Image Indexing Using Prosemantic Features

Gianluigi Ciocca; Dept. of Informatics, Systems and Communication - Università degli Studi di Milano-Bicocca; Milan; Italy
Claudio Cusano; Dept. of Electrical, Computer and Biomedical Engineering - Università degli Studi di Pavia; Pavia; Italy
Simone Santini; Escuela Politécnica Superior - Universidad Autónoma de Madrid; Madrid; Spain
Raimondo Schettini; Dept. of Informatics, Systems and Communication - Università degli Studi di Milano-Bicocca; Milan; Italy

Abstract

We present here, an image description approach based on prosemantic features. The images are firstly represented by a set of low-level features related to their structure and color distribution. Those descriptions are fed to a battery of image classifiers trained to evaluate the membership of the images with respect to a set of 14 overlapping classes. Prosemantic features are obtained by packing together the scores. In this paper we will show how prosemantic features outperform traditional low-level features in a variety of tasks. One is content-based retrieval: we included prosemantic features into the framework of the QuickLook^2 image retrieval system. Target search experiments show that the use of prosemantic features, combined with the relevance feedback mechanism of QuickLook^2, allows for a more successful and quick retrieval of the query images with respect to low-level features. Moreover, we will show the effectiveness of our features for the browsing and visualization of the results obtained from image search engines.

Introduction

The gap between the pictorial features and the image's semantics makes it difficult for purely content-based retrieval systems to obtain satisfactory results. To overcome the necessity of manually describing the images content, many systems are essentially based on low-level image features that are directly and automatically computed from the images themselves. However, the use of low-level features can't overcome the gap between the content and the semantic of the images [19]. Some systems explicitly extract and embed in the retrieval process semantic information about the image content by exploiting automatic classification techniques [13]. These techniques can then be employed to automatically annotate the image content by keywords, which are then used in the retrieval process. If the underlying annotation is reliable, text-based image retrieval can be semantically more meaningful than other indexing approaches [10].

These annotation approaches can be considered as crisp annotation: if an image is annotated with a given label then the image expresses that concept or belong to that class. One of the first works that try to bring semantic information under the same model vector paradigm used in query-by-example systems is [20]. Semantic information is learned directly from the image content and forms a vector of semantic weights. Each weight is associated to a concept and is derived from the confidence score obtained by a support vector machine trained to recognize that concept. Retrieval in the semantic space corresponds to performing a similarity comparison between two model vectors using the L_2 measure. A similar approach is followed in [16].

Following a similar paradigm, we designed an approach to CBIR based on the information provided by several image classifiers. One of the main problems in integrating automatic image classification into a content-based retrieval system is the choice of classes. It is very hard to identify a set of categories that are representative of the majority of the pictures and that can be used to reliably approximate their semantics. Moreover, state of the art image classification systems are far from perfect and, consequently, their use in image retrieval requires a high degree of tolerance with respect to misclassification errors.

To circumvent these problems, we did not exploit the classifiers to obtain a "crisp" semantic description of the images (e.g. "sunset on the beach"), but rather to provide a rich description of visual content that correlates low-level features to prototypical scenes (e.g. "image with an edge distribution that can easily be found in seaside scenes"). In our approach, this level of description is provided by a set of *prosemantic* features. These features are obtained by training several image classifiers so designed that their output can be interpreted as membership values of an image in the class that they embody.

In this paper we will discuss the design and implementation of prosemantic features and we will summarize the results obtained in image retrieval, browsing, and visualization. We will conclude with an overview of our current and future work on this topic.

Prosemantic features

Figure 1 shows the process of the prosemantic features extraction. Prosemantic features extraction begins by describing the images with a suitable set of "low-level" features. As low level features we considered: color mean and standard deviation of the values of the LUV color channels on 9 image subregions, global color histogram in the RGB color space, statistics about the direction of edges and the descriptors associated to the Scale Invariant Feature Transform [15]. More details on how these feature have been computed can be found in [4, 5].

In order to provide a semantically meaningful information about the content of the images, each feature is used as input to an array of 14 soft classifier, trained to recognize partially overlapping classes. We selected a set of 14 classes: animals, city, close-up, desert, flowers, forest, indoor, mountain, night, people, rural, sea, street, and sunset. Some classes describe the image at a scene level (city, close-up, desert, forest, indoor, mountain, night, rural, sea, street, sunset) other describe the main subject of the picture (animals, flowers, people). The set of classes is not meant to be exhaustive, or to be able to characterize the content of the images with sufficient specificity for our purposes. Our intent,
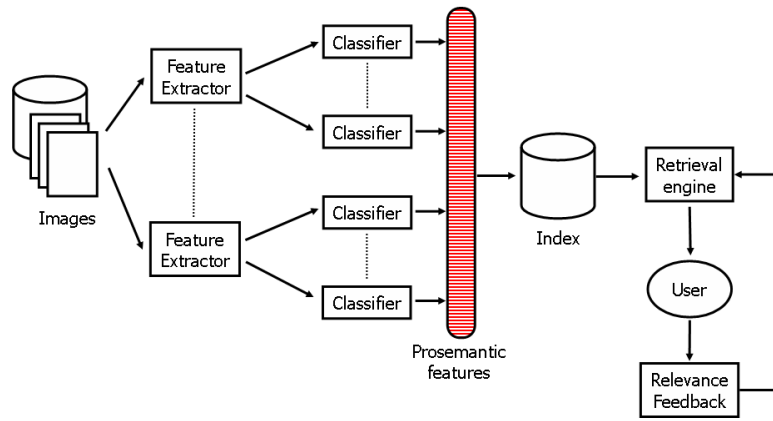
Figure 1: Prosemantic features extraction.

here, was to select a variegated set of concepts providing a wide range of low-level descriptions of typical scenes. The fact that the categories are overlapping is a practical choice reinforced by an intuition based, in turn, on an analogy. As a matter of praxis, it would be problematic, next to impossible in fact, to find a reasonably extended collection of categories that show no overlap, not in the least because the very concept of "semantic overlap" is all but well defined, and can be used, at best, as a generic regulative principle.

In order to collect suitable training samples for the classifiers, we queried various image search engines on the web with several keywords related to the classes, and downloaded the resulting pictures. The images were then manually inspected in order to remove those that did not belong to the classes as well as low quality images. For each class, a set of negative examples was also selected by taking pictures from the other classes. Since the classes may overlap, a manual inspection was needed to verify that all the selected images were actually negative examples.

For each combination of low-level feature and class, a Support Vector Machine (SVM) with a Gaussian kernel has been trained. There are two parameters that need to be tuned (the cost parameter C and the scale of the Gaussian kernel γ), and they have been selected by maximizing the cross validation performance of the resulting classifier.

At the end of training, we have a distinct SVM for each feature and for each class. Given a class c and a new image Q, represented by the feature vector $\mathbf{x}_Q$, the SVM provides a class-membership score:

$$s(\mathbf{x}_Q, c) = b^{(c)} + \sum_{I \in T^{(c)}} \alpha_I^{(c)} y_I^{(c)} k(\mathbf{x}_I, \mathbf{x}_Q), \tag{1}$$

with the kernel k defined as

$$k(\mathbf{x}_I, \mathbf{x}_Q) = \exp\left(-\gamma^{(c)} \|\mathbf{x}_I - \mathbf{x}_Q\|^2\right), \tag{2}$$

where $T^{(c)}$ is the training set for class c, $\mathbf{x}_I$ denotes the feature vectors computed on the image I, $y_I^{(c)}$ is the label in $\{-1, +1\}$ which indicates whether I is a positive or a negative example, $b^{(c)}$ and $\alpha_I^{(c)}$ are the parameters determined by the training procedure, and $\gamma^{(c)}$ is the scale parameter of the kernel. The score is expected to be positive when the image belongs to the class c, and negative otherwise. Packing together the 56 scores we obtain a compact vector of prosemantic features.

Image retrieval

When embedded into an image retrieval system, prosemantic features demonstrated remarkable capabilities. In particular, we used the retrieval functionalities of the QuickLook[2] system which is based on a relevance feedback mechanism [4]. By exploiting the statistical analysis of the image feature distributions of the retrieved items the user has judged relevant, or not relevant, the system is able to identify what features the user has taken into account (and to what extent) in formulating his judgment. The use of this information is twofold: to reformulate the user's query and to modify the image similarity measure.

The query is computed from the feature values that mostly agree with the user selection, while the outliers are removed from the computation (*query refinement*). Let R_+ be the set of relevant images, and $\vec{x}_I(k)$ be the k-th value of a feature vector of image I, the components of the query Q are computed from each feature vector as the average of the elements in the following sets:

$$\{\vec{x}_I(k): |\vec{x}_I(k) - v_k| \leq 3\sigma_k, I \in R_+\}, \tag{3}$$

where v_k and and σ_k are the average of the k-th values of the feature in R_+. Image dissimilarity is assessed using a weighed sum of dissimilarities between image features. The influence of relevant features is accentuated while the influence of non-relevant features is damped (*feature reweighing*) by analyzing the relevant and non-relevant images. Let R_- the set of non relevant images. The weight for a given feature is computed as:

$$w = \frac{1}{\varepsilon + \mu_+} - \alpha \frac{1}{\varepsilon + \mu_*}, \tag{4}$$

where ε and α are positive constants, μ_+ is the average of the dissimilarities computed on the feature between each pair of images in R_+, and μ_* the average of the dissimilarities computed on the feature between each image in R_+ and each image in R_-. Negative weights are set to 0.

Quantitative evaluation

A user study has been conducted to evaluate the performance of our prosemantic features against the corresponding low-level ones. For our purpose, we substituted the original features in the QuickLook[2] system with ours and asked 40 subjects to perform ten target search retrieval sessions. All subjects came from the computer science department of the University of Milan - Bicocca: four of them have a background on image processing or
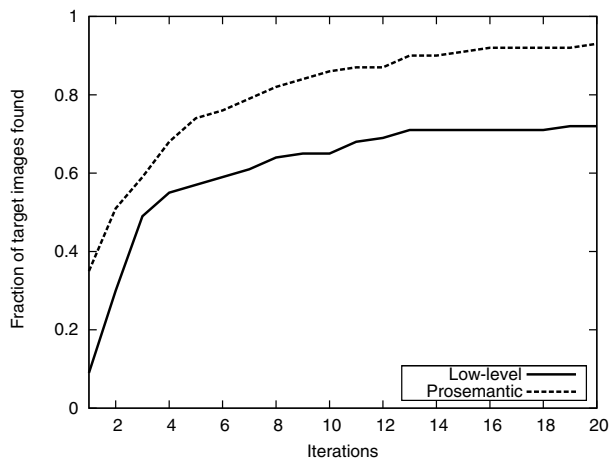
Figure 2: Fraction of images successfully retrieved as a function of the number of iterations.

computer vision (two Ph.D. students and two post-doctoral fellows), the other 16 are graduate (three) or undergraduate (13) students. Each subject was constrained to retrieve the target image by selecting any number of relevant and not relevant images within the top 60 retrieved images. The subjects were oblivious to what kind of features they were currently using.

At the end of the user study, each target image was searched ten times for each of the four different sets of features considered: prosemantic, preclassification (i.e. low-level features used to build prosemantic features), the default features of QuickLook[2], and a combination of QuickLook[2] and prosemantic features.

The dataset used consists of 1875 images taken from the Benchathlon dataset [14]. The dataset includes typical consumer photographs showing a very different distribution of concepts with respect to the dataset used to train the classifiers. For instance, very often the image would fall in the "people" class, while very few images can be considered as belonging to the "desert" or "flowers" classes. The target images have been randomly selected. Other 60 images have been randomly selected to compose the page from which the users started all their searches.

The outcome of the 200 searches clearly demonstrates the effectiveness of prosemantic features with respect to low-level features. Using the prosemantic features, only seven times were the users not able to retrieve the target images within the limit of 20 retrieval operations. By contrast the limit has been exceeded 49 times in the case of low-level features. Figure 2 shows the cumulative success rate for the two sets of features as a function of the number of iterations. The plot shows how prosemantic features allows the retrieval of more target images and with less iterations with respect to the alternatives considered (random browsing has been added as a baseline).

Qualitative evaluation

A qualitative evaluation of the proposed prosemantic features "on-the-field" has been carried out with the contribution of Fratelli Alinari Photo Archive[1] that supplied us a subset of images extracted from their extensive photo archives [9]. The dataset of images they supplied is composed of 3,000 images (b/w

---

[1] http://www.alinari.com/

and color), mainly regarding cityscape, landscape, art, painting and sculptures. This dataset has been taken in the center of Italy (Tuscany region) from locations of high cultural interest such as Firenze, Pisa, and San Giminiano among others. The images depict sculptures, palaces, plazas, and various other artifacts taken from different perspectives and sometimes under different illuminations. Some images represent a single object captured from a distance, while other images of the same subject have been taken much more closer (close-up). In some instances, whole panoramas of the surroundings have been acquired as well.

This dataset has been selected in order to evaluate how QuickLook[2] can perform with a specific genre, with a mix of contemporary color and historical b/w images, and to test if and how it can find images of specific objects or part of them.

The experiments have been conducted at the Alinari Archives where a copy of QuickLook[2] has been installed. Users have been asked to perform queries without supplying them with a specific task: they were free to choose the type and aim of their searches. For the purpose of the experiments, image examples are selected only from the first page of the retrieved results.

Different users tested the system for a few days. A questionnaire in two parts was administered to these users in order to collect their impression about the system on the overall and on its functionalities. The first part was inspired by the System usability Scale (SUS) questionnaire developed by John Brooke at DEC (Digital Equipment Corporation) [2]. It is composed of statements related to different aspects of the experience, and the subjects were asked to express their agreement or disagreement with a score taken from a Likert scale of 5 numerical values: 1 expressing strong disagreement with the statement, 5 expressing strong agreement and 3 expressing a neutral answer. The second part of the questionnaire focuses more on the functionalities of the QuickLook[2] system and was administered with the same modalities.

The results of the questionnaire are reported in Table **??**. The score given by the users are summarized by majority vote. The results show that on the overall the system perform well. The retrieval capabilities of QuickLook[2] are judged positively as well as its efficiency and efficacy "The system is robust, fast to manage, and speedy in the query mechanism". The relevance feedback mechanism coupled with the prosemantic features is efficient and is able to retrieve satisfactory results in few iterations and requiring the users to select a moderate amount of image examples. Retrieval by examples is still considered a plus in a retrieval engine that cope mainly with images. With respect to the users' experiences the weakest point of the system is the graphical user interface. Although the selection of the images is quite intuitive, the other components of the user interface has been rated poorly.

Browsing and Visualization

Several works addressed the problem of visualizing sets of images on the basis of the associated metadata [11, 21, 3]. However, metadata are already taken into account by the search engines while visual content is usually ignored. Therefore it may be considered as an additional source of information which can be exploited to provide a more convenient and efficient way to browse the results of the queries.

Content-based approaches have been used for the visualization of large database of images [17, 18, 1]. Since indexing is

Table 1: The usability questionnaire administered to the QuickLook[2] users. The numerical scale goes from 'strongly disagree' to 'strongly agree' (1 → 'strongly disagree', 2 → 'disagree', 3 → 'neutral', 4 → 'agree', 5 → 'strongly agree').

#	Statement	1	2	3	4	5
1.	I think that I would like to use this system frequently					⋆
2.	I found the system unnecessarily complex			⋆		
3.	I thought the system was easy to use		⋆			
4.	I would need the support of a technical person to be able to use this system					⋆
5.	I found the various functions in this system were well integrated				⋆	
6.	I thought there was too much inconsistency in this system		⋆			
7.	I would imagine that most people would learn to use the system very quickly			⋆		
8.	I found the system very cumbersome to use		⋆			
9.	I felt very confident using the system				⋆	
10.	I needed to learn a lot of things before I could get going with this system		⋆			
11.	The queries are executed rapidly					⋆
12.	Too many iterations are required to obtain acceptable results		⋆			
13.	The system allows effective category searches				⋆	
14.	The system allows effective target searches				⋆	
15.	The system is useful for browsing images					⋆
16.	The system is useful for the retrieval of images					⋆
17.	The system user interface is easy to understand	⋆				
18.	The relevance feedback mechanism is too complicated	⋆				
19.	I think that retrieval by image examples is useless	⋆				
20.	Too many positive examples must be selected		⋆			
21.	Too many negative examples must be selected		⋆			

performed off-line, very powerful methods can be used to extract rich descriptions of the content of the images, and to analyze their distribution in the feature space.

We used prosemantic features to implement a real-time, content-based prototypical system which queries an image search engine, and displays the results in such a way that the user can understand "at a glance" the composition of the set of retrieved images [6]. The small dimensionality of the feature space allows for an efficient analysis of the distribution of the images. As a result, the whole visualization process can be performed in real time.

Images are displayed on a plane in such a way that similar images (either visually or semantically) are grouped in the same region of the plane. Images are not allowed to cover each over: those which would overlap are recursively grouped into clusters which can be explored independently.

The execution of each query is composed of three major steps:

- query processing;
- image description;
- thumbnail placement and visualization.

Each step is performed asynchronously on each image, which is shown to the user as soon as its processing is finished.

In the query processing step the text entered by the user is sent to an image search engine such as Google images, flickr®... The search engine returns a set of images whose content is described by a set of prosemantic features.

To reduce the computational load and allow a fast visualization, SIFT descriptors have been removed from the set of features thus obtaining a reduced 42-dimensional prosemantic space. Unfortunately, computer monitors are limited to two dimensions only. Therefore, we need to further reduce the dimensionality of the space. For this purpose, we decided to adopt a standard, widely used technique: Principal Component Analysis (PCA).

We investigated two different approaches: in the first the PCA basis is computed on-line on the prosemantic features extracted on the downloaded images; in the second approach the basis has been precomputed on a large dataset of images (here, we used the training set of the Pascal VOC 2007 challenge [12]).

The visualization in the subspace of the first two principal components allows showing to the users the result of their queries "at a glance". However, this visualization strategy is far from optimal: some regions of the space can result too crowded with an unwanted overlapping of the images which can also be completely hidden under their neighbors.

In order to simplify the presentation of the images we applied a quantization to the transformed prosemantic space. Neighbor images are recursively grouped in clusters, and only a representative example of the cluster is shown to the user (who can select the cluster to explore its content). The quantization process works as follows:

- a set of seed points is generated according to a hexagonal grid which uniformly subdivides the plane of the first two principal components;
- images are grouped into clusters according to the nearest seed point;
- for each cluster, the nearest image is selected as representative;
- the representative images are shown at the coordinates of the corresponding seed point.

Note that some seed point may correspond to an empty cluster; in this case an empty space will be left on the screen.

Figure 3 shows the search results of flickr® for the query "apple", sorted by decreasing relevance. Figure 4 shows the first 200 images visualized following our strategy. The downloaded images can be placed in different semantic categories: images representing fruits have been placed on the top while images with a technological theme have been placed on the bottom part of the screen.
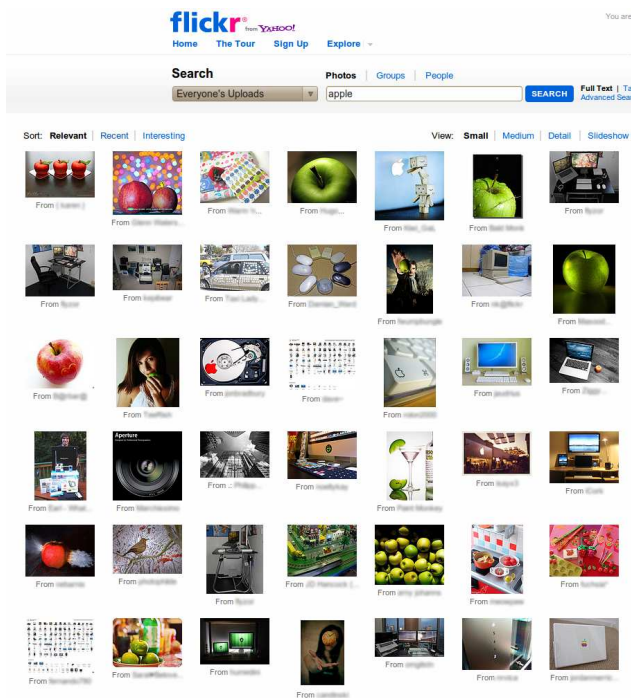
Figure 3: Results for the query "apple" submitted to the flickr® service.

Figure 4: Visualization of the 200 images which correspond to the query "apple". Overlapping images have been grouped into clusters (singleton clusters are displayed without the indication of their size).

Further investigations

In this section, we briefly illustrate two other applications of the prosemantic features.

Unsupervised categorization, often done through the use of clustering algorithms, is one of the most powerful techniques available to the designer of image management systems, as it allows categorization with no other information than that contained in the data themselves. Grouping images into semantically homogeneous classes is often a sine qua non for efficiently processing, structuring, querying, and browsing large collections of images. Prosemantic features have been tested for the task of unsupervised image categorization. We performed comparison with different state-of-the-art clustering algorithms on various standard data sets and against supervised features as well as low-level fea-

tures. Results showed that prosemantic features exhibit a remarkable expressiveness which allows to effectively group images into the categories defined by the data sets authors. Detailed results can be found in [8].

Image databases may include several nearly duplicated images that contain more or less the same information. If one of these images is very relevant with respect to the query, it is likely that all of them will be, and that the result set will be composed of very similar images. Although formally relevant, each of these images adds little information to what one already has with just one of them. To address this issue, the information retrieval community introduced the concepts of diversity and novelty. Diversity is the notion that allows the result set to deal with queries which can have several interpretations. Given an interpretation of the query there may be different aspects in which the user may be interested. An image is novel to the extent in which it covers aspects of a query not covered by other images in the result set, that is, to the extent in which images are not redundant. To maximize novelty and diversity, we implemented two algorithms based on the prosemantic features. Description of the algorithms can be found in [7].

Conclusions

We presented here, an image description approach based on prosemantic features. According to the experimental results, prosemantic features outperform traditional low-level features in a variety of tasks, including image retrieval, browsing and visualization. One of the most common task for visual features is represented by image classification. Experiments showed that prosemantic features can be used to obtain remarkable performance in supervised image classifications and unsupervised image classification. In fact, they outperformed low-level features on a variety of standard benchmarks and other reference high level features.

On the basis of the results obtained we can conclude that prosemantic features can be successfully used for different tasks. Currently, these features have been heuristically defined. In future work we will take advantage of the insights provided by the results obtained so far in order to investigate how to identify the categories that form the base for the prosemantic features as well as how many categories are required to obtain good results.

Acknowledgments

References

[1] R. Achanta, A. Shaji, P. Fua and S. Süsstrunk, Image summaries using database saliency, ACM SIGGRAPH ASIA 2009 Posters, pg. 1. (2009).

[2] J. Brooke, SUS: A Quick and Dirty Usability Scale, In P. W. Jordan, B. Thomas, B. A. Weerdmeester, and I. L. McClelland, Usability Evaluation in Industry. Taylor & Francis, London, 1996.

[3] M. Chang and J. Leggett, Collection understanding through

streaming collage, Proc. of the Information Visualization Interfaces for Retrieval and Analysis (IVARA) Workshop, (2003).

[4] G. Ciocca, C. Cusano, S. Santini, and R. Schettini, Halfway through the semantic gap: prosemantic features for image retrieval, Information Sciences, 181, 4943 (2011).

[5] G. Ciocca, C. Cusano, S. Santini, and R. Schettini, Prosemantic features for content-based image retrieval, LNCS on Adaptive Multimedia Retrieval. Understanding Media and Adapting to the User, pg. 87. (2011).

[6] G. Ciocca, C. Cusano, S. Santini, and R. Schettini, Spatially organized visualization of image query results, Proc. of SPIE Multimedia Content Access: Algorithms and Systems V, pg. 788116. (2011).

[7] G. Ciocca, C. Cusano, S. Santini, and R. Schettini, Prosemantic image retrieval, LNCS Computer Vision – ECCV 2012. Workshops and Demonstrations, pg. 643. (2012).

[8] G. Ciocca, C. Cusano, S. Santini, and R. Schettini, On the use of supervised features for unsupervised image categorization: an evaluation, Computer Vision and Image Understanding, in print, (2014).

[9] G. Ciocca, C. Cusano, S. Santini, R. Schettini, A. de Polo, and F. Tavanti, Searching through photographic databases with quicklook, Multimedia on Mobile Devices 2012; and Multimedia Content Access: Algorithms and Systems VI, pg. 83040V. (2012).

[10] R. Datta, J. Li, and J.Z. Wang, Content-based image retrieval: approaches and trends of the new age, Proc. of the 7th ACM SIGMM international workshop on Multimedia information retrieval, pg. 253. (2005).

[11] M. Dontcheva, M. Agrawala, and M. Cohen, Metadata visualization for image browsing, In ACM Symposium on User Interface Software and Technology, (2005).

[12] M. Everingham, L. Van Gool, C. Williams, J. Winn and A. Zisserman, The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results, http://www.pascal-network.org.

[13] J.Fan, Y. Gao, H. Luo, and G. Xu, Automatic image annotation by using concept-sensitive salient objects for image content representation, Proc. of the 27th annual international ACM SIGIR conference on Research and development in information retrieval, pg. 361, (2004).

[14] N.J. Gunther and G. Beretta, A benchmark for image retrieval using distributed systems over the internet: BIRDS-I, Technical Report HPL-2000-162, HP Labs, Palo Alto, (2001).

[15] D.G. Lowe, Distinctive image features from scale-invariant keypoints, International Journal of Computer Vision, 60, 91 (2004).

[16] J. Lu, S.P. Ma, and M. Zhang, Automatic image annotation based on model space, Proc. of IEEE International Conference on Natural Language Processing and Knowledge Engineering, pg. 455, (2005).

[17] G.P. Nguyen and M. Worring, Interactive access to large image collections using similarity-based visualization, Journal of Visual Languages & Computing, 19, 203 (2008).

[18] D.S. Ryu, W.K. Chung, and H.G. Cho, PHOTOLAND: a new image layout system using spatio-temporal information in digital photos, Proc. of the 2010 ACM Symposium on Applied Computing, pg. 1884, (2010).

[19] A.W.M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain, Content-based image retrieval at the end of the early years, IEEE Transactions on Pattern Analysis and Machine Intelligence, 22, 1349 (2000).

[20] J.R. Smith, M.Naphade, and A. Natsev, Multimedia semantic indexing using model vectors, Proc. of IEEE International Conference on Multimedia and Expo, pg. 445, (2003).

[21] K. Yee, K. Swearingen, K. Li, and M. Hearst, Faceted metadata for image search and browsing, Proc. of the SIGCHI conference on Human factors in computing systems, pg. 408, (2003).

Author Biography

Gianluigi Ciocca took his degree (Laurea) in Computer Science at the University of Milan in 1998, and since then he has been a fellow at the Institute of Multimedia Information Technologies of the Italian National Research Council, where his research has focused on the development of systems for the management of image and video databases and the development of new methodologies and algorithms for automatic indexing. He is currently assistant professor in computer science at DISCo (Dipartimento di Informatica, Sistemistica e Comunicazione) of the University of Milano-Bicocca, working on video analysis and abstraction.

Claudio Cusano is assistant professor at the Dep. of Electrical, Computer and Biomedical Engineering of the University of Pavia. He took his Ph.D. in 2006 at the the University of Milano-Bicocca. Since April 2001 he has been a fellow of the the ITC Institute of the Italian National Research Council. The main topics of his current research concern 2D and 3D imaging, with a particular focus on image analysis and classification, and on face recognition.

Simone Santini received his Ph.D. from the University of California in 1998 and, since 2005, he is a professor at the Universidad Autnoma de Madrid. In the past, he has published extensively on computer vision, image and video retrieval, and multimedia semantics. His current research interests are in the use of context for multimedia search and in the development of semantic features for image description.

Raimondo Schettini is a professor at the University of Milano Bicocca (Italy). He is Vice-Director of the Department of Informatics, Systems and Communication, and head of Imaging and Vision Lab (www.ivl.disco.unimib.it). He has been associated with Italian National Research Council (CNR) since 1987 where he has leaded the Color Imaging lab from 1990 to 2002. He has been team leader in several research projects and published more than 250 refereed papers and six patents about color reproduction, and image processing, analysis and classification. Raimondo Schettini has been elected Fellow of the International Association of Pattern Recognition (IAPR) for his contributions to pattern recognition research and color image analysis.