

IMAGE REPLICA DETECTION USING R-TREES AND LINEAR DISCRIMINANT ANALYSIS

Spyros Nikolopoulos, Stefanos Zafeiriou, Panagiotis Sidiropoulos, Nikos Nikolaidis and Ioannis Pitas

Dept. of Informatics, Aristotle University of Thessaloniki, Box 451, 54124 Thessaloniki, Greece
e-mail: {nikolaid,pitas}@aiaa.csd.auth.gr

ABSTRACT

In this paper a novel system for image replica detection is presented. The system uses color-based descriptors in order to extract robust features for image representation. These features are used for indexing the images in a database using an R-Tree. When a query about whether a test image is a replica of an image in the database is submitted, the R-Tree is traversed and a set of candidate images is retrieved. Then, in order to obtain a single result and at the same time reduce the number of decision errors the system is enhanced with Linear Discriminant Analysis (LDA). The conducted experiments show that the proposed approach is very promising.

1. INTRODUCTION

Recent technological advances in the area of multimedia content distribution have resulted in a major reorganization of this trade. Valuable digital artworks can be reproduced and distributed arbitrarily without any control by the copyright holders. Thus, issues related to intellectual property rights protection arise.

Numerous systems addressing the issue of copyright protection can be found in the literature, the vast majority of them being based on watermarking. Watermarking is the technique of imperceptibly embedding information within the content of the original image [1]. Although watermarking has attracted considerable interest from both industry and academia, it bears certain deficiencies. The requirement of embedding information inside a digital image before it is made publicly available, automatically excludes digital images that are already in the public domain and need to be copyright protected. In addition, watermarking is unable to counter content leakages, when an unwatermarked copy of the original artwork is stolen.

In order to overcome these inherent watermarking deficiencies, the scientific community recently started to investigate image copyright protection and Digital rights Management from another perspective. Specifically, the problem is envisaged as an image similarity one where the system decides if a query image resembles a reference image (i.e. it is a

replica of this image). These systems are referred to as image replica recognition/detection systems [2, 3, 4, 5]. Replica detection is the process of identifying all images that have been generated from the original version through intentional or unintentional manipulations. It is assumed that the modified images maintain sufficient visual quality in order to keep their commercial value and also maintain the semantic content of the original. Severely distorted copies are of no interest for a replica detection system. The major benefit of this approach derives from the fact that no additional information should be embedded within the image content, thus making the system applicable to images that are already in the public domain.

In this paper an image replica detection system that utilizes a database of original images that can be queried with a suspect image and decide whether this image is a replica of a stored original is proposed. Images are represented by a feature vector comprising of color-based descriptors [6]. Then, we implement a multidimensional indexing structure based on R-Trees. Although substantially reduced, the probability that the R-Tree returns more than a single image as candidates for being the originals of the query is existing and prevents the system to decide unambiguously. We introduce the use of image-class information for resolving cases unsuccessfully handled by the indexing structure. Specifically, LDA (preceded by PCA) is applied in order to reformulate the solution space and yield more discriminant image representations.

2. REPLICA DETECTION SYSTEM

2.1. System Overview

The process of engineering the proposed image replica detection system can be separated to two independent phases. The first phase deals with the database organization and construction. Each time a new original copyright protected image is added into the database, the image is subjected to a series of predefined attacks (image manipulations) selected according to the system's design specifications. Feature vectors are extracted from each attacked version resulting in a feature table which contains samples from the feature space neighborhood of the original image. The latter is utilized for the calculation of an extent vector that specifies the neighborhood extent for each dimension. Finally, the original image is indexed within

This work has been partly supported by EU and Greek national funds under Operational Programme in Education and Initial Vocational Training II through the Archimedes project "Retrieval and Intellectual Property rights protection of Multidimensional Digital Signals" (04-3-001/4).

the R-Tree structure, according to the extent vector (Fig. 1).

The second phase implements the actual replica detection functionality, once the database has been organized. An arbitrary image is submitted as a query to the indexing structure. R-Tree prunes the redundant branches according to the query image feature vector and results in a set of candidate images or an empty set. The next step attempts to enhance the system performance through LDA, preceded by a PCA dimensionality reduction step. Finally, the system returns the closest image based on some similarity metric. Alternatively, the query image may be found to reside outside the formulated neighborhoods. In this case the result is an empty set. Thus, the image rejection (i.e. the decision that the query image is not a replica of the images in the database) may take place either during the R-Tree traversal or after the application of LDA. The way the queries are handled is demonstrated in Fig 2.




Fig. 1. Database organization of the image replica recognition system.




Fig. 2. Searching the database for the detection of a replica.

2.2. Feature Extraction Method

The proposed system is based on work conducted in [6] for image feature extraction. In this work a comparative evaluation of various feature extraction methods has been performed. All methods are based on the color histogram and try to benefit from its inherent resilience to a number of common manipulations and especially geometrical transforms. Dimensionality reduction is performed by quantizing the color histogram distribution with respect to a specific color palette. The comparative evaluation indicates that the best tradeoff between compaction and information capacity is accomplished when the quantization scheme utilizes the Macbeth color palette for constructing a feature vector containing 24 scalar values

[6]. It must be noted that the proposed scheme can be combined with any other feature vector and its use is not limited to the features used in this paper.

2.3. Indexing Multidimensional Feature Vectors

Optimization requirements for both the system's computational cost and performance accuracy resulted in the integration of an R-Tree for achieving efficient image retrieval [7]. An R-Tree is a height-balanced tree with index records in its leaf nodes (containing pointers to data objects). Originally, R-Trees were created to index spatial objects using their bounding boxes (BBs). For a given image query, the R-Tree returns all the records whose BBs include the image query.

Our system works under the assumption that the features of image replicas are localized around those corresponding to the original images. Therefore, the used R-Tree is constructed by associating a bounding box to every original image. These bounding boxes are defined using an extent vector. In order to determine the extent vector for each image in the database we simulate all attacks that an image may undergo when made public and we wish the system to be able to withstand. Thus, before inserting an original copyright protected image into the database, a series of predefined attacks are performed. The produced images are used for determining the extent vector.

More specifically, a feature vector is extracted from every attacked version of an image and the distances in each feature dimension between each attacked image and the original image are calculated. The maximum distance for each dimension is selected as the extent in this dimension. Thus, the aforementioned procedure derives for each original image an extent vector consisting of 24 scalar values that determines the extent of the neighborhood in each feature dimension. The concept of utilizing a set of attacked versions of an image for determining the extent vector is based on the assumption that all potential attacks will be adequately modelled by this set and thus the selected neighborhood will be sufficient for enclosing all image replicas.

It is obvious that the extent vector selection is crucial for the system behavior. In our case we have chosen an implementation where a different extent is kept for every image and for every feature vector dimension. Moreover, in order to fine-tune the system performance a constant a that is multiplied with the values of all elements in the extent vector was introduced. Changing the value of a allows for extending or shrinking the bounding boxes and thus modifying the system performance.

2.4. Applying Linear Discriminant Analysis

The fact that the R-Tree can return more than one original image as candidates for being related to the query image does not allow the system to decide unambiguously. In order to obtain a single result and at the same time reduce the number

of decision errors we propose the use of LDA for discriminant feature selection. Prior to the application of the LDA, a dimensionality reduction step using PCA is performed. By eliminating the dimensions that correspond to the smaller eigenvalues, an implicit denoising of the data is achieved. The set of participating classes in the LDA coincides with the set of classes (images) returned by the R-tree. Each of these classes is comprised of the feature vector of the original image along with the feature vectors of its attacked versions. These are the observations used for calculating the class statistics. The LDA space is trained every time a query is submitted. The result of LDA is a linear transformation \mathbf{W}_o that transforms and/or reduces the dimensionality of the image feature vectors $\mathbf{x}_k \in \mathbb{R}^{24}$ as:

$$\hat{\mathbf{x}}_k = \mathbf{W}_o^T \mathbf{x}_k. \quad (1)$$

Let the set $\mathcal{C} = \{\mathcal{C}_1, \mathcal{C}_2, \dots, \mathcal{C}_K\}$ be the set of K classes that has been retrieved after the R-tree traversal. Let also assume that each class \mathcal{C}_i (original image) has M_i feature vectors \mathbf{x}_k coming from the attacked images. In order to use the information of how the sets of attacked images are separated to different classes in the feature space, Fisher's discriminant criterion [8] is used. Let \mathbf{S}_b be the between-class and \mathbf{S}_w be the within-class scatter matrices. The goal of the linear transformation \mathbf{W}_o is to maximize the between class scatter while minimizing the within class scatter, i.e.,

$$\mathbf{W}_o = \arg \max_{\mathbf{W}} \frac{\text{tr}[\mathbf{W}^T \mathbf{S}_b \mathbf{W}]}{\text{tr}[\mathbf{W}^T \mathbf{S}_w \mathbf{W}]} \quad (2)$$

where $\text{tr}[\mathbf{R}]$ denotes the trace of matrix \mathbf{R} . The optimal projection matrix \mathbf{W}_o is the one whose columns are the eigenvectors corresponding to the largest eigenvalues of $\mathbf{S}_w^{-1} \mathbf{S}_b$. The matrix \mathbf{W}_o projects the feature vectors \mathbf{x}_k to $\min(K - 1, 24)$ dimensions and a new feature vector $\hat{\mathbf{x}}_k$ is derived.

By projecting the samples in \mathcal{C} to the newly created solution space, better separation of classes is achieved and a similarity metric is then used to find the closest class (image). Still, an extent vector is necessary for accepting or rejecting the query image as a replica. The selection of this vector is done using a procedure analogous to the one described in Section 2.3.

3. EXPERIMENTAL STUDY

3.1. Performance Metrics

In order to evaluate the performance of the proposed system when being queried with images that are replicas of the images in the database we measure the percentage of the query images that are falsely identified as non replicas (false rejection rate) as well as the percentage of the query images that are identified as replicas but are assigned to a different image in the database (misclassification rate). Moreover, in order to evaluate the performance of the proposed system when being

Table 1. Original Image Set

No.	Monument Name	Number of Images
1.	Parthenon (Greece)	343
2.	White Tower (Greece)	78
3.	Liberty Statue (USA)	71
4.	Sargada Familia (Spain)	247
5.	Lighthouse (Arbitrary Photos)	56
6.	Coliseum (Italy)	128
7.	Eiffel Tower (England)	233
8.	Pisa Tower (Italy)	78
9.	Pyramids (Arbitrary Photos)	63
10.	Sphinx (Egypt)	193
11.	Duomo (Italy)	266
12.	Big Ben (England)	476
	Total	2232

queried with images that are not replicas of the images in the database the percentage of such query images that are falsely identified as replicas (false acceptance rate) is used.

3.2. Test Set Characteristics and Performance Results

A sample of 2,232 original images were used to compose the database of copyright protected images. The images were selected so as to form 12 different content categories, each corresponding to a world famous cultural monument (Table 1). This image content selection was dictated by the high dependency between the performance of a replica detection system and the level of visual similarity among the database members. Evaluation of the system performance over a database containing groups of similar images was done in an effort to assess its behavior under the least favorable situation and introduces a sense of fairness compared to other technological solutions (e.g., watermarking) whose performance is unaffected by the image content distribution.

A training set comprised of attacked versions of the originals is involved in two different stages of the replica recognition system functional chain. Initially, it is utilized for estimating the optimal extent value for each image feature dimension, i.e., evaluating the extent vector which is used in the R-Tree, and at a second level for providing the samples for the evaluation of the linear transform \mathbf{W}_o in the LDA. Essentially, the goal of the training set is to effectively model all possible distortions that an original image may undergo.

A total number of 77 attacks per original image were applied for constructing a training set consisting of 171.864 images. The aforementioned copies were derived from 4 different attack categories by modifying the attack severity. Table 2 summarizes the attacks, which are performed by a preprocessing procedure applied each time a new original image is being inserted into the image database. The constant a mentioned in Section 2.3 was selected to be equal to 1.2 since this value was experimentally found to be optimal in terms of minimizing the errors mentioned above.

Three different query image sets were constructed in order

Table 2. Attacks on Original Images (Training Image Set).

Attack Category & Severity Range	Step	Produced Images
Jpeg compression 10-90 (quality factor)	5	17
Rotation 1°-359° (degrees)	5°	36
Resizing 0.3 - 2.0 (scale factor)	0.1	16
Cropping 99% - 50% (Remaining Portion)	5%	8
Total		77

Table 3. Attacks on Original Images (Extended Image Set used in testing).

Attack Category & Severity Range	Step	Produced Images
Jpeg 10-90 (quality factor)	1	90
Rotation 1°-359° (degrees)	1°	360
Resizing 0.3 - 2.0 (scale factor)	0.05	35
Cropping 99% - 50% (Remaining Portion)	1%	50
Total		535

to encompass all potential cases that the system may be confronted with. The first query set includes only the 2232 original images, since, before testing the system behavior against queries with distorted images we have to reassure that all original images are identified correctly. For this set, both the false rejection and the misclassification rate were 0%.

An extended set containing more attacks than the training set was utilized during the second experiment. This query set still contains attacked versions derived by the same attack categories and within the same severity range, though using more dense attack steps (Table 3). 240 images that reside in the database were randomly selected for this experiment and the 535 modified versions of each of them were used as query images (total: 128400 images). This experiment aimed at testing the appropriateness of the selected training set to adequately model the total intensity range of the specified attacks. It is obvious that, in this case, images not included in the training set are incorporated in the query set, thus introducing a sense of fairness. For the queries on this set the false rejection and misclassification rates were 0.54% and 1.28% respectively.

Finally, images outside the image database formulated the content of the third query set, since, one of the most important aspects of an image replica detection system is to recognize when the query image is not contained in the database and to return an empty set. 450 images that are not included in the database obtained from the Bridgeman Art Library were used to this end. The false acceptance rate in this case was equal to 7.33% .

4. CONCLUSIONS- FUTURE WORK

A novel system for image replica detection has been presented. Color-based features are extracted for image representation. These features are used for multidimensional in-

dexing with the use of an R-Tree. LDA, preceded by PCA, is then applied in the images returned by the R-Tree in order to obtain a single result and at the same time reduce the number of decision errors through improved class separability. Experiments show that the proposed system can be successfully used for image replica detection. Currently the system is enhanced by introducing, after the R-tree step, a module that performs image similarity computation using SIFT features [9]. Preliminary results show that this approach can significantly reduce the false acceptance rate while retaining the other errors at the same levels.

5. REFERENCES

- [1] A. Tefas, N. Nikolaidis, and I. Pitas, "Watermarking techniques for image authentication and copyright protection," in *The Handbook of Image and Video Processing, 2nd edition*, edited by Al Bovik, Elsevier, 2005.
- [2] Y. Ke, R. Sukthankar, and L. Huston, "An efficient parts-based near-duplicate and sub-image retrieval system," in *Proceedings of the 12th annual ACM international conference on Multimedia*, New York, USA, 2004, pp. 869–876.
- [3] S. Roy and E.-C. Chang, "Watermarking with retrieval systems," *ACM Multimedia Systems*, vol. 9, no. 5, pp. 433–440, March 2004.
- [4] Roy S, E.-C. Chang, and K. Natarajan, "A unified framework for resolving ambiguity in copy detection," in *Proceedings of the 13th annual ACM international conference on Multimedia*, Singapore, 2005, pp. 648–655.
- [5] A. Qamra, Y. Meng, and E.Y. Chang, "Enhanced perceptual distance functions and indexing for image replica recognition," *IEEE Transaction on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 379–391, March 2005.
- [6] M. Gavrielides, E. Sikudova, and I. Pitas, "Color-based image fingerprinting," *IEEE Transactions on Multimedia*, accepted for publication, 2005.
- [7] V. Gaede and O. Gunther, "Multidimensional access methods," *ACM Computing Surveys*, vol. 30, no. 2, pp. 170–231, 1998.
- [8] D. L. Swets and J. Weng, "Using discriminant eigenfeatures for image retrieval," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 18, no. 8, pp. 831–836, 1996.
- [9] D.G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, vol. 60, no. 2, pp. 91–110, November 2004.