Blind Verification of Digital Image Originality: A Statistical Approach

Babak Mahdian, Radim Nedbal, and Stanislav Saic

***Abstract*—Many methods for verifying the integrity of digital images employ various fingerprints associated with acquisition devices. Data on an acquisition device and fingerprints are extracted from an image and confronted with a reference data set that includes all possible fingerprints of the acquisition device. This allows us to draw a conclusion whether the digital image has been modified or not. Thus it is critical to have a sufficiently large, reliable, and true reference data set, otherwise critical miscalculations can arise. Reference data sets are extracted from image data sets that in turn are collected from unknown and nonguaranteed environments (mostly from the Internet). Since often software modifications leave no obvious traces in the image file (e.g., in metadata), it is not easy to recognize original images, from which fingerprints of acquisition devices can be extracted to form true reference data sets. This is the problem addressed in this paper. Given a database consisting of "unguaranteed" images, we introduce a statistical approach for assessing image originality by using the image file's header information (e.g., JPEG compression parameters). First a general framework is introduced. Then the framework is applied to several fingerprint types selected for image integrity verification.**

***Index Terms*—Blind verification, camera fingerprints, image forensics, image forgery detection, image originality, image trustworthiness, JPEG compression.**

I. Introduction

TRUSTWORTHINESS of digital images has an essential role in many areas, including: forensic investigation, criminal investigation, surveillance systems, intelligence services, medical imaging, and journalism. As a result, verifying the integrity of digital images and detecting the traces of tampering without using any protecting preextracted or preembedded information has become an important and hot research field of image processing [1].

A. Reference Sets for Verification of Digital Image Integrity

When verifying the integrity of digital images, one of the critical tasks is to determine if a given image is original or additional modifications have been done to the digital image or its metadata. There are several methods how to approach this problem. An effective way is to extract a certain set of features from the digital image file and match them to the corresponding camera model. (We shorten "camera model" to "camera" in the text.)

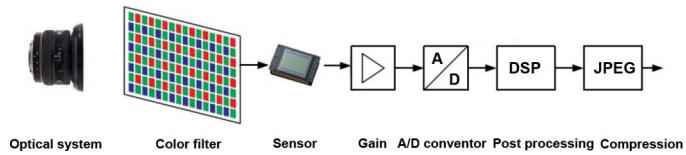

Fig. 1. Typical digital camera system.

For instance, having a digital image of resolution of 2000×1600 pixels and a claim that the digital image has been captured by a particular camera model (camera model name is found in digital image metadata), we can simply check if that particular camera device can produce digital images with such a resolution. Thus if we know that the camera always produces digital images with resolution 1200×1200, we obviously can draw a conclusion that the above claim is false. Consequently, we can conclude that the digital image has been modified and processed by software.

The above example illustrates how the image resolution can be employed as a feature (of a particular camera model) to determine the image integrity. The feature is an impression left by the camera. Generally, the traces of an impression from the friction ridges of any part of a human or other primate hand are called fingerprint. Using this analogy and for the sake of simplicity, we denote camera associated features left in digital images (e.g., metadata or JPEG compression parameters) as *camera fingerprints* . In the literature, camera fingerprints are of various types. In this paper, the term camera fingerprint refers only to the kind of features (fingerprints) that help to link the digital image to a specific camera model (e.g., Nikon Coolpix P80) with some degree of uncertainty.

A typical camera has several components (see Fig. 1) that leave fingerprints useful for integrity verification of digital images. Fingerprints left by the post processing and compression components are the most interesting for our purpose as they characterize a camera model.

Obviously, it is critical that the information about the camera fingerprints (image resolution, etc.) be true and guaranteed. Otherwise, miscalculations can arise which might have catastrophic impacts on people's lives. In our example, we considered a single digital image. In such a case, finding a corresponding camera model to evaluate fingerprints is a feasible task (though it still might be time-consuming). However, in

Manuscript received February 03, 2013; revised May 15, 2013 and July 24, 2013; accepted July 27, 2013. Date of publication August 01, 2013; date of current version August 15, 2013. This work was supported in part by the Ministry of the Interior of the Czech Republic under Project MV CR, VG20102013064, and in part by the Czech Science Foundation under Project GACR 13-28462S. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Alex ChiChung Kot.

B. Mahdian and S. Saic are with the Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, 182 08 Prague 8, Czech Republic (e-mail: mahdian@utia.cas.cz; ssaic@utia.cas.cz).

R. Nedbal is with the Fondazione Bruno Kessler, Trento 38122, Italy (e-mail: nedbal@fbk.eu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TIFS.2013.2276000

a real-life verification system, digital images from a plethora of camera models have to be verified. Thus it is desirable that guaranteed information about fingerprints (of as many as possible digital cameras) be available for a real time use.

A "perfect approach" would be based on collecting fingerprints directly from manufactures. Unfortunately, our attempt has revealed that manufacturers do not have such a type of information themselves: the data they provided were rather imprecise and noisy. Another safe approach would be based on (direct or reliable) access to tens of thousands of acquisition devices. Unfortunately, this is not feasible. Thus to collect a sufficiently large set of camera fingerprints (for as high as possible number of image acquisition devices), researchers resort to collecting large image sets from online photo sharing websites (e.g., Flickr is a popular photo sharing website providing an online API for easy access to its photos).

The problem is that photos in the reference data sets are collected from unknown environments (mostly from the Internet), where software modifications leave misleading modifications in image metadata. Consequently, it is not easy to determine which images are original and usable for extracting fingerprints of their acquisition devices. This consequence is addressed in this paper. Specifically, we introduce a statistical approach for handling *information noise* in databases consisting of "unguaranteed" images. This is a critical task for major forensics methods that require large-scale collections of reliable data. First a general framework is developed. Then the framework is demonstrated on several specific fingerprint types.

The rest of the paper is organized as follows. The next subsection gives a motivating example. Section II introduces the related work. After that, some basic concepts are presented to build up the necessary mathematical background. Section IV presents the main technical contribution—the statistical method. The following section shows experimental results demonstrating the efficiency of the method. In the last section, we discuss some subtleties revealed by the experiment, recall main properties of our approach and a highlight the main contributions.

B. Motivating Example

As aforementioned, a collection of reference camera fingerprints is usually extracted from the Internet. Therefore, it is a must to understand how image data are transferred to Internet storage places. Unfortunately, this important point is missed in the related published work.

To analyze image integrity, consider quantization tables (QTs), which encode digital images to JPEG format (see Subsect. V for more information on quantization steps and the JPEG procedure), as a fingerprint type. Indeed, QTs have been used as fingerprints in various works [2]–[4] since different image acquisition devices and software editors typically use different QTs.

Let us download one million digital images from a typical photo sharing site and extract a reference fingerprint data set. To discard nonoriginal (i.e., manipulated) images and create a reliable reference data set, photos containing obvious traces of modifications are eliminated. To further eliminate nonoriginal images, only those that form sufficiently big clusters of images with the same paired make, model, resolution, and QTs are retained and employed to extract camera fingerprints. Today, this is the most commonly used approach for denoising reference fingerprint data sets used by researchers [2], [3].

TABLE I
QTs (Luminance in Zig-Zag Order and Number of its Appearances)

QTs	cardinality of clusters
$5, 7, 7, 7, 8, 8, 810, 9, 10, 11, 12, 12, 11, \cdots$	9
$9, 9, 10, 12, 14, 12, 14, 15, 15, 14, 17, 15, 16, \cdots$	8
$2, 1, 1, 2, 1, 1, 1, 2, 2, 3, 6, 3, 3, \cdots$	4
$5, 7, 8, 8, 8, 9, 11, 14, 13, 12, 15, 17, 17, \cdots$	2

Fig. 2. Typical ways of uploading photos to photo-sharing sites (reprinted from Flickr.com).

For the sake of illustration, assume that Table I shows fingerprint values (QTs) (of a single camera) extracted from all the downloaded images. Employing the above simple denoising method, where we set the threshold for discarding clusters to five entries, we simply end up with a conclusion that QTs shown in the first two rows are true fingerprint values of the camera.

The approach seems to be rational. But the problem is that in reality only the third row belongs to the camera and all the other ones are QTs generated by software editors. The reason behind this is visualized in Fig. 2: the growing popularity of social web sites like Facebook or smart phones (iPhone, etc.) brings a number of new channels for photo upload to web sites like Flickr etc. An example of a traditional channel is a web browser. Today, there are also desktop versions of photo uploaders. Another transfer bridge is by directly using e-mail. Last but not least, there are a high number of apps allowing image upload using mobile devices.

Many of these channels modify photos during the transfer process automatically: they recompress or resize photos in order to achieve a more effective transfer speed, or they change image metadata for adding some marketing information, etc. As a result, there are large clusters of nonoriginal photos. Size of individual clusters are dependant on popularity of specific transfer tools and cameras. Taking into account the capability of these modifications to change the camera generated fingerprints without leaving any obvious traces of modification in image data (e.g., metadata), it is apparent that the illustrated information denoising method totally fails. So far, known methods for integrity verification undervalue or ignore this fact, and thus they might end up with completely incorrect results, causing critical consequences.

Hence the question of what fingerprints belong to cameras and what fingerprints are produced by software is the main topic of this paper. The proposed solution is general and applicable to any fingerprint type.

II. Related Work

In general, there are two approaches for verifying the integrity of digital images: *active* and *passive-blind* approaches.

A. Active Approaches

The area of active methods can be divided into:

- the *data hiding* approach [31] (digital watermarks [1], [25], [32] are most popular), where some secondary data are embedded into the image,
- the *digital signatures* approach [30], [34].

B. Blind Methods

In this work, we focus on blind methods [22]. In contrast to active methods, they need no prior information about the image being analyzed. For example, there are blind methods for detecting

- image splicing [24], [18],
- traces of inconsistencies in color filter array interpolation [28], [14],
- traces of geometric transformations [20], [27],
- cloning [19], [13], [11],
- computer graphics generated photos [6], [23],
- JPEG compression inconsistencies [10], [29], [21].

All these methods are most often based on the fact that forgeries can bring specific detectable statistical changes into the image.

There is a group of efficient blind methods based on the fact that each imaging device introduces specific fingerprints into the photo during the process of photo creation. Considering a typical digital camera system as shown in Fig. 1, we can notice several points in the system that are characteristic for each camera model, e.g., the demosaicking method or compression properties employed by the device. (Also, we can notice points exhibiting unique fingerprints dependent on individual devices like sensor-based noise propagated to photos). Since these fingerprints can be corrupted when a photo is digitally edited, they form efficient tools for forensic analysis of digital images.

There are a number of proposed fingerprints to verify the image integrity.

1) Image Thumbnails: Eric Kee and Hany Farid [2] have employed embedded *image thumbnails* to create camera fingerprints. These fingerprints are based on the fact that the creation of a thumbnail is modeled with a series of filtering operations, contrast adjustment, and compression, which significantly differ between camera manufacturers and photo-editing software.

2) Imaging Sensor Properties: Jessica Fridrich *et al.* [9], [5], [12] analyzed how photo-response nonuniformity (noise-like patterns caused by inhomogeneity of the silicon wafer from which the sensor is made) of imaging sensors can be used for a variety of image forensic tasks including forgery localization.

3) Demosaicking: Sevinc Bayram *et al.* [2], Mehdi Kharrazi *et al.* [16], Sevinc Bayram *et al.* [3], Ashwin Swaminathan *et al.* [31] used the traces of demosaicing to analyze photos.

4) Sensor Dust: Sensor dust characteristics (e.g., Ahmet Emir Dirik *et al.* [7]) showed that the location and shape of dust specks in front of the imaging sensor and their persistence make dust spots a useful fingerprint for digital single lens reflex cameras.

5) Quantization Tables: Hany Farid [8], [15] proposed to use the quantization tables to distinguish between original and modified photos, etc.

Also, there are methods dealing with identification of source cell-phones (e.g., Oya Celiktutan *et al.* [33] used binary similarity measures, image quality measures and higher order wavelet statistics to achieve this goal).

III. Basic Notations and Preliminaries

In this section, we introduce elementary concepts and outline JPEG compression needed for further exposition.

A. Digital Images

A *digital image* is a file consisting of

1) pixel data
2) and metadata.

Metadata can be internal or external. For example:

- ID of the user that has taken the photograph,
- properties of the camera with which the photograph has been taken,
- the size of the photograph, etc.

B. Cameras

1) Camera Fingerprints: More formally, we say that a digital image has attributes, and the image metadata are their respective values, which characterize the digital image. Essentially, some attribute values are dependent on the camera with which the digital image has been taken. As already mentioned in Section I-A, we refer to such camera associated features left in the digital image as *camera fingerprints* .

The properties that characterize an acquisition device (camera) explicitly include (cf. Section II-B), e.g.,

- its maker and model,
- the output file format,
- imaging sensor properties,
- the digital zoom interpolation method,
- the color filter array interpolation method used to encode an image etc.

Some of these properties identify a camera uniquely, and some of them can be considered as camera model fingerprints.

2) A Camera ID Vector and a Fingerprint Vector: Here, we assume a fixed tuple of properties sufficient for unique identification of any camera. We denote such a tuple by $\overrightarrow{cm}$, a *camera ID vector* . Next, we assume a tuple of camera attributes, whose values pose suitable fingerprints of most cameras on the current market. We denote such a tuple of fingerprints by $\overrightarrow{\theta}$, a *fingerprint vector*. Note that a camera leaves different fingerprints in different digital images, and accordingly, each camera ID vector is associated with a set of fingerprint vectors.

C. Data

1) A Reference Data Set: Let S denote our *reference data set* – a subset of the ternary Cartesian product $Cm \times \Theta \times U$ of sets

- Cm of camera ID vectors of all existing cameras,
- Θ of all possible fingerprint vectors,
- U of user IDs representing (all potential) camera end users.

Let each element of S, which is a tuple constituted by concatenation of a camera ID vector, a fingerprint vector, and a user ID, $\langle \overrightarrow{cm}, \overrightarrow{\theta}, u\rangle$, represent a photograph (e.g., downloaded from the Internet). We say that, in accordance with $\langle \overrightarrow{cm}, \overrightarrow{\theta}, u\rangle$, the photograph has been taken by the (camera) user u with the camera $\overrightarrow{cm}$ that has left the fingerprint vector $\overrightarrow{\theta}$.

2) Information Noise: Many photographs have been software-manipulated: a software application has changed a photograph metadata so that the changed fingerprint vector does not match the camera with which the photograph actually has been taken. This is *information noise* in S. For instance, $\langle \overrightarrow{cm}, \overrightarrow{\theta}, u\rangle$ (from S) represents a photograph taken with the camera $\overrightarrow{cm}$ by the user u. Nevertheless, it does not necessarily entail that $\overrightarrow{\theta}$ is the fingerprint of $\overrightarrow{cm}$. In fact, $\overrightarrow{cm}$ and/or $\overrightarrow{\theta}$ present vectors of values that might have been software-manipulated and thus $\overrightarrow{\theta}$ may or may not be real, genuine fingerprint vector of $\overrightarrow{cm}$. This is the noise inherent in S.

IV. A Statistical Approach for Noise Removal

In general, the question arises: Given observation represented as S, our reference data set, can we quantify the "confidence" that $\overrightarrow{\theta}$ can be (real) fingerprint vector of $\overrightarrow{cm}$, where $\overrightarrow{cm}$ and $\overrightarrow{\theta}$ are a camera ID vector and a fingerprint vector found in metadata of a digital image of interest? We show how to make a lower estimation of this confidence. Our approach is based on statistical hypothesis testing.

Now we introduce some basic terminology and notation from probability theory needed for further exposition.

A. A Null Hypothesis and a Test Statistic

In brief, we analyze information noise inherent in S statistically. Specifically, given a "testing" tuple

$$t_0 = \langle \overrightarrow{cm}_0, \overrightarrow{\theta}_0\rangle,$$

our default position is that $\overrightarrow{\theta}_0$ can't be a (real) fingerprint vector of $\overrightarrow{cm}_0$. That is, all the tuples from S containing both $\overrightarrow{cm}_0$ and $\overrightarrow{\theta}_0$ represent information noise only. Accordingly, we set out the following null hypothesis

$$H_0: \quad ''\overrightarrow{\theta}_0 \mathrm{can't be a fingerprint of}\ \overrightarrow{cm}_0\ \overrightarrow{cm}_0''$$

and introduce a *test statistic*, which, in general, is a numerical summary of S that reduces S to a set of values that can be used to perform the hypothesis test. It will be seen that the test statistic enables to estimate how unlikely S is under the H_0 assumption. In particular, we determine the upper estimation of the probability p of (possible) observing the test statistic that is at least as extreme as the test statistic that has been actually observed.

Definition 1 (Test Statistic): Let T denote the mapping

$$T : Cm \times \Theta \longrightarrow \mathbb{N}_0$$

that maps each pair $\langle \overrightarrow{cm}, \overrightarrow{\theta}\rangle$ from the binary Cartesian product $Cm \times \Theta$ to the cardinality ($\mathbb{N}_0$ denotes the set of nonnegative integers) of the set of all and only those users who, in accordance with S, have taken some image with the camera cm that leaves the fingerprint vector $\overrightarrow{\theta}$. Then the *test statistic* is defined as the image of $\langle \overrightarrow{cm}_0, \overrightarrow{\theta}_0\rangle$ under T:

$$T\left(\overrightarrow{cm}_0, \overrightarrow{\theta_0}\right) = \left|\left\{u \mid \langle \overrightarrow{cm}_0, \overrightarrow{\theta_0}, u\rangle \in S\right\}\right|.$$

□

Taking into account possible software manipulations, $\overrightarrow{cm}_0$ or $\overrightarrow{\theta}_0$ (or both) in the reality might have been changed ("damaged") in metadata of the photograph by a software application used by a user to modify the photograph. Actually, H_0 presupposes that any photograph that indicates in its metadata $\overrightarrow{cm}_0$ and $\overrightarrow{\theta_0}$ must have been software-modified. Accordingly, under H_0, $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ is the number of all the (distinct) users (captured in S) who have modified photographs with software applications that, in accordance with S, have left $\overrightarrow{cm}_0$ and $\overrightarrow{\theta_0}$ in metadata of respective photographs.

Speaking in broad terms, we conclude that the number of these users is too big to be attributed exclusively to H_0 and information noise in S if the number exceeds a specified threshold. To determine the threshold, we define the *sampling distribution* of $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$. The sampling distribution is derived from three parameters that partially capture our knowledge on S. It will be seen that this sampling distribution is the *hypergeometric distribution*.

To work with data that have desirable statistical properties wrt. $\overrightarrow{cm}_0$, $\overrightarrow{\theta_0}$, we restrict ourselves only to some specific test subset of S – tuples containing camera ID vectors from some set C. Essentially, C includes camera vectors $\overrightarrow{cm}_i$ only of those cameras (" $\overrightarrow{cm}_i$ -cameras")

a) that never produce $\overrightarrow{\theta_0}$,
b) each user has taken exactly the same number of photographs with a $\overrightarrow{cm}_i$ -camera.

a) can be ensured by an expert—at least to a certain degree of certainty. b) can be achieved by a random selection from our database: Consider photographs taken by a user u_j with a $\overrightarrow{cm}_i$ -camera for example. We have to select a certain number of such photographs arbitrarily and disregard them if the user u_j has taken to many photographs with the $\overrightarrow{cm}_i$ -camera. Both a) and b) are crucial for the statistical modeling.

Definition 2 (A Test Subset C of Cameras): C is a set of camera ID vectors such that:

i. For each camera (identified by) $\overrightarrow{cm}_i$ from C, its fingerprint vectors are (always) different from $\overrightarrow{\theta}_0$.
ii. As, in accordance with H_0, fingerprint vectors of $\overrightarrow{cm}_0$ are always different from $\overrightarrow{\theta_0}$, we include $\overrightarrow{cm}_0$ in C.
iii. There is a fixed positive integer k such that, for any pair of $\overrightarrow{cm}_i$ from C and a user ID u_j, exactly k photographs are represented as respective tuples in S. That is,

$$|U_{ij}| = k \tag{1}$$

where

$$U_{ij} = \left\{\langle \overrightarrow{cm}_i, \overrightarrow{\theta}, u_j\rangle \;\middle|\; \langle \overrightarrow{cm}_i, \overrightarrow{\theta}, u_j\rangle \in S \mathrm{and} \overrightarrow{cm}_i \in C\right\}. \tag{2}$$

B. Initial Probabilities

First we address the initial probability space, denoted as the triplet $\left(\Omega, 2^{\Omega}, P\right)$, where Ω is a *sample space*, 2^{Ω} its powerset, and $P;\ 2^{\Omega} \rightarrow [0,1]$ the *probability measure* P that takes the simple form: $P(D) = \sum_{\omega \in D} p(\omega)$ for any event D, i.e., $D \subseteq \Omega$, where $p;\ \Omega \rightarrow [0,1]$ is a uniform probability mass function, i.e., $p(\omega) = 1/|\Omega|$ for each ω from Ω.

Next, we introduce *events* A, B, U_{ij}, i.e., subsets of Ω.

- Ω represents reading a tuple (from S) with $\overrightarrow{c}\,m$ from C:

$$\Omega = \left\{ \langle \overrightarrow{cm},\ \overrightarrow{\theta}, u\rangle \ \middle|\ \langle \overrightarrow{cm},\ \overrightarrow{\theta}, u\rangle \in S \text{ and } \overrightarrow{cm} \in C \right\}. \quad (3)$$

  That is, Ω includes exactly those photographs that have been taken with some "C-camera" (a camera identified with an ID vector $\overrightarrow{cm}_i$ from C). Note that, in accordance with the third requirement on C (Def. 2), $|\Omega| = k \cdot N(C)$ where $N(C) = \left|\left\{\langle \overrightarrow{cm}, u\rangle \ \middle|\ \langle \mathrm{cm},\ \overrightarrow{\theta}, u\rangle \in \Omega\right\}\right|$.
- The event A represents reading a tuple (from S) with the camera ID vector $\overrightarrow{cm}_0$:

$$A = \left\{ \langle \overrightarrow{cm}_0,\ \overrightarrow{\theta}, u\rangle \ \middle|\ \langle \overrightarrow{cm}_0,\ \overrightarrow{\theta}, u\rangle \in \Omega \right\}. \quad (4)$$

  That is, A includes exactly those photographs that have been taken with the $\overrightarrow{cm}_0$ -camera. Note that, in accordance with the third requirement on C, $|A| = k \cdot G\left(\overrightarrow{cm}_0\right)$ where $G\left(\overrightarrow{cm}_0\right) = \left|\left\{\langle \overrightarrow{cm}_0, u\rangle \ \middle|\ \langle \overrightarrow{cm}_0,\ \overrightarrow{\theta}, u\rangle \in \Omega\right\}\right|$.
- The event B represents reading a tuple (from S) with $\overrightarrow{cm}$ from C and the fingerprint vector $\overrightarrow{\theta_0}$:

$$B = \left\{ \langle \overrightarrow{cm}, \overrightarrow{\theta_0}, u\rangle \ \middle|\ \langle \overrightarrow{cm}, \overrightarrow{\theta_0}, u\rangle \in \Omega \right\}.$$

  That is, B includes exactly those photographs with the $\overrightarrow{\theta_0}$ fingerprint that have been taken with some C-camera. As B depends on $\overrightarrow{\theta_0}$ and C, we denote its cardinality $|B|$ as the value $n(\overrightarrow{\theta_0}, C)$ of a two variable function n:

$$|B| = n\left(\overrightarrow{\theta_0}, C\right). \quad (5)$$

- An event U_{ij} represents reading a tuple (from S) with $\overrightarrow{cm}_i$ from C and the user ID u_j:

$$U_{ij} = \left\{ \langle \overrightarrow{cm}_i,\ \overrightarrow{\theta}, u_j\rangle \ \middle|\ \langle \mathrm{cm}_i,\ \overrightarrow{\theta}, u_j\rangle \in \Omega \text{and} \overrightarrow{cm}_i \in C \right\}. \quad (6)$$

  That is, U_{ij} includes exactly those $\overrightarrow{cm}_i$ -photographs that have been taken by a user u_j . Note that this equation coincides with (2).

Most importantly, note that $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0}) = |A \cap B|$. Consequently, we need to derive the distribution of the cardinality of $A \cap B$. To this end, we use only the partial knowledge of S captured by its numerical characteristics: $G(\overrightarrow{cm}_0)$, $N(C)$, and $n(\overrightarrow{\theta_0}, C)$. First we will be concerned with the *conditional probability* $P(A|B)$ of A given B, which is the probability of A, given the occurrence of B, defined by the following equality: $P(A|B) = P(A \cap B)/P(B)$. Then we will study how the conditional probability changes after new evidence is taken into account: we will consider removing specific tuples from B (Section IV-C). This will provide us with all the knowledge we need to derive the probability model of $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ (Section IV-D).

Lemma 1 (Conditional Probability): $P(A|B)$ coincides with the *unconditional (marginal)* probability of A: $P(A|B) = P(A)$.

Proof: First observe that it follows from H_0 and the first two requirements on C (Def. 2) that

$$P(B|A) = P(B). \quad (7)$$

Then the lemma follows readily from the Bayes' theorem: $P(A|B) = P(B|A) \cdot P(A)/P(B) = P(A)$.. ■

As a result, we get:

$$P(A|B) = \sum_{\omega \in A} p(\omega) = \frac{|A|}{|\Omega|} = \frac{k \cdot G\left(\overrightarrow{cm}_0\right)}{k \cdot N(C)} = \frac{G\left(\overrightarrow{cm}_0\right)}{N(C)},$$

the probability that a tuple $\langle \overrightarrow{cm}, \overrightarrow{\theta_0}, u\rangle$ chosen arbitrarily from B also is in A, i.e., $\overrightarrow{cm} = \overrightarrow{cm}_0$.

C. Probabilities After Removing Tuples

Now we will study how the conditional probability changes after new evidence is taken into account.

1) Conditional Probability After Removing A 1-th Tuple: Suppose that we have removed a tuple $\langle \overrightarrow{cm}_i, \overrightarrow{\theta_0}, u_j\rangle$ from B. How has this affected the (conditional) probability that another tuple from B is also in A? To answer this question, we consider another pair A', B' of events, where $A' \subseteq A,\ B' \subseteq B$, and pursue the conditional probability of A' given B'. To this end, we introduce a new probability space denoted as $\left(\Omega', 2^{\Omega'}, P'\right)$ where $\Omega' \subseteq \Omega,\ A' \subseteq \Omega',\ B' \subseteq \Omega'$, and where the conditional probability $P'(A'|B')$ can be expressed easily. Again, we assign the probability measure P' the following, simple form: $P'(D) = \sum_{\omega \in D} p'(\omega)$ for any event D, i.e., $D \subseteq \Omega'$, where $p' : \Omega' \rightarrow [0,1]$ is a uniform probability mass function, i.e., $p'(\omega) = 1/|\Omega'|$ for each ω from Ω'.

To sum up,

- Ω' represents reading a tuple (from S) different from any tuple from U_{ij}:

$$\Omega' = \Omega - U_{ij}.$$

  Note that, in accordance with (3) and(6), $U_{ij} \subseteq \Omega$, and thus by (1), $|\Omega'| = |\Omega| - k$.
- The event A' represents reading a tuple (from S) with the camera ID vector $\overrightarrow{cm}_0$ but different from any tuple from U_{ij}:

$$A' = A - U_{ij}.$$

Note that either

$$\overrightarrow{cm}_\imath = \overrightarrow{cm}_0 \tag{8}$$

and thus we have $U_{\imath\jmath} \subseteq A$ by (4) and (6), which entails $|A'| = |A| - k$, due to (1), or

$$\overrightarrow{cm}_\imath \neq \overrightarrow{cm}_0 \tag{9}$$

and thus $U_{\imath\jmath} \cap A = \emptyset$, which entails $|A'| = |A|$.

- The event B' represents reading a tuple (from S) with $\overrightarrow{c}\,m$ from C and the fingerprint vector $\overrightarrow{\theta_0}$ but different from $\langle \overrightarrow{cm}_\imath, \overrightarrow{\theta_0}, u_\jmath \rangle$:

$$B' = B - \{\langle \overrightarrow{cm}_\imath, \overrightarrow{\theta_0}, u_\jmath \rangle\} = B - U_{\imath\jmath},$$
$$|B'| = |, |B| - 1.$$

The following two lemmata justify addressing conditional probability $P'(A'|B')$ in the newly defined probability space $(\Omega', 2^{\Omega'}, P')$.

Lemma 2 (Invariance of Conditional Probabilities): The conditional probability of A given B' in the probability space $(\Omega, 2^\Omega, P)$ coincides with the conditional probability of A' given B' in the probability space $(\Omega', 2^{\Omega'}, P')$:

$$P(A|B') = P'(A'|B').$$

□

Proof: The assertion of the lemma follows readily from the uniformity of respective probability mass functions p, p'. Indeed, $P(A|B') = P(A \cap B')/P(B') = \sum_{\omega \in A \cap B'} p(\omega) / \sum_{\omega \in B'} p(\omega) = |A \cap B'|/|B'| = |A' \cap B'|/|B'| = \sum_{\omega \in A' \cap B'} p'(\omega) / \sum_{\omega \in B'} p'(\omega) = P'(A' \cap B')/P'(B') = P'(A'|B')$. ■

Lemma 3 (Conditional Probability): $P'(A'|B')$ coincides with the *unconditional (marginal)* probability of A':

$$P'(A'|B') = P'(A').$$

Proof: It can be observed that it follows from H_0 and the requirements on C (Def. 2) that $P'(B'|A') = P'(B')$ if we assume that all the users whose IDs are in tuples in Ω are equally probable to modify their photographs with a software application. This is a simplifying assumption, but we argue that it doesn't poses significant accuracy damage to our model. Then the lemma follows from the following chain of equalities: $P'(A'|B') = P'(B'|A') \cdot P'(A')/P'(B') = P'(A')$. ■

As a result, we get the probability (in the initial probability space $(\Omega, 2^\Omega, P)$) that a triplet chosen at random from B' also is in A: $P(A|B') = P'(A') = \sum_{\omega \in A'} p'(\omega) = |A'|/|\Omega'|$, which is equal to $k \cdot G\left(\overrightarrow{cm}_0\right) - k/k \cdot N(C) - k = G\left(\overrightarrow{cm}_0\right) - 1/N(C) - 1$, if (8) holds, or $k \cdot G\left(\overrightarrow{cm}_0\right)/k \cdot N(C) - k = G(\mathrm{cm}_0)/N(C) - 1$, otherwise, i.e., when (9) holds.

2) Conditional Probability After Removing An ℓ -th Tuple: Repeating the above train of thoughts, it can be observed that we arrive at the following general rule describing the conditional probability of A given B from which some tuples have been removed. Suppose that

- B^ℓ is defined as B from which we have removed ℓ tuples: $B^\ell = B - \left\{\langle \overrightarrow{cm}_{\imath_1}, \overrightarrow{\theta_0}, u_{\jmath_1} \rangle, \ldots, \langle \overrightarrow{cm}_{\imath_\ell}, \overrightarrow{\theta_0}, u_{\jmath_\ell} \rangle\right\}, |B| > \ell$, and
- k values from $\overrightarrow{cm}_{\imath_1}, \ldots, \overrightarrow{cm}_{\imath_\ell}$ coincide with $\overrightarrow{cm}_0$.

Then the probability $P(A|B^\ell)$ that another triplet $\langle \overrightarrow{cm}_{\imath_{\ell+1}}, \overrightarrow{\theta_0}, u_{\jmath_{\ell+1}} \rangle$ read from B^ℓ also is included in A is

$$P\left(A|B^\ell\right) = \frac{G\left(\overrightarrow{cm}_0\right) - k}{N(C) - \ell} \tag{10}$$

provided that C fulfills the three requirements and H_0 holds.

D. A Probability Model of the Test Statistic

In this subsection, we view $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ as a random phenomenon and show how its probability model is derived from (10). Specifically, we show that $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ can be modeled as a *random variable* that follows the *hypergeometric probability distribution*. Perhaps the easiest way to see this is in terms of the *urn problem*, well-known in statistics.

The urn problem is an idealized mental exercise in which some objects of real interest—such as tuples from B are represented as colored balls in an urn. One pretends to draw (remove) one or more balls from the urn; the goal is to determine the probability of drawing one color or another, or some other properties.

We use the well known *urn model* that contains red and blue balls that are not returned to the urn once drawn. Knowing that G out of N balls in the urn are red, it is easily seen that the probability of drawing a red ball provided that k out of ℓ balls drawn from the urn are red is equal to the $G - k$ to $N - \ell$ ratio,

$$\frac{G - k}{N - \ell}\ , \tag{11}$$

i.e., the proportion of red balls remaining in the urn. In particular, note that it is well known in probability theory and statistics that the number of k red balls in a sequence of n draws from this urn (without replacement) has the hypergeometric distribution whose *probability mass function*

$$h: \mathbb{N}_0 \longrightarrow [0, 1]$$

is defined by the following rule:

$$h(k) = \frac{\binom{G}{k}\binom{N-G}{n-k}}{\binom{N}{n}}. \tag{12}$$

Consequently, observing that

- Equation (11) coincides with the rule (10) if we set

$$G = G\left(\overrightarrow{cm}_0\right), \quad N = N(C), \tag{13}$$

and view

— the k tuples $\langle \overrightarrow{cm}_0, \overrightarrow{\theta_0}, u_{\jmath_1} \rangle, \ldots, \langle \overrightarrow{cm}_0, \overrightarrow{\theta_0}, u_{\jmath_k} \rangle$ removed from $A \cap B$ as red balls

— and the other $\ell - k$ tuples, removed from $B - A$, as blue balls,

- Equation (5) holds, i.e., the cardinality of B is $n(\overrightarrow{\theta_0}, C)$,

the following theorem is clear upon reflection.

Theorem 1 (Sampling Distribution of Test Statistic): Suppose that H_0 holds, and C fulfills all the three requirements (Def. 2). Then the sampling (discrete cumulative) distribution of the test statistic $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ coincides with the hypergeometric (cumulative) distribution function

$$H : \mathbb{N}_0 \longrightarrow [0, 1]$$

defined by the following rule

$$H(x) = \sum_{k=0}^{x} h(k) \tag{14}$$

where $h(k)$ is given by (12),(13), and $n = n(\overrightarrow{\theta_0}, C)$.

E. Confidence of Correctly Rejecting the Null Hypothesis

Now we discuss an important subtlety of the three requirements on C in the above theorem. Admittedly, the first one may be hard to fulfill as we might have no prior information on cameras. Specifically, we might not know cameras that never leave the fingerprint vector $\overrightarrow{\theta_0}$ in images. In fact, even no expert might know such cameras. Accordingly, we, in general, are able to fulfill the first requirement on C only partially—to some degree of certainty less or equal to 100%. The following corollary addresses the estimate of the *p-value*

$$p = 1 - H\left(T\left(\overrightarrow{cm}_0, \overrightarrow{\theta_0}\right)\right), \tag{15}$$

which is interpreted as the probability of observing a value for the test statistic at least as extreme as $T(\overrightarrow{cm}_0, \overrightarrow{qt}_0)$, assuming that the null hypothesis H_0 is true and C fulfills all the three requirements.

Corollary 1 (p-Value): We get an upper estimation of the p-value if the first requirement on C is fulfilled only partially.

Proof: To see the assertion of the corollary, recall that, among others, the first requirement on C conditions the independence of events A and B (refer to Lemmata 1 and 3). Without the guarantee that the requirement is fulfilled, their independence can't be assumed any more. That is, the conditional probability $P(A|B)$ of A given B may not coincide with the unconditional (marginal) probability of A.

More formally, in the probability space $(\Omega, 2^{\Omega}, P)$, the conditional probability of A given B is no more than the unconditional (marginal) probability of A if H_0 and the second and third requirements on C hold:

$$P(A|B) \leq P(A). \tag{16}$$

Essentially, this claim follows from the observation that the equality, in general, can't be assumed any more in (7). Instead, observe that $P(B|A) \leq P(B)$ holds. Therefore, following the lines of the proof of Lemma 1, we get (16). Similarly, it can be observed that neither the independence of A' and B' can be assumed any more, and by the same argument as above, we get $P'(A'|B') \leq P'(A')$. Then repeating the train of thoughts as in Section IV-C, it can be seen that $P\left(A|B^{\ell}\right) \leq G\left(\overrightarrow{cm}_0\right) - k/N(C) - \ell$, i.e., the rule (10) may overvalue $P(A|B^{\ell})$. Because of this, it follows from properties of

- sampling (cumulative) distributions, which define *finite numerical* , *monotonically increasing sequences* with the greatest members equal to 1,
- the hypergeometric distribution,

that the (real) sampling (discrete cumulative) distribution $H'(x)$ of $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$, is greater or equal to $H(x)$(14) for any nonnegative integer x that is less or equal to n: $0 \leq x \leq n \Rightarrow H(x) \leq H'(x)$. Now the corollary is immediate by (15). ■

Note that rejecting H_0 entails accepting the *alternative hypothesis* , namely that $\overrightarrow{\theta_0}$ is a (possible) fingerprint that a camera $\overrightarrow{cm}_0$ may leave in a photograph metadata. As a result, the confidence of correct accepting the alternative hypothesis can be quantified by the value $1 - p$ with the well-known, rigorous interpretation: the estimate of the probability $H\left(T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})\right)$ of observing a value for the test statistic less extreme than $T(\overrightarrow{cm}_0, \overrightarrow{\theta_0})$ if the null hypothesis H_0 is true and provided that C fulfills all the three requirements (Def. 2). The presented statistical approach provides a lower estimate of this probability.

V. Basics of JPEG Compression

Since camera fingerprints employed to verify the originality of digital images in the next (experimental) part are directly related to the JPEG encoder and file format, it is necessary to briefly introduce the basic idea behind JPEG.

Every JPEG image file consists of a sequence of segments carrying information about the image, codec, producer, etc. Each segment begins with a marker having binary format 0xFF followed by a byte indicating what kind of marker it is. For instance, 0xFFD8 defines SOI (Start of image), which means the entry point of the JPEG image file. On the other hand, 0xFFD9 defines EOI (End of image), which means the ending point of the JPEG image file. Typical JPEG files contain markers defining a thumbnail image, used Huffman tables, Quantization tables (QTs), etc. Basic format of markers is shown below:

$$0xFF + Marker\ Number\ (1\ byte) \\ + Data\ size(2\ bytes) + Data\ (n\ bytes).$$

Although JPEG file can be encoded in various ways, the most common algorithm is the following one. Typically, the image is first converted from RGB to YCbCr, consisting of one luminance component (Y), and two chrominance components (Cb and Cr). Then each component is split into adjacent blocks of 8×8 pixels. After this step, each block undergoes a discrete cosine transform (DCT) resulting in 64 DCT coefficients, $F(u, v)$, for each block. In the next step, all 64 $F(u, v)$ coefficients are quantized. This is done by simply dividing each component in the frequency domain by a constant for that component and then rounding to the nearest integer. Quantization steps for each DCT frequency u and v are defined in quantization tables $QT(u, v)$.

These QTs can be found in the EXIF of JPEG file and are denoted by a marker called DQT (Define Quantization Table) beginning with 0xFF and followed by 0xDB. In the final step, entropy coding is carried out. For more detailed information on JPEG, please refer to [34].

VI. Experimental Results

Section IV introduced a statistical method for addressing the noise in reference knowledge databases. In order to demonstrate its efficiency, we need to evaluate the *power of the test*, which is the probability that the test will reject the H_0 when H_0 is false, i.e., the probability of not committing an *error of the second kind* (making a false negative decision). To this end, we applied the presented method to a collection of a large number of random digital images that are original, i.e., not modified by software.

A. Proposed Fingerprints

There are a number of fingerprints which can be used to distinguish between original and altered JPEG images. In this experiment, we chose the following JPEG-related fingerprint types:

- $F_{Markers}$– EXIF markers,
- F_{QTs}– luminance and chrominance quantization tables,
- F_{Thumb}– information on the JPEG thumbnail image:
  - — the thumbnail width and height,
  - — luminance and chrominance QTs, and Huffman tables, used in encoding the thumbnail image,
  - — chroma subsampling scales of the thumbnail image (both horizontal and vertical directions).

First of all, we tried to employ widely used libraries for extraction of JPEG related data to extract the above mentioned fingerprints. Unfortunately, we learned that these libraries are not capable of extracting so detailed and precise features from JPEG files due to high variety of JPEG file formats available on market as well as due to a number of imperfections brought into JPEG files by camera or software producers. For these reasons, we created our own JPEG forensics fingerprint reader. This reader was optimized and applied to 5 millions JPEG images of various formats in the period of 6 months.

B. Reference Image Data Set

As pointed out in previous sections, to automatically differentiate between noisy and original data, we need a large reference data set. Therefore we collected a large number of digital images from a noncontrolled image arena. Keeping at disposition a variety of popular photo-sharing servers from which photos can be downloaded, we opted for Flickr, one of the most popular photo sharing sites. We downloaded 5 million images labeled as "original". Nevertheless, as has been pointed out, Flickr, in fact, is an "uncontrolled arena:" Flickr photos are with no guarantee that they have been captured with the camera as "officially" indicated in their metadata. Indeed, Flickr has no practical reason to filter out modified images. Most often seen camera makers in our database are Apple, Canon, Casio, Eastman Kodak, Fuji, Hewlett-Packard, Nikon, Nokia, Panasonic, Pentax, Olympus, Samsung, Sony, etc.

C. Power of Test

Power of the test, is defined as $1-\beta$, where β is the probability of the error of the second type, also referred to as *false negative rate*, i.e., "failing to reject H_0" in our setting. (Please note that "failing to reject H_0" means "not to reject H_0 when H_0 is not true.")

To carry out experiments, we picked 24 cameras (see Table II). Each camera has been used to capture 100 digital images of indoor and outdoor scenes resulting in a set of 2400 digital images in total. These images are guaranteed to be original: they present our ground-truth data.

Since camera settings can directly affect fingerprint values such as, e.g., quantization tables, we have imposed no restrictions on them when capturing ground-truth digital images. All photographers producing ground-truth data were totally free to capture photos as they wished. In this way, we attempted to minimize any systematic influence on experimental test data and results. Moreover, photos downloaded from the Internet forming the reference image data set also had no restrictions on camera settings.

For every image and each of the three types of fingerprints, we have repeated a statistical test procedure with the significance level set to 1% and 5%, which is the probability of the *error of the first kind*, also referred to as *false positive rate*, i.e., mistakenly rejecting H_0 provided that H_0 is true. Thus we have obtained six sequences (columns) of results in Table II, revealing the powers of our test for respective combinations of fingerprints and significance levels. For example, the column headed $F_{Markers}[1\%]$ shows a sequence of 24 values that correspond to the power of the test based on the $F_{Markers}$ fingerprint type at the significance level 1%.

To perform the experimental test, we assumed presumably the most common and arguably also the most challenging scenario with a lay user, who has no background knowledge concerning camera properties (e.g., of a specific camera fingerprint values). Accordingly, we opted for a coarse, ignorant approach and included all the cameras from S, our reference data set, in the set C, the parameter of the sampling distribution of the test statistic. Consequently, in accordance with Corollary 1, we had to expect to obtain a rather coarse upper estimate of the p-value.

To fully understand the results shown in Table II, it is crucial to point out that knowledge represented by the reference set is limited due to the size of this database. Thus, it happens that a particular fingerprint obtained from an image being tested is not found in the reference set. Apparently, the more we want to eliminate this problem, the bigger reference image data set we have to accumulate to capture bigger knowledge. Still, considering the high variety of cameras and software packages, it is impossible to create a complete reference image data set covering all existing fingerprints. Nonetheless, there are several ways how to deal with this kind of fingerprints that have no associated data in the reference data set. In our experiments, we opted for drawing no conclusion about the originality of this kind of tested fingerprints, filtering them out from results. This happened

- 341 times for $F_{Markers}$,
- 528 times for F_{QTs},

TABLE II
POWER OF THE TEST. DATA IN EACH CELL (A NUMBER OF TIMES OF REJECTING H_0 CORRECTLY – OUT OF 100 TESTED IMAGES THAT HAVE BEEN TAKEN WITH A SHOWN CAMERA MAKER, CAMERA MODEL) ARE OBTAINED USING 2400 JPEG IMAGES ACQUIRED BY 24 DIFFERENT CAMERAS

Camera maker	Camera Model	$F_{Markers}[1\%]$	$F_{QTs}[1\%]$	$F_{Thumb}[1\%]$	$F_{Markers}[5\%]$	$F_{QTs}[5\%]$	$F_{Thumb}[5\%]$
Apple	iPhone 4	100	100	100	100	100	100
Apple	iPhone 4S	100	100	100	100	100	100
Canon	Digital Ixus 55	78	96	99	88	98	100
Canon	Digital Ixus 70	91	93	97	92	98	99
Canon	EOS 20D	80	79	89	84	86	98
Canon	EOS 30D	76	77	81	83	79	81
Canon	EOS 50D	72	75	90	83	82	90
Canon	EOS 5D Mark II	81	79	81	88	86	88
Canon	PowerShot A3000 IS	80	94	94	91	99	100
Canon	PowerShot A640	100	96	100	100	96	100
Canon	PowerShot A75	83	91	95	88	99	100
Canon	PowerShot G12	86	86	94	93	97	99
Casio	EX-Z150	82	80	91	100	80	94
Fujifilm	FinePix J50	100	92	96	100	98	98
Konica	KD-400Z	100	62	100	100	81	100
Nikon	E990	100	91	97	100	100	100
Nikon	Nikon D200	93	81	91	93	85	96
Nikon	Nikon D70	100	69	97	100	83	97
Olympus	C740UZ	100	84	87	100	98	98
Olympus	SP600UZ	100	57	100	100	100	100
Olympus	u1010,S1010	79	82	94	82	100	100
Olympus	X450,D535Z,C370Z	100	96	100	100	98	100
Panasonic	DMC-LX2	94	81	82	94	84	84
Sony	DSC-W40	100	87	100	100	89	100

- 783 times for F_{Thumb} out of 2400 tested images in each case. i.e., 7200 extracted fingerprints in total.

VII. DISCUSSION AND SUMMARY

We point out that in accordance with Corollary 1, results of our experiments are affected by the test subset C of cameras (Def. 2) in the expected fashion. For example, note the value 57 for Olympus SP600UZ in the $F_{QTs}[1\%]$ column in Table II. It says that 43 out of 100 tests on images captured by Olympus SP600UZ failed to reject H_0. This is the biggest error of the second kind that we have obtained in our experiments. Most of these tests fail to reject H_0 because the fingerprint value used in these tests is commonly produced by too many other camera models. Because of this, our coarse, ignorant approach, when all the cameras from S (our reference data set) are included in the set C (the parameter of the sampling distribution of the test statistic) is "too ignorant," resulting in fulfilling the first requirement on C only partially. Hence, by Corollary 1, we get too coarse upper estimations of 43 respective p-values for the QTs fingerprints in the tested images captured by Olympus SP600UZ: all 43 were greater than the significance level 1%. Nevertheless, none of them was greater than the significance level 5% as is documented by the corresponding 100 value in the $F_{QTs}[5\%]$ column.

Careful selection of cameras to be included in C, which can be managed by an expert or based on an appropriate heuristics (exploiting some background knowledge concerning cameras), will improve results remarkably: it will make the estimate of the p-value more accurate, which in turn will eliminate rare huge errors of the second kind. Altogether, this will result in increase of the already high power of the test, and thus we get a lower probability of failing to reject H_0 when H_0 is (really) false (the error of the second kind).

It is clear that also the definition of fingerprint vectors is crucial for the accuracy of the presented approach. Recall that we took into account three various fingerprint vectors in our experiments. Comparing values in respective columns in Table II, it is seen that respective fingerprint vectors yield different results. Most importantly, observe that a big error of the second kind associated with one fingerprint vector is usually "compensated" by very small error of the second kind associated with another fingerprint vector and vice versa. Pursuit for an ideal fingerprint vector that would yield the least error of the second kind is a very tempting challenge, which however requires bigger practical experience with the proposed method as well as further investigation and extensive testing. Therefore it is left for future research.

Our goal was to estimate confidence that a given digital image truly may have been taken by a camera indicated in the image metadata. We based ourselves on carefully selected fingerprints: we confronted them with a large database of various kinds of acquisition devices and their fingerprints, extracted from images coming from an uncontrolled environment. Handling the information noise inherent in such a reference database populated with data from "unguaranteed" sources is a complex task as cameras and pieces of software often have complex and unpredictable behavior. Indeed, many pieces of software modify images (e.g., enhance the contrast or rotate the image) without leaving any obvious traces in their JPEG file metadata. Despite this, the proposed approach proves to be extremely effective. On top of that, the approach is general and can be applied straightforwardly to other features formed by acquisition devices and software packages stored in various image file formats.

REFERENCES

[1] M. Arnold, M. Schmucker, and S. D. Wolthusen, *Techniques and Applications of Digital Watermarking and Content Protection Inc*. Norwood, MA, USA: Artech House, 2003.

[2] S. Bayram, H. T. Sencar, and N. D. Memon, "Classification of digital camera-models based on demosaicing artifacts," *Digital Investigation*, vol. 5, no. 1-2, pp. 49–59, 2008.

[3] S. Bayram, H. T. Sencar, N. D. Memon, and I. Avcibas, "Source camera identification based on CFA interpolation," in *Proc. ICIP*, 2005, vol. 3, pp. 69–72.

[4] O. Celiktutan, I. Avcibas, and B. Sankur, "Blind identification of source cell-phone model," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 553–566, Sep. 2008.

[5] M. Chen, M. Goljan, and J. Lukas, "Determining image origin and integrity using sensor noise," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 1, pp. 74–90, Mar. 2008.

[6] A. E. Dirik, S. Bayram, H. T. Sencar, and N. Memon, "New features to identify computer generated images," in *Proc. IEEE Int. Conf. Image Processing (ICIP '07)*, 2007, vol. 4, pp. 433–436.

[7] A. E. Dirik, H. T. Sencar, and N. Memon, "Digital single lens reflex camera identification from traces of sensor dust," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 539–552, Sep. 2008.

[8] H. Farid, Digital Image Ballistics From JPEG Quantization, Department of Computer Science, Dartmouth College, Tech. Rep. TR2006–583, 2006.

[9] J. Fridrich, "Digital image forensics," *IEEE Signal Process. Mag.*, vol. 2, no. 26, pp. 26–37, Mar. 2009.

[10] J. Fridrich and T. Pevny, "Detection of double-compression for applications in steganography," *IEEE Trans. Inf. Security Forensics*, vol. 3, no. 2, pp. 247–258, Jun. 2008.

[11] J. Fridrich, D. Soukal, and J. Lukas, "Detection of copy-move forgery in digital images," in *Proc. Digital Forensic Res. Workshop IEEE Comput. Soc.*, Cleveland, OH, USA, Aug. 2003, pp. 55–61.

[12] M. Goljan, J. J. Fridrich, and T. Filler, "Large scale test of sensor fingerprint camera identification," *Media Forensics Security*, vol. 12, p. 72540, 2009.

[13] H. Huang, W. Guo, and Y. Zhang, "Detection of copy-move forgery in digital images using sift algorithm," in *Proc. 2008 IEEE Pacific-Asia Workshop on Computational Intell. and Ind. Applicat. IEEE Comput. Soc. (PACIIA '08)*, Washington, DC, USA, 2008, pp. 272–276.

[14] Y. Huang and Y. Long, Demosaicking recognition with applications in digital photo authentication based on a quadratic pixel correlation model 2008, pp. 1–8.

[15] E. Kee and H. Farid, "Digital image authentication from thumbnails," *Proc. SPIE, Electronic Imaging, Media Forensics and Security XII*, vol. 12, pp. 72792–72799, 2010.

[16] M. Kharrazi, H. T. Sencar, and N. D. Memon, "Blind source camera identification," in *Proc. ICIP*, 2004, pp. 709–712.

[17] J. D. Kornblum, "Using JPEG quantization tables to identify imagery processed by software," in *Proc. Digital Forensic Workshop*, Aug. 2008, pp. 21–25.

[18] Z. Lint, R. Wang, X. Tang, and H.-Y. Shum, "Detecting doctored images using camera response normality and consistency," in *Proc. 2005 IEEE Comput. Soc. Conf. on Comput. Vision and Pattern Recognition (CVPR'05) IEEE Comput. Soc.)*, Washington, DC, USA, 2005, vol. 1, pp. 1087–1092.

[19] B. Mahdian and S. Saic, "Detection of copy-move forgery using a method based on blur moment invariants," *Forensic Sci. Int.*, vol. 171, no. 2– 3, pp. 180–189, 2007.

[20] B. Mahdian and S. Saic, "Blind authentication using periodic properties of interpolation," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 3, pp. 529–538, Sep. 2008.

[21] B. Mahdian and S. Saic, "Detecting double compressed JPEG images," in *Proc. 3rd Int. Conf. Imaging for Crime Detection and Prevention (ICDP-09)*, London, U.K., Dec. 2009.

[22] B. Mahdian and S. Saic, "A bibliography on blind methods for identifying image forgery," *Image Commun.*, vol. 25, no. 6, pp. 389–399, 2010.

[23] T.-T. Ng and S.-F. Chang, "An online system for classifying computer graphics images from natural photographs," in *Proc. SPIE Electron. Imaging*, San Jose, CA, USA, Jan. 2006.

[24] T.-T. Ng and M.-P. Tsui, "Camera response function signature for digital forensics—Part I: Theory and data selection," in *Proc. IEEE Workshop Inf. Forensics and Security*, Dec. 2009, pp. 156–160.

[25] N. Nikolaidis and I. Pitas, "Robust image watermarking in the spatial domain," *Signal Process.*, vol. 66, no. 3, pp. 385–403, May 1998.

[26] W. B. Pennebaker and J. L. Mitchell, *JPEG Still Image Data Compression Standard*. Norwell, MA, USA: Kluwer Academic Publishers, 1992.

[27] A. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of resampling," *IEEE Trans. Signal Process.*, vol. 53, no. 2, pt. 2, pp. 758–767, Feb. 2005.

[28] A. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," *IEEE Trans. Signal Process.*, vol. 53, no. 10, pp. 3948–3959, Oct. 2005.

[29] Z. Qu, W. Luo, and J. Huang, "A convolutive mixing model for shifted double JPEG compression with application to passive image authentication," in *Proc. IEEE Int. Conf. Acoust., Speech and Signal Processing*, Las Vegas, NV, USA, Apr. 2008, pp. 4244–1483.

[30] M. Schneider and S. F. Chang, "A robust content based digital signature for image authentication," in *Proc. IEEE Int. Conf. Image Processing (ICIP'96)*, Lausanne, 1996, vol. 3, pp. 227–230.

[31] H. T. Sencar, M. Ramkumar, and A. N. Akansu, *Data Hiding Fundamentals and Applications: Content Security in Digital Multimedia*. Orlando, FL, USA: Academic, 2004.

[32] S. Servetto, C. Podilchuk, and K. Ramchandran, "Capacity issues in digital image watermarking," in *Proc. Int. Conf. Image Processing*, 1998, pp. 445–449.

[33] A. Swaminathan, M. Wu, and K. J. R. Liu, "Component forensics," *IEEE Signal Process. Mag.*, vol. 26, no. 2, pp. 38–48, Mar. 2009.

[34] C.-H. Tzeng and W.-H. Tsai, "A new technique for authentication of image/video for multimedia applications," in *Proc. 2001 Workshop on Multimedia and Security ACM Press*, New York, NY, USA, 2001, pp. 23–26.

Babak Mahdian received the M.Sc. degree in computer science from the University of West Bohemia, Czech Republic, in 2004, and the Ph.D. degree in mathematical engineering from the Czech Technical University, Prague, Czech Republic, in 2008.

His current research interests include all aspects of digital image processing and pattern recognition, particularly digital image forensics, cyber-security, OCR, and multimodal HCI. He has been awarded by several prestigious national and international scientific as well as commercial awards because of his outstanding achievements in scientific and commercial fields.

Radim Nedbal received the M.Sc. degree in computer science from the Czech Technical University, Prague, Czech Republic, in 2003, and the Ph.D. degree in mathematical engineering from the Czech Technical University, Prague, Czech Republic, in 2011.

Currently, he is with the Fondazione Bruno Kessler, Trento, Italy. His current research interests include combining statistical learning methods with symbolic approaches in artificial intelligence with the aim of enhancing reasoning capabilities in the presence of noisy data.

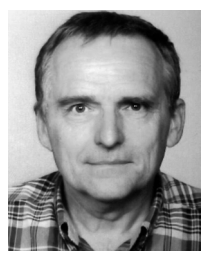

Stanislav Saic received the M.Sc. degree in physical electronics from the Czech Technical University, Prague, Czech Republic, in 1973, and the C.Sc. degree (corresponding to the Ph.D. degree) in radio electronics from the Czechoslovak Academy of Sciences, Prague, in 1980.

Since 1973, he has been with the Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Prague, where he was Head of the Department of Image Processing from 1985 to 1994. His current research interests include all aspects of digital image and signal processing, particularly Fourier transform, image filters, remote sensing, and geosciences.