Fragile Sensor Fingerprint Camera Identification

Erwin Quiring
Department of Information Systems
University of Münster, Münster, Germany
Email: erwin.quiring@wwu.de

Matthias Kirchner
Department of Electrical and Computer Engineering
Binghamton University, Binghamton, NY
Email: kirchner@binghamton.edu

***Abstract*—We study digital camera identification based on sensor noise in an adversarial environment with asymmetries. We focus on fingerprint-copy attacks, where the attacker has access to JPEG images only, while the defender may leverage uncompressed images. This leads to the notion of fragile sensor fingerprints that are only available to the defender but do not survive lossy compression. Experiments with seven different cameras suggest a highly reliable detection of the attack as long as no high-quality images are shared with the public.**

I. Introduction

Camera sensor noise is one of the most powerful device characteristics in digital image forensics, particularly suitable to uniquely identify digital cameras [1]. This quality is commonly attributed to photo-response non-uniformity (PRNU), a camera-specific unique multiplicative noise pattern caused by inevitable material imperfections and variations in the manufacturing process of sensor elements. PRNU occurs very similarly for images captured with the same camera, but differs between images from different cameras. Hence, it can serve as a "sensor fingerprint" in forensic applications.

Sensor noise fingerprints can be estimated from a number of images taken with same camera. This makes forensic schemes vulnerable to fingerprint leakage from publicly available images [2], [3]. If Eve wants to pretend that an arbitrary image was captured by Alice's camera, she can obtain a suitable camera fingerprint from some of Alice's photos and plant it on the spurious image. This *fingerprint-copy attack* has been studied to some extent in the literature, eventually leading to a counter-technique known as the triangle test [4]. The triangle test exploits small fractions of residual correlation between the images used by Eve and the bogus fingerprint. A practical disadvantage of this procedure is that Alice—when asked to prove that Eve's image was not taken by her camera—may have to test all images ever made public by her. This can easily result in a substantial computational burden in times of widely-used photo sharing platforms. At the same time, the triangle test is known to become less reliable as the number of candidate images grows. Recent works have also proposed variations of the fingerprint-copy attack that are more likely to remain undetected [5]–[7] (but do not relieve Eve of the non-trivial problem to make the spoofed image plausible otherwise [4]).

In the following, we consider a more narrow, but not unrealistic, scenario for which we will demonstrate that Alice will not need to run the triangle test to verify that an image was not captured by her camera. Specifically, we will assume that

1) Alice's camera supports taking pictures in raw or uncompressed file format;
2) Eve's goal is to make an uncompressed image look like one of Alice's uncompressed images;
3) Alice, however, has decided to share JPEG-compressed images with the public only.[1]

The first assumption is only a minor restriction as more and more modern cameras offer raw format support. Open source firmware modifications exist for a considerable number of camera models that do not support raw images natively. Also the Android platform supports raw images, for instance. The third assumption is not unrealistic either. The vast majority of images shared online is stored in JPEG format. The second assumption is where our setup differs the most from the more general scenario in prior work. We have situations in mind where Eve wants to support her malicious activity with a high-quality forgery, for instance when interacting with a court of law.

We exploit the asymmetry of assumptions 2 and 3 by focusing on the very component of the sensor fingerprint that is *fragile* to lossy JPEG compression. While Alice will always be able to provide a "full" fingerprint from uncompressed images of her own camera, Eve's fingerprint estimate from public JPEG images will comprise the component that is robust to JPEG compression but lack the fragile part. Testing for the presence of the fragile fingerprint component only, Alice can establish that Eve's image was not taken with her camera, even though the robust component alone will typically be sufficient to fool the standard detector. Before we discuss the technical details of fragile sensor fingerprints in more detail in Sect. III, Sect. II briefly recapitulates the basic concepts of sensor noise forensics. Section IV reports experimental results. Section V concludes the paper.

Our notation is as follows: vectors and matrices are set in boldface font. Operators $\odot$ and $\oslash$ denote element-wise multiplication and division; $[\boldsymbol{X}]$ denotes element-wise rounding and truncation according to the dynamic range of $\boldsymbol{X}$.

II. Background

A. Camera Identification from Sensor Noise Fingerprints

The sensor noise fingerprint of a camera can be estimated from a sufficiently large number of images $\boldsymbol{I}_1, \dots, \boldsymbol{I}_N$ captured with that camera. For each image, a denoising filter $F(\cdot)$ extracts a noise residual $\boldsymbol{W}_k = \boldsymbol{I}_k - F(\boldsymbol{I}_k)$, commonly modeled as [1]

$$\boldsymbol{W}_k = \boldsymbol{I}_k \odot \boldsymbol{K} + \boldsymbol{\Theta}_k \,. \tag{1}$$

Multiplicative factor $\boldsymbol{K}$ represents the camera-specific PRNU term, i. e., the sensor fingerprint. $\boldsymbol{\Theta}$ subsumes a variety of

[1] We ignore the possibility of (data) theft in our forensics-centered scenario.

other noise terms, modeled as i.i.d. Gaussian. The maximum likelihood estimator of $\boldsymbol{K}$ under these assumptions is [1]

$$\hat{\boldsymbol{K}} = \left(\sum_{k=1}^{N} \boldsymbol{W}_k \odot \boldsymbol{I}_k\right) \oslash \left(\sum_{k=1}^{N} (\boldsymbol{I}_k)^2\right) . \tag{2}$$

Estimates $\hat{\boldsymbol{K}}$ typically require post-processing to remove non-unique artifacts, e. g., due to demosaicing or lens distortion correction [1], [8], [9]. Camera identification then works by computing the noise residual from a query image $\boldsymbol{J}$, $\boldsymbol{W}_{\boldsymbol{J}} = \boldsymbol{J} - F(\boldsymbol{J})$, and evaluating its similarity to a camera fingerprint estimate,

$$\rho = \mathrm{sim}(\boldsymbol{W}_{\boldsymbol{J}}, \boldsymbol{J} \odot \hat{\boldsymbol{K}}) . \tag{3}$$

The literature has proposed correlation, normalized cross-correlation and peak-to-correlation energy (PCE) as suitable similarity measures.

B. Fingerprint-Copy Attack

For Eve to make an arbitrary image $\boldsymbol{J}$ look like it was taken by Alice's camera, (at least) the following steps are necessary. First, Eve needs to collect a number of images from Alice's camera to obtain an estimate of the camera fingerprint, $\hat{\boldsymbol{K}}_E$. We assume that Eve follows the procedure outlined in the section above. Then, Eve plants her estimate on the image $\boldsymbol{J}$,

$$\boldsymbol{J}' = [\boldsymbol{J} \odot (1 + \alpha \odot \hat{\boldsymbol{K}}_E)] , \tag{4}$$

adjusting the scalar fingerprint strength parameter α suitably. A good choice of α is crucial. If α is too low, the forged image will not be identified as one of Alice's images. If α is set too high, Alice's chances of running a successful triangle test may increase substantially [4], [7]. Before and after superimposing the fingerprint, Eve may apply further processing steps to make her forgery more compelling, e. g., removing the genuine camera fingerprint [10], synthesizing demosaicing artifacts [11], and removing or adding traces of JPEG compression [12].

III. Fragile Sensor Noise Fingerprint

In the remainder of the manuscript, we assume Alice and Eve to operate in the scenario described in Sect. I. Specifically, Alice computes her fingerprint estimate $\hat{\boldsymbol{K}}$ from uncompressed images. Eve's estimate $\hat{\boldsymbol{K}}_E$ originates from JPEG-compressed images of the same camera.

The main building block of JPEG compression is the discrete cosine transform (DCT) of non-overlapping 8×8 pixel blocks. Denote $\boldsymbol{D}$ the orthogonal DCT transformation matrix for one dimension, and $\boldsymbol{X}$ a block of 8×8 pixels. The corresponding DCT coefficients $\boldsymbol{Y}$ are given as

$$\boldsymbol{Y} = \boldsymbol{D}\, \boldsymbol{X}\, \boldsymbol{D}^\top . \tag{5}$$

JPEG compression encodes quantized DCT coefficients

$$\tilde{\boldsymbol{Y}} = [\boldsymbol{Y} \oslash \boldsymbol{Q}] \tag{6}$$

based on an 8×8 quantization table $\boldsymbol{Q}$, which holds a quantization factor for each of the 64 coefficients. Quantization factors generally become larger towards higher DCT modes that correspond to high-frequency details of the image, but also with lower JPEG quality. The larger the quantization factors, the more coefficients are quantized to zero.

When Eve computes a sensor noise fingerprint estimate from JPEG-compressed images, she thus faces particularly strong quantization errors in the high-frequency DCT modes. These quantization errors will distort her estimate. In the extreme case, when quantization is too strong, Eve's images are lacking high-frequency information altogether and so will the sensor noise fingerprint estimate. At the same time, Alice's fingerprint can be expected to distribute almost evenly over all DCT modes, as illustrated in Fig. 1. The two panels compare the distributions of the 8×8 block-DCT coefficients of $\hat{\boldsymbol{K}}$ and $\hat{\boldsymbol{K}}_E$. Both fingerprints were estimated from the same 25 flat field frames, taken by a Nikon D200 camera, but Eve's images were JPEG-compressed with quality factor 90.[2]

Following this line of thought, Alice can test for the presence of a *fragile sensor noise fingerprint* that represents information from the high-frequency DCT modes only. Specifically, define a mode-selective highpass filter $H_c(\cdot)$ that retains a defined set of block-DCT coefficients only, setting all other coefficients to zero. While Alice may use this filter at various stages of the identification algorithm, we found empirically that computing a similarity measure of the form

$$\mathrm{sim}(H_c(\boldsymbol{W}_{\boldsymbol{J}}), H_c(\boldsymbol{J} \odot \hat{\boldsymbol{K}})) \tag{7}$$

works particularly well. The choice of sub-bands depends on the maximum JPEG quality of the public images. For a sufficiently conservative choice, Alice can assume that the retained portion of the fingerprint is available exclusively to her. We will consider a cut-off along (minor) anti-diagonals of the DCT coefficient matrix. For auxilliary matrix $\boldsymbol{C}$ defined below,

$$\boldsymbol{C} = \begin{pmatrix} & & & & & & & 0 \\ & & & & & & 0 & 1 \\ & & & & & 0 & 1 & 2 \\ & & & & 0 & 1 & 2 & 3 \\ & & & 0 & 1 & 2 & 3 & 4 \\ & & 0 & 1 & 2 & 3 & 4 & 5 \\ & 0 & 1 & 2 & 3 & 4 & 5 & 6 \\ 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{pmatrix} , \tag{8}$$

filter H_c retains all coefficients for which $C_{ij} \geq c$, i. e., $c = 1$ refers to all DCT modes in the lower right triangle.

We emphasize that Eve, under common modelling assumptions, cannot gain knowledge about the fragile fingerprint from the low-frequency information she has access to. Without loss of generality, consider the one-dimensional DCT and its inverse,

$$\boldsymbol{y} = \boldsymbol{D}\,\boldsymbol{x} \tag{9}$$

$$\boldsymbol{x} = \boldsymbol{D}^\top \boldsymbol{y} = \boldsymbol{D}^\top \boldsymbol{D}\,\boldsymbol{x} . \tag{10}$$

By the linearity of the DCT, Eq. (9) can be decomposed into

$$\boldsymbol{y} = (\boldsymbol{H}_c \odot \boldsymbol{D})\,\boldsymbol{x} + (\boldsymbol{L}_c \odot \boldsymbol{D})\,\boldsymbol{x} \tag{11}$$

where $\boldsymbol{H}_c$ is a binary multiplicative mask that encodes the frequency selection of highpass filter $H_c(\cdot)$ and $\boldsymbol{L}_c = \boldsymbol{H}_c$ XOR 1 is the corresponding low-frequency mask. A similar decomposition for $\boldsymbol{x}$ is

$$\boldsymbol{x} = (\boldsymbol{H}_c \odot \boldsymbol{D})^\top (\boldsymbol{H}_c \odot \boldsymbol{D})\,\boldsymbol{x} + (\boldsymbol{L}_c \odot \boldsymbol{D})^\top (\boldsymbol{L}_c \odot \boldsymbol{D})\,\boldsymbol{x} \tag{12}$$

$$= \boldsymbol{D}_H\,\boldsymbol{x} + \boldsymbol{D}_L\,\boldsymbol{x} . \tag{13}$$

[2] We refer to Sect. IV for a description of our experimental setup.

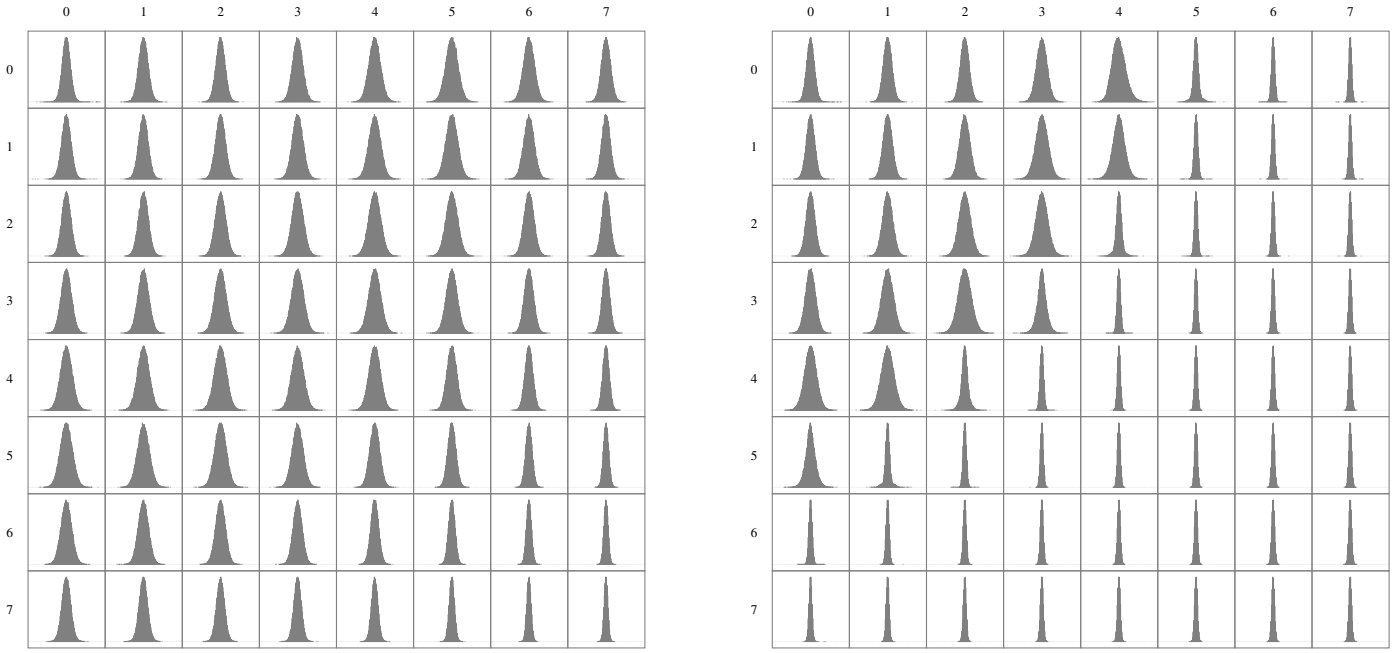

Figure 1. DCT coefficient distribution of $\hat{\boldsymbol{K}}$ (left) and $\hat{\boldsymbol{K}}_E$ (right), estimated from each 25 uncompressed and JPEG90-compressed flat field images (Nikon D200). Histograms are arranged in correspondence with DCT mode indices. All histograms are plotted on the same horizontal scale.

The first component corresponds to the fragile part of the signal. The second component represents the remaining low-frequency information. The correlation between the two components is zero. By the orthogonality of the DCT, the rows of matrices $\boldsymbol{D}_H$ and $\boldsymbol{D}_L$ are pairwise orthogonal, $\boldsymbol{D}_H^\top \boldsymbol{D}_L = \boldsymbol{0}$, and thus

$$(\boldsymbol{D}_H \boldsymbol{x})^\top (\boldsymbol{D}_L \boldsymbol{x}) = \boldsymbol{x}^\top \boldsymbol{D}_H^\top \boldsymbol{D}_L \boldsymbol{x} = 0\,. \tag{14}$$

The following section explores empirically how well Eve can estimate the fragile fingerprint component from JPEG-compressed images.

IV. Experimental Results

We examine the reliability of fragile sensor fingerprints in three different scenarios. First, we demonstrate in Sect. IV-B that fragile fingerprints are discriminative enough to distinguish between uncompressed images from Alice's camera and images from other cameras. We then switch to Eve's perspective in Sect. IV-C, quantifying the quality of fragile fingerprint estimates obtained from JPEG images. Finally, we let Eve plant spoofed JPEG fingerprints on uncompressed images from other cameras and test whether Alice can detect such fingerprint-copy attacks in Sect. IV-D.

A. Dataset and Experimental Setup

We work with the Dresden Image Database [13] and the RAISE Image Database [14], excluding images with extremely dark or saturated content. From the first dataset, we use 1442 uncompressed Adobe Lightroom images from three different camera models (two devices each, cf. Table I). Our subset of the RAISE database includes 4948 uncompressed natural images from a single Nikon D7000 camera. All images were synchronized to landscape orientation, cropped to a common size of 2000×2000 pixels, and converted to grayscale before any further processing. JPEG versions of both databases were created with the Independent JPEG Group reference library and standard quantization tables.

Table I. Natural images per camera

Database	Camera model	Camera 0	Camera 1
Dresden [13]	Nikon D70	175	188
	Nikon D70s	175	174
	Nikon D200	360	370
RAISE [14]	Nikon D7000	4948	—

Noise residuals, from both uncompressed images and JPEGs, were computed with the "classic" Wavelet denoising filter [15]. Clean sensor fingerprint estimates $\hat{\boldsymbol{K}}$ were obtained from uncompressed images for each camera according to Eq. (2), also applying the post-processing suggested in [1]. The Dresden Image Database provides 25 homogeneously lit flat field images per camera for this purpose. The fingerprint estimate of the RAISE camera was computed from 300 randomly chosen natural images. The natural images from the Dresden Image Database serve as standard benchmark set in all our tests. As all six cameras by and large gave similar results in all tested scenarios, we aggregate our outcomes over these devices. The remaining 4648 RAISE images allow us to study the effect of a large number of public images on Eve's fingerprint estimation.

B. Camera Identification

A first series of benchmark experiments focuses on the question how fragile fingerprints compare to traditional "full" sensor noise fingerprints in a typical camera identification scenario. It can be expected that fragile fingerprints will be less discriminative, as they represent only a fraction of the full sensor noise fingerprint characteristic. The PCE serves as similarity criterion, computed as in Eq. (7) for all images of each Dresden Image Database camera (true positives), and also for all

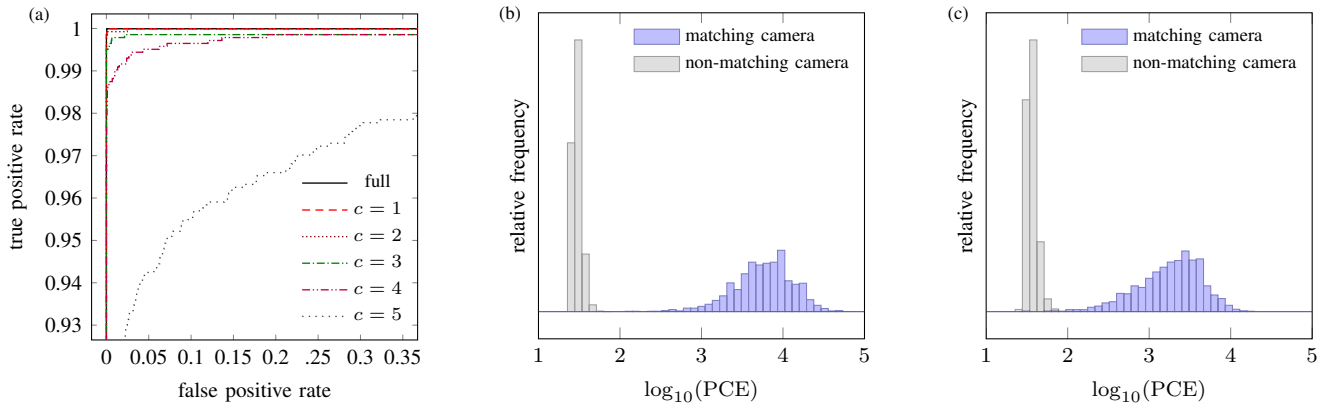

Figure 2. Camera identification with fragile sensor noise fingerprints. Results aggregated over 1442 uncompressed Dresden Image Database images from six different cameras. (a) ROC curves for different cut-off frequencies c. The "full" frequency range of standard camera identification corresponds to $c = 0$. (b) PCE distribution when considering the full frequency range. (c) PCE distribution for cut-off parameter $c = 3$.

remaining natural images from that database (true negatives). By the very nature of the fragile fingerprints, we only consider uncompressed images here.

Figure 2a depicts the aggregated ROC curves for different cut-off frequencies c, cf. Eq. (8). A curve obtained from the full frequency range (corresponding to $c = 0$) is included for reference. The results indicate that even a cut-off parameter $c = 4$ yields still relatively good separation. This corresponds to a setting where only 10 high-frequency DCT coefficients are considered. Larger numbers of coefficients guarantee almost perfect detection in our test set. Figures 2b and 2c give additional insight by comparing the underlying distributions of PCE values from standard camera identification ($c = 0$) and fragile fingerprint identification with $c = 3$. While PCE values generally decrease with each increase of the cut-off parameter, our results suggest that the separation between matching and non-matching camera images is highly reliable overall.

C. Fragile Fingerprint Estimation

A second relevant question concerns the quality of Eve's fragile fingerprint estimate from public JPEG images. We can expect that stronger JPEG compression will remove more high-frequency components of the camera sensor noise, making it less likely to obtain a good estimate of the fragile fingerprint from JPEG noise residuals.

We start by examining differences between noise residuals from uncompressed images (true positives) and JPEG images (true negatives) from all six Dresden Image Database cameras. We calculate the PCE as in Eq. (7) for both groups. The camera fingerprint estimates are always computed from uncompressed images. Figure 3 depicts aggregated ROC curves at JPEG qualities 100, 90 and 80 for different cut-off frequencies c, cf. Eq. (8). Compression with JPEG quality 100 in Fig. 3a is equivalent to simple rounding of DCT coefficients, which is why noise residuals extracted from JPEGs images are not distinguishable from those of uncompressed images. Results for JPEG qualities 90 (Fig. 3b) and 80 (Fig. 3c) indicate that stronger compression of public images is clearly to Alice's advantage because noise residuals from individual images contain less and less relevant information. These results also underline that the high robustness against JPEG compression that "classical" sensor noise camera identification is known for [1] is to a large degree due to low-frequency noise components.

Putting more emphasis on Eve's perspective, we also report the quality of fingerprint estimation [4] in terms of the correlation between Alice's fingerprint from uncompressed images, $\hat{\boldsymbol{K}}$, and Eve's fingerprint from compressed images, $\hat{\boldsymbol{K}}_{\boldsymbol{E}}$:

$$\mathrm{corr}(H_c(\hat{\boldsymbol{K}}), H_c(\hat{\boldsymbol{K}}_{\boldsymbol{E}}))\,. \tag{15}$$

In this scenario, we assume Eve to have access to N_E public JPEG images from Alice's camera. Varying the parameter N_E, we report results averaged over five randomly compiled JPEG image sets of size N_E per camera. Alice's camera-specific fingerprint was kept constant throughout all repetitions. Table II summarizes the aggregated correlations between the fingerprint estimate from uncompressed flat field images and $N_E = 150$ images from the Dresden Image Database. Table III resembles this setup for the RAISE Image Database, granting Eve access to a much larger number of 2000 and 4648 JPEG images however.[3] A closer inspection of the results suggests that Eve's fingerprint quality increases only very slowly with the number of available public JPEG images. Most importantly, we observe that the correlation values remain extremely low for suitable combinations of JPEG quality and cut-off parameter c, preventing Eve from obtaining a good estimate of Alice's fragile fingerprint even from a substantial amount of public data.

D. Fingerprint-Copy Attack

We finally consider the more realistic scenario where Eve embeds her fingerprint estimate $\hat{\boldsymbol{K}}_E$ into 100 randomly chosen uncompressed images taken by a different camera, as described in Sect. II-B. Note that we do not attempt to determine the optimal embedding strength α for each image [4], [7], but rather test over a variety of settings in the range $\{0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 1, 2, 5, 10\}$. For all parameter combinations, this process was repeated over the same image set partitions as in Sect. IV-C.

[3] Note that $N_E = 4648$ is the maximum number of available images in our setup, so that the correlation values are from a single instance of the experiment only in this case.

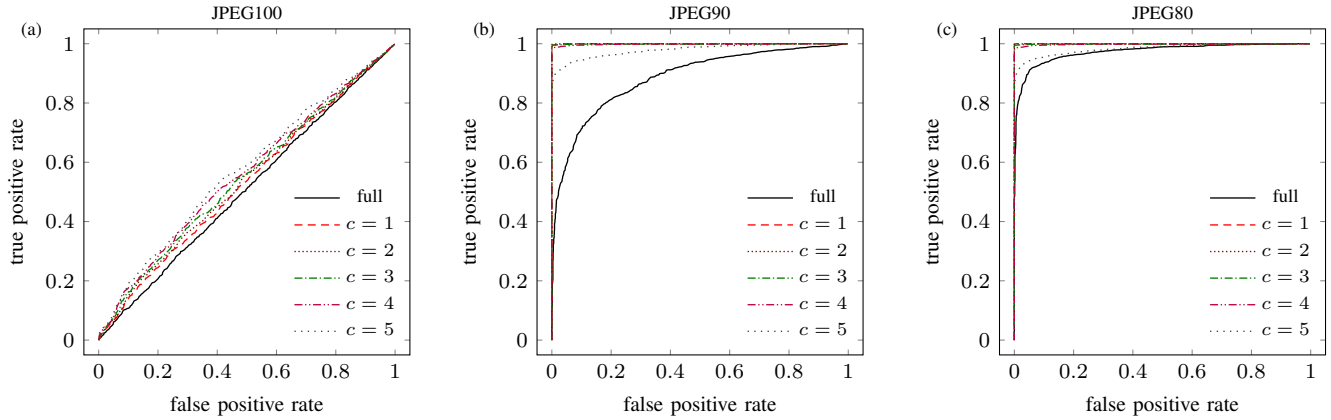

Figure 3. PCE-based distinction between noise residuals from uncompressed images and JPEG images with different JPEG qualities. ROC curves aggregated over 1442 Dresden Image Database images from six different cameras.

Figure 4 summarizes this experiment by plotting the average PCE values after the fingerprint-copy attack as a function of embedding strength for JPEG quality factors 100, 90 and 80. The upper row corresponds to fingerprints estimated from $N_E = 150$ images from the Dresden Image Database, the lower three panels depict the results for the RAISE Image Database with $N_E = 4648$. Not surprisingly, fingerprint similarity increases rapidly when Eve has high-quality JPEG images available (Figs. 4a and 4d). The situation is substantially different with stronger compression. For a JPEG quality factor of 90, only the full fingerprint yields high PCE values for small embedding strengths α, yet larger cut-offs c would prevent Eve from mounting a successful attack (Figs. 4b and 4e). For even lower JPEG qualities, no reasonable choice of α will produce Eve's desired result when $c > 0$ (Figs. 4c and 4f). The influence of parameter N_E is generally weak. A comparison of Figs. 4b and 4e suggests that a substantially larger number of public images is only to Eve's advantage for extremely strong embedding.

Overall, these results indicate that fragile fingerprints are a viable means to counter fingerprint-copy attacks in a scenario where Eve has only JPEG images available. For a JPEG quality of 90 for instance, Alice may choose $c = 3$ to ensure reliable camera identification in an adversarial environment.

Table II. QUALITY OF FINGERPRINT ESTIMATION (DRESDEN)

N_E	JPEG	c					
		full	1	2	3	4	5
150	100	0.3720	0.3484	0.3245	0.2850	0.2302	0.1607
	95	0.2522	0.0870	0.0561	0.0337	0.0160	0.0100
	90	0.1865	0.0294	0.0157	0.0058	0.0009	0.0029
	85	0.1449	0.0109	0.0029	-0.0007	-0.0022	0.0012
	80	0.1174	0.0027	-0.0014	-0.0031	-0.0027	-0.0000
	75	0.0977	-0.0012	-0.0029	-0.0030	-0.0026	-0.0005
	70	0.0851	-0.0029	-0.0037	-0.0036	-0.0030	-0.0011

Table III. QUALITY OF FINGERPRINT ESTIMATION (RAISE)

N_E	JPEG	c					
		full	1	2	3	4	5
2000	100	0.6128	0.6002	0.5565	0.4838	0.3828	0.2627
	95	0.5291	0.3752	0.2645	0.1600	0.0800	0.0524
	90	0.4381	0.1513	0.0762	0.0357	0.0169	0.0177
	85	0.3758	0.0506	0.0154	0.0055	0.0023	0.0046
	80	0.3347	0.0172	0.0009	-0.0014	-0.0008	0.0028
	75	0.3035	0.0055	-0.0041	-0.0034	-0.0007	0.0016
	70	0.2837	0.0015	-0.0053	-0.0042	-0.0015	0.0007
4648	100	0.6414	0.6322	0.5924	0.5250	0.4302	0.3123
	95	0.5704	0.4536	0.3464	0.2235	0.1173	0.0773
	90	0.4798	0.2047	0.1087	0.0520	0.0243	0.0253
	85	0.4167	0.0705	0.0231	0.0084	0.0028	0.0078
	80	0.3756	0.0249	0.0024	-0.0019	-0.0013	0.0037
	75	0.3445	0.0085	-0.0049	-0.0042	-0.0000	0.0027
	70	0.3254	0.0036	-0.0063	-0.0059	-0.0008	0.0017

V. CONCLUDING REMARKS

This paper has explored sensor fingerprints that are fragile to JPEG compression. In the broad context of counter-forensics [16] and adversary-aware signal processing [17], our work has considered a scenario where an attacker attempts to frame a victim by planting a fake fingerprint, estimated from JPEG images only, on an uncompressed image. We demonstrated how fragile fingerprints from uncompressed images provide valuable side information to the defender. In a general hypothesis testing framework, this asymmetry may resemble a situation where attacker and defender have access to training data of different quality [18]. Experimental results indicate that an attacker faces substantial problems when attempting to recover the fragile fingerprint component from JPEG images alone.

We emphasize that fragile fingerprints are not a replacement of the triangle test [4]. In fact, both approaches should rather be seen as a powerful combination, as Eve faces the following dilemma: choosing a too strong fingerprint strength, she risks being uncovered by the triangle test; too weak embedding may be an easy catch for a detector based on fragile fingerprints.

Possible extensions and research questions abound. It will be interesting to see how the proposed approach generalizes to other forms of side information, for instance also involving raw images or JPEG images of different qualities. Future work may explore strategies for Alice to make even more informed selections of DCT sub-bands, e. g., based on a set of candidate images. Finally, we expect that the incorporation of DCT coefficient distribution assumptions will contribute to a more thorough understanding of the limits of fragile fingerprints and serve as stepping stone for further applications.

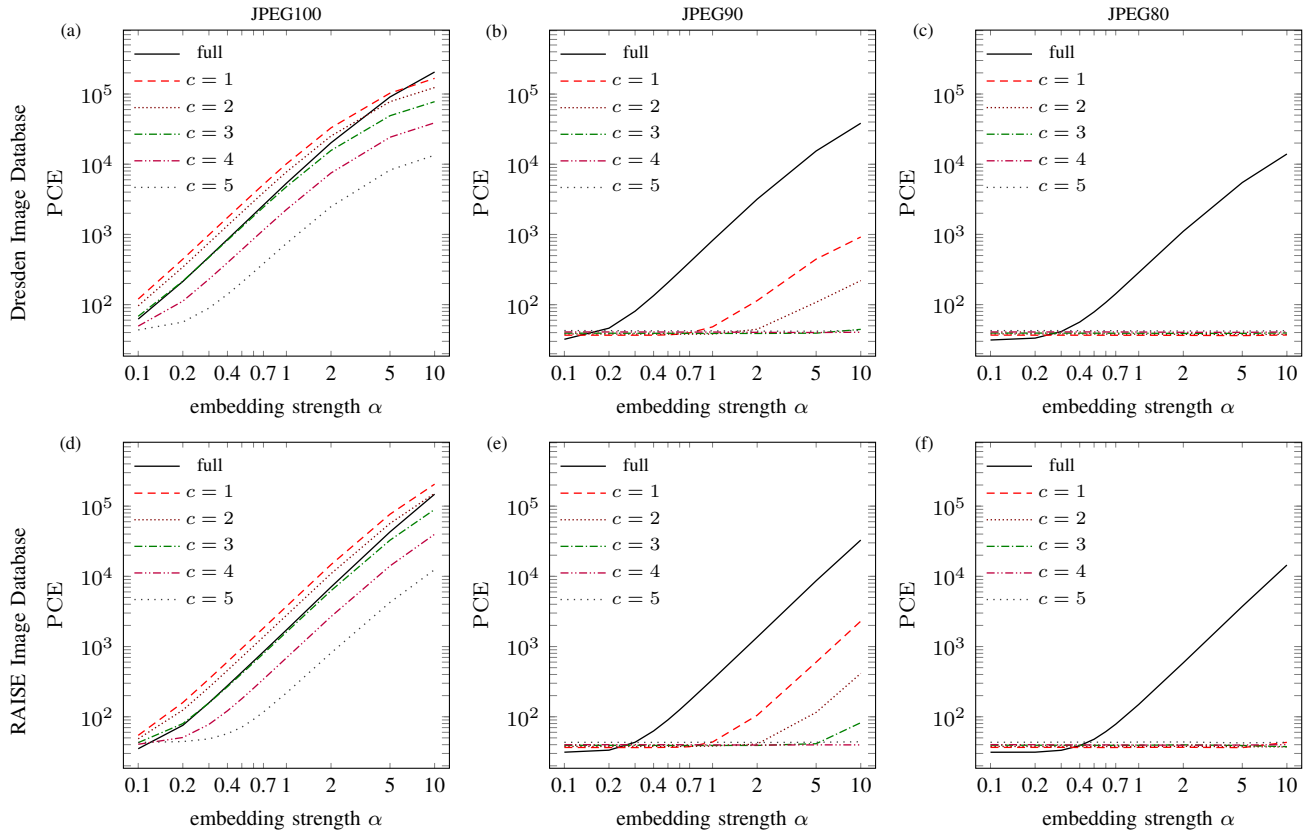

Figure 4. Fingerprint-copy attack with fragile fingerprints. Average PCE values as a function of the embedding strength α for different JPEG qualities. Results aggregated over embedding into 100 images with (a)–(c) $N_E = 150$ (Dresden Image Database) and (d)–(f) $N_E = 4648$ (RAISE Image Database).

ACKNOWLEDGMENTS

This work is supported in part by the NSF grant 1464275. The first author thanks the German Academic Exchange Service (DAAD) for financial support during his stay in Binghamton.

REFERENCES

[1] J. Fridrich, "Sensor defects in digital image forensic," in *Digital Image Forensics: There is More to a Picture Than Meets the Eye*, H. T. Sencar and N. Memon, Eds. Springer, 2013, pp. 179–218.

[2] J. Lukáš, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 205–214, 2006.

[3] T. Gloe, M. Kirchner, A. Winkler, and R. Böhme, "Can we trust digital image forensics?" in *15th International Conference on Multimedia*, 2007, pp. 78–86.

[4] M. Goljan, J. Fridrich, and M. Chen, "Defending against fingerprint-copy attack in sensor-based camera identification," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 1, pp. 227–236, 2011.

[5] R. Caldelli, I. Amerini, and A. Novi, "An analysis on attacker actions in fingerprint-copy attack in source camera identification," in *IEEE International Workshop on Information Forensics and Security*, 2011.

[6] Q. Rao, H. Li, W. Luo, and J. Huang, "Anti-forensics of the triangle test by random fingerprint-copy attack," in *Computational Visual Media Conference*, 2013.

[7] F. Marra, F. Roli, D. Cozzolino, C. Sansone, and L. Verdoliva, "Attacking the triangle test in sensor-based camera identification," in *IEEE International Conference on Image Processing*, 2014, pp. 5307–5311.

[8] M. Goljan and J. Fridrich, "Sensor-fingerprint based identification of images corrected for lens distortion," in *Media Watermarking, Security, and Forensics 2012*, ser. Proceedings of SPIE, N. Memon, A. M. Alattar, and E. J. Delp, Eds., vol. 8303, 2012, 83030H.

[9] T. Gloe, S. Pfennig, and M. Kirchner, "Unexpected artefacts in PRNU-based camera identification: A 'Dresden Image Database' case-study," in *ACM Multimedia and Security Workshop*, 2012, pp. 109–114.

[10] A. Karaküçük and A. E. Dirik, "Adaptive photo-response non-uniformity noise removal against image source attribution," *Digital Investigation*, vol. 12, pp. 66–76, 2015.

[11] M. Kirchner and R. Böhme, "Synthesis of color filter array pattern in digital images," in *Media Forensics and Security*, ser. Proceedings of SPIE, E. J. Delp, J. Dittmann, N. Memon, and P. W. Wong, Eds., vol. 7254, 2009, 72540K.

[12] M. C. Stamm and K. J. R. Liu, "Anti-forensics of digital image compression," *IEEE Transactions on Information Forensics and Security*, vol. 6, no. 3, pp. 1050–1065, 2011.

[13] T. Gloe and R. Böhme, "The Dresden Image Database for benchmarking digital image forensics," *Journal of Digital Forensic Practice*, vol. 3, no. 2–4, pp. 150–159, 2010.

[14] D.-T. Dang-Nguyen, C. Pasquini, V. Conotter, and G. Boato, "RAISE: a raw images dataset for digital image forensics," in *6th ACM Multimedia Systems Conference*, 2015, pp. 219–224.

[15] M. K. Mıhçak, I. Kozintsev, and K. Ramchandran, "Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 6, 1999, pp. 3253–3256.

[16] R. Böhme and M. Kirchner, "Counter-forensics: Attacking image forensics," in *Digital Image Forensics: There is More to a Picture Than Meets the Eye*, H. T. Sencar and N. Memon, Eds. Springer, 2013, pp. 327–366.

[17] M. Barni and F. Pérez-González, "Coping with the enemy: Advances in adversary-aware signal processing," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2013, pp. 8682–8686.

[18] M. Barni and B. Tondi, "Binary hypothesis testing game with training data," *IEEE Transactions on Information Theory*, vol. 60, no. 8, pp. 4848–4866, 2014.