

Aligning codebooks for near duplicate image detection

Sebastiano Battiato · Giovanni Maria Farinella ·
Giovanni Puglisi · Daniele Ravi

Published online: 21 April 2013
© Springer Science+Business Media New York 2013

Abstract The detection of near duplicate images in large databases, such as the ones of popular social networks, digital investigation archives, and surveillance systems, is an important task for a number of image forensics applications. In digital investigation, hashing techniques are commonly used to index large quantities of images for the detection of copies belonging to different archives. In the last few years, different image hashing techniques based on the Bags of Visual Features paradigm appeared in literature. Recently, this paradigm has been augmented by using multiple descriptors (e.g., Bags of Visual Phrases) in order to exploit the coherence between different feature spaces. In this paper we propose to further improve the Bags of Visual Phrases approach considering the coherence between feature spaces not only at the level of image representation, but also during the codebook generation phase. Also we introduce a novel image database specifically designed for the development and benchmarking of near duplicate image retrieval techniques. The dataset consists of more than 3,300 images depicting more than 500 different scenes having at least three real near duplicates. The dataset has a huge variability in terms of geometric and photometric transformations between scenes and their corresponding near duplicates. Finally, we suggest a method to compress the proposed image representation for storage purposes. Experiments show the effectiveness of the proposed near duplicate retrieval technique, which outperforms the original Bags of Visual Phrases approach.

S. Battiato · G. M. Farinella (✉) · G. Puglisi · D. Ravi
Department of Mathematics and Computer Science, Image Processing Laboratory,
University of Catania, Viale A. Doria 6, Catania 95125, Italy
e-mail: gfarinella@dmi.unict.it

S. Battiato
e-mail: battiato@dmi.unict.it

G. Puglisi
e-mail: puglisi@dmi.unict.it

D. Ravi
e-mail: ravi@dmi.unict.it

Keywords Image forensics · Near duplicate images · Image retrieval · Bags of visual words · Bags of visual phrases · Codebooks alignment

1 Introduction and motivations

Image Forensics is a science which, among the other questions, aims to answer the following one during investigation: is the image under consideration contained in a specific digital archive? The increasing use of low cost imaging devices and the availability of large databases of digital photos makes the near duplicate image retrieval (NDIR) task a common activity for a number of applications. In particular, NDIR in large databases (such as popular social networks, collections of surveillance images and videos, or digital investigation archives) is a key ingredient for different forensics activities.

During digital investigation (e.g., for copyright violation, child abuse, etc.), classic hashing techniques (e.g., MD5 [32], SHA1 [12], etc.) are commonly used to index large quantities of images in order to detect copies in different archives. However, these methods are unsuitable to find altered copies, even in case of slight modifications (e.g., near duplicates). Indeed, classic hashing techniques usually fail because just a small change in the image (even a single bit) will, with overwhelming probability, results in a completely different hash code. For example, two images depicting a scene of crime are perceptually identical under small viewpoint changes, partial occlusion, and/or low photometric distortions, but their hash code is completely different when a classic hashing approach is used to check their similarity. In order to cope with all related problems, robust hashing techniques based on image content must be developed: perceptually identical images in terms of content should have the same (or at least very similar) hash value with high probability, while perceptually different images should have independent hash values.

Most of the near duplicate detection techniques based on image content exploit the bag of visual word approach to build the image signature [2, 15, 44, 49, 50]. A problem of the bag-of-visual-word based methods is related to the ambiguity of some generated visual words [39, 42]. On the other hand, since different descriptors represent different aspects of a local region, there is no single descriptor which is superior to the others [28]. Here we propose to combine different descriptors through an alignment procedure based on clusters correspondence (i.e., number of shared local regions). The advantage of the proposed codebook alignment method is related to the enforcement of the coherence across multiple descriptors in order to capture different aspects of the considered local region (e.g., shape, texture, etc.) and hence reduce both, the visual word ambiguity and the quantization error in the visual codebook generation [15, 22]. The different aspects of a local regions are captured by the alignment during the codebook generation in the sense that the local regions falling in the intersection of two aligned clusters agree with respect to both descriptors, whereas the others agree just with one descriptor and not with the other. Taking into account such peculiarity, we split the clusters of each feature domain obtaining new codebook prototypes which consider the intersecting part of the aligned clusters, as well as the part which not intersect. Since, by using multiple descriptors there is an overhead in terms of storage of the representations of the

images, and considering that image datasets are becoming more and more popular and huge (i.e., Facebook proceeds at a rate of about 22,000 uploads per minute), we also propose a method to compress the image representation by maintaining performances in terms of near duplicate image detection accuracy.

The remainder of the paper is organized as follows: Section 2 gives a brief survey of the related work. Section 3 describes the proposed model. In Section 4 the method to compress the image descriptors is suggested. The dataset built for testing purposes is described in Section 5, whereas Section 6 details the experimental settings and reports the obtained results. Finally, conclusions and avenues for further research are given in Section 7.

2 Related works

Recently, some commercial approaches for robust content based hashing methods have been proposed for photos (PhotoDNA, <http://www.microsoftphotodna.com/>) and videos (Videntifier [23]). These techniques make use of the recent developments in the field of Near Duplicate Image (NDI) retrieval. Note that there is no agreement on the technical definition of near-duplicates (see [11] for an in-depth discussion). The definition of near duplicate depends on the degree of variability (photometric and geometric) that is considered acceptable for each particular application. Some approaches [19] consider as NDI, images obtained by slightly modifying the original ones through common transformations such as changing contrast or saturation, scaling, cropping, etc. Other techniques [15] consider as NDI, images of the same scene but with different viewpoint and illumination. A drawback in testing near duplicate retrieval approaches is that usually near duplicate images used in the experiments are synthetically generated from a set of images or correspond to different frames of a video, hence there is an high correlation in terms of visual content, and there is no variability in terms of resolution and compression. To better evaluate the different algorithms it is needed a database composed by images depicting the same scene and/or subject whose have been acquired by different cameras, with different viewpoint, luminance condition, and variability in terms of background.

In the last few years, different image hashing techniques have been proposed in literature to cope with image retrieval and near-duplicate image detection problems. Most of these techniques are based on the Bags of Visual Words paradigm (BoVW) [24, 41] to build a holistic representation of the images.

Ke et al. [19] detected near-duplicate images by employing local descriptors [28] extracted on interest points [27] to represent and match images under several transformations. They used a hash-based indexing technique to efficiently search into the image databases, and also applied an optimized storage layout to further improve efficiency. Chum et al. [9] proposed two novel image similarity measures for image indexing through local feature descriptors and enhanced min-Hash techniques. The authors of [10] introduced a method to combine visual words with geometric information to improve hashing-based image retrieval and object detection, obtaining a novel algorithm (called Geometric min-Hash) which shows significant advantages against geometrical deformations and occlusions. Cheng et al. [8] considered local dependencies among descriptors both in scale and space, and encoded not only visual appearance but also their scale and space co-occurrence. Moreover they

built SuperNodes that embody the neighbor information to speed up the retrieval execution time. Wu et al. [44] proposed a novel scheme to exploit geometrical constraints by spatially grouping features to improve the retrieval precision. Spatial verification stage to re-rank the results with the bag of-words model have been also exploited by Philbin et al. in [31]. Wang et al. [43] combined appearance-based and keypoint-based methods. The algorithm is able to extract and match keypoints from images by discarding outliers through a voting procedure based on the affine invariant ratio of normalized lengths. Later, to further validate the correspondences, the algorithm compares the color histograms of the corresponding areas which have been previously identified by the matched points. In [46], Xu et al. proposed a two stage method based on the Spatial Pyramid Matching (SPM) technique [21] and image blocks alignment through linear programming [45]. The aim of the cascade is to deal with spatial shifts and scale variation; both transformations frequently occur between frames of a video. Since there is an increasing interest in the scalability of the Bag-of-Words based near duplicate visual search paradigm, a method to parallelize the near duplicate visual search architecture to index images over multiple servers have been proposed by Rongrong et al. in [35]. Near duplicate image retrieval based on BoVW paradigm has been also exploited for the annotation of web videos as addressed by Zhao et al. in [50, 51]. Taking into account their previous work [49], the authors extract keypoints from keyframes and generate a visual dictionary by using a clustering algorithm. Each keyframe is then described by a BoVW representation. Moreover, to speed up the keyframe retrieval, inverted file indexing plus Hamming embedding is employed. A re-rank strategy based on a weak geometric consistency checking is also proposed to improve the overall performance of the system. The final similarity of a video is obtained considering both the scores of keyframes and their temporal consistency with respect to the query video. The memory usage and query-response time are two of the main issues in the retrieval task. The problem of compressing the visual codebook to better handle with storage and retrieval complexity has been studied by Rongrong et al. in [34].

In a recent study of Hu et al. [15], the BoVW paradigm has been augmented by using multiple descriptors (Bags of Visual Phrases) to exploit the coherence between different feature spaces in which local image regions are described. Specifically, to reduce the amount of false matchings in the BoVW model the authors of [15] introduced the coherent phrase model. In this model, a local image region (i.e., the patch surrounding the local interest point [27]) is described by a visual phrase of multiple descriptors instead of a visual word of a single descriptor. In the Bags of Visual Phrases approach, both feature (local regions are described by descriptors of different types) and spatial coherence (multiple descriptors are obtained from local areas at different sizes) are taken into account.

To further improve the Bags of Visual Phrases model, taking into account our preliminary work [3], we propose to exploit the coherence between feature spaces not only in the image representation, but also during the generation of codebooks. This is obtained by aligning the codebooks of different descriptors to produce a more significant quantization of the involved spaces of descriptors, which leads to a more distinctive representation. In particular, to reduce the amount of false matchings, instead of separately obtain the codebooks corresponding to the different feature spaces as proposed in [15], we generate the final codebooks taking into account the correspondence of the clusters of the involved spaces of descriptors to further enforce

feature correspondence. To properly perform tests, a new image database of near duplicate images has been built by collecting images from Flickr (<http://www.flickr.com>) and private collections. The dataset contains 3,148 images of 525 different scenes which have from 3 to 34 real near duplicates. Finally, a method to compress the image representation to be stored for near duplicate purposes is suggested. The experiments performed on the aforementioned dataset show the effectiveness of the proposed approach, which obtains a good margin of performances with respect to the approach described in [15].

3 Proposed model

Most of the image hashing techniques for near-duplicate image detection problems typically represent images through feature vectors encoding color, texture, and/or other visual cues such as corners, edges or local interest points [4, 5, 13, 17, 20, 25–28, 33]. These information are automatically extracted using several algorithms and then represented by many different local descriptors. Most of these techniques are based on the Bags of Visual Words paradigm (BoVW) [41] to build a global representation of the visual content within the images. The basic idea is to consider images as visual documents composed of repeatable and distinctive visual elements, which are comparable to the words in texts. Indeed, the BoVW originates in the text categorization community [37] where it was used to describe documents by how many words (belonging to a pre-built vocabulary) occur within them. Each word embracing a semantic meaning, has an inherent set of topics where it is used more often than others. To exploit this model in computer vision and multimedia, a vocabulary of distinctive patterns, usually called “visual words”, is built through a clustering approach from a set of local descriptors [5, 13, 17, 20, 25, 28] extracted in correspondence of interest points [4, 26, 27, 33] which have been previously detected on images of a training database. A local descriptor encodes properties of the region surrounding the interest point in the image from which have been generated. Hence, the “visual words” obtained by clustering the training set of local descriptor are used to identify properties, structures and textures present in the images whose are finally described as an unordered set (a bag) of “visual words”. Specifically, each image is represented as a normalized histogram whose bins correspond to “visual words” of the built codebook. Since the bag of visual words description is compact, it is suitable to represent huge image databases.

The proposed approach is built by taking into account the coherent phrase model introduced in [15], where the BoVW paradigm has been augmented by using multiple descriptors (called Bags of Visual Phrases Model) to exploit the coherence between different feature spaces (i.e., local descriptors) in which the local image regions corresponding to interest points are described. To further improve the Bags of Visual Phrases model, we propose to exploit the coherence between feature spaces (i.e., local descriptors) not only in the image representation step (e.g., using a two dimensional distribution of co-occurrence of visual words of codebooks corresponding to two different feature spaces), but also during the generation of codebooks. This is obtained by aligning the codebooks of different descriptors to produce a more significant quantization of the involved spaces of descriptors, which leads to a more distinctive image representation. Differently than Hu et al. [15],

we do not obtain the final codebooks corresponding to the different feature spaces separately, but we generate the final codebooks taking into account the correspondence of the clusters of the involved spaces of descriptors to further enforce feature correspondence. Specifically, the partitions obtained through the clustering procedure on each descriptor space are further analyzed with respect to the involved local regions in order to find correspondence between clusters of different features spaces. This alignment allows to further improve the Bag of Visual Phrases Model by adding the coherence of different feature spaces also during codebooks generation phase. The approach is formalized in the following.

Let I an image, and M the number of local regions extracted by making use of a local detector [27] or through dense sampling [1, 21]. Each extracted local region $r_i, i = 1, \dots, M$, is described by H different local descriptors $\phi_{ih}, h = 1, \dots, H$. Each region r_i is then associated to a set of local descriptors [28] $\phi_i = \{\phi_{i1}, \phi_{i2}, \dots, \phi_{iH}\}$. A vocabulary V_h is built for each type of local descriptor, and the different local descriptors ϕ_{ih} of a region r_i are hence associated to visual words v_h belonging to the codebook V_h as in the classic BoVW paradigm [41]. This produces a H -tuple $p_i = \{v_h | h \in [1, 2, \dots, H]\}$, called “visual phrase”, which contains visual words of different feature spaces for each $\phi_i, i = 1, \dots, M$, corresponding to the M local regions detected into the considered image I . Each image is then described by the frequency distribution of visual phrases, called “Bags of Visual Phrases”. This model, called “coherent phrase model” [15], incorporates the coherence across multiple descriptors in order to describe different aspects of the appearance of a local region detected within an image.

Our approach augments the coherent phrase model by improving the vocabulary generation step. In [15], H codebooks (one per local descriptor type) are generated separately and independently by using a classical clustering approach on each descriptor space. Then the images are described with a normalized multidimensional histogram in which each bin is related to a visual phrase (e.g., a 2-D distribution by considering two different local descriptors). The underlying rationale is that, although different descriptors encode different properties of a local region, they represent the same local region, hence the clustering, and the visual words belonging to different feature spaces, are in “some way” related. Hence, the coherence among different local descriptors should be exploited also in the vocabulary generation step. The main schema of the proposed approach is summarized in Fig. 1. First, the H different local descriptor spaces are clustered separately and K visual words (cluster centroids) are obtained for each vocabulary V_h (one visual vocabulary per local descriptor) as in the classic BoVW paradigm [41]. The relative ordering of cluster labels in all of the clustering are hence rearranged according to the first one. A $K \times K$ similarity matrix between pairs of clusters belonging to the two partitions is obtained by counting the number of elements (local image regions) they share. The Hungarian algorithm [30] is then used to find the best assignment for the cluster correspondence problem which is encoded in the computed similarity matrix. Therefore, the alignment between clusters of different partitions is thought as a classical resources assignment problem to be solved by a combinatorial optimization algorithm. We choose to exploit Hungarian algorithm since it has been successfully used in Computer Vision to solve different problems which can be seen as a resources assignment problem (e.g., cluster correspondence [36], feature matching [6]). By using the Hungarian method the alignment of the different vocabularies can be done

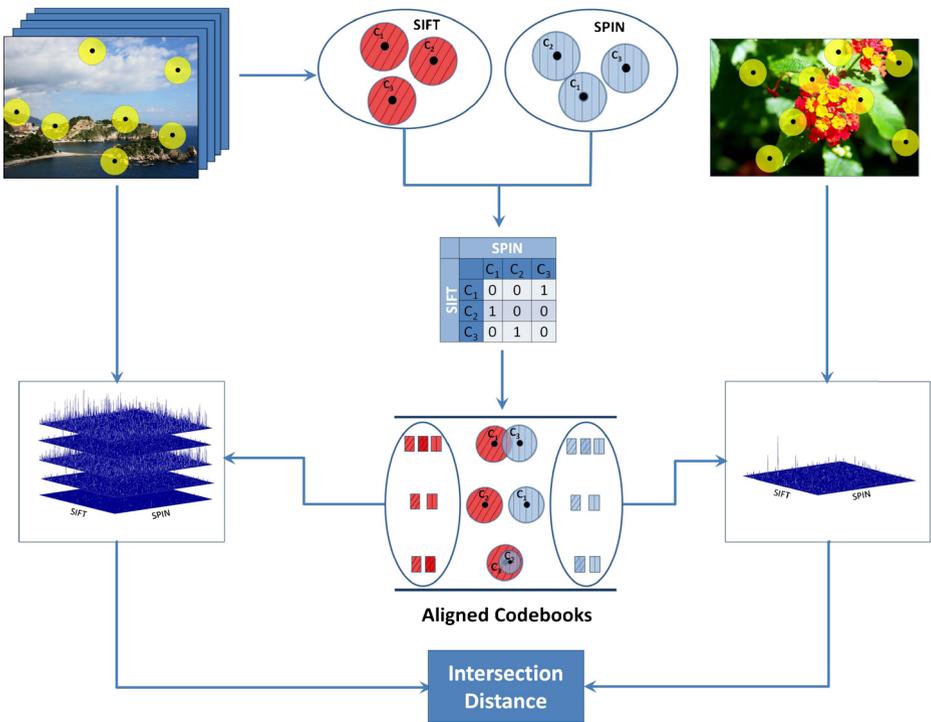


Fig. 1 Proposed Bags of Visual Phrases with codebooks alignment. First a set of keypoints are extracted from a training dataset of images by using a local detector (Hessian-Laplace in our experiments). Each local keypoints is then described by two different descriptors (SIFT [25], SPIN [17] in our experiments) and clustering is performed separately in these two feature spaces. A similarity matrix between pairs of clusters belonging to the two partitions is obtained counting the number of elements (local image regions) they share. The Hungarian algorithm is then used to find the best assignment for the cluster correspondence problem which is encoded in the similarity matrix. The obtained cluster correspondences are then used to create two novel vocabularies where visual words are generated considering the centroids relative to both common and uncommon elements between aligned clusters. The training set images are then represented by using 2D histograms of co-occurrence of visual words related to the generated codebooks. When a query is performed on the training dataset, the test image is represented by using the codebooks obtained in the training phase. Test image representation is compared with those of the training images by using the intersection distance. Finally, the training image corresponding to the lowest distance is selected as the output of the query

in $O(K^3)$ time. Despite we have used the Hungarian algorithm in our experiments, there are more efficient algorithms that can be used to solve the same problem [18].

The obtained cluster correspondences are used to create H novel vocabularies where visual words are generated considering the centroids relative to both common and uncommon elements between aligned clusters (Fig. 1). Hence three new visual words (cluster means) per descriptor space are generated from two aligned clusters considering the operations of intersection (shared local image regions belonging to the overlap among aligned clusters) and difference (local image regions belonging to the non-overlapped parts of the aligned clusters). Notice that, although Hungarian

algorithm aligns all the clusters, some of them can have no common elements (Fig. 1). If two clusters are fully separated, only two new cluster centers will be computed from individual ones. In this last case the two obtained visual words are equal to the original ones.

After building the vocabularies separately on each feature space, these are aligned (as described above) to find coherence between the different spaces based on shared keypoints. After that, each cluster of each vocabulary (which define a visual word in the considered feature space) is splitted in subclusters (defining more than one visual word if the overlap of the aligned clusters is not empty). In this way, the quantization of a descriptor space is refined by taking into account of the quantization obtained in the other feature space. So, the refinement of each vocabulary encodes also information induced from the other vocabulary. This allows to make stronger the discriminativeness of the original Bags of Visual Phrases approach [15] as empirically demonstrated by the experimental results reported in Section 6.

The algorithm described above generates a multidimensional representation of the image under analysis. In particular, starting from the original image, it extracts a set of local feature points, associates them to different descriptors and, by using a pre-computed set of vocabularies, creates the final multidimensional normalized histogram. Considering two descriptors with the associated codebooks consisting of K_1 and K_2 elements respectively, the final image representation is a matrix ($2D$ normalized histogram) of $K_1 \times K_2$ values.¹

4 Image representation compression

The compactness of the image representation impacts both in terms of memory storage and computational complexity of the near duplicate detection task [34]. The cost per single image query becomes a critical feature of the overall system as the number of the images stored in the dataset increases. It is then extremely useful to study some approximations of the original representation able to reduce the amount of data to be stored and used during retrieval, without considerably reduce the performance of the overall system. The analysis of the $2D$ histogram representations of the training images shows that the $K_1 \times K_2$ matrices are pretty sparse (see Section 6), hence only a limited set of elements (visual phrases) are actually used to describe the image content. Based on this analysis, we propose a simple and effective compression technique. The most representative and discriminative T visual phrases (i.e., T bins of the $K_1 \times K_2$ matrix, with $T \ll K_1 \times K_2$), together with their IDs (i.e., the number of row and column they belong into the matrix), can be selected to represent the image under analysis. This selection considers all the images of the training dataset on which image queries are performed.

To sum up, when a query is performed on the selected training dataset considering a generic test image I , the following steps are performed:

- i) generate the multidimensional histogram Ψ_I of the image I ;
- ii) for each image J of the training dataset, select its matrix coordinates C_J relative to its most representative and discriminative T visual phrases;

¹Note that at this stage other encoding methods can be used starting from the aligned vocabulary [7].

- iii) select the elements of the histogram Ψ_I at the coordinates C_J ;
- iv) for each image J compute the similarity between the compact representation of images I and J ;
- v) provide as output of the query the image \hat{J} belonging to the training dataset with the lowest distance from the image I .

It is worth noting that the effectiveness of the proposed approximation depends on the number of selected bins T . To obtain a satisfactory improvement in terms of memory storage and computational load this number should be considerably lower than $K_1 \times K_2$, where K_h is the dimension of the vocabulary V_h . On the other hand few visual phrases (i.e., bins) could be not able to properly discriminate the images belonging to the dataset. A smart selection strategy of the best T bins can be then useful in finding a good trade-off between compression and retrieval performance. Specifically, we employ the statistical measure TF-IDF (term frequency-inverse document frequency) [38] for the selection of the most representative and discriminative visual phrases (i.e., to select the best T bins within the representation matrix). In this way, the importance of the bin is not only related to its frequency in the image representation but also consider the frequency of the bin (i.e., the discriminativeness of the visual phrase composed by a pair of descriptors) with respect to the entire training dataset. In other words, for each image of the training dataset, we select the most representative and discriminative T visual phrases (e.g., bins) as indicated by the TF-IDF measure. During a comparison of an query image I with an image of the training dataset J only the T visual phrases of the image J which have been selected taking into account the TF-IDF measure are considered.

5 The experimental datasets

An image I is considered a near-duplicate of another image J if its content is “similar”, according to some defined similarity measures, to the image J . So, the definition of a near duplicate image changes accordingly with the allowed photometric and geometric variations. As in [9], we consider an image I a near-duplicate of another image J if it contains the same scene of J with possibly different photometric and/or geometric variations (e.g., viewpoints changes, illumination and color variations, partial scene, occlusion, different compression and camera acquisition, etc.). The problem addressed here is hence the one to enumerate all the near duplicates of a given query image in a dataset.

In order to test and compare different algorithms for near duplicate image retrieval, a representative dataset should be used. Despite different datasets have been employed in literature for testing purposes, most of them are synthetic² [19] or obtained taking into account keyframes of videos [15]. Although synthetic

²We consider a dataset as synthetic when the near duplicates are generated from a set of images (or frames of videos) by using transformations typically available on image manipulation software (e.g., ImageMagick <http://www.imagemagick.org>), such as colorizing, contrast changing, cropping, despeckling, downsampling, format changing, framing, rotating, scaling, saturation changing, intensity changing, shearing. To generate near duplicates the basic transformations are usually applied changing the different involved parameters and/or making combination of them.

datasets are compliant with the definition of near duplicate given above, they aren't representative of the real variation that can be observed in real near duplicate images (see Fig. 2). On the other hand, datasets built by collecting frames of videos contain near duplicates with no variability in terms of resolution and compression factor. The classic datasets used for image retrieval testing purposes (e.g., CBIR task), such as the one introduced in [16], are not compliant with the aim of near duplicate image retrieval, where the problem is to search for the same scene with possibly different photometric and/or geometric variations, given an image as query.

The above motivations induced us in building and using a new representative dataset for the problem under consideration. In this way we can properly test and compare the proposed augmented version of Bag of Visual Phrase model with respect to the original one [15]. Specifically, a dataset with images acquired by different cameras, in different conditions (e.g., viewpoint, scale, illumination, distance from the subjects, etc.), and high content variability (indoor, outdoor,

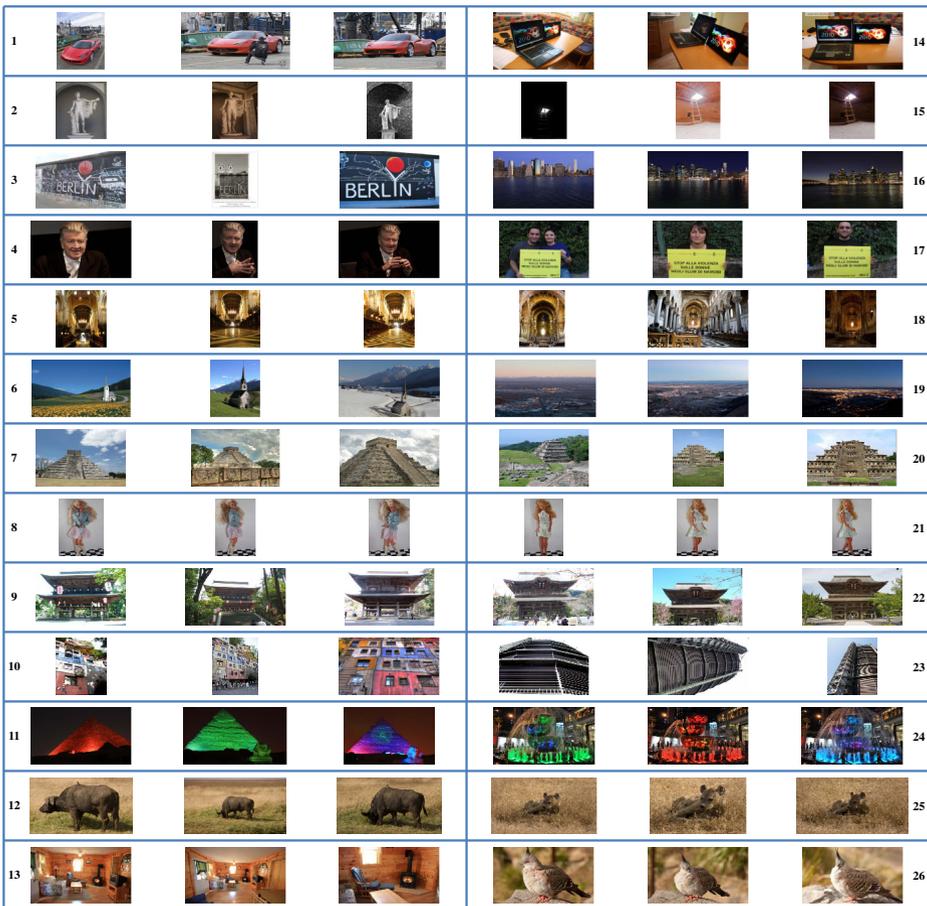


Fig. 2 Examples of 26 different scenes belonging to the considered dataset. For each scene three near duplicates are shown

object, natural scenes, etc.), has been collected from Flickr (<http://www.flickr.com>) and from private collections. To this aim, 525 different keywords (e.g., New York, Animal, Car, Church, Computer, Mountains, Landscape, etc.) have been chosen. Each keyword has been then used to retrieve images from Flickr. From the retrieved images a set of near duplicates have been hence manually sampled. Each specific set corresponding to a keyword contains from 3 to 34 near duplicates. The whole dataset contains 3,148 images. In Fig. 2 some of the images belonging to the built dataset are shown. Specifically, in the figure are reported three near duplicates of 26 different scenes. As evident by visual inspection, there is a high variability in terms of scenes (outdoor, indoor, close up objects, portraits, archeological sites, buildings, animals, open scenes, etc.) as well as a high variability in terms of geometric and photometric characteristics among near duplicates of the same scene (different point of view, luminance and color variation, zoom, rotation, background variation, etc.). Moreover, different scenes have regions with similar appearances, such as in the case of the scenes with animals (see images of the scenes with number 12 and 25 in Fig. 2) and the ones with Japanese buildings (see images of the scenes with number 9 and 22 in Fig. 2). Differently than classic content based image retrieval task in which, for instance, given an image of the scene numbered as 9 in Fig. 2 all the images of the the scene numbered as 22 are acceptable in terms of visual similarity, in the context of near duplicate image detection this become an unacceptable error. The database was hence built to properly test the challenging task under consideration.

Since near duplicate image detection techniques are usually tested on datasets used in the context of object recognition [15], we have performed tests also considering the UKBench dataset which contains a total of 10,200 images of 2,550 different objects with four near duplicate images (photometric and/or geometric variations) for each object [29].

6 Experimental results

In this section the effectiveness of the proposed approach is demonstrated through a number of experiments and comparisons. A first test, conducted on the dataset we built (see previous section), compares our method with respect to the coherent visual phrase model described in [15] and the technique proposed by Zhao et al. in [50, 51]. Note that both [15] and the classic BoVW approaches have been reimplemented at the best of our knowledge whereas the original code provided by the related authors has been used for [50, 51].

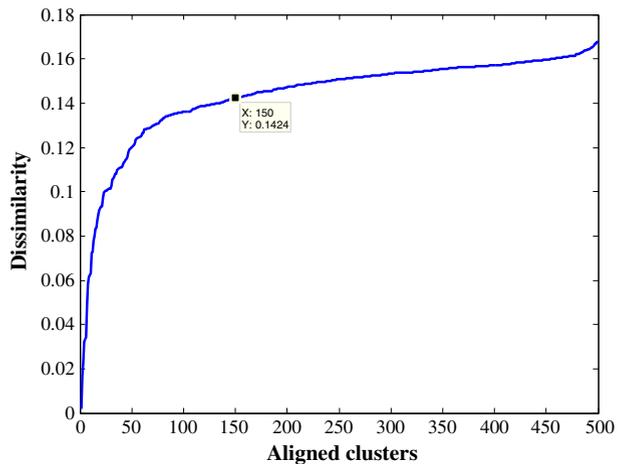
To properly evaluate the different methods, the experiments have been repeated three times. At each run the different approaches are executed on the same training and test sets. To this purpose, at each run we have built training and test sets by selecting images at random. Specifically, we have randomly selected one image per each set of near duplicates to build a training set with 525 different scenes, whereas two images per each set of near duplicates have been randomly selected to build the test set. All the parameters involved in the experiments have been learned from the corresponding training sets for each method. The results presented in the following are obtained by averaging the results of all three runs.

For every run, training images have been used for the generation of codebooks. First, local interest points have been detected (Hessian-Laplace [27]). Afterward, two

different descriptors have been extracted on each interest point: SIFT [25] and SPIN [17]. Since these descriptors are extracted considering different image properties (gradient orientation (SIFT) and intensity distribution at different distance from the center (SPIN)), they are somewhat complementary, hence can be fruitfully combined. K-means algorithm ($K = 500$ in our tests) has been then used to produce the two independent codebooks corresponding to the two involved descriptors. The two obtained partitions have been aligned with the Hungarian algorithm to generate the new codebooks (see Section 3). Finally, training images have been represented by visual phrases (with a 2D histogram) by considering the new aligned codebooks.

It is worth noting that the proposed procedure for codebook generation creates two novel vocabularies (one for each type of descriptor involved in the experiment) with a higher number of elements with respect to the original ones. Considering, as example, two codebooks of 500 elements, the alignment procedure will produce, in the worst case, two novel vocabularies of 1,500 elements for each type of descriptor. In order to reduce the dimension of the final image representation maintaining at the same time good performance, analysis and tests have been performed. In particular, useful hint can be derived from the analysis of the degree of similarity between the clusters associated in the alignment procedure performed through the Hungarian algorithm. As reported in Fig. 3, which have been obtained sorting the aligned clusters with respect to their dissimilarity, after a certain threshold, the aligned clusters cannot be considered “similar”. This means that after a given value the aligned clusters share only few keypoints (or nothing at all) and hence there is not too much coherence among these aligned clusters. The threshold imposed on cluster dissimilarity is chosen taking into account the gradient of the dissimilarity curve. At some point, the gradient of the curve starts to be very small and this fact can be used to set the threshold. Moreover, given a threshold, the number of elements of the aligned vocabulary is properly established. For example, in the case reported in Fig. 3, where a threshold which consider 150 aligned cluster is selected in correspondence of a small gradient, the final number of employed centroids is equal to $150 \times 3 + 350$ for each feature space. The cluster intersections produce 150×3 new visual words

Fig. 3 Sorted dissimilarity of aligned vocabularies. The first 150 aligned pairs of clusters can be considered “similar” in terms of shared keypoints, whereas the others are “dissimilar”



for each feature space, whereas the other not aligned 350 clusters produce 350 visual word for each feature space. So considering both, the gradient of the curve and the dimension of the final codebook, the threshold can be fixed. Taking into account the previous analysis, a more compact vocabulary can be hence generated performing the procedure for the generation of aligned codebooks (see Section 3) only for the aligned clusters having a high degree of similarity; for all other “dissimilar” clusters will be retained only the original centroids on the corresponding feature space. The analysis of the dissimilarity curves related to the different three training set considered in our tests, pointed out that the first 150 aligned pairs of clusters can be considered properly aligned (i.e., “similar” in terms of shared keypoints). In this way a final codebook of 800 visual words per descriptor (SIFT and SPIN) has been generated instead of one of 1,500. To be fair, the comparisons with the other approaches (Hu et al. [15], BoVW SIFT and BoVW SPIN) have been performed considering codebooks with 800 elements per descriptor independently generated through K-means clustering.

At each run, test images are used to perform queries on the related training dataset. Each test image is represented by a visual phrase histogram obtained considering the aligned codebooks (see Fig. 1). This representation is then used to retrieve images in the training dataset, by means of a similarity function between Bag of Phrases histograms. To cope with partial matching, we use the intersection distance τ defined as follows [14, 40]:

$$\tau(\Psi_I, \Psi_J) = \sum_{p=1}^P \min(\Psi_p(I), \Psi_p(J)) \quad (1)$$

where Ψ_I, Ψ_J are two visual phrase histograms and $\Psi_p(\cdot)$ is the p^{th} bin of the histogram. Both representation, with and without TF-IDF weighting scheme have been considered and compared.

Each query image has been associated to a list of training images. The retrieval performance has been evaluated with the probability of the successful retrieval $P(n)$ in a number of test queries [15, 47, 48, 52]:

$$P(n) = \frac{Q_n}{Q} \quad (2)$$

where Q_n is the number of successful queries according to top- n criterion, i.e., the correct NDI is among the first n retrieved images, and Q is the total number of queries.

The proposed approach has been compared with the original Bags of Phrases approach [15], with the approach proposed in [50, 51], as well as with respect to the classic BoVW approach considering both SIFT and SPIN descriptors. The obtained results are reported in Fig. 4. Both proposed strategy, with and without TF-IDF, outperforms the original Bags of Visual Phrases, the approach proposed in [50, 51], and the classic BoVW model. We also show the precision/recall values at top- $n = 1$ in Table 1. Note that the precision and recall for top- $n = 1$ are equivalent because there is only one correct match for each query. Some visual examples of the first retrieved image on a specific query are reported in Fig. 5. Specifically, for some query images reported in the first column of Fig. 5, the first retrieved images obtained with the proposed approach and the method described in [15] are shown respectively in the second and third columns of Fig. 5. The proposed method is able to detect the

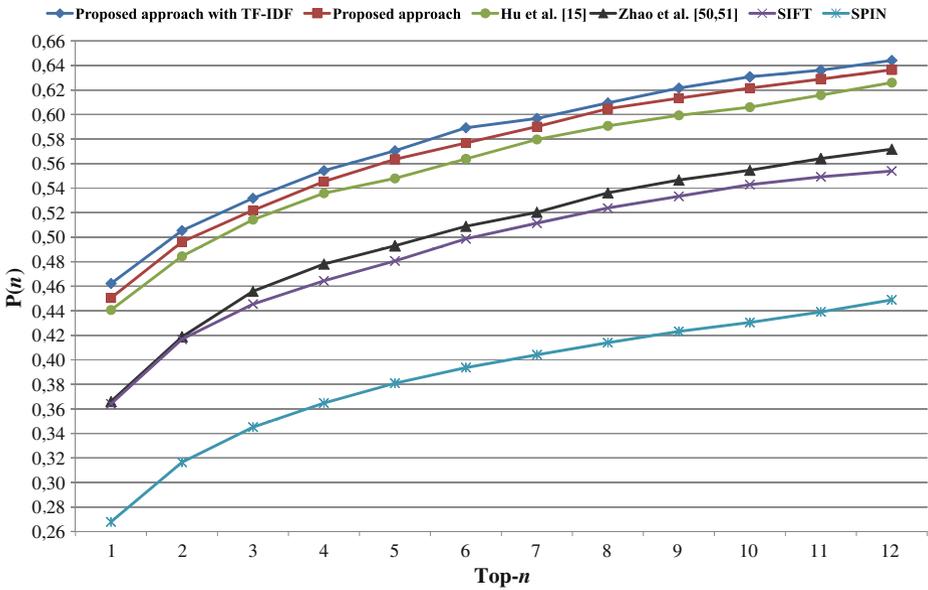


Fig. 4 Top-*n* NDI retrieval performances comparison on the proposed dataset

corresponding near duplicate within the training set, whereas the technique proposed in [15] retrieves images which aren't a near duplicate of the queries and hence fail the aim. For completeness, further visual examples in which both approaches fail, as well as some examples in which the method of Hu et al. [15] outperforms our approach are reported respectively in Figs. 6 and 7. As evident from Fig. 6, often both approaches fail in the same way (i.e., selecting the same wrong image).

As already stated in Section 4, some analysis and tests have been performed to compress the image representation in order to speed up the retrieval process and to reduce the amount of data to be stored. Although the image representation is based on a 2D histogram of 800×800 elements, only a limited number of bins are actually different from zero. Specifically, from the performed analysis we have observed that the training images are described, on average, by 1,700 non-zero elements (with a standard deviation of 1,039). This analysis motivated the compression strategy described in Section 4. Taking into account the number of non-zero elements it is possible to guess the number of bins to be used in the image representation. Several tests have been performed to validate the proposed compression strategy and to

Table 1 Precision/recall values on the proposed dataset

Method	Precision/recall
Proposed approach with TF-IDF	0.4622
Proposed approach	0.4505
Hu et al. [15]	0.4406
Zhao et al. [50, 51]	0.3660
SIFT	0.3641
SPIN	0.2679

Fig. 5 Some visual examples of the first retrieved image on a specific query. In these examples the proposed approach outperforms the method proposed by Hu et al. [15]

Query Image	First Retrieved Image	
	Proposed approach	Hu et al. [15]

Fig. 6 Some visual examples of the first retrieved image on a specific query. In these examples both, the proposed approaches and the one described in [15] fail

Query Image	First Retrieved Image	
	Proposed approach	Hu et al. [15]
		
		
		
		
		
		
		
		
		
		
		

Fig. 7 Some visual examples of the first retrieved image on a specific query. In these examples the method proposed by Hu et al. [15] outperforms the proposed approach

Query Image	First Retrieved Image	
	Proposed approach	Hu et al. [15]
		
		
		
		

find a good trade-off between compression and retrieval performance. As shown by Fig. 8, the retrieval performance increases at increasing of the elements used for image representation. Moreover, the results obtained considering 1,600 elements are comparable with the ones of the proposed approach without compression in which the overall 800×800 elements are involved. The considered number of bins (i.e., 1,600) is very close to the number of the non-zero element computed during the aforementioned analysis (i.e., 1,700). In this way we are able to obtain a compact image representation without sacrificing the retrieval performance.

It is worth noting that the built dataset is really challenging; in some cases even an human observer could have some difficulties in finding the correct near duplicate image (e.g., compare the scenes marked with number 9 and 22 in Fig. 2). Moreover, in the described near duplicate image retrieval system each of the 525 classes are described by only one image, a design choice that could limit the overall performance of the proposed method, but is realistic, for instance, in the context of forensic science where investigators have only one image example of a criminal scene.

To further confirm the effectiveness of the proposed approach, additional experiments have been performed on the UKBench dataset [29]. This dataset, usually used for object recognition tasks, contains 10,200 images of 2,550 different objects. Specifically, there are four near duplicate images with photometric and/or geometric variations for each object. In our test, the training dataset has been built randomly selecting one image per class. The remaining images have been then used for testing purposes. The test has been repeated three times and the final results are obtained

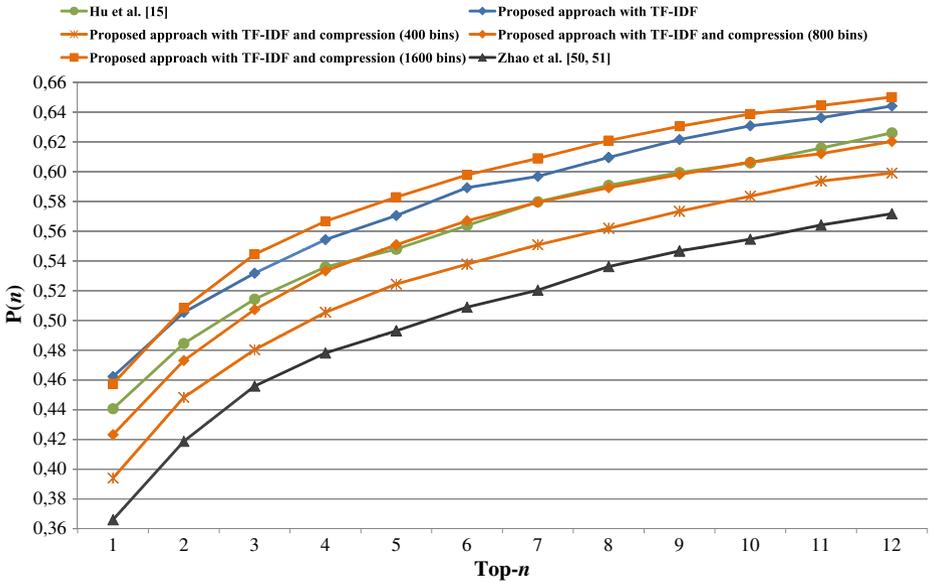


Fig. 8 Top-*n* NDI retrieval performances of proposed approach with compression on the built dataset. Results are reported at varying of the number of elements involved into the image representation

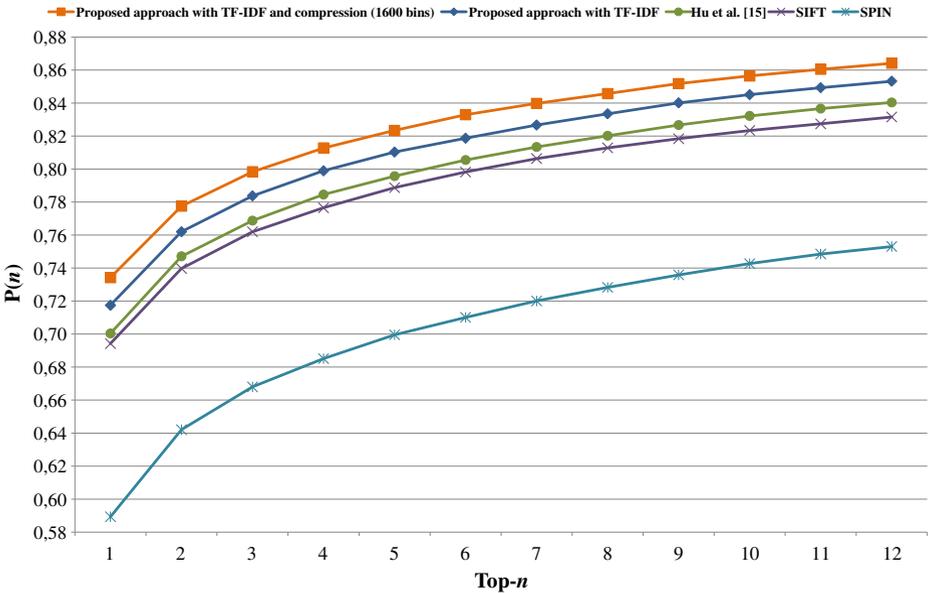


Fig. 9 Top-*n* NDI retrieval performances of the proposed approach with compression on the UKBench dataset

Table 2 Precision/recall values on the UKBench dataset

Method	Precision/recall
Proposed approach with TF-IDF (1,600 bins)	0.7342
Hu et al. [15]	0.7003

by averaging the results obtained on each test. As can be easily seen from Fig. 9 and Table 2, also considering this dataset, the proposed approach obtains satisfactory results. Also in this case the proposed approach outperforms the original Bag of Phrases approach [15] obtaining a good margin in terms of performances. The approach proposed in [50, 51] results worst than the original Bag of Phrases method and it is not reported in Fig. 9.

As pointed out in Section 3, in terms of computational complexity the proposed approach has an additional cost due to the alignment of clusters during the codebook generation. On the other hand, this allows a richer description of the images which is reflected in the increasing of the performances with respect to the original Bag of Visual Phrases paradigm [15]. Moreover, the alignment procedure is performed just once during the vocabulary generation and it does not affect the retrieval step in terms of extra costs. Considering the computational complexity during the retrieval task, the description compression proposed in Section 4 helps to reduce both, space and time with respect to the original paradigm [15] by maintaining the performances of the proposed codebooks alignment framework. Moreover, regarding the retrieval task, the proposed technique is comparable with the one proposed in [50, 51] in terms of computational complexity. Indeed, considering vocabularies with size K for the different descriptors, the mapping of each image with M local regions to the related visual vocabularies has computational complexity $O(MK)$. The time needed to build the visual phrases distribution is $O(M)$, whereas the compression of the image representation described in Section 4 takes time $O(T)$. Finally, the similarity between the query and an image belonging to the training dataset has computational complexity $O(T)$. Hence the overall computational complexity to represent and check a query image with respect to an image in the training dataset is $O(MK)+O(M)+O(T)$. It is worth noting that by employing the compression strategy the complexity in terms of computational power as well as the one related to the memory usage, have been considerably reduced. Specifically, considering a simple 2D matrix representation (without optimized data structure such as sparse matrix) the complexity of the comparisons is $O(K^2)$ instead of $O(T)$ where $K^2 \gg T$. Hence the final complexity of the algorithm without compression is $O(MK)+O(M)+O(K^2)$.

7 Conclusions and future works

In this paper we proposed an improvement of the coherent phrase model (Bags of Phrases) originally proposed in [15]. The main contribution of the presented approach is in augmenting the original paradigm by exploiting coherence between different feature spaces also during the codebook generation step. This is achieved through alignment of the feature space partitions obtained from independent clustering. Moreover, a method based on TF-IDF statistical measure to compress the proposed image representation for storage purposes is suggested. Experiments show the effectiveness of the described method on both, a novel and challenging

near duplicate image database and a classic benchmark one. Future works will be devoted to extend the proposed alignment methodology to consider multiple ($h > 2$) types of descriptors.

Acknowledgements Part of this work has been performed in the project PANORAMA, co-funded by grants from Belgium, Italy, France, the Netherlands, the United Kingdom, and the ENIAC Joint Undertaking. The authors would like to thank Giuseppe Claudio Guarnera, Tony Meccio and Rosetta Rizzo who have given some help at the beginning of this work.

References

1. Battiato S, Farinella GM, Gallo G, Ravi D (2010) Exploiting textons distributions on spatial hierarchy for scene classification. *EURASIP J Image Video Process* Article ID 919367:1–13. doi:10.1155/2010/919367
2. Battiato S, Farinella GM, Messina E, Puglisi G (2012) Robust image alignment for tampering detection. *IEEE Trans Inf Forensics Secur* 7(4):1105–1117
3. Battiato S, Farinella GM, Guarnera GC, Meccio T, Puglisi G, Ravi D, Rizzo R (2010) Bags of phrases with codebooks alignment for near duplicate image detection. In: *Proceedings of the international acm workshop on multimedia in forensics, security and intelligence (MiFor 2010)*, in conjunction with international acm multimedia conference, pp 65–70
4. Bay H, Ess A, Tuytelaars T, Van Gool L (2008) Speeded-up robust features (SURF). *Int J Comput Vis Image Understand* 110(3):346–359
5. Belongie S, Malik J, Puzicha J (2002) Shape matching and object recognition using shape contexts. *IEEE Trans Pattern Anal Mach Intell* 2(4):509–522
6. Belongie S, Malik J, Puzicha J (2002) Shape matching and object recognition using shape contexts. *IEEE Trans Pattern Anal Mach Intell* 24(24):509–521
7. Chatfield K, Lempitsky V, Vedaldi A, Zisserman A (2011) The devil is in the details: an evaluation of recent feature encoding methods. In: *Proceedings of the British machine vision conference*
8. Cheng X, Hu Y, Chia L-T (2011) Exploiting local dependencies with spatial-scale space (s-cube) for near-duplicate retrieval. *Comput Vis Image Understand* 115(6):750–758
9. Chum O, Philbin J, Zisserman A (2008) Near duplicate image detection: min-hash and tf-idf weighting. In: *Proceeding of BMVC*
10. Chum O, Perdoch M, Matas J (2009) Geometric min-hashing: finding a (thick) needle in a haystack. In: *IEEE computer society conference on computer vision and pattern recognition*, pp 17–24
11. De Oliveira R, Cherubini M, Oliver N (2010) Looking at near-duplicate videos from a human-centric perspective. *ACM Trans Multimedia Comput Commun Appl* 6(3):15:1–15:22
12. Eastlake D, Jones P (2001) RFC 3174. <http://tools.ietf.org/html/rfc3174>
13. Freeman W, Adelson E (1991) The design and use of steerable filters. *IEEE Trans Pattern Anal Mach Intell* 13(9):891–906
14. Grauman K, Darrell T (2005) The pyramid match kernel: discriminative classification with sets of image features. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp 1458–1465
15. Hu Y, Cheng X, Chia L-T, Xie X, Rajan D, Tan A-H (2009) Coherent phrase model for efficient image near-duplicate retrieval. *IEEE Trans Multimedia* 11(8):1434–1445
16. Huiskes MJ, Lew MS (2008) The MIR Flickr retrieval evaluation. In: *MIR '08: proceedings of the 2008 ACM International conference on multimedia information retrieval*. ACM, New York, NY
17. Johnson AE, Hebert M (1999) Using spin images for efficient object recognition in cluttered 3d scenes. *IEEE Trans Pattern Anal Mach Intell* 21(5):433–449
18. Jonker R, Volgenant A (1987) A shortest augmenting path algorithm for dense and sparse linear assignment problems. *Computing* 38(4):325–340
19. Ke Y, Sukthankar R, Huston L (2004) Efficient near-duplicate detection and sub-image retrieval. In: *Proceeding of ACM multimedia*, pp 869–876
20. Koenderink J, van Doorn A (1987) Representation of local geometry in the visual system. *Biol Cybern* 55:367–375

21. Lazebnik S, Schmid C, Ponce J (2006) Beyond bags of features: spatial pyramid matching for recognizing natural scene categories. In: Proceedings of the 2006 IEEE computer society conference on Computer Vision and Pattern Recognition, CVPR '06, pp 2169–2178
22. Lazebnik S, Raginsky M (2009) Supervised learning of quantizer codebooks by information loss minimization. *IEEE Trans Pattern Anal Mach Intell* 31(7):1294–1309
23. Lejsek H, ÁðormóÁrðsdóttir H, Ásmundsson F, DaÁrason K, Jóhannsson ÁÁð, Jónsson BÁð, Amsaleg L (2010) Identifier forensic: large-scale video identification in practice. In: Proceeding of ACM workshop on multimedia in forensics, security and intelligence, pp 1–6
24. Leung T, Malik JJ (1999) Recognizing surfaces using three-dimensional textons. In: Proceedings of the IEEE international conference on computer vision, pp 1010–1017
25. Lowe D (2004) Distinctive image features from scale-invariant keypoints. *Int J Comput Vis* 60(2):91–110
26. Matas J, Chum O, Urban M, Pajdla T (2002) Robust wide-baseline stereo from maximally stable extremal regions. In: Proceedings of the British machine vision conference, pp 384–393
27. Mikolajczyk K, Schmid C (2004) Scale & affine invariant interest point detectors. *Int J Comput Vis (IJCV)* 60(1):63–86
28. Mikolajczyk K, Schmid C (2005) A performance evaluation of local descriptors. *IEEE Trans Pattern Anal Mach Intell (PAMI)* 27(10):1615–1630
29. Nistèr D, Stewènius H (2006) Scalable recognition with a vocabulary tree. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR), pp 2161–2168
30. Papadimitriou CH, Steiglitz K (1982) Combinatorial optimization: algorithms and complexity. Prentice-Hall, Inc
31. Philbin J, Chum O, Isard M, Sivic J, Zisserman A (2007) Object retrieval with large vocabularies and fast spatial matching. In: Proceedings of the International conference on computer vision and pattern recognition
32. Rivest RL (1992) RFC 1321. <http://tools.ietf.org/html/rfc1321>
33. Rosten E, Drummond T (2006) Machine learning for high-speed corner detection. In: Proceedings of the European conference on computer vision, pp 430–443
34. Rongrong J, Hongxun Y, Wei L, Xiaoshuai S, Tian TQ (2012) Task-dependent visual-codebook compression. *IEEE Trans Image Process* 21(4):2282–2293
35. Rongrong J, Duan L-Y, Chen J, Xie L, Yao H, Gao W (2013) Learning to distribute vocabulary indexing for scalable visual search. *IEEE Trans Multimedia* 15(1):153–166
36. Saffari A, Bischof H (2007) Clustering in a boosting framework. In: Computer vision winter workshop, pp 75–82
37. Salton G, McGill M (1983) Introduction to modern information retrieval. McGraw-Hill
38. Salton G, Buckley C (1988) Term-weighting approaches in automatic text retrieval. *Inf Process Manage* 24(5):513–523
39. Sivic J, Russell BC, Efros AA, Zisserman A, Freeman WT (2005) Discovering object categories in image collections. In: Proceedings of the international conference on computer vision
40. Swain MJ, Ballard DH (1991) Color indexing. *Int J Comput Vis* 7(1):11–32
41. Szeliski R (2010) Computer vision: algorithms and applications. Springer Available at <http://szeliski.org/Book>
42. van Gemert LC, Veenman CJ, Smeulders AWM, Geusebroek JM (2010) Visual word ambiguity. *IEEE Trans Pattern Anal Mach Intell* 32(7):1271–1283
43. Wang Y, Hou Z, Leman K (2011) Keypoint-based near-duplicate images detection using affine invariant feature and color matching In: Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP 2011), pp 1209–1212
44. Wu Z, Ke Q, Isard M, Sun J (2009) Bundling features for large scale partial-duplicate web image search. In: Proceedings of the international conference on computer vision and pattern recognition, pp 25–32
45. Xu D, Chang S-F (2007) Visual event recognition in news video using kernel methods with multi-level temporal alignment. In: Proceeding of IEEE international conference on computer vision and pattern recognition
46. Xu D, Cham TJ, Yan S, Duan L, Chang S-F (2010) Near duplicate identification with spatially aligned pyramid matching. *IEEE Trans Circuits Syst Video Technol (TCSVT)* 20(8):1068–1079
47. Zhang D-Q, Chang S-F (2004) Detecting image near-duplicate by stochastic attributed relational graph matching with learning. In: Proceedings of the ACM multimedia conference, pp 877–884
48. Zhao W-L, Ngo C-W, Tan H-K, Wu X (2007) Near-duplicate keyframe identification with interest point matching and pattern learning. *IEEE Trans Multimedia* 9(5):1037–1048

49. Zhao W-L, Ngo C-W (2009) Scale-rotation invariant pattern entropy for keypoint-based near-duplicate detection. *IEEE Trans Image Process* 18(2):412–423
50. Zhao WL, Wu X, Ngo CW (2010) On the annotation of web videos by efficient near-duplicate search. *IEEE Trans Multimedia* 12(5):448–461
51. Zhao W-L, Wu X, Ngo C-W (2011) SOTU: a toolkit for efficient near-duplicate image/video & retrieval/detection. Manual for SOTU Version 1.06. <http://www.cs.cityu.edu.hk/~wzhao2/sotu.htm>
52. Zhu J, Hoi SC, Lyu MR, Yan S (2008) Near-duplicate keyframe retrieval by nonrigid image matching. In: *Proceedings of the ACM multimedia conference*, pp 41–50



Sebastiano Battiato (M'04–SM'06) received his degree in computer science (*summa cum laude*) in 1995 from University of Catania and his Ph.D. in computer science and applied mathematics from University of Naples in 1999. From 1999 to 2003 he was the leader of the “Imaging” team at STMicroelectronics in Catania. He joined the Department of Mathematics and Computer Science at the University of Catania as assistant professor in 2004 and became associate professor in the same department in 2011. His research interests include image enhancement and processing, image coding, camera imaging technology and multimedia forensics. He has edited four books and co-authored more than 150 papers in international journals, conference proceedings and book chapters. He is a co-inventor of about 15 international patents, reviewer for several international journals, and he has been regularly a member of numerous international conference committees. Prof. Battiato has participated in many international and national research projects. Chair of several international events (IWCV2012, ECCV2012, VISAPP 2012–2013–2014, ICIAP 2011, ACM MiFor 2010–2011, SPIE EI Digital Photography 2011–2012–2013, etc.). He is an associate editor of the *IEEE Transactions on Circuits and System for Video Technology* and of the *SPIE Journal of Electronic Imaging*. Guest editor of the following special issues: “Emerging Methods for Color Image and Video Quality Enhancement” published on *EURASIP Journal on Image and Video Processing* (2010) and “Multimedia in Forensics, Security and Intelligence” published on *IEEE Multimedia Magazine* (2012). He is the recipient of the 2011 Best Associate Editor Award of the *IEEE Transactions on Circuits and Systems for Video Technology*. He is director (and co-founder) of the International Computer Vision Summer School (ICVSS), Sicily, Italy. He is a senior member of the IEEE.



Giovanni Maria Farinella (M'11) received the M.S. degree in computer science (*egregia cum laude*) from the University of Catania, Italy, in 2004, and the Ph.D. degree in Computer Science in 2008. He joined the Image Processing Laboratory (IPLAB) at the Department of Mathematics and Computer Science, University of Catania as a Contract Researcher in 2008. He is a Contract professor of Computer Science at the University of Catania (since 2008) and leads the course on Computer Vision at the Academy of Arts of Catania (since 2004). His research interests lie in the fields of Computer Vision, Pattern Recognition and Machine Learning. He has edited three books and coauthored more than 60 papers in international journals, conference proceedings and book chapters. He is also co-inventor of three international patents. He serves as a reviewer and on the programme committee for major international journals and international conferences. Giovanni Maria Farinella founded (in 2006) and currently directs the International Computer Vision Summer School (ICVSS—www.dmi.unict.it/icvss).



Giovanni Puglisi was born in Acireale, Italy, in 1980. He received his degree in Computer Science Engineering (*summa cum laude*) from Catania University in 2005 and his Ph.D. in Computer Science in 2009. He is currently contract researcher at the Department of Mathematics and Computer Science. His research interests include video stabilization and raster-to-vector conversion techniques. He is the author of several papers on these activities.



Daniele Ravi was born in Sant'Agata di Militello (ME), Italy, in 1983. He received the Master Degree in Computer Science (summa cum laude) in 2007 from University of Catania. From 2008 to 2010 he worked at STMicroelectronics (Advanced System Technology Imaging Group) as Consultant. He is currently Ph.D student at the Department of Mathematics and Computer Science, University of Catania, Italy. His interests lie in the fields of computer vision, image analysis, pattern recognition and machine learning for embedded devices.