

Image Forgery Localization via Fine-Grained Analysis of CFA Artifacts

Original

Image Forgery Localization via Fine-Grained Analysis of CFA Artifacts / P., Ferrara; Bianchi, Tiziano; A., De Rosa; A., Piva. - In: IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY. - ISSN 1556-6013. - 7:5(2012), pp. 1566-1577. [10.1109/TIFS.2012.2202227]

Availability:

This version is available at: 11583/2505936 since:

Publisher:

IEEE - INST ELECTRICAL ELECTRONICS ENGINEERS INC

Published

DOI:10.1109/TIFS.2012.2202227

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Image Forgery Localization via Fine-Grained Analysis of CFA Artifacts

P. Ferrara, T. Bianchi *Member, IEEE*, A. De Rosa, and A. Piva *Senior Member, IEEE*

Abstract—In this paper, a forensic tool able to discriminate between original and forged regions in an image captured by a digital camera is presented. We make the assumption that the image is acquired using a Color Filter Array, and that tampering removes the artifacts due to the demosaicing algorithm. The proposed method is based on a new feature measuring the presence of demosaicing artifacts at a local level, and on a new statistical model allowing to derive the tampering probability of each 2×2 image block without requiring to know a priori the position of the forged region. Experimental results on different cameras equipped with different demosaicing algorithms demonstrate both the validity of the theoretical model and the effectiveness of our scheme.

Index Terms—Image forensics, CFA artifacts, digital camera demosaicing, tampering probability map, forgery localization.

I. INTRODUCTION

Image forensics is a multidisciplinary science aiming at acquiring important information on the history of digital images, including the acquisition chain, the coding process, and the editing operators. The extraction of such data can be exploited for different purposes, one of the most interesting is the verification of the trustworthiness of digital data. Image forensic techniques work on the assumption that digital forgeries, although visually imperceptible, alter the underlying statistics of an image. These statistical properties can be interpreted as *digital fingerprints* characterizing the image life-cycle, during its acquisition and any successive processing. One of the tasks of image forensics is then to verify the presence or the absence of such digital fingerprints, similar to intrinsic watermarks, in order to uncover traces of tampering.

As a first basic application of the above principle, the presence/absence of forensic fingerprints can be verified on the whole image (or a given suspected region, as a sort of sub-image), thus providing information about the authenticity of the entire image (or the entire region). However, a more sophisticated result would be a sort of map indicating for each image pixel (or small image block) its trustworthiness: in this case no manual choice of suspected regions would be necessary. Currently, several fingerprints have been studied for acquiring information on an image at a global level, but only

few examples of tools that provide a fine-grained localization of forgery within a digital image have been proposed, in particular for double JPEG compression artifacts detection [1]–[5]. In many cases a sufficiently large portion of the image (e.g. a $B \times B$ block, with $B \geq 100$) is needed for a reliable statistical analysis of the chosen feature, so even if the image is processed block-wise only a coarse grained localization of tampering is possible.

In this paper, we focus our attention on the fine grained forgery localization problem, assuming to have no information on the position of possibly manipulated pixels. Among the numerous fingerprints considered in image forensic literature [6], [7], we consider the traces left by the *interpolation* process. Image interpolation is the process of estimating values at new pixel locations by using known values at neighbouring locations. During the image life cycle, there are two main phases in which interpolation is applied:

- Acquisition processing, to obtain the 3 color channels (red, green, and blue). The light is filtered by the *Color Filter Array* (CFA) before reaching the sensor (CCD or CMOS), so that for each pixel only one particular color is gathered. Thus, starting from a single layer containing a mosaic of red, green, and blue pixels, the missing pixel values for the three color layers are obtained by applying the interpolation process, also referred to as *demosaicing*.
- Geometric transformations, to obtain a transformed image. When scaling (shrinking and zooming), rotation, translation, shearing, are applied to an image, pixels within the to-be-transformed image are relocated to a new lattice, and new intensity values have to be assigned to such positions by means of interpolation of the known values, also referred to as *resampling* operation.

The artifacts left in the image by the interpolation process can be analyzed to reveal image forgery. Ideally, an image coming from a digital camera, in the absence of any successive processing, will show demosaicing artifacts on every group of pixels corresponding to a CFA element. On the contrary, demosaicing inconsistencies between different parts of the image, as well as resampling artifacts in all or part of the analyzed image, will put image integrity in doubt.

Our effort is focused on the study of demosaicing artifacts at a local level: by means of a local analysis of such traces we aim at localizing image forgeries whenever the presence of CFA interpolation is not present. Obviously our approach is based on the hypothesis that unmodified images coming from a digital camera are characterized by the presence of CFA demosaicing artifacts. Starting from such an assumption, we propose a new feature that measures the presence/absence

Copyright (c) 2012 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

P. Ferrara is with the National Institute of Optics (INO), Firenze, Italy (e-mail: pasquale.ferrara@ino.it). T. Bianchi, A. De Rosa, and A. Piva are with the CNIT Research Unit and Department of Electronics and Telecommunications, University of Florence, Firenze, Italy (e-mail: tiziano.bianchi@unifi.it, alessia.derosa@unifi.it, alessandro.piva@unifi.it).

This work was partially supported by the REWIND Project funded by the Future and Emerging Technologies (FET) programme within the 7FP of the European Commission, under FET-Open grant number: 268478.

of these artifacts even at the smallest 2×2 block level, thus providing as final output a forgery map indicating with fine localization the probability of the image to be manipulated.

The paper is organized as follows. In Section II, we will provide a brief overview of previous works considering the fingerprints left by the CFA and the interpolation process, highlighting if and how the localization problem is taken into account by the methods proposed so far. In Section III we will present a statistical model for describing the presence of CFA, and starting from it we will propose the new forgery localization algorithm and describe the overall system in Section IV. In Section V, firstly the proposed model will be validated through a set of experiments, and secondly the detection capability of the proposed forgery localization algorithm will be investigated. Some conclusions will end the paper in Section VI.

II. RELATED WORK

Previous works considering CFA demosaicing as the to be analyzed fingerprint can be divided in two main classes, i) algorithms aiming at estimating the parameters of the color interpolation algorithm, and ii) algorithms aiming at evaluating the presence/absence of demosaicing traces. Whereas the second class focuses on forgery detection (inconsistencies in the CFA interpolation reveal the presence of forged regions), algorithms within the first class are mostly intended to classify different source cameras, though sometimes they can also be used to detect tampering.

As to the first class, Swaminathan et al. in [8] propose a method for camera identification by the estimation of the CFA pattern and interpolation kernel; while in [9] the same authors exploit the inconsistencies among the estimated demosaicing parameters as proof of tampering. Cao and Kot in [10] aim at estimating the demosaicing formulas, employing a partial second-order image derivative correlation model, in order to classify different demosaicing algorithms. In [11], Bayram et al. detect and classify traces of demosaicing by jointly analyzing features coming from two previous works (see [12] and [13] below), in order to identify the source camera model. In [14], Fan et al. propose a neural network framework for recognizing the demosaicing algorithms in raw CFA images, and use it for digital photo authentication.

Regarding the detection of demosaicing traces, Popescu and Farid propose an approach for detecting the interpolation artifacts left on digital images by resampling [15] and demosaicing [12] processes. In their approach, the Expectation-Maximization algorithm is applied to estimate the interpolation kernel parameters, and a probability map is achieved that for each pixel provides its probability to be correlated to neighbouring pixels. The presence of interpolated pixels results in the periodicity of the map that is clearly visible in the Fourier domain. Such an analysis can be applied to a given image region, however a minimum size is needed for assuring the accuracy of the results: authors tested their algorithms on 256×256 and 512×512 sized areas.

Gallagher in [13] observed that the variance of the second derivative of an interpolated signal is periodic: he thus looked

for the periodicity in the second derivative of the overall image by analyzing its Fourier transform. Successively, for detecting traces of demosaicing, Gallagher and Chen proposed in [16] to apply Fourier analysis to the image after high pass filtering, for capturing the presence of periodicity in the variance of interpolated/acquired coefficients. The procedure has been tested only up to 64×64 image blocks, whereas a variation yielding a pixel-by-pixel tampering map is based on a 256-point discrete Fourier transform computed on a sliding window, thus lacking resolution.

In [17] by Dirik and Memon, the sensor noise power of the analyzed image is taken into account: its change across the image (i.e. its difference value for interpolated and acquired pixels) is considered for demonstrating the presence of demosaicked pixels. In the above paper, a block based CFA detection was also proposed, however the features proposed therein have to be computed on 96×96 blocks, thus permitting only a coarse grained localization of tampering.

Demosaicing can also be detected using methods which analyze generic resampling artifacts. In this area, Kirchner in [18], [19] consider an approach similar to [15], by observing that the actual prediction weights of the resampling filter are not necessary for revealing periodic artifacts, thus simplifying the analysis, however experimental results consider only 512×512 images. Mahdian and Saic in [20] consider the derivatives of the interpolated image and apply the method to suspected windows of size at least 64×64 , while in [21] they adopt the spectral correlation function, but only considering 512×512 sized images. Finally, in [22] Vazquez-Padin et al. demonstrate that the interpolated image is an almost cyclostationary process, with a period depending on the resampling factor. However, the authors use image blocks of size 128×128 pixels for the analysis, which only permits a coarse forgery localization.

III. CFA MODELING

During the CFA interpolation process, the estimation of the values in the new lattice based on the known values can be locally approximated as a filtering process through an interpolation kernel periodically applied to the original image to achieve the resulting image. Thus, the identification of artifacts due to CFA demosaicing can be seen as a particular case of the detection of interpolation artifacts, that has been deeply studied in these last years, as exposed in Section II.

In [18], Kirchner demonstrated that for a resampled stationary and non-constant signal $s(x)$, with $x \in \mathbb{Z}$, the variance of the residue of a linear predictor $\text{Var}[e(x)]$ is periodic with a period equal to the original sampling rate. Hence, if we consider the signal resampled according to an integer interpolation factor r , we have $\text{Var}[e(x)] = \text{Var}[e(x+r)]$, since the original sampling period corresponds to r samples of the resampled signal.

For the case of CFA demosaicing, if we consider a single dimension, the general result presented in [18] turns into $\text{Var}[e(x)] = \text{Var}[e(x+2)]$, that is the variance of the prediction error assumes only two possible values, one for the odd positions and another one for the even positions. In more

detail, considering for example the interpolation of the green color channel $G(x)$ in a particular row of the image, the acquired signal $s_A(x)$ is

$$s_A(x) = \begin{cases} G(x) & x \text{ even} \\ 0 & x \text{ odd} \end{cases} \quad (1)$$

If we consider a simplified demosaicing model, the resulting signal $s_R(x)$, composed by the acquired component $s_A(x)$ and by the interpolated component, takes values:

$$s_R(x) = \begin{cases} s_A(x) = G(x) & x \text{ even} \\ \sum_u h_u s_A(x+u) & x \text{ odd} \end{cases} \quad (2)$$

where h_u represents the interpolation kernel. In the above model, we assume that each color channel is independently interpolated using a linear filter and that original sensor samples are not modified by the interpolation process¹. In practice, since only odd values of u contribute to the above summation, we will restrict our attention to the case $h_u = 0$ for u odd. The prediction error is then defined as $e(x) = s_R(x) - s_P(x)$, with:

$$s_P(x) = \sum_u k_u s_R(x+u) \quad (3)$$

the predicted signal, and k_u the prediction kernel. Hence:

$$e(x) = \begin{cases} G(x) - \sum_u k_u s_R(x+u) & x \text{ even} \\ \sum_u h_u s_A(x+u) - \sum_u k_u s_R(x+u) & x \text{ odd} \end{cases} \quad (4)$$

By assuming to use the same kernel for the interpolation and the prediction (i.e. $h_u = k_u$), the prediction error in odd positions is identically zero, while in the even positions takes values different from zero. Hence, in such an ideal case, $\text{var}[e(x)]$ is expected to be zero in the positions corresponding to the demosaicked signal, and different from zero in the positions corresponding to the acquired signal.

In general, the exact interpolation coefficients may not be known, however we can assume that $k_u = 0$ for u odd. Moreover, we can also assume $\sum_u k_u = \sum_u h_u = 1$, which usually holds for common interpolation kernels. In this case, equation (4) above can be rewritten as

$$e(x) = \begin{cases} G(x) - \sum_u k_u \sum_v h_v G(x+u+v) & x \text{ even} \\ \sum_u (h_u - k_u) G(x+u) & x \text{ odd} \end{cases} \quad (5)$$

By assuming the acquired signal samples to be independent and identically distributed (i.i.d.) with mean μ_G and variance σ_G^2 , the mean of the prediction error can be evaluated as

$$E[e(x)] = \begin{cases} \mu_G - \mu_G \sum_u k_u \sum_v h_v = 0 & x \text{ even} \\ \mu_G (\sum_u h_u - \sum_u k_u) = 0 & x \text{ odd} \end{cases} \quad (6)$$

whereas the variance of the prediction error is

$$\begin{aligned} \text{Var}[e(x)] &= \text{Var} \left[\left(1 - \sum_u k_u h_{-u} \right) G(x) \right. \\ &\quad \left. + \sum_{t \neq 0} \left(\sum_u k_u h_{t-u} \right) G(x+t) \right] \\ &= \sigma_G^2 \left[\left(1 - \sum_u k_u h_{-u} \right)^2 + \sum_{t \neq 0} \left(\sum_u k_u h_{t-u} \right)^2 \right] \end{aligned} \quad (7)$$

¹The first assumption is often not verified in practice, however the second one usually holds since most practical demosaicing algorithms do not change the value of acquired pixels.




Fig. 1. (a) the Bayer's filter mosaic; (b) the quincunx lattice \mathcal{A} for the acquired green channels and the complementary quincunx lattice \mathcal{I} for the interpolated green channels.

for x even and

$$\text{Var}[e(x)] = \text{Var} \left[\sum_u (h_u - k_u) G(x+u) \right] = \sigma_G^2 \sum_u (h_u - k_u)^2 \quad (8)$$

for x odd. According to the above model, the prediction error has zero mean and variance proportional to the variance of the acquired signal. However, when the prediction kernel is close to the interpolation kernel, the variance of prediction error will be much higher at the positions of the acquired pixels than at the positions of interpolated pixels.

Leaving the ideal conditions, the acquired signal will be only locally i.i.d. and its variance only locally stationary: thus σ_G^2 has to be computed on small parts of the signal and consequently $\text{var}[e(x)]$ will assume different values depending on the specific signal content. Also, additive noise may be present on pixel values due to rounding and truncation effects. Nevertheless, we can still expect the variance of $e(x)$ to be higher at the positions of acquired pixels.

IV. PROPOSED ALGORITHM

In order to extend the previous analysis to the bidimensional case, without loss of generality we will consider as specific CFA the most frequently used Bayer's filter mosaic, a 2×2 array having red and green filters for one row and green and blue filters for the other (see Fig. 1(a)). Furthermore, we will consider only the green channel; since the green channel is up-sampled by a factor 2, for a generic square block we have the same number of samples (and the same estimation reliability) for both classes of pixels (either acquired or interpolated).

By focusing on the green channel, the even/odd positions (i.e. acquired/interpolated samples) of the one-dimensional case turn into the quincunx lattice \mathcal{A} for the acquired green values and the complementary quincunx lattice \mathcal{I} for the interpolated green values (see Fig. 1(b)). Similar to the one-dimensional case, we assume that in the presence of CFA interpolation the variance of the prediction error on lattice \mathcal{A} is higher than the variance of the prediction error on lattice \mathcal{I} , and in both cases it is content dependent. On the contrary, when no demosaicing has been applied, the variance of the prediction error assumes similar values on the two lattices. Hence, in order to detect the presence/absence of demosaicing artifacts, it is possible to evaluate the imbalance between the variance of the prediction error in the two different lattices.

A. Proposed feature

Let us suppose that $s(x, y)$, with $(x, y) \in \mathbb{Z}^2$, is an observed image. The prediction error can be obtained as:

$$e(x, y) = s(x, y) - \sum_{u,v \neq 0} k_{u,v} s(x+u, y+v) \quad (9)$$

where $k_{u,v}$ is a bidimensional prediction filter. In the ideal case, $k_{u,v} = h_{u,v} \forall (u, v)$ where $h_{u,v}$ is the interpolation kernel of the demosaicing algorithm. In general, we can assume that $k_{u,v} \neq h_{u,v}$, since the in-camera demosaicing algorithm is usually unknown.

Because of the local stationarity of the residue, the variance of the prediction error $e(x, y)$ is *locally* estimated pixel-by-pixel for each position (demosaicked or acquired) only from a neighborhood of interpolated (\mathcal{I}) or acquired (\mathcal{A}) pixels respectively. In this work, we assume to know the spatial pattern of the CFA (for example the Bayer CFA). This hypothesis is not a serious constraint, because it is reasonable to suppose either to know the CFA pattern or to estimate it by adopting a proper estimation algorithm [8].

By assuming that the local stationarity of prediction error is valid in a $(2K+1) \times (2K+1)$ window, it is possible to define the local weighted variance of the prediction error as:

$$\sigma_e^2(x, y) = \frac{1}{c} \left[\left(\sum_{i,j=-K}^K \alpha_{ij} e^2(x+i, y+j) \right) - (\mu_e)^2 \right] \quad (10)$$

where α_{ij} are suitable weights, $\mu_e = \sum_{i,j=-K}^K \alpha_{ij} e(x+i, y+j)$ is a local weighted mean of the prediction error and $c = 1 - \sum_{i,j=-K}^K \alpha_{ij}^2$ is a scale factor that makes the estimator unbiased, i.e., $E[\sigma_e^2(x, y)] = \text{var}[e(x, y)]$, for each pixel class. The weights α_{ij} are obtained as $\alpha_{ij} = \alpha'_{ij} / \sum_{i,j} \alpha'_{ij}$ where

$$\alpha'_{ij} = \begin{cases} W(i, j) & \text{if } e(x+i, x+j) \text{ belongs to} \\ & \text{the same class of } e(x, y) \\ 0 & \text{otherwise} \end{cases}$$

and $W(i, j)$ is a $(2K+1) \times (2K+1)$ Gaussian window with standard deviation $K/2$.

Given a $N \times N$ image, we analyze it by considering $B \times B$ non-overlapping blocks, where B is related to the period of Bayer's filter mosaic: the smallest period (and block dimension) is $(2, 2)$, but also multiples can be adopted. The generic block in position (k, l) is denoted as $\mathcal{B}_{k,l}$ with $k, l = 0, \dots, \frac{N}{B} - 1$. Each block is composed by disjoint sets of acquired and interpolated pixels, indicated as $\mathcal{B}_{A_{k,l}}$ and $\mathcal{B}_{I_{k,l}}$, respectively. We then define the feature \mathbf{L} :

$$\mathbf{L}(k, l) = \log \left[\frac{GM_A(k, l)}{GM_I(k, l)} \right] \quad (11)$$

where $GM_A(k, l)$ is the *geometric mean* of the variance of prediction errors at acquired pixel positions, defined as:

$$GM_A(k, l) = \left[\prod_{i,j \in \mathcal{B}_{A_{k,l}}} \sigma_e^2(i, j) \right]^{\frac{1}{|\mathcal{B}_{A_{k,l}}|}} \quad (12)$$

whereas $GM_I(k, l)$ is similarly defined for the interpolated pixels.

The proposed feature \mathbf{L} allows us to evaluate the imbalance between the local variance of prediction errors when an image is demosaicked: indeed, in this case the local variance of the prediction error of acquired pixels is higher than that of interpolated pixels and thus the expected value of $\mathbf{L}(k, l)$ is a nonzero positive amount. On the other hand, if an image is not demosaicked, this difference between the variance of prediction errors of acquired and interpolated pixels disappears, since the content can be assumed to present locally the same statistical properties, and the expected value of $\mathbf{L}(k, l)$ is zero. Our inference will be based on these two key observations.

Let us now suppose that a demosaicked image has been tampered by introducing a new content, and that in order to make this forgery more realistic, some processing (blurring, shearing, rotation, compression, etc.) has been likely applied to the added content, thus destroying the demosaicing traces on the forged region. The proposed feature $\mathbf{L}(k, l)$ will assume inconsistent values within the tampered image: in some regions (the untampered ones) it will be significantly greater than zero, while in other regions (the tampered ones) the feature will be close to zero. We can thus employ these inconsistencies to finely localize forgeries.

In some respects, the proposed feature is conceptually similar to the approach in [16], where the variance is approximated using the average of absolute values. However, an important difference is that the technique of [16] requires a Fourier analysis, thus limiting the resolution of the method when aiming at the fine-grained localization of CFA artifacts. Moreover, the proposed feature can be described using a very convenient statistical model, described in the following, which allows us to associate to each block a probability of being manipulated.

B. Feature modeling

By using a *Bayesian approach*, for each block $\mathcal{B}_{k,l}$ it is possible to derive the probability that CFA artifacts are present/absent conditioned on the observed values of $\mathbf{L}(k, l)$.

Let M_1 and M_2 be the hypotheses of presence and absence of CFA artifacts, respectively. In order to have a simple and tractable model, we assume that $\mathbf{L}(k, l)$ is Gaussian distributed under both hypotheses and for any possible size B of the blocks $\mathcal{B}_{k,l}$. For a fixed B , we can characterize our feature using the following *conditional probability density functions*:

$$Pr\{\mathbf{L}(k, l) | M_1\} = \mathcal{N}(\mu_1, \sigma_1^2) \quad (13)$$

with $\mu_1 > 0$, and

$$Pr\{\mathbf{L}(k, l) | M_2\} = \mathcal{N}(0, \sigma_2^2). \quad (14)$$

The above densities hold $\forall k, l = 0, \dots, \frac{N}{B} - 1$, i.e., we assume that the parameters of the two conditional pdfs do not change over the considered image, such that they can be globally estimated.

If a demosaicked image contains some tampered regions in which CFA artifacts have been destroyed (as it may occur in a common splicing operation), both hypotheses M_1 and M_2 are present, therefore $\mathbf{L}(k, l)$ can be modeled as a mixture of Gaussian distributions. The first component, with $\mu_1 > 0$, is




Fig. 2. The work flow of our algorithm.

due to regions in which CFA artifacts are present, whereas the second component, with $\mu_2 = 0$, is due to tampered regions in which CFA artifacts have been removed². In order to estimate simultaneously the parameters of the proposed Gaussian Mixture Model (GMM), we employ the *Expectation-Maximization (EM) algorithm* [23]. This is a standard iterative algorithm that estimates the mean and the variance of the component distributions by maximizing the expected value of a *complete log-likelihood function* with respect to the distribution parameters. In our case, the EM algorithm is used to estimate only μ_1 , σ_1 , and σ_2 , since we assume $\mu_2 = 0$.

C. Map generation

The final aim we point at is to achieve a map indicating for each $B \times B$ block $\mathcal{B}_{k,l}$ its probability to be original/tampered, based on its probability to contain or not CFA artifacts. Starting from equations (13) and (14) and assuming a-priori probabilities $Pr\{M_1\} = Pr\{M_2\} = 1/2$, we obtain the *posterior probability* of being an original block. By exploiting Bayes' Theorem and relying on the observed feature $\mathbf{L}(k, l)$ for each $\mathcal{B}_{k,l}$ block, we achieve:

$$Pr\{M_1|\mathbf{L}(k, l)\} = \frac{Pr\{\mathbf{L}(k, l)|M_1\}}{Pr\{\mathbf{L}(k, l)|M_1\} + Pr\{\mathbf{L}(k, l)|M_2\}} \quad (15)$$

which can be expressed as:

$$Pr\{M_1|\mathbf{L}(k, l)\} = \frac{1}{1 + \mathcal{L}(\mathbf{L}(k, l))} \quad (16)$$

where \mathcal{L} is the *likelihood ratio* of $\mathbf{L}(k, l)$ defined as:

$$\mathcal{L}(\mathbf{L}(k, l)) = \frac{Pr\{\mathbf{L}(k, l)|M_2\}}{Pr\{\mathbf{L}(k, l)|M_1\}}. \quad (17)$$

Let us note that equations (16) and (17) have the same statistical information. Applying equation (17) to each block of an image, we obtain a *likelihood map* (LM), where each pixel of the map is the likelihood ratio associated to a $B \times B$ block.

These maps are usually noisy because they associate a probability (or a likelihood ratio) value to a single realization of $\mathbf{L}(k, l)$, which is very noisy itself. In order to denoise these maps, we can cumulate feature values on larger blocks whose size is $C \times C$, where $C = n \cdot B$ with $n \in \mathbb{Z}^+$. Assuming blocks

to be conditionally independent given either M_1 or M_2 , the accumulated likelihood ratio is obtained as:

$$\mathcal{L}_{cum}(\mathbf{L}(k', l')) = \frac{\prod_{k,l} Pr\{\mathbf{L}(k, l)|M_2\}}{\prod_{k,l} Pr\{\mathbf{L}(k, l)|M_1\}}. \quad (18)$$

In order to further improve the localization performance, we note that in a realistic forged image the manipulated areas are usually connected regions, due to the image semantic content. These connected regions can be highlighted by applying to the map a simple low-pass spatial filter, like a mean filter or a median filter. For better numerical stability, such filters are applied to the logarithm of the likelihood map.

D. Overall system

In Fig. 2 we show the overall system that, given a suspected image, produces the corresponding forgery map: each pixel in the forgery map indicates for each $C \times C$ image block its probability to contain CFA artifacts, so that low values in the output map correspond to likely forged areas.

As a first step, the green channel is extracted from the image, and the prediction error is computed. Because in-camera processing algorithms are usually unknown, a fixed predictor is used: the choice of the adopted predictor will be discussed and validated in Section V. The weighted local variance is then estimated and the feature $\mathbf{L}(k, l)$ is obtained for each $B \times B$ block. The GMM parameters are globally estimated exploiting the EM algorithm and used for the generation of the forgery map. When $C = B$ the forgery map is generated using the likelihood ratios in (17), whereas for $C > B$ we use the cumulated likelihood map in (18). Optionally, the intermediate log-likelihood map can be filtered using either a mean filter or a median filter.

V. EXPERIMENTAL RESULTS

The results presented in this section have been obtained on a dataset consisting of 400 original color images, in TIFF uncompressed format, coming from 4 different cameras (100 images for each camera): Canon EOS 450D, Nikon D50, Nikon D90, Nikon D7000. All cameras are equipped with a Bayer CFA, thus respecting our requirement that authentic images come from a camera leaving demosaicing traces, but the in-camera demosaicing algorithms of such devices are unknown. Each image was cropped to 512×512 pixels, maintaining the original Bayer pattern, which is assumed to

²The above model may not be accurate in the case of copy-move forgeries exhibiting a nonaligned CFA pattern, since these areas will result in negative values of $\mathbf{L}(k, l)$. However, this is only a small subset of the possible forgeries and it does not appear reasonable to complicate the model to cope with this particular case.

be known³. We will refer to such a dataset as the *original dataset*.

A. Model Validation

The first step was to verify the assumption of Gaussian distribution on $\mathbf{L}(k, l)$, both in the presence and in the absence of CFA artifacts. To this end, starting from 100 images selected from the *original dataset*, we have created two datasets satisfying the M_1 (presence of CFA) and M_2 (absence of CFA) hypotheses. To create the dataset corresponding to M_1 , the original images have been sampled according to the Bayer CFA pattern and then re-interpolated using four possible demosaicing algorithms, namely bilinear, bicubic, gradient-based and median (see [12] for more details on such interpolation algorithms). This allowed us to know the interpolation kernel on the whole image, and then to exactly predict the interpolated values with the four different predictors (we refer to these cases as 'ideal'). To create the dataset corresponding to M_2 , each color channel of the original images has been upsampled by a factor two, blurred with a 7×7 median filter, and downsampled by a factor two, thus removing all CFA artifacts. Features are then computed using again the four predictors as before.

Moving towards realistic conditions, we also computed the value of $\mathbf{L}(k, l)$ under the M_1 hypothesis on the *original dataset* of 400 TIFF uncompressed images interpolated using their unknown in-camera demosaicing algorithms, and applying bilinear, bicubic, gradient-based and median predictors.

We verified the approximate Gaussian distribution of the features for all the classes described so far, i.e.: absence of CFA, presence of CFA with known interpolation kernel, and the four sets of cameras with unknown CFA demosaicing algorithms; for each of these six classes, the features have been computed with the four different interpolation algorithms (bilinear, bicubic, gradient-based, median) setting $B = 8$. The approximately Gaussian behavior of the features has been verified by fitting them with a generalized Gaussian distribution (GGD), given by

$$p(\mathbf{L}) = \frac{1}{Z} e^{-(|\mathbf{L}-\mu|/\eta)^\nu} \quad (19)$$

where μ is a location parameter (mean), η is a scale parameter, ν is a shape parameter, and Z is a normalization factor so that $p(\mathbf{L})$ integrates to one. The Gaussian distribution is a particular case of the GGD for $\nu = 2$. Since our conjecture is that the Gaussian assumption holds for a single image, but not necessarily over the whole dataset, the shape parameter has been independently estimated for each image using the Mallat's method [24]. In Table I we report the median value of the estimated shape parameters for the six classes and the four interpolation algorithms. The values indicate a reasonable fit of the proposed model. Interestingly, the model appears more fitting in the presence of CFA artifacts, and when the predictor is matched to the actual interpolation algorithm.

Furthermore, we plot the mean value of the features in order to verify how features in M_1 hypothesis can be discriminated

³The correct CFA configuration has been verified by inspecting the technical specifications of the raw image format.

TABLE I
MEDIAN VALUE OF THE GGD SHAPE PARAMETERS ESTIMATED FROM THE DISTRIBUTION OF THE FEATURE $\mathbf{L}(k, l)$ FOR EACH IMAGE, CONSIDERING DIFFERENT PREDICTORS ON DIFFERENT DATASETS.

	bilinear	bicubic	gradient-based	median
No CFA	1.589	1.558	1.672	1.812
Ideal	2.168	2.134	2.049	2.016
Canon EOS	2.001	1.908	1.897	1.962
Nikon D50	1.736	1.797	1.834	1.814
Nikon D7000	2.206	2.066	1.729	1.899
Nikon D90	1.998	1.924	1.667	1.927

by features in M_2 hypothesis, both in ideal and in realistic cases. In Fig. 3, we show the results for the ideal case in absence of CFA (first row) and presence of known CFA (second row). In Fig. 4, we show the 16 histograms of the mean values of $\mathbf{L}(k, l)$: along each row we have histograms referring to the same camera, from top to bottom, Canon EOS 450D, Nikon D50, Nikon D90, Nikon D7000. For both the Figures along each column we have histograms referring to the same predictor, from left to right, bilinear (red), bicubic (blue), gradient-based (green), median (violet).

Globally, the above results confirm that the proposed features has zero mean under the M_2 hypothesis and mean greater than zero under the M_1 hypothesis. The histograms also highlight that the four predictors have different behaviors. The median predictor does not seem well suited to detect CFA artifacts, since it produces values of $\mathbf{L}(k, l)$ closer to zero than the other predictors, irrespective of the camera.

B. Detection Performance Validation

In this section, the detection capability of the proposed forgery localization algorithm is investigated. Firstly, the behavior with respect to different predictors is analyzed. Then, in order to characterize the algorithm performance in different conditions, a particular predictor is chosen – the bilinear – and the results are evaluated considering different scenarios, different forgery sizes, and different choices of algorithm parameters.

1) *Experimental Methodology*: The considered scenarios correspond to nine different datasets derived from the *original dataset*: a first group of four datasets include uncompressed images obtained by applying bilinear, bicubic, gradient-based, and median demosaicing (as described in the previous section), representing the ideal case; a second group of five datasets include uncompressed images obtained using the demosaicing algorithm of the respective four cameras and JPEG compressed images obtained from the previous images using four different quality factors: 100%, 95%, 90% and 85%. The idea underlying this choice is to verify the performance on sets of images that completely satisfy the requirements of the proposed model as well as on more realistic images.

For each dataset, forgery has been simulated by applying to the central region of the image the procedure for removing CFA artifacts described in the previous section. As to the size of the forgery, we considered tampered regions of 128×128 , 64×64 , and 32×32 pixels. In the case of JPEG datasets, CFA removal has been simulated before JPEG compression.




Fig. 3. Distribution of the average value of $\mathbf{L}(k, l)$ on an image, feature evaluated on 8×8 blocks, in the absence of CFA artifacts (top row) and when the predictor is the same as the demosaicing algorithm (bottom row), using different predictors: from left to right, bilinear (red), bicubic (blue), gradient-based (green), median (violet).

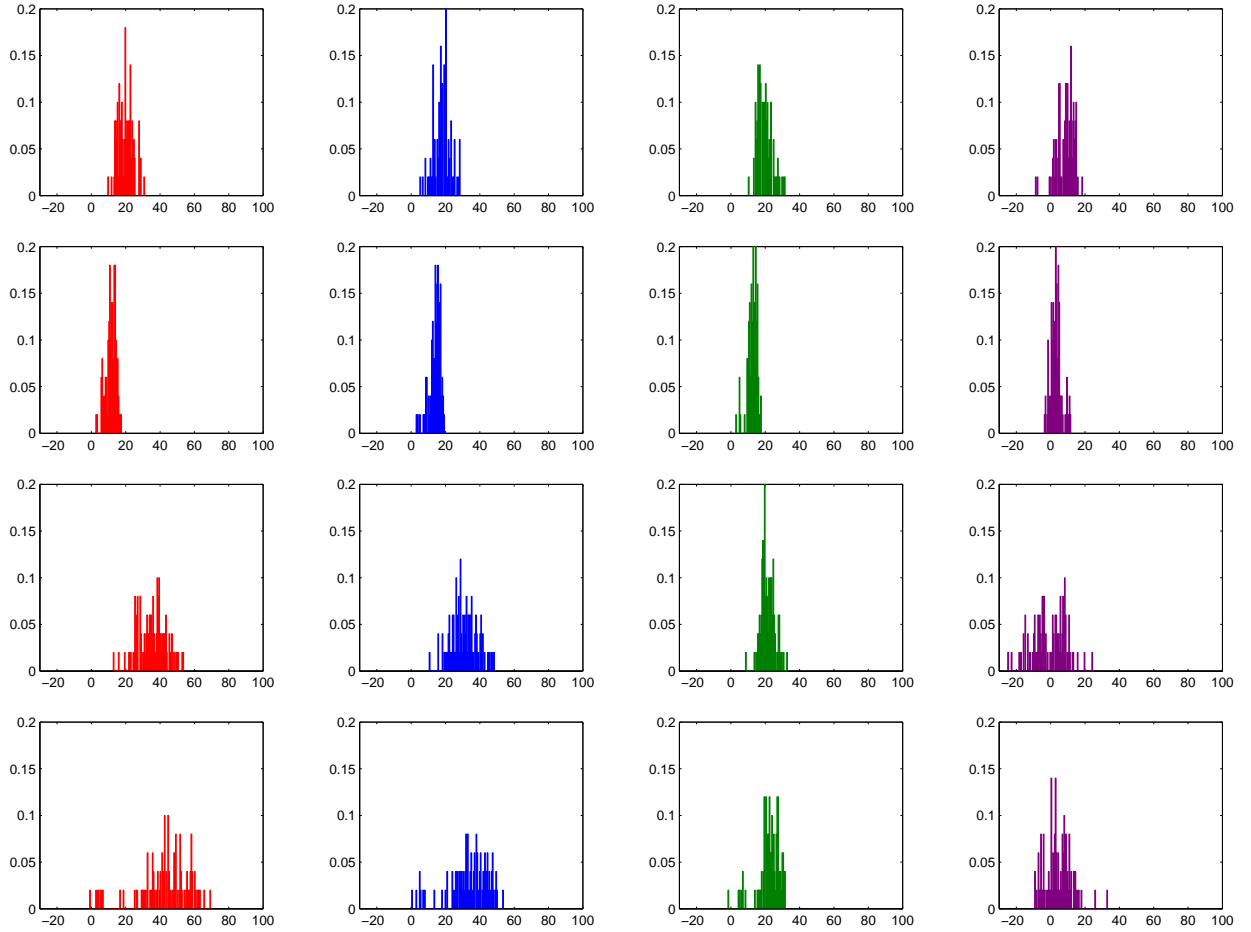


Fig. 4. Distribution of the average value of $\mathbf{L}(k, l)$ on an image, feature evaluated on 8×8 blocks, with unknown in-camera demosaicing algorithms and using different predictors: along each row we have histograms referring to the same camera, from top to bottom, Canon EOS 450D, Nikon D50, Nikon D7000, Nikon D90; along each column we have histograms referring to the same predictor, from left to right, bilinear (red), bicubic (blue), gradient-based (green), median (violet).

The analysis has been carried out under different *resolutions* and *filtering* of the *likelihood map*. Concerning the *resolution*, in order to permit a fine-grained localization of the tampered regions, we chose to compute the proposed metric \mathbf{L} starting from 2×2 blocks ($B = 2$), the smallest possible size to detect CFA artifacts. Different resolutions, equivalent to 4×4 blocks and 8×8 blocks, can be obtained in two ways: the first one is to define our features on larger blocks (e.g. $B = 4$ or $B = 8$). The second way is to compute the proposed metric on 2×2 or 4×4 blocks, and then to cumulate the posterior probabilities according to (18) on $C \times C$ blocks ($C = 8$). Concerning the *filtering* of the *likelihood map*, three cases were considered: no filtering at all, 5×5 weighted average filtering using a Gaussian window, and 5×5 bidimensional median filtering. In all cases, filtering is applied on *log likelihood maps* to avoid numerical problems.

As to the EM algorithm, we initialized μ_1 and σ_1^2 to the mean and variance of the observed features, $\sigma_2^2 = \sigma_1^2/10$, and $\alpha = 0.5$. Convergence was assumed if the increase of the likelihood function with respect to the previous iteration was less than 10^{-3} or after 500 iterations.

The performance of the proposed algorithm has been measured by the *true positive rate* (R_{TP}), measuring the fraction of tampered blocks correctly detected as forgery, and the *false positive rate* (R_{FP}), measuring the fraction of unchanged blocks erroneously detected as forgery. If we assume N_{R1} the amount of blocks in the untampered region R_1 , N_{R2} the amount of blocks in the forged region R_2 , N_{mR1} the amount of blocks detected as tampered in region R_1 and N_{mR2} the amount of blocks detected as tampered in region R_2 , we have:

$$R_{TP} = \frac{N_{mR2}}{N_{R2}}; \quad (20)$$

$$R_{FP} = \frac{N_{mR1}}{N_{R1}}. \quad (21)$$

The overall performance of the detector is evaluated by plotting its *receiver operating characteristic* (ROC) curve, obtained by thresholding the output maps (i.e. the *cumulated and filtered likelihood maps*) using a varying threshold value and recording the corresponding values of R_{TP} and R_{FP} . Finally, the *area under the ROC curve* (AUC) is used as a scalar parameter to describe detector capabilities: an AUC close to one indicates good detection performance, whereas an AUC close to 0.5 indicates that the detector has no better performance than choosing at random.

2) *Results*: In Fig. 5(a), we show the detection performance on the four ideal datasets, where for each datasets we use a predictor matched to the demosaicing algorithm, whereas in Fig. 5(b), we show the detection performance on the dataset using in-camera demosaicing when different predictors are applied. For each test a 128×128 tampered region has been considered. Detection results are averaged over the four different cameras. As to the resolution of the likelihood map, we have $B = C = 8$. The results show that when the predictor matches the demosaicing algorithm the performance is nearly optimal, irrespective of the used predictor, whereas in the presence of a realistic and unknown demosaicing




Fig. 5. ROC curves considering images from the *original dataset* with 128×128 tampered regions. Features are computed on 8×8 blocks: (a) ideal case: the 400 original images have been sampled according to the Bayer CFA pattern and then re-interpolated using the four chosen interpolation algorithms; results from all the 400 images are aggregated for each of the four predictors and the behaviour is shown separately; for the sake of readability, we show a zoom of the ROC curves for $R_{TP} > 0.8$ and $R_{FP} < 0.2$; AUC values are: bicubic 0.9975, bilinear 0.9845, gradient-based 0.9975, median 0.9954; (b) real case: the 400 original images coming from the 4 cameras with unknown demosaicing algorithms; results from all the 400 images are aggregated for each of the four predictors and the behaviour is shown separately.

algorithm the best average performance is obtained using the bilinear predictor. It is worth noting that in the latter case the performance of the median predictor is far worse than that of the other predictors, which is in accordance with the histograms in Fig. 4.

The following results show the detection performance, averaged over the four cameras, when using the bilinear predictor and different choices of algorithm parameters. In Fig. 6 we report the AUC values obtained using different *likelihood map* resolutions without filtering the likelihood map, under six different scenarios and considering different sizes of the tampered area. In all cases, the best performance is obtained when the exact interpolation kernel is known (in this case bilinear). Note also that the ability to localize forged regions sensibly decreases when the JPEG compression quality is below 95%. This is due to the low-pass behavior of JPEG compression, which drastically attenuates high frequency signals, such as the prediction error. With a quality factor 85%, our algorithm is unable to discriminate between the presence and the absence of CFA artifacts.

By comparing the different curves, we observe that defining our features on larger blocks makes our model more robust. These better performances are obtained at the expense of map resolution. However, in realistic conditions forgery sizes less than 8 pixels are unusual. It is also worth noting that computing the features on 2×2 or 4×4 blocks and cumulating the probabilities on 8×8 block yields slightly worse results than directly computing the features on 8×8 blocks. Lastly, the performance of the proposed detector appears similar for different forgery sizes, even though smaller tampered areas are more difficult to detect due to the reduced number of tampered blocks which decreases the reliability of the GMM estimation.

In Fig. 7, we compare the performance of the proposed detector using the most favorable combination of parameters, namely 8×8 resolution without cumulation, with the performance of the algorithms proposed by Dirik and Memon in [17] (DM) and by Gallagher and Chen in [16], namely the block-




Fig. 6. Effects of the *Likelihood Map* resolution on the AUC values. We consider TIFF images with bilinear interpolation (I) and TIFF images with in-camera demosaicing (II). These latter images are then compressed in JPEG format with quality at 100% (III), 95% (IV), 90% (V) and 85% (VI). Different forgery sizes are investigated: (a) 32×32 pixels; (b) 64×64 pixels; (c) 128×128 pixels.

wise version (GC-B) and the version based on local statistics (GC-L). For a fair comparison, the DM and GC-B algorithms have been applied on 8×8 blocks, whereas the features of GC-L algorithm have been computed using 7×7 local averaging and 16-point discrete Fourier transform. The proposed feature clearly outperforms the previous approaches, demonstrating far better localization capabilities. It is also evident that the performance of all CFA-based methods degrades similarly in the presence of JPEG compression when such methods are used to localize CFA artifacts at a fine-grained resolution.

We also investigated the use of filtering on the *likelihood map*. In Figure 8, the AUC values are shown in the absence or presence of either mean or median filtering, using 8×8 -features. The size of the tampered region is 128×128 pixels. We can see that filtering improves performances, except in the ideal case, where the effects of the loss of resolution on the edges of the tampered region is predominant, and that median filtering gives better results than mean filtering.




Fig. 7. Comparison between the proposed algorithm and the algorithms by Dirik and Memon (DM) [17] and by Gallagher and Chen (GC-B and GC-L) [16]. We consider TIFF images with bilinear interpolation (I) and TIFF images with in-camera demosaicing (II). These latter images are then compressed in JPEG format with quality at 100% (III), 95% (IV), 90% (V) and 85% (VI). The features are computed on 8×8 blocks. Tampered region is 128×128 pixels.




Fig. 8. Effects of *Likelihood Map* filtering on the AUC values. We consider TIFF images with bilinear interpolation (I) and TIFF images with in-camera demosaicing (II). These latter images are then compressed in JPEG format with quality at 100% (III), 95% (IV), 90% (V) and 85% (VI). The features are computed on 8×8 blocks. Tampered region is 128×128 pixels.

C. Examples

In this section, some examples of forgery localization are shown on realistically tampered images. In all the cases, the corresponding *forgery maps* have been obtained by computing features on 8×8 blocks ($C = B = 8$), using the bilinear predictor and applying median filtering on the *log likelihood map*.

In Fig. 9 a copy-move forgery on an image acquired with a Nikon D90 is shown. Both the original image, in Fig. 9(a), and the tampered copy, in Fig. 9(b), are saved in TIFF uncompressed format. The flower in the upper-left corner has been pasted disaligning the CFA pattern, whereas the flower in the upper right corner has been pasted maintaining the same CFA pattern. In Figs. 9(c)-(f) we show the forgery maps obtained with the proposed algorithm and the DM, GC-B, and GC-L algorithms, respectively. Even if the case of copy-move forgery does not perfectly fit the proposed model, since in the case of misaligned CFA artifacts the expected value of \mathbf{L} is less than zero, the proposed algorithm correctly localizes the flower in the upper-left corner, whereas it is not able to localize the flower in the upper-right corner. This is not surprising, since the proposed method gives higher likelihood values for positive values of the feature and reveals local inconsistencies of the CFA artifacts even when $\mathbf{L} < 0$. As to the other algorithms, only the GC-B is able to localize the




Fig. 9. Example of a copy-move forgery in an image with CFA artifacts. The resulting image is saved in TIFF format: (a) original image acquired by the Nikon D90 camera; (b) tampered image; forgery maps obtained with the proposed (c), DM (d), GC-B (e), and GC-L (f) algorithms. Bright areas indicate high probability of presence of CFA artifacts (unchanged blocks), whereas dark areas indicate low probability of presence of CFA artifacts (tampered blocks).

upper-left flower. Moreover, some false alarms are present in the case of saturated white regions, in which CFA artifacts are not detectable.

Very often, to make the forgery more convincing some image processing operations, like smoothing, filtering, stretching, rotating, etc., are applied. These operations, removing CFA artifacts from the tampered regions, make easier the forgery localization. In Fig. 10, we show an example where a tampering is done by splicing a geometrically transformed image onto an image taken by a Nikon D90 camera. In Figs. 10(c)-(n) we show forgery maps obtained with different algorithms, from top to bottom, the proposed algorithm, DM, GC-B, and GC-L algorithms, assuming that the tampered image was saved in JPEG format with quality, from left to right, 100%, 95%, and 90%. As can be seen, the forged region can be correctly detected in high quality images, but false alarms increase abruptly when the quality of JPEG compression decreases, because lossy compression tends to delete CFA artifacts. On this example, DM algorithm appears less effective than the other algorithms.

The inspection of the forgery maps in Figs. 9-10 suggests that the proposed method is less effective in the presence of either almost flat areas or sharp edges. In the first case, the prediction error is almost zero irrespective of the presence of CFA artifacts, so that this appears as an intrinsic limit of the method. In the second case, this can be ascribed to the signal adaptive and possibly non-linear behavior of realistic

in-camera demosaicing algorithms. At least in theory, such effects could be eliminated by using some prior knowledge regarding in-camera CFA interpolation, which should yield results very close to the ideal behavior shown in Fig. 5. An alternative approach could be that of reverse engineering the CFA interpolation algorithm, for example using methods such as in [8] to take into account a signal adaptive behavior. However, in the presence of heavily manipulated images this approach is likely to produce a biased estimate and must be handled with care.

VI. CONCLUSIONS

In this work, a forensic algorithm to localize forged regions in a digital image without any a-priori knowledge about the location of the possibly tampered areas has been presented. Considering the CFA demosaicing artifacts as a digital fingerprint, we proposed a new feature measuring the presence of demosaicing artifacts even at the smallest 2×2 block level; by interpreting the local absence of CFA artifacts as an evidence of tampering, the proposed scheme provides as output a forgery map indicating the probability of each block to be trustworthy.

The validity of the proposed system has been demonstrated by computing the ROC curve of a forgery detector based on thresholding the probability map, considering different scenarios and different algorithm parameters, and comparing the performance with the approaches in [17] and [16]. The




Fig. 10. Example of a forgery in which a processed content (statue) is pasted on an image with CFA artifacts: (a) original image; (b) tampered image; (c)-(n) *forgery maps* obtained after saving in JPEG format with quality, from left to right, 100%, 95% and 90%: (c)-(e) proposed algorithm; (f)-(h) DM algorithm; (i)-(k) GC-B algorithm; (l)-(n) GC-L algorithm.

results show that by a proper parameter configuration CFA artifacts are usually reliably localized even at 8×8 block resolution. Results are also confirmed by tests carried out on realistic forgeries.

The fine-grained localization of tampered regions using CFA artifacts is the main contribution of this work, since in previous approaches either the area to be investigated has to be manually selected, or automatic block processing obtains poor detection performance when forced to reveal CFA artifacts at a fine-grained scale. The results show that the proposed

algorithm can be a valid tool for detecting and localizing forgeries in images acquired by a digital camera. However, it should be remarked that the detection performance is strongly affected by JPEG compression, limiting the applicability to scenarios in which the image under test is either uncompressed or compressed with high quality factors. Moreover, the present method may not be directly applicable to cameras using a super CCD [25].

Test on realistically tampered images demonstrate that, due to the presence of uniform or very sharp regions, automatic

detection may give a remarkable false positive rate. Therefore, in order to limit the incidence of false positives human interpretation of the *forgery maps* is still required. Future work will be then devoted to the study of segmentation algorithms that, by taking into account the local content characteristics, allow to produce a final map with reduced false positives.

REFERENCES

- [1] Z. Lin, J. He, X. Tang, and C.-K. Tang, "Fast, automatic and fine-grained tampered JPEG image detection via DCT coefficient analysis," *Pattern Recognition 2009 - Elsevier*, pp. 2492–2501, 2009.
- [2] T. Bianchi, A. D. Rosa, and A. Piva, "Improved DCT coefficient analysis for forgery localization in JPEG images," in *Proc. of ICASSP 2011*, Prague, Czech Republic, May 2011, pp. 2444–2447.
- [3] T. Bianchi and A. Piva, "Analysis of non-aligned double JPEG artifacts for the localization of image forgeries," in *Proc. of WIFS 2011*, Foz do Iguaçu, Brazil, Nov.-Dec. 2011.
- [4] W. Wang, J. Dong, and T. Tan, "Exploring DCT coefficient quantization effect for image tampering localization," in *Proc. of WIFS 2011*, Foz do Iguaçu, Brazil, Nov.-Dec. 2011.
- [5] T. Bianchi and A. Piva, "Image forgery localization via block-grained analysis of JPEG artifacts," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 3, pp. 1003–1017, June 2012.
- [6] H. Farid, "Image forgery detection – a survey," *IEEE Signal Processing Mag.*, vol. 2, no. 26, pp. 16–25, Mar. 2009.
- [7] J. A. Redi, W. Taktak, and J.-L. Dugelay, "Digital image forensics: a booklet for beginners," *Multimedia Tools and Applications*, vol. 51, no. 1, pp. 133–162, 2011.
- [8] A. Swaminathan, M. Wu, and K. R. Liu, "Nonintrusive component forensics of visual sensors using output images," *IEEE Trans. Inf. Forensics Security*, vol. 2, no. 1, pp. 91–106, Mar. 2007.
- [9] A. Swaminathan, M. Wu, and K. J. R. Liu, "Digital image forensics via intrinsic fingerprints," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 1, pp. 101–117, 2008.
- [10] H. Cao and A. Kot, "Accurate detection of demosaicing regularity for digital image forensics," *IEEE Transactions on Information Forensics and Security*, vol. 4, no. 4, pp. 899–910, 2009.
- [11] S. Bayram, H. T. Sencar, and N. Memon, "Classification of digital camera-models based on demosaicing artifacts," *Digital Investigation*, vol. 5, pp. 46–59, 2008.
- [12] A. C. Popescu and H. Farid, "Exposing digital forgeries in color filter array interpolated images," *IEEE Trans. Signal Proc.*, pp. 3948 – 3959, Oct. 2005.
- [13] A. C. Gallagher, "Detection of linear and cubic interpolation in JPEG compressed images," *Computer and Robot Vision, Canadian Conference*, vol. 0, pp. 65–72, 2005.
- [14] N. Fan, C. Jin, and Y. Huang, "A pixel-based digital photo authentication framework via demosaicking inter-pixel correlation," in *11th ACM Multimedia and Security Workshop (MM&Sec '09)*, 2009, pp. 125–129.
- [15] A. C. Popescu and H. Farid, "Exposing digital forgeries by detecting traces of resampling," *IEEE Trans. Signal Proc.*, pp. 758–767, Feb. 2005.
- [16] A. C. Gallagher and T. Chen, "Image authentication by detecting traces of demosaicing," in *IEEE Computer Vision and Pattern Recognition Workshops (CVPRW 2008)*, 2008, pp. 1–8.
- [17] A. E. Dirik and N. Memon, "Image tamper detection based on demosaicing artifacts," in *16th IEEE Int. Conf. on Image Processing (ICIP '09)*, 2009, pp. 1497–1500.
- [18] M. Kirchner, "Fast and reliable resampling detection by spectral analysis of fixed linear prediction residue," in *10th ACM Multimedia and Security Workshop (MM&Sec '08)*, 2008, pp. 11–20.
- [19] M. Kirchner and T. Gloe, "On resampling detection in re-compressed images," in *First IEEE International Workshop on Information Forensics and Security*, 2009, December 2009, pp. 21–25.
- [20] B. Mahdian and S. Saic, "Blind authentication using periodic properties of interpolation," *IEEE Transactions on Information Forensics and Security*, vol. 3, no. 3, pp. 529–538, 2008.
- [21] —, "A cyclostationarity analysis applied to image forensics," in *Proceedings of the 2009 IEEE Workshop on Applications of Computer Vision*, Snowbird, UT, USA, 2009, pp. 389–399.
- [22] D. Vazquez-Padin, C. Mosquera, and F. Perez-Gonzalez, "Two-dimensional statistical test for the presence of almost cyclostationarity on images," *17th IEEE Int. Conf. on Image Processing (ICIP 2010)*, pp. 1745 – 1748, September 2010.
- [23] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society: Series B* 39, pp. 1–38, 1977.
- [24] S. Mallat, "A theory for multiresolution signal decomposition: the wavelet representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674 –693, July 1989.
- [25] http://www.fujifilm.com/products/digital_cameras/topics/2008/0922_01.html.



Pasquale Ferrara was born in Figline Valdarno, Italy, 1985. He received the Laurea degree (cum laude) in Telecommunication Engineering in April 2011 from the University of Florence.

Since June 2011, he is a Research Assistant at National Institute of Optics, involved in "Multispectral Imaging Diagnostics of Skin Tumors" (MIDST) project. His main research interests are in the field of image processing for medical and forensic applications, and automatic 3-D texture mapping.



Tiziano Bianchi (S'03-M'05) was born in Prato, Italy, in 1976. He received the M.Sc. degree (Laurea) in electronic engineering and the Ph.D. degree in information and telecommunication engineering from the University of Florence, Italy, in 2001 and 2005, respectively.

Since March 2005, he is with the Department of Electronics and Telecommunications, University of Florence as a Research Assistant. His research interests have involved processing of SAR images, signal processing in communications, multicarrier modulation techniques, and ultra-wideband systems. Current research topics include multimedia security technologies and signal processing in the encrypted domain.



Alessia De Rosa received the electronic engineering degree and the Ph.D. degree in informatics and telecommunications from the University of Florence (Italy) in 1998 and 2002 respectively.

Since 2002, she has worked at the University of Florence as a Research Assistant with the Department of Electronics and Telecommunications. Her main research interests are in the field of image processing and protection, including digital watermarking, human perception models for watermarking and quality assessment, image processing for Cultural Heritage applications and image forensics. She holds an Italian patent regarding digital watermarking.



Alessandro Piva (M'04-SM'10) received the Ph.D. degree in computer science and telecommunications engineering from the University of Florence in 1999.

From 2002 to 2004, he was a Research Scientist at the National Inter-University Consortium for Telecommunications. He is at present Assistant Professor at the University of Florence, Firenze, Italy. His current research interests are the technologies for multimedia content security, and image processing techniques for the Cultural Heritage field. He is coauthor of more than 100 papers published in international journals and conference proceedings.

Dr. Piva holds three Italian patents and an international one regarding watermarking. He serves as Associate Editor of the *IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY*, of the *EURASIP Journal on Information Security*, and of the *LNCS Transactions on Data Hiding and Multimedia Security*.