

# On the Security and Applicability of Fragile Camera Fingerprints

Erwin Quiring<sup>1</sup>, Matthias Kirchner<sup>2</sup>, and Konrad Rieck<sup>1</sup>

<sup>1</sup> TU Braunschweig, Germany

<sup>2</sup> Binghamton University, USA

**Abstract.** Camera sensor noise is one of the most reliable device characteristics in digital image forensics, enabling the unique linkage of images to digital cameras. This so-called camera fingerprint gives rise to different applications, such as image forensics and authentication. However, if images are publicly available, an adversary can estimate the fingerprint from her victim and plant it into spurious images. The concept of fragile camera fingerprints addresses this attack by exploiting asymmetries in data access: While the camera owner will always have access to a full fingerprint from uncompressed images, the adversary has typically access to compressed images and thus only to a truncated fingerprint. The security of this defense, however, has not been systematically explored yet. This paper provides the first comprehensive analysis of fragile camera fingerprints under attack. A series of theoretical and practical tests demonstrate that fragile camera fingerprints allow a reliable device identification for common compression levels in an adversarial environment.

**Keywords:** Fragile Camera Fingerprint · PRNU · Authentication.

## 1 Introduction

Minimal, inevitable manufacturing imperfections of digital camera sensors lead to the photo-response non-uniformity (PRNU) signal, a highly unique and reliably detectable camera device characteristic [8]. Similar to a robust digital watermark, the PRNU signal is unnoticeably present in any image taken by the same camera, but differs between images from different cameras. These properties make the PRNU a natural *camera fingerprint*. It has found widespread applications in forensics to attribute digital images to their source camera [8]. Recent works have also proposed to use the PRNU as a means to link mobile device authentication to inherent hardware characteristics of the mobile device [2, 26].

In practice, however, these use cases face the problem of *fingerprint copy-attacks* [10, 19]. If Alice shares images from her camera with the public, Mallory can estimate Alice’s fingerprint, plant it into her images, and pretend that an arbitrary image was captured by Alice’s camera. The so-called *triangle test* [13] detects such attacks *ex post*, but it potentially requires an exhaustive search over all public images shared by Alice. A proactive defense based on the notion of *fragile camera fingerprints* has recently been proposed by Quiring and Kirchner [23] for

scenarios that warrant camera identification from high-quality (uncompressed) images. Here, the camera owner Alice can exploit an asymmetry in the quality of accessible data by only sharing JPEG-compressed images with the public while retaining her uncompressed images private. As a result, she will always be able to provide her full fingerprint from high-quality images when asked to do so. In contrast, Mallory’s estimate of Alice’s fingerprint from public JPEG images will only contain the part that is robust to lossy JPEG compression while lacking the fragile component. A test for the presence of the fragile fingerprint will then prevent Mallory from making an uncompressed image look like one of Alice’s uncompressed images.

In forensics applications, fragile camera fingerprints are of particular relevance to the prevention of fingerprint-copy attacks in support of high-quality image forgeries, which may otherwise convey a false sense of trustworthiness [4]. Equally important, fragile fingerprints are currently the only scalable approach to establish mobile device authentication based on physical camera characteristics that mitigates fingerprint leakage from public images: conducting the triangle test [13] on every authentication attempt is computationally infeasible, and an alternative proposal for a targeted fingerprint-copy attack detector by Ba et al. [2] can be defeated by an adversary with two cameras.

The practical applicability of fragile camera fingerprints in these security-related scenarios crucially depends on their robustness against attacks. As Quiring and Kirchner’s work [23] only provided preliminary results in this regard, this paper sets out to deliver a thorough and more comprehensive security analysis. Specifically, we examine the amount of information that Mallory can estimate, recover and exploit in a series of theoretical and empirical considerations. First, we analytically derive an upper bound on the correlation between Alice’s and Mallory’s fingerprint estimates with respect to the JPEG quality of publicly shared images Mallory has access to. Second, to test for dependencies beyond linear correlation, a kernel statistical test is used to assess whether Alice’s fragile fingerprint is statistically independent of Mallory’s fingerprint. Third, we demonstrate that practical attempts to recover quantized JPEG coefficients from potentially remaining dependencies do not increase Mallory’s ability to mount successful attacks. Fourth, we test the resistance of fragile fingerprints against practical fingerprint-copy attacks. We finally illustrate that fragile fingerprints and the triangle test are a powerful combination in forensics applications.

The rest of this paper is organized as follows. Section 2 reviews the background of sensor noise forensics before Section 3 discusses fragile fingerprints and their possible applications. Section 4 provides a comprehensive analysis of Mallory’s attack surface, while Section 5 reports on experiments around the applicability of fragile fingerprints. Section 6 concludes the paper.

## 2 Background

Before introducing fragile fingerprints, we give a short primer on camera identification, the possible fingerprint-copy attack, and the triangle test as defense.

Throughout our work, the *notation* is as follows: vectors and matrices are set in boldface font. Operations on vectors and matrices are point-wise if not stated otherwise; the operator  $\bullet$  denotes matrix multiplication.

## 2.1 Camera Identification from Sensor Noise Fingerprints

Due to sensor element manufacturing imperfections, each camera image does not only contain the original noise-free image content  $\mathbf{I}^0$ , but also the PRNU  $\mathbf{K}$  as a camera-specific, multiplicative noise factor. A common simplified model of the image capturing process assumes the final image  $\mathbf{I}$  to take the form [8]

$$\mathbf{I} = \mathbf{I}^0 + \mathbf{I}^0 \mathbf{K} + \mathbf{\Gamma}, \quad (1)$$

where  $\mathbf{\Gamma}$  reflects a variety of other additive noise terms. Due to its multiplicative nature, the PRNU is not present in images with dark scene contents (i. e.,  $\mathbf{I}^0 \approx \mathbf{0}$ ). Extensive experiments have demonstrated that the PRNU factor  $\mathbf{K}$  represents a unique and robust camera fingerprint [14] that can be estimated from a number of images  $\mathbf{I}_1, \dots, \mathbf{I}_N$  taken with a given camera of interest. The standard approach utilizes a denoising filter  $F(\cdot)$  and models noise residuals  $\mathbf{W}_k = \mathbf{I}_k - F(\mathbf{I}_k)$  as [8]

$$\mathbf{W}_k = \mathbf{I}_k \mathbf{K} + \mathbf{\Theta}_k. \quad (2)$$

Modeling noise  $\mathbf{\Theta}$  subsumes  $\mathbf{\Gamma}$  and residues of the image content due to inherent imperfections of the denoising filter in separating image content from noise. Adopting an i.i.d. Gaussian noise assumption for  $\mathbf{\Theta}$ , the maximum likelihood (ML) estimator of  $\mathbf{K}$  is [8]

$$\hat{\mathbf{K}} = \left( \sum_{k=1}^N \mathbf{W}_k \mathbf{I}_k \right) \cdot \left( \sum_{k=1}^N (\mathbf{I}_k)^2 \right)^{-1}. \quad (3)$$

A more simple estimator takes the pixel-wise average of the noise residuals [19]. A post-processing step is recommended to clean  $\hat{\mathbf{K}}$  from so-called non-unique artifacts, e. g., due to demosaicing or lens distortion correction [8, 11, 12]. Given a query image  $\mathbf{J}$  of unknown provenance, camera identification then works by computing the residual  $\mathbf{W}_J = \mathbf{J} - F(\mathbf{J})$ , and evaluating its similarity to a camera fingerprint estimate against a set threshold  $\tau$ ,

$$\phi_{\mathbf{W}_J, \mathbf{J}\hat{\mathbf{K}}} = \text{sim}(\mathbf{W}_J, \mathbf{J}\hat{\mathbf{K}}) \geq \tau. \quad (4)$$

Suitable similarity measures for this task are normalized correlation or peak-to-correlation energy (PCE) [8, 19].

## 2.2 Fingerprint-Copy Attack

Following the procedure described in Section 2.1, Mallory may obtain an estimate of Alice's camera fingerprint from a set of  $N_E$  publicly available images. Denoting

this estimate  $\hat{\mathbf{K}}_E$ , Mallory can then attempt to make an arbitrary image  $\mathbf{J}$  look as if it was captured by Alice’s camera. The multiplicative nature of PRNU suggests a *fingerprint copy attack* of the form [19]

$$\mathbf{J}' = [\mathbf{J}(1 + \alpha\hat{\mathbf{K}}_E)], \quad (5)$$

with  $\alpha > 0$  being the scalar fingerprint strength parameter. Attacks of this type have been demonstrated to be effective, in the sense that they can successfully mislead a camera identification algorithm in the form of Equation (4). However, the attack’s success generally depends on a good choice of  $\alpha$ : too low values mean that the bogus image  $\mathbf{J}'$  may not be assigned to Alice’s camera; a too strong embedding will make the image appear suspicious [13, 20]. In practical scenarios, Mallory may have to apply further processing to make her forgery more compelling, e. g., removing the genuine camera fingerprint [7, 16], synthesizing demosaicing artifacts [17], and removing or adding traces of JPEG compression [25].

### 2.3 Triangle Test

Under realistic assumptions, it is impossible to prevent Mallory from forcing a high similarity score in Equation (4) for arbitrary images from a foreign camera. Yet Alice can utilize that noise residuals computed with practical denoising filters will always contain remnants of image content to establish that image  $\mathbf{J}'$  underwent a fingerprint-copy attack [13]. The key observation here is that the already existing similarity between a noise residual  $\mathbf{W}_I$  from an image  $\mathbf{I}$  taken with Alice’s camera and the noise residual  $\mathbf{W}_{J'}$  due to a common PRNU term will be slightly increased by some shared residual image content, if  $\mathbf{I}$  contributed to Mallory’s fingerprint estimate  $\hat{\mathbf{K}}_E$ . Alice can thus test which of her public images have been used by Mallory to mount the attack by evaluating whether the similarity of their noise residuals  $\mathbf{W}_I$  with  $\mathbf{W}_{J'}$  is suspiciously large.

Because the additional correlation imposed by shared image content is generally rather weak and also varies with macroscopic image characteristics, Goljan et al. [13] propose a *triangle test* to calibrate the test statistic. Specifically, the test does not only consider the observed correlation  $\nu_{\mathbf{W}_I, \mathbf{W}_{J'}}$  between residuals  $\mathbf{W}_I$  and  $\mathbf{W}_{J'}$ , but it also employs a correlation predictor to estimate the correlation  $\tilde{\nu}_{\mathbf{W}_I, \mathbf{W}_{J'}}$  between  $\mathbf{W}_I$  and  $\mathbf{W}_{J'}$  if image  $\mathbf{I}$  had *not* participated in the computation of  $\hat{\mathbf{K}}_E$ . This predictor takes the correlation between Alice’s own fingerprint and both  $\mathbf{W}_I$  and  $\mathbf{W}_{J'}$  into account—hence the name triangle test. Assuming a linear relationship between the observed and the predicted correlation, the proposed test then evaluates

$$\nu_{\mathbf{W}_I, \mathbf{W}_{J'}} - \theta\tilde{\nu}_{\mathbf{W}_I, \mathbf{W}_{J'}} - \mu \geq t \quad (6)$$

for a suitably chosen threshold  $t$ . The parameters  $\theta$  and  $\mu$  are estimated from a set of *safe* images, for which it can be guaranteed that they have not been used by Mallory. We refer to Goljan et al. [13] for a detailed exposition of the correlation predictor and the parameter estimators. The test statistic in Equation (6) is

expected to have zero mean when two noise residuals share only a common PRNU term. A larger difference indicates an additional shared term, possibly due to image  $I$ 's involvement in the attack. Observe that Alice may have to test *all images* ever made public by her as part of a comprehensive defense. We finally point out that a number of fingerprint-copy attack variations have been proposed recently that are reportedly less likely to be exposed by the triangle test [e.g. 20].


### 3 Fragile Camera Fingerprint

As a novel and proactive defense against fingerprint-copy attacks, Quiring and Kirchner [23] introduce the notion of fragile camera fingerprints that vanish under lossy JPEG compression. The idea is based on two mild assumptions: 1) Alice's device supports capturing images in uncompressed format, which is true nowadays for many devices operating under mobile platforms, such as iOS and Android; 2) Alice only shares JPEG images with the public, which is already today's quasi-standard for image online storage and sharing. When combined, these two assumptions allow Alice to effectively exploit an asymmetry in the quality of data access. With full access to her camera, Alice is always in the position to present a fingerprint estimate  $\hat{\mathbf{K}}$  from uncompressed images while Mallory is restricted to estimate  $\hat{\mathbf{K}}_E$  from JPEG-compressed images.

On a technical level, the concept of fragile camera fingerprints exploits the lossy nature of JPEG compression. JPEG maps each non-overlapping  $8 \times 8$  pixel block in an image to  $8 \times 8$  discrete cosine transform (DCT) coefficients. Each of the 64 coefficients quantifies the influence of a particular frequency subband and will be quantized based on an  $8 \times 8$  quantization table with quantization factors for the 64 DCT subbands. Larger quantization factors mean that the DCT coefficients in the corresponding subband are more likely to be quantized to zero. Quantization factors generally increase with decreasing JPEG quality and grow towards the bottom right corner of the quantization table to suppress high-frequency image details more aggressively.

In consequence, Mallory's camera fingerprint estimate from JPEG-compressed images will be strongly distorted in the high-frequency DCT subbands due to larger quantization errors. If the quantization is too strong, Mallory's images will lack high-frequency content altogether and so will her fingerprint estimate. In other words, her estimate only comprises the fingerprint component that is robust to JPEG compression. A fingerprint estimate from uncompressed images is in turn distributed almost evenly over all subbands [23]. Hence, Alice has access to a *fragile camera fingerprint*, computed from the high-frequency subbands only.

To obtain the fragile part, it is instructive to define a mode-selective highpass filter  $H_c(\cdot)$ . Based on a binary multiplicative mask  $\mathbf{H}_c = [h_{i,j}], 1 \leq i, j \leq 8$ , the filter retains a defined set of DCT subbands and sets all other subband coefficients to zero. Alice should choose  $\mathbf{H}_c$  depending on the maximum JPEG quality of her published images. For a sufficiently conservative choice, she can assume that the retained subbands are available exclusively to her. Quiring and Kirchner [23] propose to parameterize  $\mathbf{H}_c$  with a cut-off along the  $(-c)$ -th anti-diagonal of the



**Fig. 1.** Fragile fingerprint computation based on a subband-selective filter  $H_1(X)$ : Each pixel block  $X$  is mapped to its DCT representation  $Y$ , element-wise multiplied by a binary mask  $H_1$ , and transformed back to the spatial domain to give  $\tilde{X}$ .

DCT coefficient matrix,

$$h_{i,j} = [(i + j - 8 - c) > 0], \quad (7)$$

where  $[\cdot]$  denotes the Iverson bracket. Figure 1 summarizes the internal steps of  $H_c(\cdot)$  for cut-off parameter  $c = 1$ , which retains all DCT subbands in the lower right triangle. Equipped with  $H_c$ , a refined similarity test of the form

$$\phi_{\mathbf{W}_J, \mathbf{J}\hat{\mathbf{K}}}(c) = \text{sim} \left( H_c(\mathbf{W}_J), H_c(\mathbf{J}\hat{\mathbf{K}}) \right) \geq \tau. \quad (8)$$

then establishes camera identification from fragile fingerprints, which is of particular relevance in the following two application scenarios.

*Digital Image Forensics.* Testing for the presence of camera fingerprints facilitates device identification and image manipulation detection in forensic applications [8]. Fragile camera fingerprints can benefit scenarios that warrant the analysis of high-quality images, for instance when uncompressed images ought to be presented as a source of particularly high trustworthiness. In this case, Alice can establish that a spurious image was not captured by her camera by having kept her uncompressed images private. She presents the fragile fingerprint when needed, potentially in combination with cryptographic safeguards [21] or in some form of zero-knowledge proof to further secure her fragile fingerprint from leakage.

*Mobile Device Authentication.* Camera fingerprints have been proposed as building blocks for augmenting mobile device authentication schemes with physical hardware characteristics [2, 5, 26]. Yet like with device signatures from other types of hardware sensors, the vulnerability to fingerprint-copy/spoofing attacks is of particular concern [1]. To the best of our knowledge, the concept of fragile fingerprints is the only existing approach that would address the problem in a proactive and scalable manner. Performing a triangle test upon every authentication attempt is computationally infeasible, let alone that Alice may object to the idea of sharing all her images with the service she wants to authenticate to.

Ba et al. [2] attempt to work around that issue in their authentication protocol by requiring the user to take two images of different visual codes during the identification phase. Similar to the triangle test, the reasoning is that two spoofed images will correlate implausibly strongly as they do not only share Alice's

fingerprint but also the one from Mallory’s camera. This measure can be easily circumvented by an attacker who uses two different devices to take the respective pictures however. In this way, Mallory prevents the additionally shared signal.

In summary, fragile camera fingerprints enable novel applications in image forensics and mobile device authentication. Their applicability, however, depends on the robustness against attacks. In the remainder of this work, we thus perform a comprehensive security analysis.

## 4 Security Analysis

A secure application of fragile camera fingerprints demands that Mallory cannot estimate the fragile fingerprint from JPEG-compressed images. We guide our analysis along the following three questions:

- (Q1) Can we bound the quality of Mallory’s fingerprint estimate  $\hat{K}_E$ ?
- (Q2) Can Mallory improve her fingerprint estimate by exploiting the quantized high-frequency or robust low-frequency information?
- (Q3) Can Mallory perform a successful fingerprint-copy attack?

### 4.1 Datasets and Experimental Setup

Where empirical tests are warranted, we adopt the setup described by Quiring and Kirchner [23]. The dataset consists of images from the Dresden Image Database [9] (DDB) and the RAISE Image Database [6], cf. Table 1. In particular, we use 25 homogeneously lit flat field images of each DDB camera to obtain uncompressed fingerprint estimates. 1442 natural images serve as benchmark data. We present aggregated results over the six cameras in the following, as all gave similar results. The RAISE database only provides natural images. We randomly select 300 images for fingerprint computation, leaving us with 4648 images for a benchmark set that facilitates the study of attacks where Mallory has access to a large number of public images. Note that the usage of 300 natural images for fingerprint estimation can be attributed to the heterogenous content of natural images. In an authentication scenario, a user can be asked to take a much smaller number of suitable images (e. g. a white wall) without nuisance image content. If not stated otherwise, we use the standard Wavelet denoiser to obtain noise residuals [19] and the ML formulation in Equation (3) to estimate fingerprints.

**Table 1.** Number of images per test set and camera.

Database	Camera model	Fingerprint estimate		Benchmark data	
		Camera 0	Camera 1	Camera 0	Camera 1
Dresden [9]	Nikon D70	25	25	175	188
	Nikon D70s	25	25	175	174
	Nikon D200	25	25	360	370
RAISE [6]	Nikon D7000	300	—	4648	—

In order to also guide our evaluation at the practically used JPEG quality, we collected over 1.4 Million JPEG images from Twitter, Instagram, Imgur, Deviantart and Flickr. The average JPEG quality is 83.5 with a standard deviation of 9.2. This, for instance, fits to recommendations from Flickr, Wikimedia and the official Android documentation that recommend qualities less than 90.

#### 4.2 (Q1) Analytical Quality of Fingerprint Estimation

Our first objective is to establish a bound on the quality of Mallory’s fingerprint estimate from Alice’s camera irrespective of a concrete image data set to reflect Mallory’s chances of performing a successful fingerprint-copy attack. In particular, we adopt the notion of *quality of fingerprint estimation* by Goljan et al. [13]. We derive an analytical expression for the correlation between Alice’s fingerprint from uncompressed images,  $\hat{\mathbf{K}}$ , and Mallory’s fingerprint from compressed images,  $\hat{\mathbf{K}}_E$ ,

$$\text{cor}(H_c(\hat{\mathbf{K}}), H_c(\hat{\mathbf{K}}_E)). \quad (9)$$

This quantity can be seen as a simplified version of the similarity measure in Eq. (8) for images taken under ideal conditions, e.g. homogeneously lit. As we focus on high-frequency subbands only, less image content disturbs the fingerprint calculation. Thus,  $H_c(\mathbf{J}\hat{\mathbf{K}})$  resembles  $H_c(\hat{\mathbf{K}})$ , as well  $H_c(\mathbf{W}_J) \approx H_c(\hat{\mathbf{K}}_E)$ .

We make three assumptions to simplify the calculation. First, Mallory computes her fingerprint from the same, but compressed, image set that Alice uses for her uncompressed estimate. This will yield a loose upper bound for cases where Mallory obtains a different JPEG-compressed image set that Alice has not used. Moreover, we apply a simple fingerprint estimator that takes the pixel-wise average of noise residuals. Finally, we assume a negligible correlation between individual DCT subbands and across images. The first and the second assumption imply that the fingerprint calculation can be modeled as pixel-wise averaging. Denote  $\mathbf{X}_i$  the  $i$ -th uncompressed image from Alice’s camera and  $\tilde{\mathbf{X}}_i$  its JPEG-compressed version to rewrite Equation (9) as

$$\text{cor}(H_c(\hat{\mathbf{K}}), H_c(\hat{\mathbf{K}}_E)) = \text{cor}\left(\sum_i H_c(\mathbf{X}_i), \sum_i H_c(\tilde{\mathbf{X}}_i)\right). \quad (10)$$

Appendix A establishes that the *sample correlation coefficient* based on Equation (10),  $r(c)$ , can be computed in the DCT domain directly. We write  $\mathbf{Y}_i$  and  $\tilde{\mathbf{Y}}_i$  for the DCT representations of  $\mathbf{X}_i$  and  $\tilde{\mathbf{X}}_i$ , respectively, to obtain

$$r(c) \cong \text{cor}\left(\sum_i (\mathbf{H}_c \mathbf{Y}_i), \sum_i (\mathbf{H}_c \tilde{\mathbf{Y}}_i)\right). \quad (11)$$

The coefficient is parametrized by the cut-off  $c$  from Equation (7). High-pass filter  $H_c$  is now made explicit through the DCT mask  $\mathbf{H}_c$  (see Section 3), yielding a convenient formulation to compute the sample correlation coefficient between Alice’s and Mallory’s fingerprint directly in the DCT domain. This formulation thus allows us to use known statistical distribution models for DCT coefficients.

We continue to derive the *population correlation coefficient* by assuming a Laplacian distribution for the AC DCT coefficients [24]. Specifically, denote  $Y_{i,s}$



the random variable representing the  $s$ -th subband of the  $i$ -th uncompressed image. Equivalently, denote  $\tilde{Y}_{i,s}$  the respective quantized counterpart to reflect the effect of JPEG compression on  $Y_{i,s}$ . Appendix B establishes the general relation between the two random variables in terms of their covariance, which can be expressed solely on the basis of the distribution of the uncompressed variable  $Y_{i,s}$ . We highlight this by defining  $\text{Cov}^+(Y_{i,s}) = \text{Cov}(Y_{i,s}, \tilde{Y}_{i,s})$ . A similar derivation for the variance yields  $\text{Var}^+(Y_{i,s}) = \text{Var}(\tilde{Y}_{i,s})$ . Appendix B shows how aggregating these quantities over the various subbands  $s \in S_c$  as specified by filter  $H_c$  of various images leads to the following formulation for the population correlation coefficient:

$$\rho(c) \cong \frac{\sum_i \sum_{s \in S_c} \text{Cov}^+(Y_{i,s})}{\sqrt{\sum_i \sum_{s \in S_c} \text{Var}(Y_{i,s})} \sqrt{\sum_i \sum_{s \in S_c} \text{Var}^+(Y_{i,s})}}. \quad (12)$$

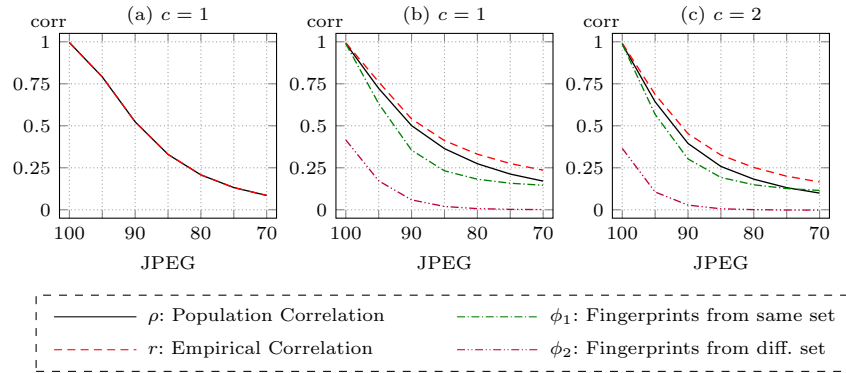
This equation is a first step towards an analytical understanding of the impact of JPEG-induced quantization on the ability to estimate fragile camera fingerprints. Specifically, it allows Alice to deduce the expected correlation of Mallory’s fingerprint with her fingerprint based on general DCT distribution assumptions. Note that the derived correlation is computed under the assumption of a strong attacker: Mallory bases her fingerprint estimation on the same image set as Alice; her images only differ in that they are JPEG-compressed. Consequently, Equation (12) can serve as upper bound for the more realistic scenario when Mallory has only access to a different image set. Alice, as camera owner, will always be able to create new images for her fingerprint. The next section demonstrates the validity of our analytical derivation under practical conditions for both scenarios.

### 4.3 (Q1) Empirical Quality of Fingerprint Estimation

We start with the quality of Mallory’s fingerprint estimate when Alice and Mallory operate on the same image set, and then continue with different image sets.

*Same image sets.* In a first step, we compute the population correlation coefficient  $\rho$  from Equation (12) on a set of 250 synthetic images. Each image follows a zero-mean Laplacian distribution with a randomly generated scale parameter. This allows us to examine  $\rho$  on idealized conditions. We compare  $\rho$  with its empirical counterpart,  $r$ , as given in Equation (11). Figure 2a shows that the two derived quantities are consistent under varying JPEG compression levels.

In the next experiment, we use natural images from the Nikon D7000. We give Mallory access to  $N_E = 250$  JPEG-compressed images, derived from the same set that Alice uses for her fingerprint. Varying the JPEG quality and cut-off parameter  $c$ , we compute  $\rho$  and  $r$ . The computation of  $\rho$  involved a standard maximum likelihood estimator to obtain the Laplace scale parameter for each DCT subband per image. For benchmark purposes, we also include the sample correlation coefficient  $\phi_1$  between Alice’s and Mallory’s fingerprint, both calculated with the ML formulation in Equation (3). We repeat the experiments five times and report averaged results for  $\rho$ ,  $r$  and  $\phi_1$  in Figures 2b–2c for  $c \in \{1, 2\}$ .



**Fig. 2.** Quality of fingerprint estimation. Results from (a)  $N_E = 250$  synthetic images and (b)-(c)  $N_E = 250$  natural Nikon D7000 images.

The curves resemble each other reasonably well, with  $r$  generally predicting a slightly higher fingerprint quality than  $\rho$  due to the implied independence assumptions in the latter. As  $c$  increases,  $\rho$  slowly approaches  $\phi_1$ . This indicates that the analytically derived  $\rho$  is a good approximation of Mallory’s fingerprint quality under idealized conditions particularly in high-frequency DCT subbands.

*Different image sets.* In a more realistic scenario, we assume Alice and Mallory to work on different image sets. Figures 2b and 2c thus include  $\phi_2$ , the corresponding sample correlation coefficient between Alice’s and Mallory’s fingerprints as obtained with the ML estimator, averaged over five randomly compiled JPEG image sets of size  $N_E = 250$  that Alice has not used for computing  $\hat{\mathbf{K}}$ . Alice’s camera-specific fingerprint from 300 images was kept constant throughout all repetitions. Not surprisingly, the population correlation coefficient  $\rho$  is a loose upper bound to the observed correlation  $\phi_2$  when Mallory operates on a different image set:  $\phi_2$  approaches zero quickly with increasing  $c$  and decreasing JPEG quality. Appendix C gives additional insights by reporting the correlations for a much larger number of JPEG images,  $N_E$ . Mallory’s fingerprint quality increases only slowly with the number of available JPEG images. For a suitable combination of JPEG quality and cut-off parameter  $c$ , the correlation remains extremely small. As a result, less restrictive quality and cut-off parameters are possible compared to the contrived situation where Alice and Mallory access the same images.

*Analysis summary.* Overall, strong guarantees for a scenario where Mallory has access to JPEG-compressed versions of the very images Alice used for her fingerprint are possible for JPEG qualities of 70 or smaller. In a more realistic scenario with different image sets, a secure operation is already possible with JPEG quality factors 90 or lower. For JPEG 85—the average quality factor on various image platforms (see Section 4.1)—Alice may choose  $c \geq 3$  to ensure a reliable identification in an adversarial environment.

#### 4.4 (Q2) Independence Test

The previous section has analyzed the correlation between Alice’s and Mallory’s high-frequency fingerprint estimates—deriving first bounds when Alice’s fingerprint remains private. We continue with this analysis under the scenario of different image sets in the following sections.

Quiring and Kirchner have shown [23] that the high-frequency pixel part kept by  $\mathbf{H}_c$  is uncorrelated to the complementary low-frequency part kept by  $\mathbf{L}_c = \mathbf{H}_c \text{ XOR } 1$ . Consequently, Mallory cannot exploit linear dependencies between her robust low-frequency fingerprint and Alice’s fragile fingerprint. However, correlation does not cover all modes of dependence. In this section, we thus examine if Mallory can exploit non-linear dependencies and conduct a kernel statistical test of independence. In particular, we choose the Hilbert-Schmidt independence criterion<sup>3</sup> (HSIC)[15]. In simplified terms, this test maps the possibly non-linear dependencies to a linear space where independence is tested. The test is consistent in the sense that the level of *alpha* controls the type I error (detects dependence although independence is true) while the type II errors goes to zero for an increasing sample size [15].

We consider the following two scenarios. First, we test if Alice’s high frequency fingerprint is independent to Mallory’s high-frequency fingerprint from JPEG-compressed images:

$$\mathcal{H}_0 : H_c(\hat{\mathbf{K}}) \perp H_c(\hat{\mathbf{K}}_E). \quad (13)$$

Equivalently, the second scenario tests if Alice’s high frequency fingerprint is independent to Mallory’s full fingerprint from JPEG-compressed images:

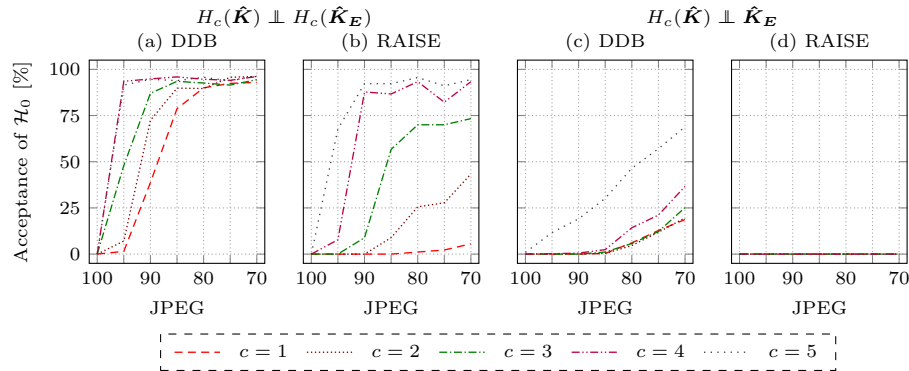
$$\mathcal{H}_0 : H_c(\hat{\mathbf{K}}) \perp \hat{\mathbf{K}}_E. \quad (14)$$

We grant Mallory access to  $N_E = 150$  images of each DDB camera and  $N_E = 1000$  RAISE images. We aggregate results over ten randomly compiled sets of size  $N_E$ .

For both scenarios, Figure 3 depicts the observed  $\mathcal{H}_0$  acceptance rates, i. e. the percentage of cases for which we cannot detect a measurable dependence between the two quantities under test. This rate increases with lower JPEG qualities or larger cut-off parameters  $c$ . In the first scenario, the test statistic suggests independence at a considerable rate for  $c = 4$  and JPEG quality 90 for all cameras. Interestingly, the second scenario—where Alice’s high frequency fingerprint is tested against Mallory’s full fingerprint—is characterized by a lower rate of independence. A comparison of both scenarios thus suggests that remaining dependencies may result from the low-frequency part. This dependency would have to be non-linear, as the low- and high-frequency signal are not correlated to each other. As we show in the next section, it is unclear how Mallory can exploit these potentially remaining dependencies in practice, however.

We surmise that the notably less conclusive results on the RAISE data are due to non-trivial remnants of image content in the noise residuals. In

<sup>3</sup> We use a Gaussian kernel, an alpha value of 0.05, and split the images into  $320 \times 320$  pixel blocks with varying offsets to keep the sample size manageable.



**Fig. 3.** Kernel statistical test of independence. Plots (a) and (b) depict the first scenario; Plots (c) and (d) the second scenario for both databases.

contrast to the Dresden database, Alice’s fingerprint is here calculated from natural instead of homogeneously lit images, raising the bar for establishing independence considerably.

*Analysis summary.* The chosen HSIC test establishes statistical independence for suitable JPEG and cut-off parameters, which gives a strong evidence that Mallory cannot exploit any dependence to recover Alice’s fingerprint. Considering the high-frequency signals, Alice may choose  $c \geq 3$  for quality factor 85.

#### 4.5 (Q2) DCT Recovery

In the next experiment, we examine if Mallory can exploit remaining dependencies to recover DCT coefficients. Although a DCT coefficient that is quantized to zero does not reveal information about the fingerprint, non-zero coefficients may leak information at least with their sign. By averaging enough images, Mallory may thus obtain a coarse fingerprint estimate. We test below if Mallory can improve her fingerprint by recovering DCT coefficients that were quantized to zero.

We adapt the systematic approach by Li et al. [18], since it is in principle also applicable to the recovery of high-frequency DCT coefficients. The recovery is a linear optimization problem with the objective to minimize the spatial distance of neighboring pixels within and across the  $8 \times 8$  pixel blocks from JPEG compression. The first constraint is that the recovered pixel values must correspond to their DCT coefficients. Second, DCT coefficients that should not be recovered are fixed. Finally, the pixel and DCT coefficients have to be within their dynamic range. The optimization problem can be summarized as

$$\min \sum_{l,l'} |\mathbf{X}(l) - \mathbf{X}(l')| \quad (15)$$

$$\text{s.t. } \mathbf{X} - \mathbf{D}^\top \bullet \mathbf{Y} \bullet \mathbf{D} = \mathbf{0}, \quad (16)$$

$$\mathbf{Y}(s) = \mathbf{Y}^*(s), \quad (17)$$

$$\mathbf{X}(l) \in [x_{\min}, x_{\max}], \mathbf{Y}(s) \in [y_{\min}, y_{\max}], \quad (18)$$

**Table 2.** Contingency table of DCT recovery from 50 Nikon D70 images

		Fraction Predicted Sign					
		JPEG 100			JPEG 95		
		neg	zero	pos	neg	zero	pos
True Sign	neg	0.08	0.09	0.05	0.11	0.15	0.09
	zero	0.16	0.08	0.16	0.07	0.11	0.07
	pos	0.05	0.09	0.24	0.09	0.15	0.17

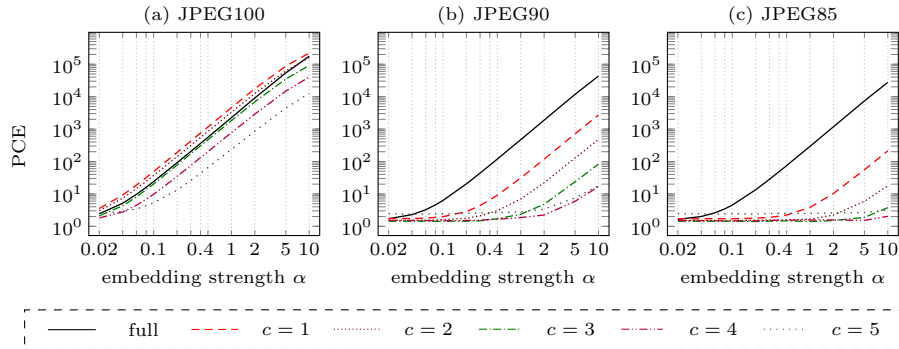
where  $l$  and  $l'$  are the indices of neighboring pixels in the spatial domain,  $\mathbf{D}$  denotes the DCT transformation matrix, and  $s$  is the index of a DCT subband. The second constraint fixes with  $\mathbf{Y}^*$  all DCT coefficients that are not part of the subbands retained by filter  $H_c$  or are non-zero in the subbands retained by  $H_c$ . As a result, we recover only zero-valued DCT coefficients that  $H_c$  retains. For each image, and for each  $8 \times 8$  pixel block, we set up such an optimization problem and include its direct neighboring blocks.

We report results for 50 images from a Nikon D70 over the JPEG qualities 100 and 95 as well as the cut-off frequency  $c = 1$ . The performance does not change considerably for smaller JPEG qualities or larger cut-off frequencies and thus are omitted. Table 2 depicts a contingency table that summarizes the frequency of correctly predicted signs. This is the case when the sign of the predicted DCT coefficient equals the sign from the corresponding original uncompressed image or both the predicted and uncompressed coefficient lie in the zero range  $[-0.25, 0.25]$ . Even for JPEG quality 100, the recovery cannot reliably predict the sign. The correct distinction drops further for a smaller JPEG quality and tends towards a random classifier. In each case, the correlation to Alice’s fragile fingerprint decreases when Mallory uses the recovered images for her estimate. In contrast, the recovery of low-frequency subbands is successful with an average recovery rate of 70%. However, only the correlation to Alice’s low-frequency fingerprint increases in our experiments.

*Analysis summary.* The recovery of the correct sign is partly possible for low-frequency subbands, while the recovery of high-frequency subbands is already difficult for JPEG quality 100. Mallory can thus not improve her estimate of Alice’s fragile fingerprint through a DCT recovery.

#### 4.6 (Q3) Fingerprint-Copy Attack

We finally consider a realistic fingerprint-copy attack where Mallory plants her calculated fingerprint estimate  $\tilde{\mathbf{K}}_E$  from Alice’s camera into 100 randomly chosen uncompressed images taken by a different camera (see Section 2.2). Figure 4 depicts the average PCE values with respect to the embedding strength for varying JPEG qualities. We present results only for the Nikon D7000 from the RAISE database with  $N_E = 4648$ . This allows us to depict the effect when Mallory uses a large number of public images. We refer to Quiring and Kirchner [23] for results from the DDB, which are similar to the results reported here.



**Fig. 4.** Fingerprint-copy attack with fragile fingerprints. Average PCE values as a function of the embedding strength  $\alpha$  with  $N_E = 4648$  (Nikon D7000) for different JPEG qualities.

As expected, high-quality JPEG 100 images enable Mallory to perform a successful attack due to the negligible quantization (Fig. 4(a)). The situation is substantially different with stronger compression. For JPEG 90, only the full fingerprint gives high PCE values for small embedding strengths. Yet, larger cut-offs demand extremely strong embeddings to achieve high PCEs. For JPEG quality 85, no choice of  $\alpha$  will produce Mallory’s desired result with  $c \geq 3$ .

*Analysis summary.* Fragile fingerprints allow a secure identification starting from JPEG 90 and lower. In accordance to our results from previous sections, no choice of  $\alpha$  will allow an attack with  $c \geq 3$  for quality factor 85.

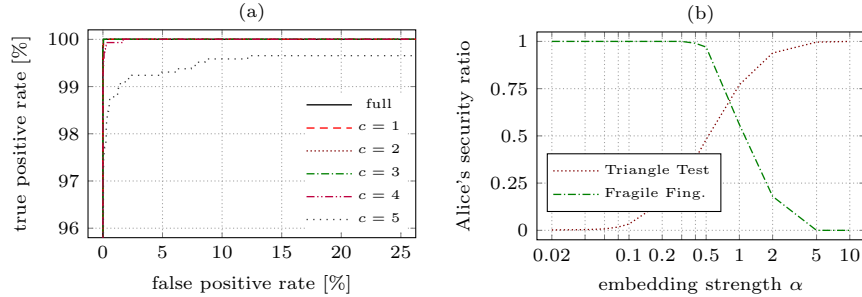
## 5 Application Analysis

We finally examine the application of fragile sensor noise fingerprints. First, we verify that they are still discriminative enough to distinguish different cameras. Second, we compare them with the triangle test against fingerprint-copy attacks.

### 5.1 Camera Identification

In the following, we show that fragile fingerprints allow a reliable camera identification compared to traditional *full* camera fingerprints. We only consider uncompressed images here by the very nature of fragile fingerprints. The PCE is used as similarity measure for images of each camera (true positives) and all remaining natural images from the Dresden Image Database (true negatives). Figure 5a shows the ROC curves for different cut-off frequencies  $c$ —aggregated over 1442 images from the six DDB cameras. The full frequency range is included for comparison.

Although a fragile fingerprint with  $c = 1$  employs only 28 DCT coefficients in each block, it achieves the same detection performance as the full fingerprint with



**Fig. 5.** Applications. Plot (a) shows the camera identification of uncompressed images. Plot (b) depicts the defense performance against fingerprint-copy attack using the triangle test and fragile fingerprints (JPEG quality 90,  $c = 1$ ).

64 coefficients. An almost perfect detection is possible with  $c \leq 4$  for the Dresden database. The results for the Nikon D7000 camera are comparable for  $c \leq 3$ . We contribute this smaller choice of  $c$  to a more perturbed fingerprint estimate of this camera—due to more image content in the respective noise residuals.

*Analysis summary.* Fragile fingerprints allow a reliable camera identification. Together with our security analysis, for a common JPEG quality factor of 85, Alice can choose  $c = 3$  to achieve both a reliable camera identification and attack resistance.

## 5.2 Comparison with Triangle Test

While the triangle test cannot be recommended for authentication, it is a reasonable defense in digital image forensics. Our final experiment highlights its powerful combination with fragile fingerprints in forensic applications against fingerprint-copy attacks. Our previous results underline that remaining fingerprint information after quantization are usable for large embedding strengths with too small cut-off parameters ( $c \leq 2$ ). However, the triangle test shows its strengths exactly in these cases, as the additional residual image content from the forgery process emerges more clearly with a larger embedding strength [13].

In the following, we assume Mallory to have access to  $N_E = 150$  public JPEG images with quality factor 90 from Alice’s Nikon D200 camera. Mallory embeds her spoofed fingerprint into an uncompressed image from another camera while varying  $\alpha$  as defined in the previous section. On the defender side, the linear parameters  $\theta$  and  $\mu$  of the test statistic are estimated from 200 images that Mallory has not used, cf. Eq. (6). We set the threshold  $t$  such that the false alarm probability is  $10^{-3}$ . Finally, Alice reports her security ratio: the percentage of images that Mallory has used and that are correctly marked as those. We repeat the process over 100 randomly chosen uncompressed images where Mallory embeds her spoofed fingerprint. Figure 5b depicts the averaged security ratio.

For reference, we include the corresponding results with the fragile fingerprint approach. We focus on  $c = 1$  and JPEG quality 90, where Mallory obtains considerably high PCE values with a fingerprint-copy attack (see Figure 4(b)). In particular, we estimate the distribution of the PCE values from uncompressed images with a Gaussian kernel density estimator. The PCE threshold under which an image is not assumed as one from Alice’s camera is set such that the false positive probability is  $10^{-3}$ . Alice’s security ratio expresses the percentage of Mallory’s images that do not exceed the PCE threshold and thus are correctly identified as being not from Alice’s camera.

Figure 5b emphasizes that both approaches are a powerful combination when Mallory has just access to JPEG images. At the point where Mallory starts to circumvent fragile fingerprints, the triangle test already detects more than 50% of images that are involved in Mallory’s attack; usually enough to raise suspicion that Mallory has forged the image under investigation. In summary, Mallory faces the following dilemma: A too strong fingerprint strength is likely to be uncovered by the triangle test; with a too weak embedding, Mallory’s forged image will not be identified as one of Alice’s images. By using the triangle test in addition, Alice can even use smaller cut-off values for her fragile fingerprint.

## 6 Conclusion

This paper contributes to a thorough understanding of fragile camera fingerprints by providing a comprehensive security analysis. In multiple tests, we confirm that Mallory cannot estimate Alice’s camera fingerprint from JPEG-compressed images with common compression levels. Our analysis thus motivates the usage of fragile fingerprints in various applications, such as authentication or digital image forensics. Finally, we note that the concept of fragile fingerprints effectively demonstrates how asymmetries in the quality of accessible data can be exploited. In the context of recent unification attempts between related research disciplines [3, 22], this may foster novel strategies in adversarial machine learning or signal processing.

## Acknowledgments

The authors gratefully acknowledge funding from Deutsche Forschungsgemeinschaft (DFG) under the project RI 2469/3-1, from the German Federal Ministry of Education and Research (BMBF) under the project FIDI (FKZ 16KIS0786K), and by the NSF grant 1464275. The first author also thanks the German Academic Exchange Service (DAAD) for financial support during his stay in Binghamton.

## References

1. Alaca, F., van Oorschot, P.C.: Device fingerprinting for augmenting web authentication: Classification and analysis of methods. In: Annual Conference on Computer Security Applications (ACSAC) (2016)



2. Ba, Z., Piao, S., Fu, X., Koutsonikolas, D., Mohaisen, A., Ren, K.: ABC: Enabling smartphone authentication with built-in camera. In: Proc. of Network and Distributed System Security Symposium (NDSS) (2018)
3. Barni, M., Pérez-González, F.: Coping with the enemy: Advances in adversary-aware signal processing. In: IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP) (2013)
4. Böhme, R., Freiling, F.C., Gloe, T., Kirchner, M.: Multimedia forensics is not computer forensics. In: Geradts, Z.J.M.H., Franke, K.Y., Veenman, C.J. (eds.) Computational Forensics, IWCF, Lecture Notes in Computer Science, vol. 5718. Springer (2009)
5. Bojinov, H., Michalevsky, Y., Nakibly, G., Boneh, D.: Mobile device identification via sensor fingerprinting. CoRR abs/1408.1416 (2014)
6. Dang-Nguyen, D.T., Pasquini, C., Conotter, V., Boato, G.: RAISE: a raw images dataset for digital image forensics. In: 6th ACM Multimedia Systems Conference (2015)
7. Entrieri, J., Kirchner, M.: Patch-based desynchronization of digital camera sensor fingerprints. In: IS&T Electronic Imaging: Media Watermarking, Security, and Forensics (2016)
8. Fridrich, J.: Sensor defects in digital image forensic. In: Sencar, H.T., Memon, N. (eds.) Digital Image Forensics: There is More to a Picture Than Meets the Eye, pp. 179–218. Springer (2013)
9. Gloe, T., Böhme, R.: The Dresden Image Database for benchmarking digital image forensics. *Journal of Digital Forensic Practice* 3(2–4) (2010)
10. Gloe, T., Kirchner, M., Winkler, A., Böhme, R.: Can we trust digital image forensics? In: 15th International Conference on Multimedia (2007)
11. Gloe, T., Pfennig, S., Kirchner, M.: Unexpected artefacts in PRNU-based camera identification: A 'Dresden Image Database' case-study. In: ACM Multimedia and Security Workshop (2012)
12. Goljan, M., Fridrich, J.: Sensor-fingerprint based identification of images corrected for lens distortion. In: Memon, N., Alattar, A.M., Delp, E.J. (eds.) Media Watermarking, Security, and Forensics. Proceedings of SPIE, vol. 8303, p. 83030H (2012)
13. Goljan, M., Fridrich, J., Chen, M.: Defending against fingerprint-copy attack in sensor-based camera identification. *IEEE Transactions on Information Forensics and Security (TIFS)* 6(1) (2011)
14. Goljan, M., Fridrich, J., Filler, T.: Large scale test of sensor fingerprint camera identification. In: Delp, E.J., Dittmann, J., Memon, N., Wong, P.W. (eds.) Media Forensics and Security. Proceedings of SPIE, vol. 7254, p. 72540I (2009)
15. Gretton, A., Fukumizu, K., Teo, C.H., Song, L., Schölkopf, B., Smola, A.J.: A kernel statistical test of independence. In: Advances in Neural Information Processing Systems (NIPS) (2008)
16. Karaküçük, A., Dirik, A.E.: Adaptive photo-response non-uniformity noise removal against image source attribution. *Digital Investigation* 12 (2015)
17. Kirchner, M., Böhme, R.: Synthesis of color filter array pattern in digital images. In: Delp, E.J., Dittmann, J., Memon, N., Wong, P.W. (eds.) Media Forensics and Security. Proceedings of SPIE, vol. 7254, p. 72540K (2009)

18. Li, S., Karrenbauer, A., Saupe, D., Kuo, C.C.J.: Recovering missing coefficients in DCT-transformed images. In: IEEE International Conference on Image Processing (ICIP) (2011)
19. Lukáš, J., Fridrich, J., Goljan, M.: Digital camera identification from sensor pattern noise. IEEE Transactions on Information Forensics and Security (TIFS) 1(2) (2006)
20. Marra, F., Roli, F., Cozzolino, D., Sansone, C., Verdoliva, L.: Attacking the triangle test in sensor-based camera identification. In: IEEE International Conference on Image Processing (ICIP) (2014)
21. Mohanty, M., Zhang, M., Asghar, M.R., Russello, G.: e-PRNU: Encrypted domain PRNU-based camera attribution for preserving privacy. IEEE Transactions on Dependable and Secure Computing (TDSC) (2019)
22. Quiring, E., Arp, D., Rieck, K.: Forgotten siblings: Unifying attacks on machine learning and digital watermarking. In: IEEE European Symposium on Security and Privacy (EuroS&P) (2018)
23. Quiring, E., Kirchner, M.: Fragile sensor fingerprint camera identification. In: IEEE International Workshop on Information Forensics and Security (WIFS) (2015)
24. Reininger, R.C., Gibson, J.D.: Distributions of the two-dimensional DCT coefficients for images. IEEE Transactions on Communications 31(6) (1983)
25. Stamm, M.C., Liu, K.J.R.: Anti-forensics of digital image compression. IEEE Transactions on Information Forensics and Security (TIFS) 6(3) (2011)
26. Valsesia, D., Coluccia, G., Bianchi, T., Magli, E.: User authentication via PRNU-based physical unclonable functions. IEEE Transactions on Information Forensics and Security (TIFS) 12(8) (2017)
27. Widrow, B., Kollár, I.: Quantization Noise. Cambridge University Press (2008)

## A Sample Correlation Coefficient

The objective is to compute Pearson’s sample correlation coefficient between two images  $\mathbf{u}$  and  $\mathbf{v}$  equivalently in the DCT domain. Without loss of generality, we focus on an  $8 \times 8$  pixel block, so that the correlation is given as

$$r = \frac{n \sum u(l)v(l) - \sum u(l) \sum v(l)}{n \sqrt{\sum u(l)^2 - (\sum u(l))^2} \sqrt{\sum v(l)^2 - (\sum v(l))^2}} \quad (19)$$

where  $u(l)$  and  $v(l)$  are the pixel values. The total number of pixel values or DCT coefficients is given by  $n$ , thus for one block  $n = 64$ . To obtain the same correlation value just with the DCT representation  $\mathbf{U}$  and  $\mathbf{V}$  of both images, we

use the following identities between spatial pixels and DCT coefficients:

$$n \sum u(l)v(l) = n \sum U(l)V(l) \quad (20)$$

$$\sum u(l) \sum v(l) = n^2 \bar{\mathbf{u}} \bar{\mathbf{v}} = n U(0)V(0) \quad (21)$$

$$n \sum u(l)^2 = n \sum U(l)^2 \quad (22)$$

$$\left( \sum u(l) \right)^2 = (n \bar{\mathbf{u}})^2 = n U(0)^2. \quad (23)$$

Incorporating these identities in Eq. (19) and canceling  $n$ , we obtain the following correlation equation for one  $8 \times 8$  pixel block:

$$r = \frac{\sum \mathbf{U}(l)\mathbf{V}(l) - \mathbf{U}(0)\mathbf{V}(0)}{\sqrt{\sum (\mathbf{U}(l)^2) - \mathbf{U}(0)^2} \sqrt{\sum (\mathbf{V}(l)^2) - \mathbf{V}(0)^2}} \quad (24)$$

The generalization over all image blocks yields the same result. If we now just focus on AC coefficients, the DC coefficient  $U(0)$  and  $V(0)$  become zero. As the AC coefficient's mean goes to zero, Eq. (24) corresponds to Eq. (19). In other words, we can directly feed the AC DCT coefficients into the standard Pearson correlation equation.

## B Population Correlation Coefficient

Given a quantizer and uniform step size  $q$ , we denote by  $U$  an uncompressed image as random variable and by  $V$  its quantized output,  $V = \lfloor U/q + 0.5 \rfloor \cdot q$ . The objective is to compute the population correlation coefficient between  $U$  and  $V$ :

$$\rho = \frac{\text{Cov}(U, V)}{\sqrt{\text{Var}(U)} \sqrt{\text{Var}(V)}}. \quad (25)$$

The following general relations between the random variable  $U$  and output  $V$  can be established when  $U$  is assumed to have a symmetrical and zero-mean pdf  $f_U(x)$  with characteristic function  $\Phi_U(x)$  [27]:

$$\begin{aligned} \text{Var}(V) = \text{Var}(U) + \frac{q^2}{12} + \frac{q^2}{\pi^2} \sum_{k=1}^{\infty} \Phi_U \left( \frac{2\pi k}{q} \right) \cdot \frac{(-1)^k}{k^2} \\ + \frac{2q}{\pi} \sum_{k=1}^{\infty} \Phi_U' \left( \frac{2\pi k}{q} \right) \frac{(-1)^{k+1}}{k} \end{aligned} \quad (26)$$

$$\text{Cov}(U, V) = \text{Var}(U) + \frac{q}{\pi} \sum_{k=1}^{\infty} \Phi_U' \left( \frac{2\pi k}{q} \right) \frac{(-1)^{k+1}}{k} \quad (27)$$

For a zero-mean Laplacian distribution with parameter  $\lambda$ , the characteristic function is given as:

$$\Phi_U(x) = \frac{\lambda^2}{x^2 + \lambda^2}. \quad (28)$$

To highlight that the covariance and variance terms are only based on variable  $U$ , we write  $\text{Cov}(U, V) = \text{Cov}^+(U)$  and  $\text{Var}(V) = \text{Var}^+(U)$ .

In the next step, to determine the fingerprint quality, we need to calculate the correlation after averaging the uncompressed images and their compressed counterparts, respectively:

$$\text{cor}(\sum_i U_i, \sum_i V_i) . \quad (29)$$

We start with the distribution on one subband and denote by  $U_{i,s}$  and  $V_{i,s}$  the  $s$ -th subband of the  $i$ -th image and its compressed version. The aggregation over various subbands follows from the linear property of the covariance and the assumption of uncorrelated DCT subbands:

$$\text{Cov}(\sum_s U_{i,s}, \sum_s V_{i,s}) = \sum_s \text{Cov}(U_{i,s}, V_{i,s}) . \quad (30)$$

This is also possible for the variance of a sum of random variables. Finally, we assume the images to be uncorrelated to average the covariance over all images:

$$\text{Cov}(\sum_i U_i, \sum_i V_i) = \sum_i \sum_s \text{Cov}(U_{i,s}, V_{i,s}) . \quad (31)$$

Taking all together, the population correlation coefficient is given as

$$\rho(\sum_i U_i, \sum_i V_i) = \frac{\sum_i \sum_s \text{Cov}^+(U_{i,s})}{\sqrt{\sum_i \sum_s \text{Var}(U_{i,s})} \sqrt{\sum_i \sum_s \text{Var}^+(U_{i,s})}} . \quad (32)$$

## C Empirical Quality of Fingerprint Estimation

For the different image set scenario, Table 3 shows the correlations for  $N_E = 2000$  and  $N_E = 4648$  JPEG images.  $N_E = 4648$  is the maximum number of available images in our setup, so that its values are from a single instance of the experiment.

**Table 3.** Quality of fingerprint estimation (RAISE)

$N_E$	JPEG	$c$					
		full	1	2	3	4	5
2000	100	0.6241	0.6070	0.5613	0.4867	0.3858	0.2654
	95	0.5484	0.3853	0.2694	0.1622	0.0824	0.0551
	90	0.4633	0.1588	0.0793	0.0375	0.0181	0.0185
	85	0.4030	0.0550	0.0178	0.0071	0.0031	0.0065
	80	0.3619	0.0195	0.0026	-0.0003	0.0002	0.0045
	75	0.3301	0.0070	-0.0029	-0.0024	0.0004	0.0023
	70	0.3093	0.0035	-0.0042	-0.0036	-0.0001	0.0017
4648	100	0.6526	0.6387	0.5968	0.5279	0.4330	0.3151
	95	0.5890	0.4654	0.3531	0.2269	0.1208	0.0808
	90	0.5054	0.2162	0.1130	0.0543	0.0264	0.0277
	85	0.4451	0.0757	0.0248	0.0099	0.0040	0.0094
	80	0.4045	0.0272	0.0033	-0.0008	-0.0005	0.0047
	75	0.3732	0.0100	-0.0044	-0.0035	0.0006	0.0033
	70	0.3535	0.0053	-0.0058	-0.0053	-0.0003	0.0022