANG and IRMA database for the color histograms and the X  $\times$  32 thumbnails. On the one hand, for the WANG database, the color histograms perform very well for error rate and mean average precision; in contrast, the image thumbnails perform poorly. On the other hand, the effect is reversed for the IRMA database: here, the color histograms perform poorly and the image thumbnails outstandingly well. It can be observed that the performance increase (resp. decrease) is in the same magnitude for mean average precision and error rate. Thus, it can be seen that a feature that performs well for the task of classification on a certain dataset, it will most probably be a good choice for retrieval of images from that dataset, too.

## 7 Conclusion

We have discussed a large variety of features for image retrieval and a setup of five freely available databases that can be used to quantitatively compare these features. From the experiments conducted it can be deduced, which features perform well on which kind of task and which do not. In contrast to other papers, we consider tasks from different domains jointly and directly compare and analyze which features are suitable for which task.

Which features are suitable for which task in CBIR? The main question addressed in this paper, which features are suitable for which task in image retrieval, has been thoroughly investigated:

One clear finding is that color histograms, often cited as a baseline in CBIR, clearly are a reasonably good baseline for general color photographs. However, approaches using local image descriptors outperform color histograms in various tasks but usually at the cost of much higher computational costs. If the images are from a restricted domain, as they are in the IRMA and in the ZuBuD task, other methods should be considered as a baseline, e.g., a simple nearest neighbor classifier using thumbnails of the images.

Furthermore, it has been shown that, despite more than 30 years in research on texture descriptors, still none of the texture features presented can convey a complete description of the texture properties of an image. Therefore a combination of different texture features will usually lead to best results.

It should be noted that for specialized tasks, such as finding images that show certain objects, better methods exist today that can learn models of particular objects from a set of training data. However, these approaches are computationally far more expensive and always require relatively large amounts of training data.

Although the selection of features tested was not completely exhaustive, the selection was wide and the methods presented can easily be applied to other features to compare them to the features presented here. On one hand, the presented descriptors were selected such that features presented many years ago, such as color histograms (Swain and Ballard 1991), Tamura texture features (Tamura et al. 1978), Gabor features, and spatial autocorrelation features (Haralick et al. 1973), as well as very recent features such as SIFT descriptors (Lowe 2004) and patches (Deselaers et al. 2005) were compared. On the other hand, the features were selected such that descriptors accounting for color, texture, and (partly) shape, as well as local and global descriptors were covered. We also included a subset of the standardized MPEG7 features.

All features have been thoroughly examined experimentally on a set of five databases. All of these databases are freely available and pointers to their location are given in this paper. This allows researchers to compare the findings from this work with other features that were not covered here or which will be presented in future. The databases chosen are representative for four different tasks in which CBIR plays an important role.

Which features are correlated and how can features be combined? We conducted a correlation analysis of the features considered showing which features have similar properties and which do not. The outcomes of this method can be used as an intuitive help to finding suitable combinations of features for certain tasks. In contrast to other methods for feature combination, the method presented here does not rely on training data/relevance judgements to find a suitable set of features. In particular, it will tell you which features are not worth combining because they produce correlated distance results. The method is not a fully automatic feature selection method but the process of selecting features is demonstrated for one of the tasks with promising results. However, the focus of this paper is not to

combine several features as this would exceed the scope and a variety of known methods have covered this aspect, e.g., (Yavlinski et al. 2004; Kittler 1998; Heesch and Rüger 2002).

Another conclusion we have drawn from this work is that the intuitive assumption that classification of images and CBIR are strongly connected is justified. Both tasks are strongly related to the concept of similarity which can be measured best if suitable features are available. In this paper, we have evaluated this assumption quantitatively by considering four different domains and analyzing the classification error rate for classification and the mean average precision for CBIR. It was clearly shown empirically that features that perform well for classification also perform well for CBIR and vice versa. This strong connection allows us to take advantage of knowledge obtained in either classification or CBIR for the other respective task. For example, in the medical domain much research has been done to classify whether an image shows a pathological case or not, likely some of the knowledge obtained in these studies can be transferred to the CBIR domain to help retrieving images from a picture archiving system.

*Future Work*. Future work in CBIR certainly includes finding new and better image descriptors and methods to combine these appropriately. Furthermore, the achievements in object detection and recognition will certainly find their way into the CBIR domain and a shift towards methods that automatically learn about the semantics of images is imaginable. First steps into this direction can be seen in (Nowak et al. 2007) where a method is presented that learns how to compare never seen objects and presents an image similarity measurement which works on the object level. Methods for automatic image annotation are also related to CBIR and the automatic generation of textual labels for images allows to use textual information retrieval techniques to retrieve images.

**Acknowledgements** This work was partially funded by the DFG (Deutsche Forschungsgemeinschaft) under contract NE-572/6. The authors would like to thank Gyuri Dorkó (formerly with INRIA Rhône-Alpes) for providing his SIFT feature extraction software and the authors of the MPEG7 XM reference implementation.

## References

- Antani, S., Kasturi, R., & Jain, R. (2002). A survey on the use of pattern recognition methods for abstraction, indexing and retrieval of images and video. *Pattern Recognition*, 35, 945–965.
- Armitage, L. H., & Enser, P. G. (1997). Analysis of user need in image archives. Journal of Information Science, 23(4), 287–299.
- Bloehdorn, S., Petridis, K., Saathoff, C., Simou, N., Tzouvaras, V., Avrithis, Y., et al. (2005). Semantic annotation of images and videos for multimedia analysis. In *European Semantic Web Conference* (*ESWC 05*). Heraklian, Greece.
- Bober, M. (2001). MPEG-7 Visual Shape Descriptors. IEEE Trans on Circuits and Systems for Video Technology, 11(6), 716–719.
- Carson, C., Belongie, S., Greenspan, H., & Malik, J. (2002). Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8), 1026–1038.
- Clough, P., Mueller, H., Deselaers, T., Grubinger, M., Lehmann, T., Jensen, J., et al. (2006). The CLEF 2005 cross-language image retrieval track. In Accessing Multilingual Information Repositories, 6th Workshop of the Cross-Language Evaluation Forum, CLEF 2005, vol. 4022 of Lecture Notes in Computer Science (pp. 535–557). Vienna, Austria.
- Clough, P., Müller, H., & Sanderson, M. (2004). The CLEF Cross Language Image Retrieval Track (ImageCLEF) 2004. In *Fifth Workshop of the Cross-Language Evaluation Forum (CLEF 2004), vol. 3491* of LNCS (pp. 597–613).
- Datta, R., Li, J., & Wang, J. Z. (2005). Content-based image retrieval—approaches and trends of the new age. In ACM Intl. Workshop on Multimedia Information Retrieval, ACM Multimedia. Singapore.

- de Vries, A. P., & Westerveld, T. A. (2004). comparison of continuous vs. discrete image models for probabilistic image and video retrieval. In Proc. International Conference on Image Processing (pp. 2387–2390). Singapore.
- Deselaers, T. (2003). Features for image retrieval. Master's thesis. Aachen: Human Language Technology and Pattern Recognition Group, RWTH Aachen University.
- Deselaers, T., Hegerath, A., Keysers, D., & Ney, H. (2006). Sparse patch-histograms for object classification in cluttered images. In DAGM 2006, Pattern Recognition, 27th DAGM Symposium, vol. 4174 of Lecture Notes in Computer Science (pp. 202–211). Berlin.
- Deselaers, T., Weyand, T., Keysers, D., Macherey, W., & Ney, H. (2006). FIRE in ImageCLEF 2005: Combining content-based image retrieval with textual information retrieval. In Accessing Multilingual Information Repositories, 6th Workshop of the Cross-Language Evaluation Forum, CLEF 2005, vol. 4022 of Lecture Notes in Computer Science (pp. 652–661). Vienna, Austria.
- Deselaers, T., Weyand, T., & Ney, H. (2007). Image retrieval and annotation using maximum entropy. In C. Peters, P. Clough, F. Gey, J. Karlgren, B. Magnini, D. Oard, et al. (Eds.), *Evaluation of Multilingual* and Multi-modal Information Retrieval—Seventh Workshop of the Cross-Language Evaluation Forum, CLEF 2000a, vol. 4730 of Lecture Notes in Computer Series (pp. 725–734). Alicante.
- Deselaers, T., Keysers, D., & Ney, H. (2004) Features for image retrieval—a quantitative comparison. In DAGM 2004, Pattern Recognition, 26th DAGM Symposium, vol. 3175 of Lecture Notes in Computer Science (pp. 228–236). Tübingen, Germany.
- Deselaers, T., Keysers, D., & Ney, H. (2004). Classification error rate for quantitative evaluation of contentbased image retrieval systems. In *International Conference on Pattern Recognition 2004 (ICPR 2004)* (Vol. 2, pp. 505–508). Cambridge.
- Deselaers, T., Keysers, D., & Ney, H. (2005). FIRE— Flexible Image Retrieval Engine: ImageCLEF 2004 evaluation. In Multilingual Information Access for Text, Speech and Images – Fifth Workshop of the Cross-Language Evaluation Forum, CLEF 2004, vol. 3491 of Lecture Notes in Computer Science (pp. 688–698). Bath: Springer.
- Deselaers, T., Keysers, D., & Ney, H. (2005). Discriminative training for object recognition using image patches. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 05)* (Vol. 2, pp. 157–162). San Diego.
- Deselaers, T., Rybach, D., Dreuw, P., Keysers, D., & Ney, H. (2005). Face-based image retrieval-one step toward object-based image retrieval. In H. Müller & A. Hanbury (Eds.), *MUSCLE/ImageCLEF Workshop on Image and Video Retrieval Evaluation* (pp. 25–32). Vienna.
- Di Sciascio, E., Donini, F. M., & Mongiello, M. (2002). Structured knowledge representation for image retrieval. *Journal of Artificial Intelligence Research*, 16, 209–257.
- Dorkó, G. (2006) Selection of discriminative regions and local descriptors for generic. Object Class Recognition. Ph.D. thesis. Institut National Polytechnique de Grenoble.
- Eidenberger, H. (2003). How good are the visual MPEG-7 features? In Proceedings SPIE Visual Communications and Image Processing Conference (Vol. 5150, pp. 476–488). Lugano.
- Faloutsos, C., Barber, R., Flickner, M., Hafner, J., Niblack, W., & Petkovic, D., et al. (1994). Efficient and effective querying by image content. *Journal of Intelligent Information Systems*, 3(3/4), 231–262.
- Fei-Fei, L., & Perona, P. (2005). A bayesian hierarchical model for learning natural scene categories. In IEEE Conference on Computer Vision and Pattern Recognition (Vol. 2, pp. 524–531). IEEE: San Diego
- Fergus, R., Perona, P., & Zissermann, A. (2003). Object class recognition by unsupervised scale-invariant learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 03)* (pp. 264–271). Blacksburg.
- Fergus, R., Perona, P., & Zisserman, A. (2005). A sparse object category model for efficient learning and exhaustive recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 05)* (Vol. 2, pp. 380–389). IEEE: San Diego.
- Forsyth, D. A., & Ponce, J. (2002) Computer Vision: A Modern Approach (pp. 599-619). Prentice Hall.
- Gu, Z. Q., Duncan, C. N., Renshaw, E., Mugglestone, M. A., Cowan, C. F. N., & Grant, P. M. (1989). Comparison of techniques for measuring cloud texture in remotely sensed satellite meteorological image data. *Radar and Signal Processing*, 136(5), 236–248.
- Haberäcker, P. (1995). Praxis der Digitalen Bildverarbeitung und Mustererkennung. München, Wien: Carl Hanser Verlag.
- Hand, D., Manila, H., & Smyth, P. (2001). Principles of data mining. Cambridge: MIT Press.
- Haralick, R. M., Shanmugam, B., & Dinstein, I. (1973). Texture Features for image classification. IEEE Transactions on Systems, Man, and Cybernetics, 3(6), 610–621.

- Heesch, D., & Rüger, S. (2002). Combining features for content-based sketch retrieval—a comparative evaluation of retrieval performance. In *European Colloquium on Information Retrieval Research, vol.* 2291 of LNCS (pp. 41–52). Glasgow, Scotland.
- Heesch, D., & Rüger, S. (2003). Performance boosting with three mouse clicks—relevance feedback for CBIR. In *European Conference on Information Retrieval Research. No. 2633 in LNCS* (pp. 363–376). Pisa: Springer Verlag.
- Iqbal, Q., & Aggarwal. J. (2002). CIRES: A system for content-based retrieval in digital image libraries. In International Conference on Control, Automation, Robotics and Vision (pp. 205–210). Singapore.
- Jain, S. (2004). Fast image retrieval using local features: Improving approximate search employing seedgrow approach. Master's thesis. INPG, Grenoble.
- Keysers, D., Deselaers, T., Gollan, C., & Ney, H. (2007). Deformation models for image recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(8), 1422–1435.
- Kittler, J. (1998). On combining classifiers. IEEE Transactions on Pattern Analysis and Machine Intelligence, 20(3), 226–239.
- Lehmann, T. M., Güld, M. -O., Deselaers, T., Keysers, D., Schubert, H., Spitzer, K., et al. (2005). Automatic categorization of medical images for content-based retrieval and data mining. *Computerized Medical Imaging and Graphics*, 29(2), 143–155.
- Lew, M. S., Sebe, N., Djeraba, C., & Jain, R. (2006). Content-based multimedia information retrieval: State of the art and challenges. ACM Transactions on Multimedia Computing, Communications and Applications, 2(1), 1–19.
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2), 91–110.
- MacArthur, S. D., Brodley, C. E., & Shyu, C. -R. (2000). Relevance feedback decision trees in contentbased image retrieval. In *Content-based access of image and video libraries* (pp. 68–72). IEEE: Hilton Head Island, SC.
- Manjunath, B., Ohm, J. -R., Vasudevan, V. V., & Yamada, A. (2001). Color and texture descriptors. IEEE Trans Circuits and Systems for Video Technology, 11(6), 703–715.
- Marée, R., Geurts, P., Piater, J., & Wehenkel, L. (2005) Random subwindows for robust image classification. In *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 34–40).
- Markkula, M., & Sormunen, E. (1998). Searching for photos—journalists' practices in pictorial IR. In Electronic Workshops in Computing—Hallenge of Image Retrieval (pp. 1–13). Newcastle.
- Meghini, C., Sebastiani, F., & Straccia, U. (2001). A model of multimedia information retrieval. *Journal of the ACM*, 48(5), 909–970.
- Mikolajczyk, K., Tuytelaars, T., Schmid, C., Zisserman, A., Matas, J., Schaffalitzky, F., et al. (2005). A comparison of affine region detectors. *International Journal of Computer Vision* 65(1/2).
- Müller, H., Müller, W., Squire, D. M., Marchand-Maillet, S., & Pun, T. (2000). Learning features weights from user behavior in Content-Based Image Retrieval. In S. Simoff & O. Zaiane (Eds.), ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (Workshop on Multimedia Data Mining MDM/KDD2000). Boston.
- Müller, H., Müller, W., Marchand-Maillet, S., & Squire, D. M. (2000). Strategies for positive and negative relevance feedback in image retrieval. In *International Conference on Pattern Recognition*, vol. 1 of *Computer Vision and Image Analysis* (pp. 1043–1046). Barcelona.
- Müller, H., Müller, W., Squire, D. M., Marchand-Maillet, S., & Pun, T. (2001). Performance evaluation in content-based image retrieval: overview and proposals. In H. Bunke & X. Jiang (Eds.), *Pattern Recognition Letters (Special Issue on Image and Video Indexing)* 22(5), 593–601.
- Müller, H., Marchand-Maillet, S., & Pun, T. (2002) The truth about corel—evaluation in image retrieval. In Proceedings of The Challenge of Image and Video Retrieval (CIVR2002), vol. 2383 of LNCS (pp. 38–49). London.
- Müller, H., Michoux, N., Bandon, D., & Geissbuhler A. (2004). A review of content-based image retrieval systems in medical applications-clinical benefits and future directions. *International Journal of Medical Informatics* 73(1), 1–23.
- Najjar, M., Ambroise, C., & Cocquerez, J. -P. (2003). Feature selection for semi supervised learning applied to image retrieval. In *ICIP 2003* (Vol. 3, pp. 559–562). Barcelona.
- Nölle, M. (2003). Distribution distance measures applied to 3-D object recognition—a case study. In DAGM 2003, Pattern Recognition, 25th DAGM Symposium, vol. 2781 of Lecture Notes in Computer Science (pp. 84–91). Magdeburg: Springer Verlag.
- Nowak, E., & Jurie, F. (2007). Learning visual similarity measures for comparing never seen objects. In *CVPR 2007*. Minneapolis.

- Obdrzalek, S., & Matas, J. (2003). Image retrieval using local compact DCT-Based representation. In DAGM 2003, Pattern Recognition, 25th DAGM Symposium, vol. 2781 of Lecture Notes in Computer Science (pp. 490–497). Magdeburg, Germany: Springer Verlag.
- Ohm, J.-R. (2001). The MPEG-7 visual description framework—concepts, accuracy and applications. In CAIP 2001. No. 2124 in LNCS (pp. 2–10).
- Opelt, A., Pinz, A., Fussenegger, M., & Auer, P. (2006). Generic object recognition with boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(3), 416–431.
- Paredes, R., Perez-Cortes, J., Juan, A., & Vidal, E. (2001). Local representations and a direct voting scheme for face recognition. In *Workshop on Pattern Recognition in Information Systems* (pp. 71–79). Setúbal, Portugal.
- Park, M., Jin, J. S., & Wilson, L. S. (2002). Fast content-based image retrieval using quasi-gabor filter and reduction of image feature. In *Southwest Symposium on Image Analysis and Interpretation* (pp. 178– 182). Santa Fe.
- Pentland, A., Picard, R., & Sclaroff, S. (1996). Photobook: Content-based manipulation of image databases. International Journal of Computer Vision, 18(3), 233–254.
- Puzicha, J., Rubner, Y., Tomasi, C., & Buhmann, J. (1999). Empirical evaluation of dissimilarity measures for color and texture. In *International Conference on Computer Vision* (Vol. 2, pp. 1165–1173). Corfu, Greece.
- Rubner, Y., Tomasi, C., & Guibas, L. J. (1998). A metric for distributions with applications to image databases. In *International Conference on Computer Vision* (pp. 59–66). Bombay.
- Rui, Y., Huang, T., & Chang, S. (1999). Image retrieval: Current techniques, promising directions and open issues. Journal of Visual Communication and Image Representation, 10(4), 39–62.
- Schaefer, G. (2004). CVPIC colour/shape histograms for compressed domain image retrieval. In DAGM 2004. vol. 3175 of LNCS (pp. 424–431). Tübingen, Germany.
- Schaefer, G., & Stich, M. (2004) UCID-An uncompressed colour image database. In Proc. SPIE Storage and Retrieval Methods and Applications for Multimedia (pp. 472–480). San Jose.
- Schmid, C., & Mohr, R. (1997). Local grayvalue invariants for image retrieval. IEEE Transactions on Pattern Analysis & Machine Intelligence, 19(5), 530–534.
- Shao, H., Svoboda, T., & van Gool, L.(2003a). ZuBuD–Zurich Buildings Database for image based recognition. Computer Vision Lab, Swiss Federal Institute of Technology, Switzerland. Zurich, Switzerland.
- Shao, H., Svoboda, T., Tuytelaars, T., & Gool, L. V. (2003b). HPAT indexing for fast object/scene recognition based on local appearance. In *Conference on Image and Video Retrieval. vol. 2728 of LNCS* (pp. 71–80). Urbana-Champaign: Springer Verlag.
- Shirahatti, N. V., & Barnard, K. (2005). Evaluating image retrieval. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR 05) (Vol. 1, pp. 955–961). IEEE: San Diego.
- Siggelkow, S. (2002). Feature histograms for content-based image retrieval. Ph.D. thesis. University of Freiburg, Institute for Computer Science. Freiburg, Germany.
- Siggelkow, S., Schael, M., & Burkhardt, H. (2001). SIMBA—Search IMages By Appearance. In DAGM 2001, Pattern Recognition, 23rd DAGM Symposium, vol. 2191 of Lecture Notes in Computer Science (pp. 9–17). Munich: Springer Verlag.
- Smeulders, A. W. M., Worring, M., Santini, S., Gupta, A., & Jain, R. (2000). Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(12), 1349–1380.
- Smith, J. R., & Chang, S. -F. (1996). Tools and techniques for color image retrieval. In SPIE Storage and Retrieval for Image and Video Databases (Vol. 2670, pp. 426–437).
- Squire, D. M., Müller, W., Müller, H., & Raki, J. (1999) Content-Based query of image databases, inspirations from text retrieval: Inverted files, frequency-based weights and relevance feedback. In *Scandinavian Conference on Image Analysis* (pp. 143–149). Kangerlussuaq.
- Sun, Y., Zhang, H., Zhang, L., & Li, M. (2002). MyPhotos-A system for home photo management and processing. In ACM Multimedia Conference (pp. 81–82). Juan-les-Pins.
- Swain, M. J., & Ballard, D. H. (1991). Color indexing. International Journal of Computer Vision, 7(1), 11–32.
- Tamura, H., Mori, S., & Yamawaki, T. (1978). Textural features corresponding to visual perception. IEEE Transaction on Systems, Man, and Cybernetics, 8(6), 460–472.
- Vailaya, A., Figueiredo, M. A. T., Jain, A. K., & Zhang, H. -J. (2001). Image classification for content-based indexing. *IEEE Transactions on Image Processing*, 10(1), 117–130.
- van Gool, L., Tuytelaars, T., & Turina, A. (2001). Local features for image retrieval. In R. C. Veltkamp, H. Burkhardt, H.-P. Kriegel (Eds), *State-of-the-art in content-based image and video retrieval* (pp. 21–41). Kluwer Academic Publishers.

- Vasconcelos, N., & Vasconcelos, M. (2004). Scalable discriminant feature selection for image retrieval and recognition. In CVPR 2004. 2 (pp. 770–775). Washington.
- Wang, J. Z., Li, J., & Wiederhold, G. (2001). SIMPLIcity: Semantics-sensitive integrated matching for picture LIbraries. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(9), 947–963.
- Yang, Z., & Kuo, C. (1999). Survey on image content analysis, indexing, and retrieval techniques and status report of MPEG-7. *Tamkang Journal of Science and Engineering*, 3(2), 101–118.
- Yavlinski, A., Pickering, M. J., Heesch, D., & Rüger, S. (2004). A comparative study of evidence combination strategies. In *IEEE International Conference on Acoustics, Speech, and Signal Processing* (ICASSP 2004) (Vol. 3, pp. 1040–1043). Montreal, Canada.
- Zahedi, M., Keysers, D., Deselaers, T., & Ney, H. (2005). Combination of tangent distance and an image distortion model for appearance-based sign language recognition. In DAGM 2005, Pattern Recognition, 26th DAGM Symposium, vol. 3663 of Lecture Notes in Computer Science (pp. 401–408). Vienna, Austria.