Available online at www.sciencedirect.com
ScienceDirect
Signal Processing: *Image Communication* 21 (2006) 688–703
SIGNAL PROCESSING: IMAGE COMMUNICATION
www.elsevier.com/locate/image

Adaptive image replica detection based on support vector classifiers

Yannick Maret*, Frédéric Dufaux, Touradj Ebrahimi

Ecole Polytechnique Fédérale de Lausanne, EPFL-STI-ITS-LTS1, Yannick MARET, ELD241, Station 11, Institut de Traitement des Signaux, CH-1015 Lausanne, Switzerland

Received 8 November 2005; received in revised form 27 June 2006; accepted 12 July 2006

Abstract

This paper presents a system for image replica detection. The idea behind the proposed approach is to adapt a system for detecting the replica of a specific reference image. The system is then able to classify test images as replicas of the reference image or as unrelated images. More precisely, the test procedure is as follows. A set of features is extracted from a test image, representing texture, colour and grey-level characteristics. These features are then feed into a preprocessing step, which is fine-tuned to the reference image. Finally, the resulting features are entered to a support vector classifier that determines if the test image is a replica of the reference image. Experimental results show the effectiveness of the proposed system. Target applications include search for copyright infringement (e.g. variations of copyrighted images) and known illicit content (e.g. paedophile images known to the police).

Keywords: Image replica detection; Features extraction; Support vector machine; Dimensionality reduction; Copyright infringement detection

1. Introduction

In this paper, we propose a system to detect image replicas. By replica, we refer not only to a bit exact copy of a given reference image, but also to modified versions of the image after minor manipulations, malicious or not, as long as these manipulations do not change the perceptual meaning of the image content. In particular, replicas include all variants of the reference image obtained after common image processing manipulations such as compression, filtering, and adjustments of contrast, saturation or colours.

The proposed image replica detection system can be applied to *detect copyright infringement* by identifying variations of a given copyrighted image. Another application is to *discover known illicit content* such as child pornography images known to the police.

The problem of image replica detection is a particular subset of the more general problem of content-based search and retrieval of multimedia content. In recent years, multimedia search and retrieval have been the subject of extensive research

*Corresponding author. Tel.: +41 21 693 4621; fax: +41 21 693 7600.
E-mail addresses: yannick.maret@epfl.ch (Y. Maret), frederic.dufaux@epfl.ch (F. Dufaux), touradj.ebrahimi@epfl.ch (T. Ebrahimi).

0923-5965/$ - see front matter
doi:10.1016/j.image.2006.07.001

works and standardisation activities (MPEG-7 [16,29] and more recently JPSearch [17]). However, the specific problem of image replica detection has so far been the focus of fewer research efforts.

Two approaches to detect image replicas are *watermarking* [12] and *robust fingerprinting* [21,33,36]. Watermarking techniques [12] consist in embedding a signature in the reference image before dissemination. Replicas of the reference image can subsequently be detected by verifying the presence of the watermark. This class of techniques typically achieves high efficiency for the correct classification of replicas and non-replicas. However, it requires to modify the reference image, namely to embed a signature, prior to its distribution. Unfortunately, this is not always possible. For instance, the method is not applicable to already disseminated copyrighted content or in the case of illicit content. Robust fingerprinting techniques [21,33,36] analyse the reference image in order to extract a signature associated with the image content. Replicas are then identified whenever their signatures are close to that of the reference. This class of techniques is often based on a single feature, for example characteristic points of the Radon transform [21], log-mapping of the Radon transform [33], or intra-scale variances of wavelet coefficients [36]. While it is usually robust, computationally efficient, and suitable for fast database indexing and retrieval, it often performs poorly for the accurate classification of replicas and non-replicas.

More recently, techniques for image replica detection have been described in [18,30]. Ke et al. [18] propose a method based on the extraction of features, referred to as key points (KPs), which are stable in a scale-space representation. An image is typically represented by thousands of KPs. Test images are then classified as replicas or non-replicas using local sensitive hashing to match their KPs to those of the reference image. More specifically, no distance is directly computed, but it is rather the number of matching KPs which quantifies the similarity between two images. While this approach achieves very good performance for replica detection, it requires a computationally complex features extraction step. Qamra et al. [30] propose a different method based on the computation of a perceptual distance function (DPF). More precisely, a DPF is generated for each pair of reference and unknown image, to measure the similarity between the two. The main idea of the approach is to activate different features for different image pairs. Hence, only the most similar features are taken into account to compute the distance. While this method achieves good performance, it is inferior to [18].

In this paper, we introduce a new approach for image replica detection based on our earlier works [25–27]. The idea behind our approach is to adapt a system for detecting the replica of a specific reference image. The system is then able to classify test images as replicas of the reference image or as unrelated images. More precisely, the test procedure is as follows. A set of 162 features is extracted from a test image, representing texture, colour and grey-level characteristics. These features are then feed into a preprocessing step, which is fine-tuned to the reference image. First, the extracted features are weighted by comparing the proportion of pixels contributing to each feature in the test image to the corresponding one in the reference image. Second, the dimensionality of the features space is reduced to keep only a subset of features relevant to replica detection of the specific reference image. In the final step, the resulting features are entered to a support vector classifier that determines if the test image is a replica of the reference image.

Simulation results show the effectiveness of the proposed system. For instance, for an average false negative rate of 8%, one achieves a fixed false positive rate of 1×10^{-4}. Indeed, the proposed technique significantly outperforms DPF [30], at low false positive rates, even if we use fewer features. Although the achieved performance is not as good when compared to KPs [18], a speed up in terms of computational complexity, in the range of one to two orders of magnitudes, is achieved.

The paper is structured as follows. We present an overview of the proposed replica detection system in Section 2, and a thorough description of the various algorithmic steps in Section 3. In order to evaluate the performance of the proposed system, an evaluation methodology is defined in Section 4, and experimental results are reported in Section 5. Section 6 discusses applications of the proposed algorithm. Finally, conclusions are drawn in Section 7.

2. Overview and preliminary remarks

The notations used throughout this paper are first detailed. Subscripts in greek letters index vector elements. Subscripts in roman letters index vectors (or scalars). Training patterns (or examples) are denoted as $\mathbf{x}_i$, with $i = 1, \dots, m$ where m is the total