

# Secure Detection of Image Manipulation by Means of Random Feature Selection

Zhipeng Chen<sup>1</sup>, Benedetta Tondi<sup>2</sup>, *Member, IEEE*, Xiaolong Li<sup>1</sup>, Rongrong Ni,  
Yao Zhao<sup>1</sup>, *Senior Member, IEEE*, and Mauro Barni<sup>2</sup>, *Fellow, IEEE*

**Abstract**—We address the problem of data-driven image manipulation detection in the presence of an attacker with limited knowledge about the detector. Specifically, we assume that the attacker knows the architecture of the detector, the training data, and the class of features  $\mathcal{V}$  the detector can rely on. In order to get an advantage in his race of arms with the attacker, the analyst designs the detector by relying on a subset of features chosen at random in  $\mathcal{V}$ . Given its ignorance about the exact feature set, the adversary attacks a version of the detector based on the entire feature set. In this way, the effectiveness of the attack diminishes since there is no guarantee that attacking a detector working in the full feature space will result in a successful attack against the reduced-feature detector. We theoretically prove that, thanks to random feature selection, the security of the detector significantly increases at the expense of a negligible loss of performance in the absence of attacks. We also provide an experimental validation of the proposed procedure by focusing on the detection of two specific kinds of image manipulations, namely adaptive histogram equalization and median filtering. The experiments confirm the gain in security at the expense of a negligible loss of performance in the absence of attacks.

**Index Terms**—Adversarial signal processing, adversarial machine learning, image manipulation detection, feature selection, multimedia forensics and counter-forensics, secure classification, randomization-based adversarial detection.

Manuscript received January 25, 2018; revised October 5, 2018 and January 27, 2019; accepted February 9, 2019. Date of publication February 26, 2019; date of current version June 5, 2019. This work was supported in part by the National Key Research and Development of China under Grant 2016YFB0800404, in part by the National Science Foundation of China under Grant 61572052, Grant 61672090, and Grant U1736213, in part by the Fundamental Research Funds for the Central Universities under Grant 2018JBZ001, in part by the Key Research and Development of Hebei Province of China under Grant 17210332, and in part by Defense Advanced Research Projects Agency (DARPA) and Air Force Research Laboratory (AFRL) under Grant FA8750-16-2-0173. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Patrizio Campisi. (*Corresponding author: Yao Zhao.*)

Z. Chen is with the Institute of Information Science, Beijing Jiaotong University, Beijing 100044, China, also with the Beijing Key Laboratory of Advanced Information Science and Network Technology, Beijing Jiaotong University, Beijing 100044, China, and also with the Department of Computer Science, Tangshan Normal University, Tangshan 063000, China (e-mail: chzhpeng@hotmail.com).

B. Tondi and M. Barni are with the Department of Information Engineering and Mathematics, University of Siena, 53100 Siena, Italy (e-mail: benedettatondi@gmail.com; barni@di.unisi.it).

X. Li, R. Ni, and Y. Zhao are with the Institute of Information Science, Beijing Jiaotong University, Beijing 100044, China, and also with the Beijing Key Laboratory of Advanced Information Science and Network Technology, Beijing Jiaotong University, Beijing 100044, China (e-mail: lixl@bjtu.edu.cn; rmi@bjtu.edu.cn; yzhao@bjtu.edu.cn).

Digital Object Identifier 10.1109/TIFS.2019.2901826

## I. INTRODUCTION

DEVELOPING secure image forensic tools, capable of granting good performance even in the presence of an adversary aiming at impeding the forensic analysis, turns out to be a difficult task, given the weakness of the traces the forensic analysis relies on [1]. As a matter of fact, a number of Counter-Forensics (CF) tools have been developed, whose application hinders a correct image forensic analysis [2]. Early CF techniques were rather simple, as they consisted in the application of some basic processing operators like noise dithering, recompression, resampling or filtering [3]–[5]. Though often successful, the application of general post-processing operators, sometimes referred to as laundering, does not guarantee that the forensic traces are completely erased and hence does not necessarily result in the failure of the forensic analysis.

When the attacker has enough information about the forensic algorithm, much more effective CF techniques can be devised. By following the taxonomy introduced in [6], we say that we are in a Perfect Knowledge (PK) scenario, when the attacker has complete information about the forensic algorithm used by the analyst. In the PK case, very powerful CF techniques can be developed allowing the attacker to prevent a correct analysis by introducing a limited distortion into the attacked image. Generally speaking, the attacker needs *only* to solve an optimisation problem looking for the image which is in *some sense* closest to the image under attack and for which the output of the forensic analysis is the wrong one. Even if such an optimisation problem may not be always easy to solve, the exact knowledge of the decision function allows the application of powerful techniques either in closed form [7]–[9], or by relying on gradient-descent mechanisms [10], [11].

In many cases, the attacker has only a Limited Knowledge about the forensic algorithm [6]. Let us consider, for example, the case of a machine-learning-based detector looking for the traces left within an image by a particular processing algorithm. The attacker may know only the type of detector used by the analyst, e.g. a Support Vector Machine (SVM) or a Neural Network, and the feature space wherein the analysis is carried out, but he may not have access to the training data. In this case, the attacker can build a surrogate version of the detector by using its own training data, and carry out the attack on the surrogate detector, hoping that the attack will also work on the detector used by the analyst [6], [10], [11]. In other cases, the attacker may know only the feature space

used by the detector. In such a situation, he may resort to a so-called universal CF strategy capable of defeating any detector working in the given feature space [12]. In most cases, the attack works by modifying the attacked-image so that its representation in the feature space is as close as possible to the representation of an image chosen in a dataset of images belonging to the desired class (e.g. non-compressed or pristine images) [10]. In [13], for instance, the attack works by bringing the histogram of the attacked image as close as possible to that of an image belonging to a reference dataset of pristine images, by solving an optimal transport problem. In [14], a similar strategy is applied in the DCT domain to attack any double JPEG detector relying on the first order statistics of block DCT coefficients.

Several anti-CF techniques have also been developed. The most common approach consists in looking for the traces left by the CF tools, and develop new forensic algorithms explicitly thought to expose images subjected to specific CF techniques. The search for CF traces can be carried out by relying on new features explicitly designed for this target as in [15]–[18], or, from a more general perspective, by using the same features of the original forensic technique and design an adversary-aware version of the classifier, as in [19] and [20]. In the latter case, it is recommendable to adopt a large feature space allowing enough flexibility to distinguish original and tampered images as well as images processed with the CF operator. If we assume that the attacker knows that the traces left by the CF tools may themselves be subjected to a forensics analysis, we fall in a situation wherein CF and anti-CF techniques are iteratively developed in a never-ending loop, whose final outcome can hardly be foreseen [21]. Some attempts to cast the above race of arms between the forensic analyst and the attacker by resorting to game theory have been made in [22] and [23]. In some cases, it is also possible to predict who between the attacker and the analyst is going to win the game according to the distortion that the attacker may introduce to impede the forensic analysis [24]. Yet in other works, the structure of the detector is designed in such a way to make CF harder. In [25], the output of multiple classifiers is exploited to design an ensemble classifier exhibiting improved resilience against adversarial attempts to induce a detection error. In [26], the robustness of a two-class classifier and the inherent superior security of one-class classifiers are exploited to design a *one and a half class* detector, that is proven to provide an extra degree of robustness under adversarial conditions. Other approaches to improve the security of machine-learning classifiers are described in [27]–[29], for applications outside the realm of image forensics. Despite all the above attempts, however, when the attacker knows the feature space used by the analyst, very powerful CF strategies can be designed whose effectiveness is only partially mitigated by the adoption of anti-CF countermeasures.

In order to restore the possibility of a sound forensic analysis in an adversarial setting, and give the analyst an advantage in his race of arms with the attacker, in this work, we propose to randomise the selection of the feature space wherein the analysis is carried out. To be specific, let us assume that to achieve his goal - hereafter deciding between

two hypotheses  $H_0$  and  $H_1$  about the processing history of the inspected image - the analyst may rely on a large set  $\mathcal{V}$  of, possibly dependent, features. The number of features used for the analysis may be in the order of several hundreds or even thousands; for instance, they may correspond to the SPAM features described in [30] or the rich feature set introduced in [31]. In most cases, the use of all the features in  $\mathcal{V}$  is not necessary and good results can be achieved even by using a small subset of  $\mathcal{V}$ . Our proposal to *secure* the forensic analysis is to randomise it by choosing a random subset of  $\mathcal{V}$  - call it  $\mathcal{V}_r$  - and let the analysis rely on  $\mathcal{V}_r$  only; in a certain sense, the *randomisation* of the feature space can be regarded as a secret key used by the analyst to improve the security of the analysis. Given its ignorance about the exact feature set used by the analyst, a possibility for the attacker is to attack the entire feature set  $\mathcal{V}$ . As we will show throughout the paper, with both theoretical and experimental results, not only attacking the entire set  $\mathcal{V}$  increases the complexity of the attack, but it also diminishes its effectiveness, since there is no guarantee that attacking a detector working in the full feature space will also result in a successful attack against a detector working in the reduced set  $\mathcal{V}_r$ .

The rest of this paper is organised as follows. In Section II, we revise prior works using randomisation to improve the security of image forensic techniques and more in general that of any detector or classifier. In Section III, we give a rigorous formulation of image manipulation detection via random feature selection, and analyse the theoretical performance of the random detector under simplified, yet reasonable, assumptions. In Section IV, we exemplify the general strategy introduced in Section III by developing an SVM detector based on randomised feature selection within a restricted subset of SPAM features, designed to detect two different kinds of image manipulations, namely adaptive histogram equalization and median filtering. In Section V, we analyse the security of the detectors described in Section IV against targeted attacks carried out in the feature and the pixel domains. As it will be evident from the experimental analysis, random feature selection considerably increases the strength required for a successful attack at the expense of a negligible performance loss in the absence of attacks. Finally, in Section VI, we draw our conclusions and highlight some directions for future research.

## II. RELATED WORKS

The use of randomisation to improve the security of a detector or a classifier is not an absolute novelty since it has already been proposed in several security-oriented applications.

Early attempts to use randomisation as a countermeasure against attacks were focusing on probing or oracle attacks, i.e. attacks that repeatedly query the detector in order to get information about it and then use such an information to build an input signal that evades the detection<sup>1</sup> (see for instance [32] for an example related to one-bit watermarking and [28] for the use of randomisation in the context of machine learning). In all these works, the outcome of the detector is

<sup>1</sup>Probing attacks are usually carried out when no any a-priori knowledge about the classifier is available to the attacker.

randomised by letting the output be chosen at random for points in the proximity of the detection boundary. In general, boundary randomisation only increases the effort necessary to the attacker to enter (or exit) the detection region; in addition, it also causes a loss of performance in the absence of attacks, that is, the robustness of the system decreases.

Other strategies exploiting randomness to prevent the attacker from gathering information about the detector have been adopted for spam filtering, intrusion detection [25] and multimedia authentication [33]. A rather common approach consists in the use of randomisation in conjunction with multiple classifiers. In [34], for instance, randomness pertains to the selection of the training samples of the individual classifiers, while in [35] is associated to the selection of the features used by the classifiers (hereafter referred to as random subspace selection), each of which is trained on the entire training set. Another randomisation strategy, used in conjunction with multiple classifiers, has been proposed in [36] and experimentally evaluated for spam filtering applications. Specifically, an additional source of randomness is introduced in the choice of the weights of the filtering modules of the individual classifiers. Random subspace selection has also been adopted in steganalysis [37], though with a different goal, that is to reduce the problems encountered when working with extremely high-dimensional feature spaces.

The use of randomization for security purposes is also common in multimedia hashing for authentication and content identification [33], [38]. In these works, *random projections* are often employed to elude attacks: specifically, a secret key is used to generate a random matrix which is then employed to generate the hash bits of the content, by first projecting the to-be-authenticated signal on the directions identified by the rows of the matrix and then comparing the absolute value of the projections against a threshold. Despite the apparent similarity, the use of random projections for multimedia hashing differs substantially from the technique proposed in this paper. In multimedia hashing applications, the random projection is applied directly in the pixel or in a transformed domain, and is possibly followed by the use of channel coding to improve robustness against noise. The kind of traces we are looking for in multimedia forensics applications, however, are so weak that a completely random choice of the feature space would not work. For this reason, in our system, randomisation is applied within a set of features explicitly designed to work in a specific multimedia forensics scenario. As a matter of fact, the system proposed in this paper can be seen as the application of the random projection method directly in the feature space, with the projection matrix designed in such way to contain exactly one non-zero in each row, with the additional computational advantage that only the selected features need to be calculated, while a full projection matrix would require the computation of the entire feature set.

The use of feature selection for security purposes has also been proposed in [39]. Although the idea of resorting to a reduced feature set is similar to our proposal, the set up considered in [39] differs considerably from the one adopted in this paper. In [39], in fact, the authors search for the best reduced feature set against an attacker with perfect knowledge

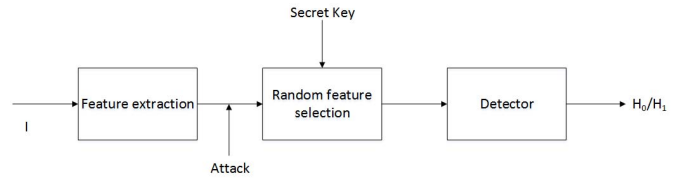


Fig. 1. Security model for the RFS detector.

about the detector, i.e. an attacker who knows the choice of the feature subset made by the defender. This is different from the scenario considered here, where feature randomization works as a kind of secret key.

### III. SECURE DETECTION BY RANDOM FEATURE SELECTION

In this section, we first describe the security assumptions behind our work, then we give a rigorous definition of binary detection based on Random Feature Selection (RFS) and provide a theoretical analysis to evaluate the security vs robustness trade-off under a simple statistical model. Though derived under simplified assumptions, the theoretical analysis is an insightful one since it provides useful insights on the impact that the statistics of the host features has on the security of the randomised detector. In addition, it permits to analyse the dependence of classification accuracy on the number of selected features both in the presence and in the absence of attacks. Even if the paper focuses on RFS, the theoretical framework is a general one and can also be used to analyse other kinds of (linear) feature randomisation.

#### A. Security Model

The security model adopted in this paper is depicted in Fig. 1. As shown in the picture, we assume that the attack is carried out directly in the feature space. This is not a problem when an invertible relationship exists between the pixel and the feature domain, as it is the case, for instance, of detectors based on block DCT coefficients or their histogram [12], [14]. In other cases, however, mapping back the attack into the pixel domain may not be easy, or even possible, since there is no guarantee that the attacked feature vector is a feasible one, i.e., that an image exists whose features are equal to those resulting from the attack. Nevertheless, analysing the performance of randomized feature detection in this scenario, which is undoubtedly more favourable for the attacker, provides interesting insights about the security improvement achieved by RFS. We also assume that the attacker is interested in inducing a missed detection error, while minimising the distortion introduced in the image as a consequence of the attack.

With regard to the knowledge available to the attacker, we assume that he knows everything with the exception of the subset of features used by the detector. This includes knowledge of the detector architecture, the full feature set  $\mathcal{V}$ , and the training set (in case of a detector based on machine learning). We also assume that it is not possible for the attacker

to query the detector to observe the result of the detection on one or more specific inputs. This means that the attacker can not run oracle or mimicry attacks on the reduced feature detector to infer information about the subset of features used by the detector [40], [41].

In order to assess the security of the RFS approach when faced with attacks carried out in the pixel domain, in section V-C we show the results of some experiments where an RFS detector trained to detect adaptive histogram equalization and median filtering is subject to an attack working in the pixel domain (see section IV-C for a description of the attack).

### B. Problem Formulation

Let  $\mathbf{v}$  be an  $n$ -long column vector with the image features the detector relies on. The detector aims at distinguishing between two hypotheses:  $H_0$  - “the image is manipulated”, and  $H_1$  - “the image is original”. We assume that the probability density function of  $\mathbf{v}$  under the two hypotheses is as follows:

$$\begin{aligned} H_0 : \mathbf{v} &\simeq \mathcal{N}(\mathbf{u}, \Sigma_0) \\ H_1 : \mathbf{v} &\simeq \mathcal{N}(-\mathbf{u}, \Sigma_1), \end{aligned} \quad (1)$$

where  $\mathcal{N}(\mathbf{u}, \Sigma_0)$  (res.,  $\mathcal{N}(-\mathbf{u}, \Sigma_1)$ ), is a multivariate Gaussian distribution with mean  $\mathbf{u}$  (res.,  $-\mathbf{u}$ ), and covariance matrix  $\Sigma_0$  (res.,  $\Sigma_1$ ). Note that assuming that the mean vectors under the two hypotheses are one the opposite of the other does not cause any loss of generality. In fact, if this is not the case, we can always apply a translation of the feature vector, for which such an assumption holds.

The derivation of the optimal detector, under the assumption that the a-priori probabilities of  $H_0$  and  $H_1$  are equal, passes through the computation of the log-likelihood probabilities of observing  $\mathbf{v}$  under  $H_0$  and  $H_1$ . In the following, we will carry out our analysis by assuming that  $\Sigma_0 = \Sigma_1 = \Sigma$ . In this case, the optimum detector is a correlation-based detector [42], according to which the detector decides for  $H_0$  if:

$$\rho = (\Sigma^{-1}\mathbf{u})^T \mathbf{v} > 0. \quad (2)$$

In the rest of the section, we will refer to the detector defined by equation (2) as the optimum *full-feature* detector.

In order to improve the security of the system, the detector randomises the decision strategy by relying on a reduced feature vector  $\mathbf{v}_r = S\mathbf{v}$  where  $S$  is a random  $k \times n$ -dimensional matrix. Several different randomisation strategies can be adopted according to the form of  $S$ . For the RFS detector proposed in this paper,  $S$  is a matrix whose entries are all zeros except for a single element of each row which is equal to one. In addition, all nonzero entries are located in different columns. This corresponds to form  $\mathbf{v}_r$  by selecting at random  $k$  elements of  $\mathbf{v}$ . Another possibility consists in generating all the elements of  $S$  independently according to a given distribution (e.g. Gaussian) and normalising the entries so that the Euclidean norm of each row is equal to one. In this way, the randomised detector relies on the projections of the vector  $\mathbf{v}$  on  $k$  random directions (Random Projection - RP).

### C. Theoretical Analysis

In this section, we analyse the trade-off between security and robustness by evaluating the performance of the randomised detector with and without attacks. As we will see, by lowering  $k$ , the security of the detector increases at the price of a loss of performance in the absence of attacks, with a better trade-off reached by the RFS detector.

1) *Performance in the Absence of Attacks (Robustness)*: As highlighted in the previous section, the sufficient statistics for the optimum full-feature detector is given by  $\rho = (\Sigma^{-1}\mathbf{u})^T \mathbf{v}$ . The performance of the full-feature detector depend on the statistics of  $\rho$ . Due to the normality of  $\mathbf{v}$ ,  $\rho$  is also normally distributed with mean and variance under  $H_0$  given by:

$$\begin{aligned} E[\rho|H_0] &= (\Sigma^{-1}\mathbf{u})^T \mathbf{u} = \mathbf{u}^T \Sigma^{-1}\mathbf{u} \\ \text{var}[\rho|H_0] &= (\Sigma^{-1}\mathbf{u})^T \Sigma \Sigma^{-1}\mathbf{u} = \mathbf{u}^T \Sigma^{-1}\mathbf{u}. \end{aligned} \quad (3)$$

Similar values are obtained under  $H_1$ , by replacing  $\mathbf{u}$  with  $-\mathbf{u}$ . The error probability of the detector is related to the  $z$ -value of the normal distribution, which is equal to:

$$z = \frac{\mathbf{u}^T \Sigma^{-1}\mathbf{u}}{(\mathbf{u}^T \Sigma^{-1}\mathbf{u})^{1/2}} = (\mathbf{u}^T \Sigma^{-1}\mathbf{u})^{1/2}. \quad (4)$$

Note that  $z$  is always positive since  $\Sigma$  is a positive-definite matrix and that higher values of  $z$  correspond to a lower error probability. Due to the symmetry of the problem, the two error probabilities, i.e., the probability of deciding in favour of  $H_0$  when  $H_1$  holds and the reverse probability of deciding for  $H_1$  when  $H_0$  holds, have the same value; hence, in the following, we will generally refer to the error probability as  $P_e$ .

In the case of randomised detection, the feature vector is:  $\mathbf{v}_r = S \cdot \mathbf{v}$ . For a given  $S$ , the statistics of the observations under the two hypotheses are as follows:

$$\begin{aligned} H_0 : \mathbf{v}_r &\simeq N(\mathbf{u}_r, \Sigma_r) \\ H_1 : \mathbf{v}_r &\simeq N(-\mathbf{u}_r, \Sigma_r), \end{aligned} \quad (5)$$

where we let  $\mathbf{u}_r = S\mathbf{u}$  and  $\Sigma_r = S\Sigma S^T$ . The optimum detector now decides for  $H_0$  when  $\rho_r = (\Sigma_r^{-1}\mathbf{u}_r)^T \mathbf{v}_r > 0$ . As for the full detector,  $\rho_r$  is a Gaussian r.v. with mean and variance (under  $H_0$ ) given by:

$$\begin{aligned} E[\rho_r|H_0] &= \mathbf{u}_r^T \Sigma_r^{-1}\mathbf{u}_r \\ \text{var}[\rho_r|H_0] &= \mathbf{u}_r^T \Sigma_r^{-1}\mathbf{u}_r. \end{aligned} \quad (6)$$

Once again the error probability depends on the  $z$ -value of the randomized detector, which is  $z_r = (\mathbf{u}_r^T \Sigma_r^{-1}\mathbf{u}_r)^{1/2}$ . By introducing the factor

$$\eta = \frac{\mathbf{u}_r^T \Sigma_r^{-1}\mathbf{u}_r}{\mathbf{u}^T \Sigma^{-1}\mathbf{u}}, \quad (7)$$

the  $z$ -value of the randomised detector can be related to that of the full-feature detector as follows:

$$z_r = \sqrt{\eta}z. \quad (8)$$

Given that  $\eta$  is always lower than one and decreases when  $k$  decreases, equations (7) and (8) determine the loss of performance due to the use of the randomised reduced-feature

detector instead of the full-feature one. Since in this case the errors are due to the natural variability of the observed features, the loss of performance can be regarded as a loss of robustness.

2) *Performance Under Attack (Security)*: Given that the attacker does not know the subset of features used by the randomised detector and, in principle, he does not even know about the randomization-based defence mechanism, we can assume that he keeps attacking the full-feature detector (an alternative strategy is considered in Section V-D). Without loss of generality, we will assume that the attacker takes a sequence generated under  $H_0$  and modifies it in such a way that the detector decides in favour of  $H_1$ . The optimum attack is the one that succeeds in inducing a decision error while minimising the distortion introduced within  $\mathbf{v}$ . Such an attack is obtained by moving the vector  $\mathbf{v}$  orthogonally to the decision boundary until  $\rho = 0$ , leading to:

$$\mathbf{v}^* = \mathbf{v} - \alpha \frac{(\Sigma^{-1}\mathbf{u})^T \mathbf{v}}{\|\Sigma^{-1}\mathbf{u}\|^2} \cdot \Sigma^{-1}\mathbf{u}, \quad (9)$$

where  $\mathbf{v}^*$  is the attacked feature vector and  $\alpha$  is a parameter controlling the strength of the attack: with  $\alpha = 1$  the attacked vector is moved exactly on the decision boundary, however the attacker may decide to use a larger  $\alpha$  (introducing a larger distortion) so to move the attacked vector more inside the wrong decision region, hence increasing the probability that the attack is successful also against the randomised detector.

By construction, when applied against the full feature detector the above attack is always successful. We now investigate the effect of the attack defined in (9), when the analyst uses a randomised detector based on a reduced set of features. We start by observing that, as a consequence of the attack, the value of  $\rho_r$  evaluated by the randomised detector is:

$$\rho_r = \mathbf{u}_r^T \Sigma_r^{-1} \mathbf{v}_r - \alpha \frac{\mathbf{u}_r^T \Sigma_r^{-1} \mathbf{w}_r}{\|\mathbf{w}\|^2} \mathbf{w}^T \mathbf{v}, \quad (10)$$

where we let  $\mathbf{w} = \Sigma^{-1}\mathbf{u}$ ,  $\mathbf{w}_r = S\mathbf{w}$ , and  $\theta = \alpha(\mathbf{u}_r^T \Sigma_r^{-1} \mathbf{w}_r / \|\mathbf{w}\|^2)$ . The statistics of  $\rho_r$  are given by the following lemma.

*Lemma 1*: By letting  $y = \mathbf{u}_r^T \Sigma_r^{-1} \mathbf{u}_r$ , and  $x = \mathbf{u}^T \Sigma^{-1} \mathbf{u}$ , we have that under  $H_0$ :

$$E[\rho_r] = y - \theta x, \quad \text{var}[\rho_r] = y + \theta^2 x - 2\theta y. \quad (11)$$

*Proof*: See the appendix.  $\square$

It is worth observing that Lemma 1 holds under the assumption that the attacker applies equation (9) even when the full-feature detector decides for  $H_1$ , that is when  $\rho < 0$ . In general this will not be the case, since when the detector already makes an error due to the presence of noise, the attacker has no interest to modify  $\mathbf{v}$ . In fact, in this case, the result of the application of equation (9) would be the correction of the error made by the detector. Given that deriving the statistics of  $\rho_r$  by taking into account that the attack is present only when  $\rho > 0$  is an intractable problem, we assume that equation (9) always holds. The results obtained in this way provide a good approximation of the real performance of the system when the error probability of the full-feature detector

in the absence of attacks is negligible, i.e., when  $z$  is much larger than 1.

The  $z$ -value of the randomised detector under attack can then be written as

$$z_{att} = z \frac{\eta - \theta}{\sqrt{\eta + \theta^2 - 2\eta\theta}}, \quad (12)$$

where  $\eta$ , defined as in (7), is equal to  $y/x$ . Note that, due to the attack,  $z_{att}$  can also take negative values thus resulting in a large error probability.

The interpretation of equation (12) is rather difficult since  $\eta$  and  $\theta$  depend on  $\mathbf{v}$  in a complicated way. In the next section we will use numerical simulations to get more insight into the performance predicted by (12). Here we observe that the expression of  $z_{att}$  can be simplified considerably when  $\Sigma = \sigma^2 I$ , that is when the features are independent and have all the same variance, and the rows of  $S$  are orthogonal, as with RFS. In this case, it is easy to see that:

$$\theta = \alpha \frac{\|\mathbf{u}_r\|^2}{\|\mathbf{u}\|^2}, \quad \eta = \frac{\|\mathbf{u}_r\|^2}{\|\mathbf{u}\|^2}, \quad (13)$$

and hence  $\theta = \alpha\eta$ . Equation (12) can then be rewritten as:

$$z_{att} = z \frac{\eta(1 - \alpha)}{\sqrt{\eta + \alpha^2\eta^2 - 2\alpha\eta^2}}. \quad (14)$$

From equations (13) and (14), we see that the error probability under attack depends only on  $\eta$ , that is the ratio of the norm of the vector with the mean value of the reduced set of features and the norm of the full-feature mean vector. Clearly, when the number  $k$  of features used by the randomised detector decreases, the value of  $\eta$  decreases as well. Given that the numerator in (14) is either null or negative, and given that the quantity

$$\frac{\eta}{\sqrt{\eta + \alpha^2\eta^2 - 2\alpha\eta^2}} = \frac{1}{\sqrt{\frac{1}{\eta} + \alpha^2 - 2\alpha}} \quad (15)$$

increases when  $\eta$  decreases, the error probability under attack decreases with  $k$ . In fact, if  $\eta$  approaches zero, i.e.,  $k = 1$ ,  $z_{att}$  will be close to 0, and the error probability under attack tends to 0.5. In other words, the probability that an attack against the full feature detector is also affective against a reduced detector based on one feature only is 0.5 (the improved security comes at the price of a reduced effectiveness in the absence of attacks, as stated by equation (8)). As we will see in the next section, even better results are obtained when the features are not independent. Expectedly, it is also easy to see that when  $\eta = 1$ , that is when all the features are used,  $z_{att} = -z$ , hence resulting in a very large error probability.<sup>2</sup>

#### D. Numerical Results

In this section, we use numerical analysis to study the performance of the randomised detector both in the presence and in the absence of attacks as predicted by equations (8) and (12).

<sup>2</sup>In fact, when  $\eta = 1$  the error probability should be equal to 1. This is not the case in the present analysis due to the assumption we made that the attacker always applies equation (9), even when  $\rho < 0$ .

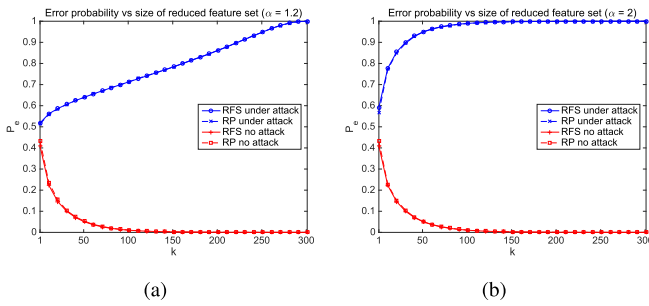


Fig. 2. Missed detection error probability of the randomised feature detector with and without attacks in the case of i.i.d. features. The plots have been obtained by letting  $z = 4$ ,  $n = 300$ ,  $\alpha = 1.2$  (a),  $\alpha = 2$  (b), and averaging over 200 random choices of the matrix  $S$ .

We start with the simple case of i.i.d. features. The performance predicted by the theory are reported in Fig. 2, where we show the dependence on  $k$  of the missed detection error probability both in the presence (upper curves) and in the absence (lower curves) of attacks.<sup>3</sup> The plots have been obtained by letting  $n = 300$ ,  $z = 4$ , and averaging over 500 random choices of the matrix  $S$ . We set  $\alpha = 1.2$  for the leftmost plot and  $\alpha = 2$  for the rightmost. As it can be seen, no particular difference can be noticed between the RFS and RP detectors. The security of the randomised detector increases for lower values of  $k$ , while the performance in the absence attacks decreases. As predicted by equation (15), when the number of features used by the detector tends to 1, the error probability in the presence of attacks tends to 0.5, thus showing that the attack designed to defeat the full-feature detector fails almost half of the times when the reduced feature detector is used. Expectedly, the attack is more successful for larger  $\alpha$ .

We now consider the more general case of dependent features. To do so, we set the statistics of the host features as follows: i) constant mean vector  $\mathbf{u}$  (we also run some simulations with a randomly generated mean vector obtaining very similar results), ii) random covariance matrix constructed by first generating a diagonal matrix with uniformly distributed random diagonal entries, and then randomly rotating the diagonal matrix so to obtain dependent features. We observe that in this way the features have different variances, however, especially after the random rotation, the difference among the variances is not big. This agrees with a practical setup in which the detector relies on normalised features. To force a different variance among features, we considered an additional setup in which the variance of the features is scaled after the random rotation of the covariance matrix. We did so by randomly generating a vector of scaling factors uniformly distributed between 0 and 1 and applying the square root of the scaling factors to the features. We used the square root of the random scaling factors to limit the scale differences among the features. In the following we refer to the first case as dependent normalised features and to the second as dependent features.

<sup>3</sup>Due to the symmetry of the problem, in the absence of attacks, the missed detection probability is equal to the overall error probability.

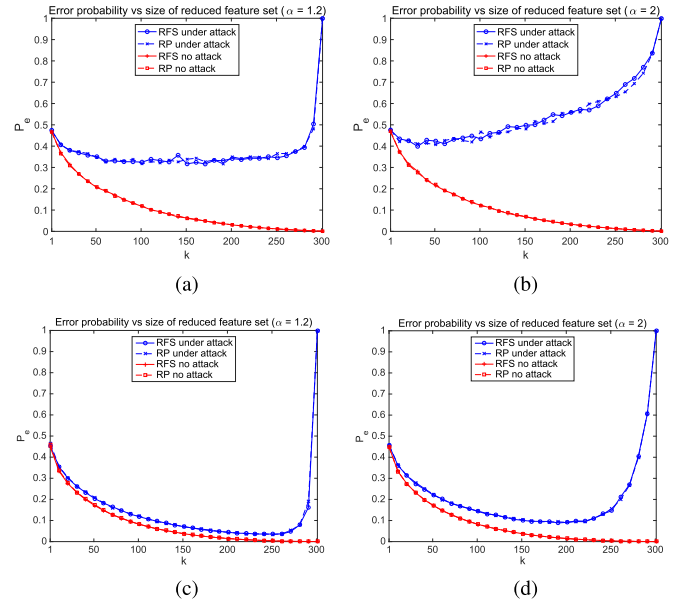


Fig. 3. Missed detection error probability of the randomised feature detector with and without attacks, in the case of dependent features. The plots have been obtained by letting  $n = 300$ ,  $\alpha = 1.2$  (a)-(c), and  $\alpha = 2$ , (b)-(d). In (a) and (b) the features variance is not normalised while (c) and (d) refer to the case of normalised features. In all cases the error probability has been obtained by averaging over 500 random choices of the host feature statistics and the matrix  $S$ . The average value of  $z$  was: (a)  $z = 4.9$ , (b)  $z = 5$ , (c)  $z = 4.8$ , (d)  $z = 4.9$ .

Alike in the i.i.d. case, we let  $n = 300$  and averaged the results over 500 repetitions, each time by randomly generating a new covariance matrix and a new matrix  $S$ . Even in this case, we considered two different values of  $\alpha$ , namely:  $\alpha = 1.2$  and  $\alpha = 2$ . Fig. 3 reports the results that we have obtained. Due to the randomness involved in the generation of the feature statistics, we could not control the exact value of  $z$ , the values resulting from each experiment are reported in the caption of the figure. As in Fig. 2, no particular difference is observed between the RFS and RP cases, however, the overall behaviour of the detector is completely different with respect to the i.i.d. case. The error probability under attack now decreases more rapidly when the number of features used by the detector is reduced, thus indicating a high security level. After a certain point, however, the error probability increases again approaching 0.5 when  $k$  tends to 1. Such a behaviour can be interpreted as a loss of robustness rather than a loss of security. In fact, the error probability in the absence of attacks exhibits a similar increase when  $k$  decreases, indicating that the detector is not able to distinguish between  $H_0$  and  $H_1$  by relying on few features only. Of course, such a problem has an impact also on the error probability in the presence of attacks. Overall, the reduced detector performs better in the case of dependent features, however, higher values of  $k$  must be used with respect to the i.i.d. case. We also observe that a significantly higher security level is obtained for the case of normalised features (cases (c) and (d) in Fig. 3). In addition, increasing the value of  $\alpha$  does not have a great impact on the success rate of the attack, especially in the case of normalized features.

In order to explain the plots reported in Fig. 3, noticeably their difference with respect to Fig. 2, we need to analyse in more detail the reason behind the security improvement achieved through random feature selection. We will do so by referring to the RFS case, however the same considerations can be applied to the RP case. We start by observing that the optimal attack given in the rightmost part of equation (9) can be decomposed in two parts: a scaling factor  $a$  and a direction  $\mathbf{e}_{att}$ :

$$a = \alpha \frac{(\Sigma^{-1}\mathbf{u})^T \mathbf{v}}{\|\Sigma^{-1}\mathbf{u}\|}$$

$$\mathbf{e}_{att} = \frac{\Sigma^{-1}\mathbf{u}}{\|\Sigma^{-1}\mathbf{u}\|}. \quad (16)$$

The direction  $\mathbf{e}_{att}$  ensures the optimality of the attack, since it is chosen so to be orthogonal to the boundary of the decision regions. The scaling factor  $a$  is chosen in such a way to move the attacked feature vector exactly on the boundary of the decision regions ( $\alpha = 1$ ) or inside the target region ( $\alpha > 1$ ). With the randomized feature detector, the attack vector is projected onto a subspace with lower dimensionality and both the scaling factor and the direction of the projected attack vector are no more optimal. The non-optimality of the scaling factor (*scale mismatch*) implies that sometimes the magnitude of the projected attack vector is too small thus failing to induce a decision error. Such a probability is augmented by the non-optimality of the direction of the projected attack vector, since there is no guarantee that such a direction is orthogonal to the decision boundary of the RFS detector. In order to understand when and to which extent the *angle mismatch* between the optimal attack direction and the direction of the projected attack vector affects the security of the RFS detector, let us consider the expression of such directions. The direction orthogonal to the decision boundary of the RFS detector is:

$$\mathbf{e}_{RFS} = \frac{\Sigma_r^{-1}\mathbf{u}_r}{\|\Sigma_r^{-1}\mathbf{u}_r\|}, \quad (17)$$

while the projection of  $\mathbf{e}_{att}$  on the reduced feature space is:

$$\mathbf{e}_{att,r} = \frac{S\Sigma^{-1}\mathbf{u}}{\|S\Sigma^{-1}\mathbf{u}\|}. \quad (18)$$

In the case of i.i.d. feature, it is easy to see that the above directions coincide given that  $\Sigma_r^{-1} = I_{k \times k}$ ,  $\Sigma^{-1} = I_{n \times n}$  and  $SS^T = I_{k \times k}$ . This is not the case with dependent features. To understand the importance of the angle mismatch in this case, we randomly generated 1000 pairs of covariance matrixes and random selection matrices  $S$ . Then we plot the histogram of the angle mismatch. The results we got are reported in Fig. 4, for both the cases of non-normalized and normalized features. The figure shows the results for  $k = 50, 150$  and  $250$  (similar results hold for other values of  $k$ ). Upon inspection of the figure, we can see that the angle mismatch is much larger in the case of normalized features, thus explaining the higher security performance of the randomized detector in this case. For small values of  $k$  the mismatch is so large that the attack vector is almost orthogonal to the optimal direction, thus nullifying the effect of the attack. Such an observation

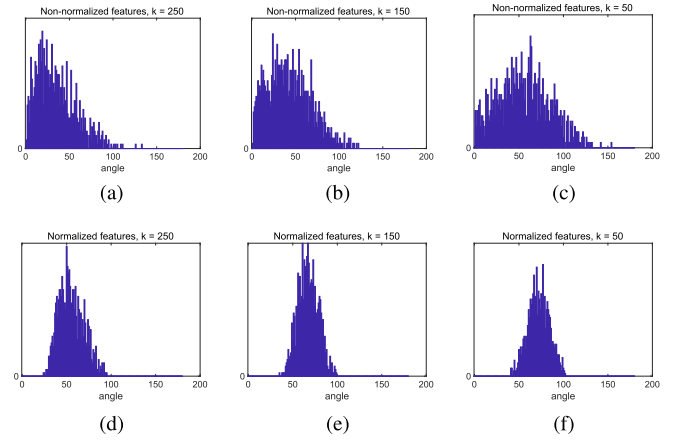


Fig. 4. Angle mismatch between the projection of the attack vector on the feature space of the RFS detector and the optimal attack direction for the RFS detector. Plots (a)-(c) refer to the case of non-normalized features, while plots (d)-(f) correspond to the case of normalized features.

also explains why in this case increasing  $\alpha$  does not increase the effectiveness of the attack, in fact, the increased value of the magnitude of the attack is *wasted* since the direction of the attack is a wrong one.

Given that the particular form of the matrix  $S$  does not have a significant impact on the performance of the reduced detector, in the following we focus on the RFS case only. An advantage of such an approach is its lower computational complexity. In the RP cases, in fact, the detector must compute the entire feature vector and then choose only  $k$  linear combinations. In the RFS case, instead, the detector can compute only the features that it intends to use, avoiding to calculate the non-selected features; this may allow a significant reduction of the complexity, especially for large values of  $n$  and small values of  $k$ .

#### IV. APPLICATION TO IMAGE MANIPULATION DETECTION

The theoretical analysis given in the previous section suggests that a detector based on a randomised subset of features provides a better security with respect to a full-feature detector. The applicability of such an idea to real world applications, however, requires great care, since the assumptions behind the theoretical analysis are ideal ones and are rarely met in practice. Since the goal of this paper is to improve the security of image forensic techniques against counter-forensic attacks, in this section, we introduce an SVM-based detector based on random feature selection and apply it to two particular image forensic problems, namely the detection of adaptive histogram equalization and the detection of median filtering. The full feature space consists of a subset of SPAM features [30], however the SVM detector is trained by relying only on a random subset of the full feature set. As we will see, the loss of performance of the RFS SVM detector in the absence of attacks is very limited, even for rather small values of  $k$ .

We also introduce two attacks aiming at deceiving the SVM detector. Both attacks are based on gradient descent, the first one works in the feature domain, while the second operates directly in the pixel domain. The attacks are very powerful

since they were able to prevent a correct detection in all the test images. In particular, the attack operating in the pixel domain is a very practical one since it does not require to map back the attack from the feature to the pixel-domain and can prevent a correct detection by introducing a limited distortion in the attacked image. In Section V, we will use these attacks to demonstrate the improved security ensured by the RFS detector.

#### A. RFS SVM-Based Detection of Image Manipulations

Residual-based features, originally devised for steganalysis applications [30], [31], have been used with success in many image forensic applications, including forgery detection [43], detection of pixel-domain enhancement, spatial filtering, resampling and lossy compression [44]. In particular, in this paper, we consider the SPAM feature set [30]. Since we carried out all our tests on grey-level images, we assume that these features are computed directly on grey-level pixel values, or on the luminance component derived from the RGB colour bands. Feature computation consists of three steps. In the first step, residual values are computed; specifically, the difference 2-D arrays are evaluated along horizontal, vertical and diagonal directions. In the second step, the residual values are truncated so that their maximum absolute value is equal to  $T$ . Finally, the co-occurrence matrices are computed. Depending on the value of  $T$  and the order of the co-occurrences considered in the computation, different sets of features with different sizes are obtained. We use second-order SPAM features with  $T = 3$ , for a dimensionality of the feature space of 686.

Based on the SPAM features, we built an SVM detector aiming at revealing Adaptive Histogram Equalization (AHE) and Median Filtering (MF). With regard to AHE, we considered the contrast-limited algorithm (CLAHE) implemented by Matlab function *adapt\_histeq* with *cliplimit* = 0.02. Some sample images manipulated in this way are shown in the second column in Fig. 5. With regard to MF, we considered window sizes  $3 \times 3$ ,  $5 \times 5$  and  $7 \times 7$ , (MF3, MF5, MF7).

To train and test the detectors, we used a set of 2000 images from the RAISE-2k dataset [45]. Specifically, we used 1400 images for training and 600 images for testing. To speed up the experiments, the images were downsampled by factor 4 and converted to grayscale. The SVM models are built by using the tools provided by the LibSVM library [46]. The RBF kernel is used for all the SVMs. The results of the tests are shown in the first row of Table II. All the four detectors got a 100% accuracy on the test data, thus confirming the excellent capabilities of the SVM trained with the SPAM features to detect global image manipulations like median filtering and histogram equalization.

To double-check that the SVM model does not overfit to the training data, Table I shows the number of support vectors for the SVM detector trained on the full feature set. These numbers are very low compared to the number of examples considered for training, which is 2800 (1400 per class). More specifically, the ratio between the number of support vectors and training examples is around 0.04 - 0.05 (which can be taken as a good indication of the upper bound of the generalisation error [47]).

TABLE I  
NUMBERS OF SUPPORT VECTORS OF FULL-FEATURE SVM MODEL

	AHE	MF3	MF5	MF7
<b>Num. of support vectors</b>	62	59	71	69

TABLE II  
ERROR PROBABILITY OF SPAM-BASED SVM DETECTORS  
IN THE ABSENCE AND PRESENCE OF ATTACKS ( $\varepsilon = 0.5$ )

	AHE	MF3	MF5	MF7
<b>Manipulated</b>	0%	0%	0%	0%
<b>Attack in feature domain</b>	100%	100%	100%	100%
<b>Attack in pixel domain</b>	100%	100%	100%	100%

We also carried out some tests with a linear SVM, however we found that the linear model does not discriminate well the two classes when the number of features decreases to less than 100-200, thus preventing the application of RFS with small values of  $k$ .

#### B. Attack in the Feature Domain

In this section, we describe the feature domain attack we have implemented against the RFS SVM detectors.

The attack in the feature domain has been built by following the system described in [6]. Specifically, given a feature vector  $\mathbf{v}$  and a discriminant function  $g(\cdot)$ , we assume that the detector decides that  $\mathbf{v}$  belongs to a manipulated image if  $g(\mathbf{v}) > 0$ , and to an original image otherwise. In this setup, the optimally attacked vector  $\mathbf{v}^*$  is determined by solving the following minimization problem:

$$\mathbf{v}^* = \arg \min_{\mathbf{v}': g(\mathbf{v}') \leq -v} d(\mathbf{v}, \mathbf{v}'), \quad (19)$$

where  $d(\cdot, \cdot)$  is a suitable distortion measure and  $v$  is a safe margin, which, similarly to the parameter  $\alpha$  in Section III, permits to move the attacked vector more or less inside the acceptance region. For an SVM detector, the discriminant function can be written as:

$$g(\mathbf{v}) = \sum_i \alpha_i y_i k(\mathbf{v}, \mathbf{v}_i) + b, \quad (20)$$

where  $\alpha_i$  and  $y_i$  are, respectively, the support value and the label of the  $i$ -th support vector  $\mathbf{v}_i$ , and where  $k(\cdot)$  is the kernel function. In our implementation, the minimization problem is solved by using a gradient descent algorithm, where the gradient at each iteration is computed as:

$$\nabla g(\mathbf{v}) = \sum_i \alpha_i y_i \nabla k(\mathbf{v}, \mathbf{v}_i). \quad (21)$$

The discrimination function and the corresponding gradient for different kernels are reported in Table III.

As a matter of fact, in our implementation we used the probabilistic output of the SVM rather than  $g(\mathbf{v})$ . The probabilistic output is built by mapping  $g(\mathbf{v})$  into the  $[0, 1]$  range and by letting the value  $g(\mathbf{v}) = 0$  correspond to a probabilistic output equal to 0.5. By indicating the probabilistic output of the SVM with  $p(\mathbf{v})$ , the minimization problem in (19) can be rewritten as:

$$\mathbf{v}^* = \arg \min_{\mathbf{v}': p(\mathbf{v}') \leq \varepsilon} d(\mathbf{v}, \mathbf{v}'), \quad (22)$$

where  $\varepsilon = 0.5$  (equivalent to  $v = 0$ ) corresponds to an attacked feature vector lying on the decision boundary and values of



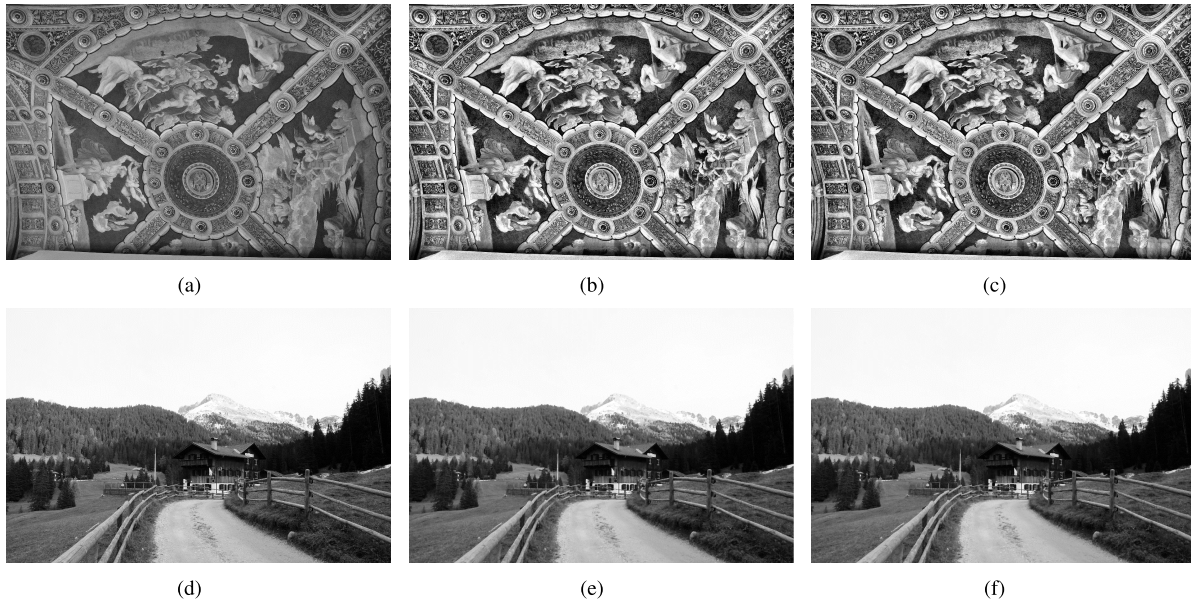


Fig. 5. Examples of manipulated and attacked images: (a) and (d) show the original images; (b) the image manipulated by AHE; (e) the image manipulated by MF3; (c), and (f) the images after a pixel-domain attack with  $\varepsilon = 0.5$ .

TABLE III  
DISCRIMINATION FUNCTIONS AND CORRESPONDING GRADIENT FOR DIFFERENT SVM KERNELS

	$g(\mathbf{v})$	$\nabla g(\mathbf{v})$
<b>Linear kernel</b>	$k(\mathbf{v}, \mathbf{v}_i) = \mathbf{v}^T \mathbf{v}_i$	$\sum_i a_i y_i \mathbf{v}_i$
<b>Polynomial kernel</b>	$k(\mathbf{v}, \mathbf{v}_i) = (\mathbf{v}^T \mathbf{v}_i + c)^p$	$\sum_i a_i y_i p (\mathbf{v}^T \mathbf{v}_i + c)^{p-1} \mathbf{v}_i$
<b>RBF kernel</b>	$k(\mathbf{v}, \mathbf{v}_i) = \exp(-\gamma \ \mathbf{v} - \mathbf{v}_i\ ^2)$	$-\sum_i 2\gamma a_i y_i \exp(-\gamma \ \mathbf{v} - \mathbf{v}_i\ ^2) (\mathbf{v} - \mathbf{v}_i)$

TABLE IV  
AVERAGE SNR (DB) COMPUTED ON THE FEATURES OF 600 ATTACKED IMAGES. SNR IS DEFINED AS THE RATIO BETWEEN THE ENERGY OF THE FEATURE VECTOR BEFORE THE ATTACK AND THE ENERGY OF THE DISTORTION

	$\varepsilon = 0.5$	$\varepsilon = 0.3$	$\varepsilon = 0.1$
<b>AHE</b>	26.01	25.09	23.88
<b>MF3</b>	18.59	17.53	16.1
<b>MF5</b>	19.19	18.15	16.76
<b>MF7</b>	19.15	18.12	16.72

$\varepsilon < 0.5$  introduce a safe margin bringing the attacked vector deeper inside the wrong detection region.

We attacked the full-feature SVM detectors for AHE and MF by applying the gradient descent attack in the feature domain; even with  $\varepsilon = 0.5$ , the attack was able to deceive all the detectors, as reported in the second row of Table II. The distortion introduced by the attack for  $\varepsilon = 0.5, 0.3$  and  $0.1$  is given in Table IV. As expected the distortion increases for smaller values of  $\varepsilon$ . We observe that the distortion values reported in the table do not give an immediate indication of the distortion of the attacked image, due to the difficulties of mapping back the attacked feature vector into the pixel domain.

### C. Attack in the Pixel Domain

In a realistic scenario, the attacker will carry out his attack in the pixel domain. Similarly to equation (22), the goal

of the attack in the pixel domain is to solve the following optimisation problem:

$$I^* = \arg \min_{I': p(f(I')) \leq \varepsilon} d(I, I'). \quad (23)$$

where  $f(\cdot)$  is the feature extraction function (i.e.  $f(I) = \mathbf{v}$ ) and, as before,  $p$  indicates the probabilistic output of the SVM. Due to the complicated form of  $f(\cdot)$  and to the necessity of preserving the integer nature of pixel values, gradient descent can not be applied directly to solve (23). For this reason we implemented the attack as described in [11], by setting to 20% the percentage of pixels modified at each iteration of the attack (see [11] for more details).

The results of the attack on the same dataset used in Section IV-A are reported in the last row of Table II. Only the results for  $\varepsilon = 0.5$  are shown for simplicity. The results show that the attack always succeeds for all the detectors. The average distortion introduced within the attacked images is reported in Table V, while some examples of attacked images are given in the third column of Fig. 5.

## V. SECURITY OF RFS-BASED MANIPULATION DETECTION AGAINST FEATURE-DOMAIN AND PIXEL-DOMAIN ATTACKS

In this section, we show the improved security provided by RFS by testing the adaptive histogram equalization and median filtering detectors introduced in Section IV-A against the attacks presented in Sections IV-B and IV-C.

TABLE V  
AVERAGE PSNR (dB) OF THE 600 ATTACKED IMAGES

	$\varepsilon = 0.5$	$\varepsilon = 0.3$	$\varepsilon = 0.1$
<b>AHE</b>	54.49	53.09	51.96
<b>MF3</b>	56.77	56.43	54.17
<b>MF5</b>	56.79	55.51	53.37
<b>MF7</b>	56.43	54.73	52.63

### A. Experimental Methodology

In our experiments, we used the same setting described in Section IV-A. The SVM models were trained by using a randomly selected subset of SPAM features extracted from the training set. We considered two versions of the detectors. The first version relies on SPAM features as they are, with no normalisation. The second version is built by normalising the SPAM features before feeding them to the SVM. Normalization is applied both during the training and testing phases, and is achieved by dividing each feature by its standard deviation, estimated on the training set. The two versions of the detector correspond to the two scenarios considered in Sect. III-D (see Fig. 3). The processed images were attacked by using the two methods presented in Section IV-B and IV-C. In both cases, we assumed that the attacker has access to a version of the full-feature SVM detectors trained on the same dataset used by the analyst. The performance of the RFS SVM detectors were then evaluated on both the attacked and non-attacked images for different values of  $k$ . We repeated the experiments 100 times, each time using a different matrix  $S$ . Various stopping conditions, namely  $\varepsilon = \{0.1, 0.3, 0.5\}$ , were considered for both the attacks. For the dimension of the reduced feature set, we considered values of  $k$  in the range  $\{1, 3, 5, 10, 20, 50, 100, 200, 300, 400, 500, 600, 686\}$ . The value of  $\gamma$  (see Table III) was obtained by means of 5-fold cross-validation carried out on the training set. To evaluate the performance in the absence of attacks, we considered both false alarms and missed detections, while for the security in the presence of attacks, we considered only the missed detection probability, since in our setup the goal of the attacker is to induce a missed detection event.

### B. Security vs Robustness Tradeoff: Feature Domain

In this section, we describe the performance of the RFS detectors against the feature-domain attack described in Section IV-B. As we already pointed out, this is an ideal situation for the attacker, since in real applications the attacker usually does not have access to the feature domain and the inverse-mapping from the feature domain to the pixel domain is often a difficult task. As we will see, despite the setup is more favourable to the attacker than in the case of a pixel domain attack, the RFS detector exhibits a good security. Fig. 6 shows the error probability of the RFS detector as a function of  $k$  with and without attacks for the MF3 and AHE detectors. The plots shown in parts (a) and (b) refer to the case of non-normalized features, while the results shown in parts (c) and (d) have been obtained by normalising the feature vector. Despite some unavoidable differences, the overall behaviour predicted by the theoretical analysis is

confirmed. To start with, the experiments confirm that reducing the dimension of the feature set does not impact much the performance of the detector in the absence of attacks. In fact, even with  $k = 1$ , the probability of correct detection is still around 0.8 for AHE and 0.7 for MF3 for the case of non-normalized features (Fig. 6 (a) and (b)). Similar results hold in the normalized feature case (Fig. 6 (c) and (d)). In the presence of attacks, the missed detection probability of the RFS detectors is significantly lower than that of the full-feature detector ( $k = 686$  in the figure). For the non-normalized case, we observe that when the stopping condition of the attack is equal to 0.5, the missed detection probability drops even for rather large values of  $k$ , while for  $\varepsilon = 0.3$  and  $\varepsilon = 0.1$ , smaller values of  $k$  are needed to make the missed detection probability drop. With regard to the detector based on normalised features, the experiments confirm the results predicted by the theory, since feature normalization results in a significantly more secure detector. In fact, for the normalized feature case we observe that the missed detection probability immediately drops when  $k$  is below a critical very large value of  $k$ , which depends on the stopping condition of the attack. In particular, for the AHE case, the missed detection error drops below around  $k = 400$  for  $\varepsilon = 0.1$ ,  $k = 500$  for  $\varepsilon = 0.3$  and  $k = 600$  for  $\varepsilon = 0.5$ . For the MF3 case, the situation is even more favorable. First of all we must mention that in this case the stopping conditions were set to  $\varepsilon = 0.5$  and 0.45 only. The reason for such a choice is that reaching  $\varepsilon = 0.45$  already requires a large number of iterations and a very strong attack. In fact, the SNR of the attacked feature vector for  $\varepsilon = 0.45$  is 4.19 (while it was 7.45 for the AHE case with  $\varepsilon = 0.1$ ).

The results in Fig. 6 clearly show that a suitable value of  $k$  can be found where the error probability in the absence of attacks (both false alarm and missed detection) is still negligible and the missed detection probability under attack is significantly smaller than 1. The results obtained with the MF5 and MF7 detectors are shown in Fig. 7 for the case of non-normalized features. We see that, with respect to the MF3 case, a lower  $k$  is necessary to get a small missed detection probability under attacks, especially for the MF7 case. The motivation for this behaviour is that, in these cases, the separation between the two classes with respect to the intraclass feature variation is very large, thus diminishing the impact of the scale mismatch on the effectiveness of the attack carried out on the full feature set.

### C. Security vs Robustness Tradeoff: Pixel Domain

In this section, we focus on the security of the RFS detectors against the pixel domain attack introduced in [11] and briefly described in Section IV-C. Before going on with the discussion of the results, we stress that attacking the images directly in the pixel domain requires a very high computational effort,<sup>4</sup>

<sup>4</sup>With the attack method in [11], attacking a single image on a 4-core PC, running at 3.3 GHz equipped with 16 GB RAM requires around two hours on the average. By considering that each curve in the figures shown in this section requires attacking 600 images (and that the effectiveness of the attack must be evaluated for all the images, for all values of  $k$  and 100 different selection matrices  $S$  for each  $k$ ) each curve takes more than one month to be computed.

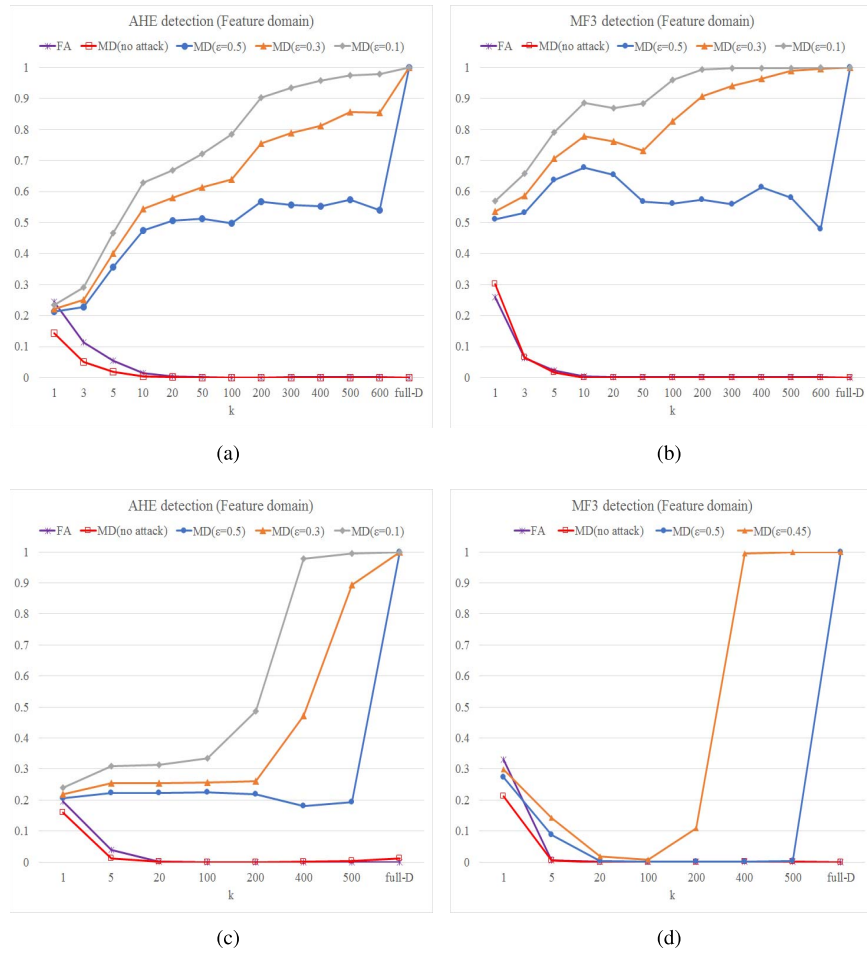


Fig. 6. Error probability of the randomised feature detector without attacks and under a feature-domain attack, in the case of *AHE* and *MF3* detection. (a) and (b) refer to the case of non-normalized features, while (c) and (d) have been obtained by normalizing the features before feeding them to the detector. The missed detection (MD) probability is reported for both the non-attacked and attacked images, by letting  $\epsilon = \{0.5; 0.3; 0.1\}$  (in (a), (b) and (c)), and  $\epsilon = \{0.5; 0.45\}$  (in (d)). The false alarm (FA) is also reported. The plots have been obtained averaging over 100 random choices of the matrix  $S$ .

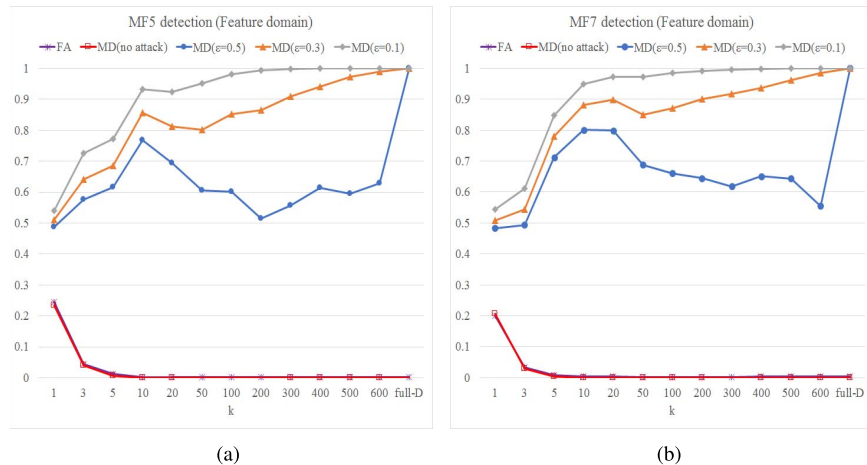


Fig. 7. Error probability of the randomised feature detector without attacks and under a feature-domain attack, in the case of *MF5* (a) and *MF7* (b) detection with non-normalized features. The MD probability is reported for both the non-attacked and attacked images, by letting  $\epsilon = \{0.5; 0.3; 0.1\}$ . The FA is also reported. The plots have been obtained averaging over 100 random choices of the matrix  $S$ . The curves of the FA and the MD for the no-attack case are overlapped.

so we limited our analysis to a lower number of cases, which, however, are sufficient to show the good performance achieved by the RFS detectors.

In Fig. 8, we show the error probability of the RFS detectors based on non-normalized features. The overall trend of the error probability is similar to that observed for the

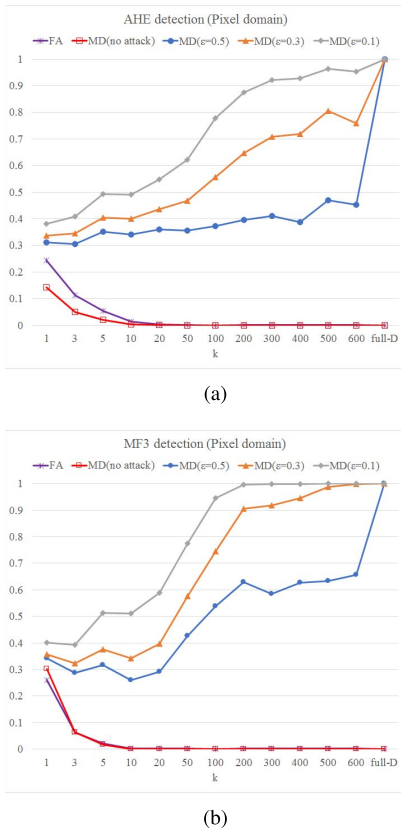


Fig. 8. Error probability of the randomised feature detector under pixel-domain attack, in the case of *AHE* (a) and *MF3* (b) detection with non-normalized features. The plots have been obtained by attacking with  $\epsilon = \{0.5, 0.3, 0.1\}$  and averaging over 100 random choices of the matrix  $S$ . The FA and the MD probabilities in the presence and absence of attacks are reported.

feature-domain attack, however the error probability is now much smaller, confirming that the possibility of carrying out the attack in the feature domain, would represent a significant advantage for the attacker. By letting  $k = 20$ , for instance, the missed detection probability in the presence of attacks ranges from 0.36 to 0.55 in the case of *AHE* and from 0.29 to 0.59 in the case of *MF3*, while the missed detection probability without attacks is equal to 0.013 and 0, respectively. As for the feature domain attack, the false alarm probability remains low (around 0) even for very small values of  $k$ , e.g.  $k = 10$ .

We also carried out some experiments with the *AHE-SVM* detector relying on normalized features. The results we have got are reported in Fig. 9. The general trend of the error probability agrees with that predicted by theory, and confirms that the RFS detector based on normalized features is considerably more secure than the detector based on non-normalized features.

#### D. Other Attack Strategies

Throughout the paper we have assumed that the attacker designs his attack by targeting the full feature detector. While this is a reasonable approach, if we assume that the attack is aware of the defence mechanism other attack strategies are possible. By assuming that the value of  $k$  is publicly available, the attacker could decide to attack a subset of  $k$  features at

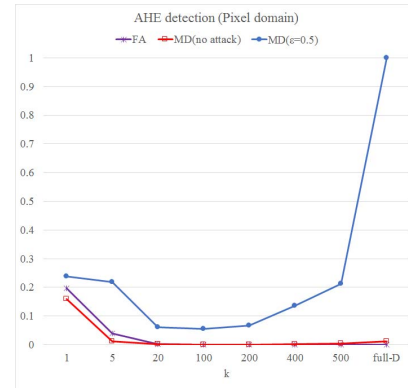


Fig. 9. Error probability of the RFS detector in the *AHE* case with normalized features under a pixel domain attack. The MD probability is reported for both the non-attacked and attacked images, with  $\epsilon = 0.5$ . The FA is also reported. The plots have been obtained averaging over 100 random choices of the matrix  $S$ .

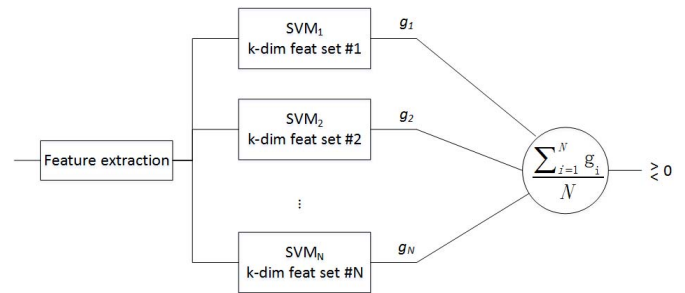


Fig. 10. Block diagram of the alternative attack strategy described in Sect. V-D. The number of reduced features  $k$  adopted by the real classifier is assumed to be known to the attacker;  $g_i$  denotes the discriminant function of the SVM trained on the  $i$ -th feature set.

random, or he could decide to attack the  $k$  most important features chosen based on principal component analysis. While we leave a thorough exploration of alternative attack strategies to a future work, in this section we evaluate the security of the randomized feature detector when the attack targets an *expected* version of the classifier, computed by averaging  $N$  random predictions obtained by  $N$  different classifiers trained on  $N$  different random subset of  $k$  features. Specifically, we consider the attack setup depicted in Fig. 10. This strategy resembles the Expectation Over Transformation (EOT) attack which has been recently proposed against randomized neural network classifiers [48]. As suggested by its name, such method computes the gradient over the expected transformation of the input. In our case, the transformation consists in the selection of the  $k$  features.

For sake of brevity, we implemented the attack depicted in the figure by operating in the feature domain and by considering the case of non-normalized features. With regard to the number of classifiers used to build the attack, we considered the cases  $N = 50$ , and  $N = 100$ . The results we have got are reported in Fig. 11. Upon inspection of the plots reported in the figure, we see that the effectiveness of the new attack increases with  $N$  and it is comparable to that of the attack carried on the full feature detector (indeed always worse than that in the case of median filtering), thus confirming the security improvement achieved by the RFS detector, even in the presence of more elaborated attacks.

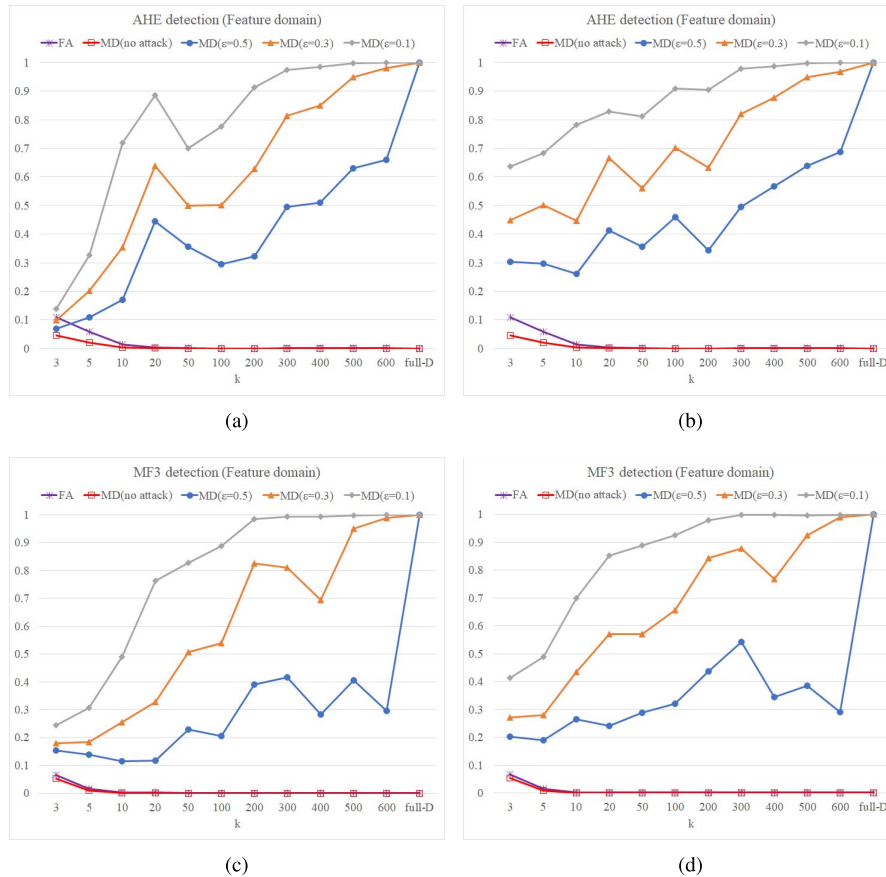


Fig. 11. Error probability of the randomised feature detector without attacks and under a feature-domain attack obtained by averaging 50 (cases (a) and (c)) and 100 (cases (b) and (d)) SVM detectors, in the case of AHE ((a) and (b)) and MF3 ((c) and (d)) with non-normalized features. The MD probability is reported for both the non-attacked and attacked images, by letting  $\varepsilon = \{0.5; 0.3; 0.1\}$ . The FA is also reported. The plots have been obtained averaging over 100 random choices of the matrix  $S$ .

## VI. CONCLUSIONS

The use of machine learning tools in the context of image forensics is endangered by the relative ease with which such tools can be deceived by an informed attacker. In this paper, we have considered a particular instantiation of the above problem, wherein the attacker knows the architecture of the detector, the data used to train it and the class of features the detector relies on. Specifically, we introduced a mechanism whereby the analyst designs the detector by randomly choosing the features used by the detector from an initial large set  $\mathcal{V}$  also known by the attacker. We first presented some theoretical results, obtained by relying on a simplified model, suggesting that random feature selection permits to greatly enhance security at the expense of a slight loss of performance in the absence of attacks. Then, we applied our strategy to the detection of two specific image manipulations, namely adaptive histogram equalization and median filtering, by resorting to SVM classification, based on the SPAM feature set. The security analysis, carried out by attacking the two detectors both in the feature and the pixel domain, confirms that security can indeed be improved by means of random feature selection, with a gain that is even more significant than that predicted by the simplified theoretical analysis. In fact, the probability of a successful attack drops from nearly 1 to

less than 0.4 in the realistic case that the attack is carried out in the pixel domain. We remark that, while an error probability lower than 0.4 would be preferable, this may be already enough to discourage the attacker in applications wherein ensuring that the attack is successful is a vital requirement for the attacker.

This work is only a first attempt to exploit randomisation, noticeably feature space randomisation, to restore the credibility of the forensic analysis in adversarial settings. From a theoretical perspective, more accurate models could be used to further reduce the gap between the analysis carried out in Section III and the conditions encountered in real applications. From a practical point of view, the use of random feature selection with detectors other than SVMs could be explored, together with the adoption of much larger feature sets, e.g., the entire set of rich features [31]. Another interesting research direction consists in the extension of our approach to counter attacks against detectors based on deep learning, specifically convolutional neural networks. In such a case, in fact, the features used by the detector are not chosen by the analyst, since they are determined by the network during the training phase, hence calling for the adoption of other forms of randomisation (see [49] and [50] for some preliminary works in this direction).

## APPENDIX

*Proof of Lemma 1.*

Let  $\rho_r = \rho_{r,1} - \theta\rho_{r,2}$ , where

$$\rho_{r,1} = \mathbf{u}_r^T \Sigma_r^{-1} \mathbf{v}_r, \quad (24)$$

$$\rho_{r,2} = \mathbf{w}^T \mathbf{v}. \quad (25)$$

By observing that  $\theta$  does not depend on  $\mathbf{v}$  and hence it is a fixed value for a given  $S$ , the statistics of  $\rho_r$  under attack depend only on the statistics of  $\rho_{r,1}$  and  $\rho_{r,2}$ , which are two Gaussian random variables. We now prove that the statistics of  $\rho_{r,1}$  and  $\rho_{r,2}$  are given as follows:

$$\begin{aligned} E[\rho_{r,1}] &= y, & E[\rho_{r,1}^2] &= y + y^2, & \text{var}[\rho_{r,1}] &= y, \\ E[\rho_{r,2}] &= x, & E[\rho_{r,2}^2] &= x + x^2, & \text{var}[\rho_{r,2}] &= x, \end{aligned} \quad (26)$$

and

$$\begin{aligned} E[\rho_{r,1}\rho_{r,2}] &= xy + y, \\ \text{cov}[\rho_{r,1}\rho_{r,2}] &= y \end{aligned} \quad (27)$$

To derive the above relations, we start by proving that:

$$E[\mathbf{v}\mathbf{v}^T] = \Sigma + \mathbf{u}\mathbf{u}^T, \quad E[\mathbf{v}_r\mathbf{v}_r^T] = \Sigma + \mathbf{u}_r\mathbf{u}_r^T. \quad (28)$$

To prove such properties, we observe that the element in position  $(i, j)$  of matrix  $\mathbf{v}\mathbf{v}^T$  is equal to  $v_i v_j$ . The element in position  $(i, j)$  of the covariance matrix  $\Sigma$  can be computed as:

$$\Sigma = E[(v_i - u_i)(v_j - u_j)] = E[v_i v_j] - u_i u_j. \quad (29)$$

Thus,  $E[v_i v_j] = \Sigma + u_i u_j$  and  $E[\mathbf{v}\mathbf{v}^T] = \Sigma + \mathbf{u}\mathbf{u}^T$ . In the same way

$$E[\mathbf{v}_r\mathbf{v}_r^T] = \Sigma_r + \mathbf{u}_r\mathbf{u}_r^T. \quad (30)$$

The expectation of  $\rho_{r,1}$  is clearly equal to:

$$E[\rho_{r,1}] = E[\mathbf{u}_r^T \Sigma_r^{-1} \mathbf{v}_r] = \mathbf{u}_r^T \Sigma_r^{-1} \mathbf{u}_r = y, \quad (31)$$

hence, based on (28), we have:

$$\begin{aligned} E[\rho_{r,1}^2] &= E[\mathbf{u}_r^T \Sigma_r^{-1} \mathbf{v}_r \mathbf{v}_r^T \Sigma_r^{-1} \mathbf{u}_r] \\ &= \mathbf{u}_r^T \Sigma_r^{-1} E[\mathbf{v}_r \mathbf{v}_r^T] \Sigma_r^{-1} \mathbf{u}_r \\ &= \mathbf{u}_r^T \Sigma_r^{-1} (\Sigma_r + \mathbf{u}_r \mathbf{u}_r^T) \Sigma_r^{-1} \mathbf{u}_r \\ &= y + y^2. \end{aligned} \quad (32)$$

In addition, the variance of  $\rho_{r,1}$  boils down to:

$$\text{var}[\rho_{r,1}] = E[\rho_{r,1}^2] - E[\rho_{r,1}]^2 = y. \quad (33)$$

In a similar manner, we can prove that:

$$\begin{aligned} E[\rho_{r,2}] &= x, \\ E[\rho_{r,2}^2] &= x + x^2, \\ \text{var}[\rho_{r,2}] &= x. \end{aligned} \quad (34)$$

Finally, by exploiting again the properties in (28), we can write:

$$\begin{aligned} E[\rho_{r,1}\rho_{r,2}] &= E[\mathbf{u}_r^T \Sigma_r^{-1} \mathbf{v}_r \mathbf{v}^T \Sigma^{-1} \mathbf{u}] \\ &= E[\mathbf{u}_r^T \Sigma_r^{-1} S \mathbf{v}\mathbf{v}^T \Sigma^{-1} \mathbf{u}] \\ &= \mathbf{u}_r^T \Sigma_r^{-1} S E[\mathbf{v}\mathbf{v}^T] \Sigma^{-1} \mathbf{u} \end{aligned}$$

$$\begin{aligned} &= \mathbf{u}_r^T \Sigma_r^{-1} S (\Sigma + \mathbf{u}\mathbf{u}^T) \Sigma^{-1} \mathbf{u} \\ &= y + xy, \end{aligned} \quad (35)$$

and then immediately:

$$\text{cov}[\rho_{r,1}\rho_{r,2}] = E[\rho_{r,1}\rho_{r,2}] - E[\rho_{r,1}]E[\rho_{r,2}] = y. \quad (36)$$

The lemma follows immediately from relations (26)-(27), by observing that

$$E[\rho] = E[\rho_{r,1}] - \theta E[\rho_{r,2}] \quad (37)$$

and

$$\text{var}[\rho] = \text{var}[\rho_{r,1}] + \theta^2 \text{var}[\rho_{r,2}] - 2\theta \text{cov}[\rho_{r,1}\rho_{r,2}]. \quad (38)$$

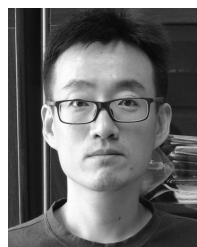
## ACKNOWLEDGMENT

The United States Government is certified to reproduce and distribute reprints for Governmental objectives notwithstanding any copyright notation thereon. The views and conclusions consist of herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Defense Advanced Research Projects Agency and Air Force Research Laboratory or U.S. Government.

## REFERENCES

- [1] T. Gloe, M. Kirchner, A. Winkler, and R. Bohme, "Can we trust digital image forensics?" in *Proc. 15th ACM Int. Conf. Multimedia*, Augsburg, Germany, Sep. 2007, pp. 78–86.
- [2] R. Böhme and M. Kirchner, "Counter-forensics: Attacking image forensics," in *Digital Image Forensics*, H. Sencar and N. Memon, Eds. Berlin, Germany: Springer, 2012.
- [3] M. Kirchner and R. Bohme, "Hiding traces of resampling in digital images," *IEEE Trans. Inf. Forensics Security*, vol. 3, no. 4, pp. 582–592, Dec. 2008.
- [4] G. Cao, Y. Zhao, R. Ni, and H. Tian, "Anti-forensics of contrast enhancement in digital images," in *Proc. 12th ACM Workshop Multimedia Secur. (MM&Sec)*. New York, NY, USA: ACM, 2010, pp. 25–34.
- [5] M. C. Stamm and K. J. R. Liu, "Anti-forensics of digital image compression," *IEEE Trans. Inf. Forensics Security*, vol. 6, no. 3, pp. 50–65, Sep. 2011.
- [6] B. Biggio *et al.*, "Evasion attacks against machine learning at test time," in *Proc. Joint Eur. Conf. Mach. Learn. Knowl. Discovery Databases*. Heidelberg, Germany: Springer, 2013, pp. 387–402.
- [7] M. Fontani and M. Barni, "Hiding traces of median filtering in digital images," in *Proc. 20th Eur. Signal Process. Conf.*, Aug. 2012, pp. 1239–1243.
- [8] P. Comesaña-Alfaro and F. Pérez-González, "Optimal counterforensics for histogram-based forensics," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2013, pp. 3048–3052.
- [9] C. Pasquini, P. Comesaña-Alfaro, and F. Pérez-González, and G. Boato, "Transportation-theoretic image counterforensics to first significant digit histogram forensics," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2014, pp. 2699–2703.
- [10] F. Marra, G. Poggi, F. Roli, C. Sansone, and L. Verdoliva, "Counterforensics in machine learning based forgery detection," *Proc. SPIE*, vol. 9409, Mar. 2015, Art. no. 94090L.
- [11] Z. Chen, B. Tondi, X. Li, R. Ni, Y. Zhao, and M. Barni, "A gradient-based pixel-domain attack against SVM detection of global image manipulations," in *Proc. IEEE Workshop Inf. Forensics Secur. (WIFS)*, Rennes, France, Dec. 2017, pp. 4–7.
- [12] M. Barni, M. Fontani, and B. Tondi, "A universal technique to hide traces of histogram-based image manipulations," in *Proc. ACM Multimedia Secur. Workshop*, Coventry, U.K., Sep. 2012, pp. 97–104.
- [13] M. Barni, M. Fontani, and B. Tondi, "A universal attack against histogram-based image forensics," *Int. J. Digit. Crime Forensics*, vol. 5, no. 3, pp. 1–18, 2013.

- [14] M. Barni, M. Fontani, and B. Tondi, "Universal counterforensics of multiple compressed JPEG images," in *Proc. Int. Workshop Digit. Watermarking*. Taipei, Taiwan: Springer, 2014, pp. 31–46.
- [15] S. Lai and R. Böhme, "Countering counter-forensics: The case of JPEG compression," in *Information Hiding*. Heidelberg, Germany: Springer, 2011, pp. 285–298.
- [16] G. Valenzise, M. Tagliasacchi, and S. Tubaro, "Revealing the traces of JPEG compression anti-forensics," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 2, pp. 335–349, Feb. 2013.
- [17] H. Zeng, T. Qin, X. Kang, and L. Liu, "Countering anti-forensics of median filtering," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2014, pp. 2704–2708.
- [18] A. D. Rosa, M. Fontani, M. Massai, A. Piva, and M. Barni, "Second-order statistics analysis to cope with contrast enhancement counter-forensics," *IEEE Signal Process. Lett.*, vol. 22, no. 8, pp. 1132–1136, Aug. 2015.
- [19] M. Barni, Z. Chen, and B. Tondi, "Adversary-aware, data-driven detection of double JPEG compression: How to make counter-forensics harder," in *Proc. IEEE Int. Workshop Inf. Forensics Secur. (WIFS)*, Dec. 2016, pp. 1–6.
- [20] M. Barni, E. Nowroozi, and B. Tondi, "Higher-order, adversary-aware, double JPEG-detection via selected training on attacked samples," in *Proc. 25th Eur. Signal Process. Conf. (EUSIPCO)*, Aug./Sep. 2017, pp. 281–285.
- [21] M. Barni and F. Pérez-González, "Coping with the enemy: Advances in adversary-aware signal processing," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Vancouver, BC, Canada, May 2013, pp. 8682–8686.
- [22] M. Barni and B. Tondi, "The source identification game: An information-theoretic perspective," *IEEE Trans. Inf. Forensics Security*, vol. 8, no. 3, pp. 450–463, Mar. 2013.
- [23] M. C. Stamm, W. S. Lin, and K. J. R. Liu, "Forensics vs. anti-forensics: A decision and game theoretic framework," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, Kyoto, Japan, Mar. 2012, pp. 1749–1752.
- [24] M. Barni and B. Tondi, "Source distinguishability under distortion-limited attack: An optimal transport perspective," *IEEE Trans. Inf. Forensics Security*, vol. 11, no. 10, pp. 2145–2159, Oct. 2016.
- [25] B. Biggio, G. Fumera, and F. Roli, "Multiple classifier systems for robust classifier design in adversarial environments," *Int. J. Mach. Learn. Cybern.*, vol. 1, nos. 1–4, pp. 27–41, 2010.
- [26] B. Biggio *et al.*, "One-and-a-half-class multiple classifier systems for secure learning against evasion attacks at test time," in *Proc. Int. Workshop Multiple Classifier Syst.* Güzburg, Germany: Springer, 2015, pp. 168–180.
- [27] D. Lowd and C. Meek, "Adversarial learning," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*. New York, NY, USA: ACM, 2005, pp. 641–647.
- [28] M. Barreno, B. Nelson, A. D. Joseph, and J. D. Tygar, "The security of machine learning," *Mach. Learn.*, vol. 81, no. 2, pp. 121–148, 2010.
- [29] B. Biggio, G. Fumera, P. Russu, L. Didaci, and F. Roli, "Adversarial biometric recognition: A review on biometric system security from the adversarial machine-learning perspective," *IEEE Signal Process. Mag.*, vol. 32, no. 5, pp. 31–41, Sep. 2015.
- [30] T. Pevný, P. Bas, and J. Fridrich, "Steganalysis by subtractive pixel adjacency matrix," *IEEE Trans. Inf. Forensics Security*, vol. 5, no. 2, pp. 215–224, Jun. 2010.
- [31] J. Fridrich and J. Kodovský, "Rich models for steganalysis of digital images," *IEEE Trans. Inf. Forensics Security*, vol. 7, no. 3, pp. 868–882, Jun. 2012.
- [32] J.-P. M. G. Linnartz and M. Van Dijk, "Analysis of the sensitivity attack against electronic watermarks in images," in *Proc. Int. Workshop Inf. Hiding*. Heidelberg, Germany: Springer, 1998, pp. 258–272.
- [33] O. Koval, S. Voloshynovskiy, F. Beekhof, and T. Pun, "Security analysis of robust perceptual hashing," *Proc. SPIE*, vol. 6819, Mar. 2008, Art. no. 681906.
- [34] L. Breiman, "Bagging predictors," *Mach. Learn.*, vol. 24, no. 2, pp. 123–140, 1996.
- [35] T. K. Ho, "The random subspace method for constructing decision forests," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 8, pp. 832–844, Aug. 1998.
- [36] B. Biggio, G. Fumera, and F. Roli, "Adversarial pattern classification using multiple classifiers and randomisation," in *Structural, Syntactic, and Statistical Pattern Recognition*. Heidelberg, Germany: Springer, 2008, pp. 500–509.
- [37] J. Kodovský and J. Fridrich, "Steganalysis in high dimensions: Fusing classifiers built on random subspaces," *Proc. SPIE*, vol. 7880, Feb. 2011, Art. no. 78800L.
- [38] R. Villán, S. Voloshynovskiy, O. Koval, F. Deguillaume, and T. Pun, "Tamper-proofing of electronic and printed text documents via robust hashing and data-hiding," *Proc. SPIE*, vol. 6505, Mar. 2007, Art. no. 65051T.
- [39] F. Zhang, P. P. K. Chan, B. Biggio, D. S. Yeung, and F. Roli, "Adversarial feature selection against evasion attacks," *IEEE Trans. Cybern.*, vol. 46, no. 3, pp. 766–777, Mar. 2016.
- [40] I. J. Cox and J.-P. M. G. Linnartz, "Public watermarks and resistance to tampering," in *Proc. Int. Conf. Image Process. (ICIP)*, 1997, pp. 26–29.
- [41] K. Wang, J. J. Parekh, and S. J. Stolfo, "Anagram: A content anomaly detector resistant to mimicry attack," in *Proc. Int. Workshop Recent Adv. Intrusion Detection*. Heidelberg, Germany: Springer, 2006, pp. 226–248.
- [42] S. M. Kay, *Fundamentals of Statistical Signal Processing: Detection Theory*, vol. II. Upper Saddle River, NJ, USA: Prentice-Hall, 1998.
- [43] D. Cozzolino, D. Gragnaniello, and L. Verdoliva, "Image forgery detection through residual-based local descriptors and block-matching," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2015, pp. 5297–5301.
- [44] H. Li, W. Luo, X. Qiu, and J. Huang, "Identification of various image operations using residual-based features," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 28, no. 1, pp. 31–45, Jan. 2018.
- [45] D.-T. Dang-Nguyen, C. Pasquini, V. Conotter, and G. Boato, "Raise: A raw images dataset for digital image forensics," in *Proc. 6th ACM Multimedia Syst. Conf. (MMSys)*. New York, NY, USA: ACM, 2015, pp. 219–224. doi: 10.1145/2713168.2713194.
- [46] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intell. Syst. Technol.*, vol. 2, no. 3, pp. 27:1–27:27, 2011.
- [47] V. N. Vapnik, *The Nature of Statistical Learning Theory*. Berlin, Germany: Springer-Verlag, 1995.
- [48] A. Athalye, N. Carlini, and D. Wagner. (2018). "Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples." [Online]. Available: <https://arxiv.org/abs/1802.00420>
- [49] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [50] N. Carlini and D. Wagner, "Adversarial examples are not easily detected: Bypassing ten detection methods," in *Proc. 10th ACM Workshop Artif. Intell. Secur.* New York, NY, USA: ACM, 2017, pp. 3–14.



**Zhipeng Chen** received the B.S. degree in computer science and technology from the Hebei University of Economics and Business, China, in 2003, and the M.E. degree in computer science and technology from the North China University of Technology, China, in 2008. He is currently pursuing the Ph.D. degree with the Institute of Information Science, Beijing Jiaotong University, Beijing, China. He is also with Tangshan Normal University, China. His research interests include multimedia signal processing, digital forensics, and data hiding.

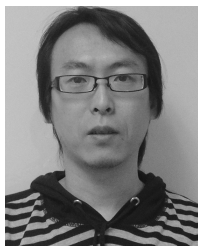


**Benedetta Tondi** (M'16) received the bachelor's and master's degree (*cum laude*) in electronics and communications engineering from the University of Siena, Siena, Italy, in 2012, and the Ph.D. degree in 2016.

From 2014 to 2015, she was a Visiting Student with the Signal Processing in Communications Group, University of Vigo. She is currently a Research Associate with the Department of Information Engineering and Mathematics, University of Siena. She is also an Assistant for the Course of

Information Theory and Coding and Multimedia Security. Her research interest focuses on the application of information theory and game theory concepts to forensics and counter-forensics analysis and more in general on the adversarial signal processing, and on the use of deep learning techniques for multimedia forensics and security.

She is a member of the IEEE Young Professionals and IEEE Signal Processing Society and a member of the National Inter-University Consortium for Telecommunications (CNIT). Since 2019, she has been a member of the IEEE Information Forensics and Security Technical Committee. She is currently involved in the American DARPA MediFor Project and Media Forensics Integrity Analysis Programme from 2016 to 2020. She was a recipient of the Best Student Paper Award at the IEEE International Workshop on Information Forensics and Security (WIFS) 2014, and the Best Paper Award at the IEEE International WIFS 2015 and at The Ninth International Conferences on Advances in Multimedia 2017. She was also a recipient of the 2017 GTTI Ph.D. Award for the Best Ph.D. Theses defended at an Italian University in the areas of communications technologies (signal processing, digital communications, and networking). Since 2018, she has been serving as an Associate Editor for the *EURASIP Journal of Information Security*. She was a Co-Organizer of a Special Session on Adversarial Multimedia Forensics at EUSIPCO 2018. She has been a designated reviewer on the technical program committee of several IEEE Workshops and International Conferences.



**Xiaolong Li** received the B.S. degree from Peking University, China, in 1999, the M.S. degree from Ecole Polytechnique, France, in 2002, and the Ph.D. degree in mathematics from ENS de Cachan (France), in 2006. He was a Post-Doctoral Fellow and then a Researcher with Peking University from 2007 to 2016. He is currently a Professor with the Institute of Information Science, Beijing Jiaotong University. His research interests are image processing and information hiding.



**Rongrong Ni** received the B.S. and Ph.D. degrees from Beijing Jiaotong University (BJTU), Beijing, China, in 1998 and 2005, respectively. Since 2005, she has been on the faculty of the School of Computer and Information Technology and the Institute of Information Science, BJTU, where she has been a Professor since 2013. Her current research interests include image processing, data hiding and digital forensics, pattern recognition, and computer vision. She is the principal investigator of four projects funded by the Natural Science Foundation of China.

She has participated in national and international projects.



**Yao Zhao** (SM'17) received the B.S. degree from the Radio Engineering Department, Fuzhou University, Fuzhou, China, in 1989, the M.E. degree from the Radio Engineering Department, Southeast University, Nanjing, China, in 1992, and the Ph.D. degree from the Institute of Information Science, Beijing Jiaotong University (BJTU), Beijing, China, in 1996. He became an Associate Professor and a Professor with BJTU in 1998 and 2001, respectively. From 2001 to 2002, he was a Senior Research Fellow with the Information and Communication

Theory Group, Faculty of Information Technology and Systems, Delft University of Technology, Delft, The Netherlands. In 2015, he visited the Swiss Federal Institute of Technology, Lausanne (EPFL), Switzerland. From 2017 to 2018, he visited the University of Southern California. He is currently the Director of the Institute of Information Science, BJTU. He is also leading several national research projects from the 973 Program, 863 Program, and the National Science Foundation of China. His current research interests include image/video coding, digital watermarking and forensics, and video analysis and understanding.

Dr. Zhao is an IET Fellow. He serves on the editorial boards of several international journals, including as an Associate Editor of the IEEE TRANSACTIONS ON CYBERNETICS, an Associate Editor of the IEEE SIGNAL PROCESSING LETTERS, the Area Editor of *Signal Processing: Image Communication*, and an Associate Editor of *Circuits, System, and Signal Processing*. He was named a Distinguished Young Scholar by the National Science Foundation of China in 2010 and was elected as a Chang Jiang Scholar of the Ministry of Education of China in 2013.



**Mauro Barni** (M'92–SM'06–F'12) graduated in electronic engineering from the University of Florence in 1991. He received the Ph.D. degree in informatics and telecommunications in 1995. During the last two decades, he has been studying the application of image processing techniques to copyright protection and authentication of multimedia, and the possibility of processing signals that have been previously encrypted without decrypting them. Lately, he has been working on theoretical and practical aspects of adversarial signal processing.

He has authored or co-authored about 300 papers published in international journals and conference proceedings and holds five patents in digital watermarking and image authentication. He has co-authored the book *Watermarking Systems Engineering: Enabling Digital Assets Security and other Applications* (Dekker, Inc., 2004). He participated in several National and European research projects on diverse topics, including computer vision, multimedia signal processing, remote sensing, digital watermarking, and IPR protection.

He is a member of EURASIP. He was a recipient of the Individual Technical Achievement Award of EURASIP for 2016. He has been the Chairman of the IEEE Information Forensic and Security Technical Committee from 2010 to 2011. He was the Technical Program Chair of ICASSP 2014. He was appointed DL of the IEEE SPS for the years 2013–2014. He was Editor-in-Chief of the IEEE TRANSACTIONS ON INFORMATION FORENSICS AND SECURITY from 2015 to 2017. He was the Funding Editor of the *EURASIP Journal on Information Security*. He has been serving as an Associate Editor of many journals, including several IEEE Transactions.