

Large Scale Test of Sensor Fingerprint Camera Identification

Miroslav Goljan, Jessica Fridrich, and Tomáš Filler

Department of Electrical and Computer Engineering

SUNY Binghamton, Binghamton, NY 13902-6000

Ph: (607) 777 5793, Fax: (607) 777 4464

Email: {mgoljan, fridrich, tfiller}@binghamton.edu

ABSTRACT

This paper presents a large scale test of camera identification from sensor fingerprints. To overcome the problem of acquiring a large number of cameras and taking the images, we utilized Flickr, an existing on-line image sharing site. In our experiment, we tested over one million images spanning 6896 individual cameras covering 150 models. The gathered data provides practical estimates of false acceptance and false rejection rates, giving us the opportunity to compare the experimental data with theoretical estimates. We also test images against a database of fingerprints, simulating thus the situation when a forensic analyst wants to find if a given image belongs to a database of already known cameras. The experimental results set a lower bound on the performance and reveal several interesting new facts about camera fingerprints and their impact on error analysis in practice. We believe that this study will be a valuable reference for forensic investigators in their effort to use this method in court.

Keywords: Camera identification, digital forensics, photo-response non-uniformity, sensor fingerprint.

1. MOTIVATION FOR EXTENSIVE TESTING

Scratches on silver-halide films and their imprints in classical photographs have been used to identify analog cameras that took the image in question. Not long after digital cameras overtook the world market with consumer and professional cameras, an effective method for identifying the source camera using sensor noise was proposed by Lukáš *et al.* in 2005. Later, the method has been refined [1], its use expanded to image forgery detection and other applications [2].

In order for the camera sensor identification method (CSI) to become an admissible evidence for establishing a link between a photograph and a camera, it is essential to provide extensive experimental verification of the theoretically estimated false alarm rate as a function of the detection threshold. To pave the way for the CSI to the courtroom, large scale tests are needed across many models and individual cameras. To this date, the performance of the CSI has only been evaluated for a rather limited number of cameras (less than 20). Tests aimed at distinguishing between cameras of the same model are even scarcer and only a handful of camera pairs were experimentally tested so far. Such tests are important because the noise component of images coming from different cameras of the same make or model exhibit weak similarities that propagate into the estimate of the sensor fingerprint [2]. While such artifacts are useful in camera model classification [5], they are highly undesirable for camera identification as they increase the false acceptance rate. To simplify the language in this paper, we will call these artifacts NUA (Non-Unique Artifact). A NUA is defined as a systematic signal that is non-unique to the individual sensor but may be shared between cameras of the same model or make or cameras with the same sensor architecture. While several countermeasures based on camera fingerprint pre-processing have been proposed [1], [6] to suppress the NUAs, their effectiveness has not been studied on a large scale.

Summarizing the motivation, the large scale tests will help answer the following questions:

- How does the performance scale to large number of camera brands and models?
- How effectively can we suppress NUA? Are there any remaining artifacts that still increase false acceptance?
- What is the overall detection rate for full size images from digital cameras in (typical) JPEG format?
- Are theoretically estimated false alarm rates a good match in reality?
- Are there any yet undiscovered issues?

In the next section, we review the camera identification algorithm as it appeared in [2]. Then, in Section 3 we describe the database on which all experiments are carried out. Section 4 contains the description of all experiments and discussion of the results. Finally, Section 5 concludes the paper.

2. PRELIMINARIES

Subtle variations among pixels in their sensitivity to light are the cause of the Photo-Response Non-Uniformity (PRNU) of both CCD and CMOS sensors [7], [8]. PRNU casts a unique pattern onto every image the camera takes. This “camera fingerprint” has been argued to be unique for each camera, [2]. The camera fingerprint can be estimated from images known to have been taken with the camera. A given digital image can be tested for the presence or absence of the fingerprint and thus shown whether or not it was taken with a specific camera.

2.1 Camera Sensor Identification based on sensor PRNU

Denoting the camera output image as \mathbf{I} and the “true scene” image that would be captured in the absence of any imperfections as \mathbf{I}_0 , the following sensor output model was established in [2] based on the model [9] (all matrix operations are understood element-wise)

$$\mathbf{I} = \mathbf{I}_0 + \mathbf{I}_0 \mathbf{K} + \boldsymbol{\Theta}, \quad (1)$$

where \mathbf{K} is the PRNU factor (sensor fingerprint) and $\boldsymbol{\Theta}$ includes all other noise components, such as dark current, shot noise, readout noise, and quantization noise [7], [8]. The fingerprint \mathbf{K} can be estimated from N images $\mathbf{I}^{(1)}, \mathbf{I}^{(2)}, \dots, \mathbf{I}^{(N)}$ taken by the camera. Let $\mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \dots, \mathbf{W}^{(N)}$, be their noise residuals obtained using a denoising filter F , $\mathbf{W}^{(i)} = \mathbf{I}^{(i)} - F(\mathbf{I}^{(i)})$, $i = 1, \dots, N$. In [2], the following maximum likelihood estimator of the PRNU factor, \mathbf{K} , was derived:

$$\hat{\mathbf{K}} = \frac{\sum_{i=1}^N \mathbf{W}^{(i)} \mathbf{I}^{(i)}}{\sum_{i=1}^N (\mathbf{I}^{(i)})^2}. \quad (2)$$

Denoting the noise residual of the image under investigation, \mathbf{I} , as \mathbf{W} , the detection of the fingerprint \mathbf{K} in \mathbf{W} can be formulated as a two-channel hypothesis testing problem [10]

$$\begin{aligned} H_0 \text{ (non-matching image): } & \mathbf{K}_1 \neq \mathbf{K}_2 \\ H_1 \text{ (matching image): } & \mathbf{K}_1 = \mathbf{K}_2 \end{aligned} \quad (3)$$

where

$$\begin{aligned} \hat{\mathbf{K}}_1 &= \mathbf{K}_1 + \boldsymbol{\Xi}_1 \\ \mathbf{W} &= \mathbf{I} \mathbf{K}_2 + \boldsymbol{\Xi}_2 \end{aligned}, \quad (4)$$

are two observables—the estimate of the camera fingerprint, $\hat{\mathbf{K}}_1$, obtained using (2) and the noise residual \mathbf{W} . Under the assumption that the image under investigation did not undergo any geometrical processing with the exception of cropping, a good approximation of the generalized likelihood ratio test [11] is the maximum of the normalized correlation ρ [12]

$$\max_{s_1, s_2} \rho(s_1, s_2; \mathbf{X}, \mathbf{Y}), \quad (5)$$

where

$$\rho(s_1, s_2; \mathbf{X}, \mathbf{Y}) = \frac{\sum_{k=1}^m \sum_{l=1}^n (\mathbf{X}[k, l] - \bar{\mathbf{X}}) (\mathbf{Y}[k + s_1, l + s_2] - \bar{\mathbf{Y}})}{\|\mathbf{X} - \bar{\mathbf{X}}\| \|\mathbf{Y} - \bar{\mathbf{Y}}\|}, \quad (6)$$

$\|\cdot\|$ is the L_2 norm, and

$$\mathbf{X} = \mathbf{I}\hat{\mathbf{K}}, \mathbf{Y} = \mathbf{W}. \quad (7)$$

The maximum in (5) is taken over all k admissible shifts between the possibly cropped image and the camera fingerprint. Denoting the image and fingerprint dimensions $m \times n$ and $m_K \times n_K$, respectively, the number of admissible shifts is

$$k = (m_K - m + 1)(n_K - n + 1).$$

Before evaluating (6), the image is padded with zeros to match the sizes of \mathbf{X} and \mathbf{Y} . The shifts $k + s_1$ and $l + s_2$ are taken modulo m and n , respectively.

The case when H_1 is rejected for an image that did originate from the camera is called *false rejection*. *False acceptance (alarm)* means accepting H_1 when the image was not taken by the camera. We denote the false rejection rate *FRR*, the false alarm rate *FAR*, and the detection rate $DR = 1 - FRR$. Following the Neyman-Pearson criterion, a bound is set on *false acceptance probability*, which determines the *detection threshold* for the test statistic (5). The FRR is obtained from experiments and depends mainly on the image content and quality, the number of images used to estimate the PRNU factor and their quality, and likely on some physical sensor parameters. Both FRR and FAR are functions of the detection threshold.

Denoting the coordinates of the peak where the maximum (5) occurs as $\mathbf{s}_{\text{peak}} = [s_1, s_2]$, the Peak to Correlation Energy ratio (PCE)

$$PCE_k = \frac{\rho(\mathbf{s}_{\text{peak}}; \mathbf{X}, \mathbf{Y})^2}{\frac{1}{mn - |\mathcal{N}|} \sum_{\mathbf{s} \in \mathcal{N}} \rho(\mathbf{s}; \mathbf{X}, \mathbf{Y})^2} = \frac{(\mathbf{X} \cdot \mathbf{Y}(\mathbf{s}_{\text{peak}}))^2}{\frac{1}{mn - |\mathcal{N}|} \sum_{\mathbf{s} \in \mathcal{N}} (\mathbf{X} \cdot \mathbf{Y}(\mathbf{s}))^2}, \quad (8)$$

is used as a measure of the peak height. Here, $\mathbf{X} \cdot \mathbf{Y}(\mathbf{s})$ is the dot product between $\mathbf{X} - \bar{\mathbf{X}}$ and $\mathbf{Y}(\mathbf{s}) - \bar{\mathbf{Y}}$ circularly shifted by vector \mathbf{s} , and \mathcal{N} is a small neighborhood around the peak (in this paper, \mathcal{N} is a square region 11×11 pixels). The normalized correlation can be replaced with the dot product (the correlation) as the norms from (6) cancel out in (8). We note that if we know a priori that the image under investigation has not been cropped, the search for the peak is not carried out and $k = 1$ in (8).

The PCE is a more stable test statistic than correlation as it is independent of the image size and has other advantages [6]. This definition is also compliant with the definition of Kumar and Hassebrook [13].

2.2 Error analysis

To set the threshold for the test statistics under the Neyman-Pearson setting, we need to determine the distribution of the test statistic (6) and (8) under H_0 . We simplify the analysis by noticing that for large data record (large number of pixels mn),

$$\frac{1}{\sqrt{mn}} \|\mathbf{X} - \bar{\mathbf{X}}\| \approx \sigma_X \quad \text{and} \quad \frac{1}{\sqrt{mn}} \|\mathbf{Y} - \bar{\mathbf{Y}}\| \approx \sigma_Y,$$

where σ_X^2, σ_Y^2 are the variances of \mathbf{X} and \mathbf{Y} , respectively. Ideally, the image noise residual $\mathbf{Y}[i, j]$ and the signal $\mathbf{X}[i, j]$ containing the estimated fingerprint should be independent. The normalized cross-correlation (6) is then well modeled as Gaussian $N(0, 1/mn)$ by virtue of the central limit theorem (for experimental verification of this modeling assumption, see [14]). This allows us to compute a theoretical decision threshold τ for PCE (8) from

$$FAR = 1 - \left(1 - Q(\sqrt{\tau})\right)^k, \quad (9)$$

where $Q(x)$ denotes the complementary cumulative distribution function of a standard normal random variable. Alternatively, for a chosen $FAR \leq \alpha$

$$\tau(\alpha, k) = \left[Q^{-1} \left(1 - (1 - \alpha)^{\frac{1}{k}} \right) \right]^2. \quad (10)$$

The threshold (10) varies over images with different sizes (because k varies), which complicates aggregating experimental data and comparing FAR from the experiment with the theory. To resolve this issue, we performed majority of our experiments without considering any search for cropping ($k = 1$), in which case $\tau = (Q^{-1}(\alpha))^2$. Note that the corresponding model for the PCE is the chi-square distribution with one degree of freedom χ_1^2 .

If \mathbf{X} and \mathbf{Y} share a weak component due to imperfect suppression of NUAs, \mathbf{X} and \mathbf{Y} are not independent and the error analysis must reflect this. We attempt to capture the dependency by writing $\mathbf{X} + a\boldsymbol{\eta}$ and $\mathbf{Y} + b\boldsymbol{\eta}$, instead of \mathbf{X} and \mathbf{Y} , where $\boldsymbol{\eta} \sim N(0,1)$ is the shared component. In this case, simple algebra shows that the correlation (6) is a Gaussian signal with mean μ and the corresponding model for the PCE is the non-central chi-square $\chi_1^2(\lambda)$ with one degree of freedom and non-centrality parameter

$$\lambda = nm \frac{\mu^2}{1 + \mu^2}, \quad \mu^2 = \frac{a^2 b^2}{(\sigma_X^2 + a^2)(\sigma_Y^2 + b^2)}. \quad (11)$$

In most of our tests, we compute two values of PCE – one for the original image and one for the image rotated by 180 degrees (because we do not know which way the user held the camera) and take their maximum. In this case, the statistic is PCE_2 and it will be modeled as the maximum of two independent realizations of a random variable $\chi_1^2(\lambda)$. Because the PCE covers a very large range from about 0 to $\sim 10^5$, we will be working with $\log_{10} PCE_2$ instead. It can be easily shown that the probability density function of $\log_{10} PCE_2$ is

$$f_{L,\lambda}(x) = 2 \ln 10 \cdot 10^x f_\lambda(10^x) F_\lambda(10^x), \quad (12)$$

where $f_\lambda(x)$ and $F_\lambda(x)$ are the pdf and cdf of a non-central chi-square distributed random variable with non-centrality parameter λ and one degree of freedom. Note that this model contains the case of independent \mathbf{X} and \mathbf{Y} ($\lambda = 0$). Also, (12) could be used to model PCE_2 in the matched case (Part 2).

3. IMAGE DATABASE USED IN EXPERIMENTS

To properly estimate the error rates, we need an oracle that would randomly draw *digital camera images* that people can make. As it is clearly impossible to obtain such an oracle, we restrict ourselves to a large public image database at www.flickr.com. The important advantage of this source is its diversity. The distribution of image content, photographic style, composition, quality, as well as the distribution of individual camera models and brands is a reasonable approximation to our ideal oracle. The Flickr image database (further denoted as \mathcal{F}) contains millions of images, many of them in full resolution and with EXIF data containing information about the camera model, camera settings, etc. The images can be accessed by various queries, such as by owner (user), camera model, time span of uploading, etc.

There are, however, disadvantages to using an open access database compared to a controlled set of images coming from physically available source cameras. Images needed for camera fingerprint estimation are collected from each user without having control over their quality and are not even guaranteed to be coming from a single camera. There might be cases when one user switches from one camera to another of the same model or shares images with another user. Such cases will lead to a “mixed camera fingerprint” and will lower the PCE for images from both cameras. Depending on the ratio of each camera contribution to $\hat{\mathbf{K}}$, missed identification (false rejection) may occur. Images captured with digital zoom will also contribute to FRR due to pixel de-synchronization because it is computationally infeasible to search for the zoom parameter in all non-matching cases.

Another complication will occur when two or more users share one camera. In this case, we will have to find sufficient arguments to prove that a large PCE for the test image and a camera fingerprint is not a false acceptance but is, in fact, a correct positive identification (a “false false alarm”). While we attempt to resolve some selected cases, it is infeasible to do this for all cases. For this reason, we will talk about a lower bound on the performance of CSI.

A subset \mathcal{D} of the Flickr database was acquired in the following manner. User names were selected based on their activity in uploading new images. For each username u that was already on the list, all images associated with u were listed by name, EXIF header, and size. The largest image size of one camera model was assumed to be the camera’s native resolution. Once at least 50 images in landscape orientation and 10 or more in any orientation were recognized for one user u and one of his camera models c (according to camera model info in EXIF header), the images were downloaded to our directory `~flickr/c/u`. This directory structure allows us to break any test results to single cameras and trace problematic cases to the single camera user. Let us denote such a subset of \mathcal{D} as $\mathcal{D}(c,u)$. We set the maximum for each user-camera pair to 200 images, $|\mathcal{D}(c,u)| \leq 200$. The target amount of images in \mathcal{D} was one million. We stopped the downloading process at $|\mathcal{D}| > 1,000,000$. We note that all images were in the JPEG format. After deleting corrupted files, the final database consisted of $|\mathcal{D}| = 1,053,580$ JPEG images. We denote the list of all camera models as \mathcal{C} and the list of all users of the same camera model c as $\mathcal{D}(c)$.

The database \mathcal{D} contains:

- $|\mathcal{C}| = 150$ camera models,
- $\sum_{c \in \mathcal{C}} |\mathcal{D}(c)| = 6,896$ total number of user-camera image sets, which yields to slightly less than 6,896 individual cameras because some are identical,
- $\sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{D}(c)} |\mathcal{D}(c,u)| = 1,052,700$ JPEG images.

Although one million images for testing the camera identification technique is a large number from the computational perspective, it is still a very small fraction of images in \mathcal{F} . We limited the number of images from one camera-user pair to the range of 60 to 200. The lower amount of 60 was chosen to allow using $N=50$ for the camera fingerprint estimation and the rest for testing. Choosing a larger N would probably lead to a better performance of CSI. The upper bound on the number of images, 200, means that more users end up in our database. For each user-camera pair, we only selected the highest image resolution. The list of all camera models and their resolutions are listed in the appendix. A restriction to less than 8 megapixels (Mp) was set to keep the computational time reasonable. Most images in \mathcal{D} had less than 6 Mp.

4. EXPERIMENTS

The experiments start with computing the fingerprint $\hat{\mathbf{K}}$ for each user-camera pair from randomly chosen 50 images in landscape orientation (Part 1). Then, for each camera fingerprint we evaluate the PCE (8) under H_1 for the remaining matching images to provide data for estimating the FRR (Part 2). Next, we run a similar test under H_0 to evaluate the FAR. We do so in two parts. In Part 3, we estimate the FAR conditioned on the event that the tested image originated in a camera of a different model than the tested fingerprint. In Part 4, we compute the FAR conditioned on the event that the test image originated in a different camera of the same model as the tested fingerprint. Comparing the conditional error rates from Part 3 and 4 provides information about the effectiveness of NUA suppression. The overall FAR is estimated by combining the conditional error rates from Part 3 and 4 using the prior probability of the event that a randomly selected image and fingerprint came from the same camera model. Finally, in Part 5, we match a randomly chosen image from each camera model across all 6,896 fingerprints simulating thus the situation when an image is tested against a database of fingerprints.

4.1 Experimental setup and technical issues

Because the computational complexity is an important issue for such a large scale test, we implemented a fast version of the CSI method. Except for data collection, all experimental components ran on a cluster of 40 2-core AMD Opteron processors; 50 of them were devoted to this application. This configuration allowed us to complete all experiments on the database of 10^6 images in about three months.

The noise residuals, \mathbf{W} , of the test images (rotated to landscape orientation if needed) and the estimates $\hat{\mathbf{K}}$ were converted to grayscale signals. First, the fingerprint was estimated in each of the three color channels (red, blue, green). Then, the channels were combined using the common linear transformation $\text{RGB} \rightarrow \text{grayscale}$, the sample mean was subtracted from $\hat{\mathbf{K}}$, and averages of each column and each row in each of four sub-sampled 2-D signals corresponding to four types of pixels in the Bayer CFA were subtracted from all elements of $\hat{\mathbf{K}}$. This “Zero-Mean” procedure is the same as described in [6]. This procedure removes a large portion of NUAs introduced by demosaicking. To remove any residual artifacts from the estimated fingerprint, the fingerprint was finally filtered using an adaptive Wiener filter in the frequency domain (e.g., to further reduce JPEG compression artifacts or artifacts inherent to sensor on-board circuitry).

Then the PCE_1 (8) was computed. To save computation time in Part 2, we only computed PCE_2 when PCE_1 was less than 60.

4.2 Part 1: Calculation of camera fingerprints

We used 50 randomly chosen images in the landscape orientation from $\mathcal{D}(c, u)$ to estimate the camera fingerprint for each $c \in \mathcal{C}$ and $u \in \mathcal{D}(c)$. This way, we avoided the problem with unknown clockwise or counter-clockwise rotation of portrait photographs. The actual number of fingerprints obtained was smaller than the total number of camera-user image sets, because we did not compute the fingerprint whenever fewer than 50 images were available for the fingerprint computation due to deleted corrupted files. Denoting $\mathcal{D}_f(c)$ the list of users of camera model c with the camera fingerprint $\hat{\mathbf{K}}_{c,u}$, the amount of fingerprints became $N_f = \sum_{c \in \mathcal{C}} |\mathcal{D}_f(c)| = 6,827$. The sets $\mathcal{D}(c, u)$ of images for

camera-users from $\mathcal{D}_f(c)$ contain those images that participated in fingerprints, for which we use upper index ^(f), and the rest ^(v), $\mathcal{D}(c, u) = \mathcal{D}^{(f)}(c, u) \cup \mathcal{D}^{(v)}(c, u)$.

4.3 Part 2: Images matching the camera fingerprint

Testing images that match the camera fingerprint assumes the hypothesis H_1 to be true. This test determines the detection rate DR and the false rejection rate FRR for any given bound τ on PCE_2 . We note at this point that due to the fact that our image database is uncontrolled, it is possible that some of the tested images did not come from the same camera as assumed, in which case, the positive identification fails (as it should), but such cases would slightly increase our reported FRR .

In Part 2, we tested all images from $\mathcal{D}^{(f)}(c, u)$. The total number of tests in Part 2 was

$$\sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{D}_f(c)} (|\mathcal{D}^{(f)}(c, u)|) = \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{D}_f(c)} (|\mathcal{D}(c, u)| - 50) = \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{D}_f(c)} |\mathcal{D}(c, u)| - 50 \sum_{c \in \mathcal{C}} |\mathcal{D}_f(c)| = 1,041,382 - 50N_f = 700,032.$$

Images taken with digital zoom are not likely to pass this test (we can tell if a digital zoom was engaged by looking into the EXIF header). The search for the zoom factor is generally needed to identify their source camera and we refer to [14] where such search is described. Surprisingly, some digitally zoomed images were correctly identified, which is probably because we did not restrict them to fall in $\mathcal{D}^{(f)}(c, u)$. Consequently, some fingerprints may be a superposition of the camera fingerprint estimates from regular images and from up-sampled (digitally zoomed) images. Nevertheless, to clean our experiments and to be consistent with our assumption that the images are in their native resolution and uncropped, we eliminated in Part 2 all images with a positive indication in the digital zoom ratio tag in their EXIF header (about 0.2%). The normalized histogram of PCE_2 from this test is in Figure 1 left. This empirical pdf can be used to determine FRR as a function of τ (see Figure 2).




Figure 1. Histogram of $\log_{10}(PCE_2)$ for the fingerprint matching images (left) and for non-matching images (right). The right tails fall double exponentially. The solid line is the pdf (12) with $\lambda = 0$, while the dashed line shows the fit with $\lambda = 0.3$.

4.4 Part 3: Images not matching the camera, different camera model

The main purpose of this test is to determine a relation between the decision threshold τ and the FAR when the camera fingerprint is not the correct one. For each $\hat{\mathbf{K}}_{c,u}$, we randomly chose 150 images from $\bigcup_{c' \in C, c' \neq c} \mathcal{D}(c', u)$ and evaluated

PCE_2 . The total number of tests in Part 3 was $N_f \times 150 = 1,024,050$. The histogram of PCE values is shown in Figure 1 right. The FAR falls sharply as the threshold τ approaches 60 (also see Figure 2). The very good separation is further reflected in the Receiver Operating Characteristic curve (ROC) shown in Figure 3. Since the largest accounted PCE_2 was $\tau_{\max} = 57.267$, we cannot experimentally estimate the FAR for $\tau > \tau_{\max}$. For example, the detection rate $DR = 97.62\%$ and $FAR < 10^{-6}$ for $\tau = 60$. Figure 1 right also shows that the values of PCE_2 exhibit much thicker tails than model (12) predicts. We attribute this to the fact that the parameters $\sigma_x^2, \sigma_y^2, a, b$ and thus λ in (11) are not constant across images. For example, the variance of the noise residual, σ_y^2 , heavily depends on the image content. Thus, PCE_2 should be modeled as a mixture of non-central chi-square distributed random variables, where λ follows some distribution $p_\lambda(x)$, which could be determined experimentally. On the other hand, further decorrelation of the noise residuals and camera fingerprints may help close the gap between the experiment and mathematical models.




Figure 2. Left: histograms from Figure 1 overlaid in one graph. Right: relationship between error rates and the decision threshold.




Figure 3. Left: ROC of CSI for native resolution images from flickr.com (camera fingerprints estimated from 50 images). Right: the same in semi-logarithmic plot.

4.5 Part 4: Images not matching the camera, the same camera model

In the next test, we studied the situation when the tested images do not match the camera while they originate from exactly the same camera model. For each $\hat{\mathbf{K}}_{c,u}$ and all $c \in \mathcal{C}$, we randomly chose 150 images from $\bigcup_{v \neq u} \mathcal{D}(c, v)$. Here, hypothesis H_0 is supposed to be true even though we have no guarantee of it. Some pairs of users, u, v , usually family members, share one camera or they exchange pictures while having two cameras (the same model). Then H_1 is true for images in $\mathcal{D}^{(f)}(c, v)$ and the fingerprint $\hat{\mathbf{K}}_{c,u}$ although $u \neq v$. The total number of tests in Part 4 was the same as in Part 3, i.e., 1,024,050.




Figure 4. Histograms of PCE_2 for images not matching the camera fingerprint. Different camera model (left), matching camera model (right).

Comparing the histograms in Figure 4, one can see an almost perfect match in the shape and size. The only difference is in the right tail of Part 4, which contains occasional large values. We contribute them to the cases with shared cameras and/or shared pictures between users. We were able to resolve a number of these false “false alarms” as we found evidence of picture sharing and large overlaps between the users to which the fingerprints belonged. Despite our lack of

resources to resolve every above-threshold PCE value, we did not find any noticeable difference between the “same model FAR” and the overall FAR for any threshold τ less than 50.

The second and stronger argument is that any hardware based similarities between cameras of the same model would show up as a significant difference between the shapes and locations of the histograms in Part 3 where camera models match, and Part 4 where camera models do not match. We conclude that this CSI performs with the same accuracy independently of whether or not the tested cameras are of a different model or brand.

4.6 Overall error rates

The results from Part 3 and 4 need to be merged to estimate the overall error probabilities associated with the test when both the image and the camera are chosen randomly. What we need is the prior probability, p_{match} , of the image and the fingerprint coming from the same camera model. For our database \mathcal{D} , $p_{\text{match}} = 0.0523$, and the overall error probabilities are

$$\begin{aligned} FAR &= p_{\text{match}} FAR(\text{Part 4}) + (1 - p_{\text{match}}) FAR(\text{Part 3}) \\ FRR &= p_{\text{match}} FRR(\text{Part 4}) + (1 - p_{\text{match}}) FRR(\text{Part 3}). \end{aligned}$$

The results reported in Sections 4.4 and 4.5 for $\tau = 60$ lead to the following overall $FRR = 2.38\%$ and $FAR = 2.4 \times 10^{-5}$. This is the most conservative error estimate for the database \mathcal{D} , while the real FAR may turn out to be zero after verifying that every false acceptance case was a false false alarm.

4.7 Part 5: Searching for the source camera among all 6,827 camera fingerprints

Apart from testing the camera fingerprint against many images, we tested images against the database of fingerprints, a task that has never been carried out on a large scale. A list of 145 images $\mathbf{I}_1, \dots, \mathbf{I}_{145}$ was randomly selected to span across 145 different camera models. The entire database of N_f camera fingerprints was searched for the source camera for each image from the list. The total number of tests in Part 5 was $N_f \times 145 = 989,915$.

Having 6,827 PCE_2 values for each image, we plot their maximum value in Figure 5 (left) while setting the threshold based on the results in Part 3 to $\tau = 60$. Below the threshold τ (dashed lines in Figure 5) are maxima for 8 images. Four of them still identify the correct camera, however the evidence is weak. All the other maxima correctly identified the source camera out of 6,827. There was one more large PCE value that was the second largest for one of the images (Figure 5 right). After some more investigation, this “double positive” (or one case of false “false acceptance”) was proved to be a correct identification. The two fingerprints indeed belong to the same camera; we found that the two users share some specific images. A sample of a typical log PCE plot appears in Figure 6 left.




Figure 5. Maximum PCE_2 for each image (left), and the second largest PCE_2 indicating a double positive (right).




Figure 6. Left: The “no doubt” correct camera for image #1. Right: a double positive for image #143.

Then we looked closely at all 8 cases where no camera was identified as the source and pinpointed the culprit.

- I₁₅** - Digital zoom tag reads 22/10. Application of the search for scaling and cropping [14] revealed the exact digital zoom ratio 2.291.
- I₅₇** - An author of this photo post-processed most if not all the snapshots on an Apple computer. Obvious processing ranges from colorizing, reduction of color depth, to color quantization. Pixel de-synchronization operations may or may not have taken place. A large portion of images in the fingerprint estimation are dark night shots, which decreased the quality of the fingerprint for this camera-user.
- I₇₁** - All other images from the same $\mathcal{D}(c, u)$ set that are in landscape orientation are easily identifiable with $PCE > 3,000$ for all of them. No images taken in the portrait orientation match the fingerprint unless we crop out the first 7 rows of the fingerprint and the last 7 rows of the images. Images in portrait orientation are apparently taken with the active part of the sensor shifted by 7 rows when compared with landscape shots. Image **I₇₁** is one of the portrait images. The camera c is HP Photosmart R707, the only one camera model with such phenomenon.
- I₉₇** - A dark night shot, average luminance is 65. FRR for this user is 14/150, 7 out of these 14 have a low average luminance < 70 . Dark images have very weak traces of the camera fingerprint due to the multiplicativity of PRNU.
- I₁₁₁** - A poor quality fingerprint was estimated from about 1/3 of problematic images, 64/200 of them were photoshopped, 5/200 digitally zoomed. PCE was still close to the threshold. Green and blue channels in the image are saturated in a large portion of all pixels.
- I₁₁₂** - The user $FRR = 23/63$ in Part 2 with some images being a complete miss while others easily identified. We do not understand what kind of image processing was involved that prevented identification. Three images out of 23 are digitally zoomed, for other 20 including **I₁₁₂** the method failed probably after some pixel desynchronization process.
- I₁₂₈** - This image is highly compressed to 2 bits/pixel in Photoshop. Other 64/139 images of this user were also photoshopped, $FRR = 5/89$ was still not too high for such a camera-user.
- I₁₃₉** - This image is half black, half highly textured. Other dark images were the cause of the poor quality fingerprint. The user $FRR = 38/154$.

5. CONCLUSION

We presented a large scale experimental evaluation of camera identification based on sensor fingerprint. The test database of images contained over one million pictures taken by 6896 cameras covering 150 camera models and the total number of camera-fingerprint tests was 3,038,015. The experiments established upper bounds on error rates of the camera identification method: false rejection rate less than 0.0238 at false acceptance rate below 2.4×10^{-5} . In experiments, the test statistics (PCE) exhibited significantly thicker tails when compared to theoretical models. The problem was traced to the fact that the distribution of the test statistics is a complicated mixture, which could be

estimated experimentally. The experiments confirmed that error rates do not increase across cameras of the same model, which indicates that current methods aimed at removing non-unique systematic artifacts from fingerprints are effective. By inspecting detection failures, we determined that the most important factor contributing to missed detection is the quality of images used for fingerprint estimation.

One of the outcomes we were hoping to obtain was to identify camera models for which the camera identification works less reliably. Unfortunately, we could not accomplish this task because the most influential factor among cameras that we tested was the quality of the estimated fingerprint or particular habits of photographers. The variations in the images content did not allow us to establish any clear dependence between median PCE and the camera make or the sensor physical size or such.

A vast amount of data that was collected will be further analyzed and utilized. We intend to optimize the system parameters and use the data and the database \mathcal{D} as a benchmark.

6. ACKNOWLEDGEMENTS

This research was supported by an NSF award CNF-0830528.

REFERENCES

- [1] Chen, M., Fridrich, J., and Goljan, M.: “Digital Imaging Sensor Identification (Further Study),” *Proc. SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX*, vol. 6505, San Jose, California, pp. 0P–0Q, 2007.
- [2] Chen, M., Fridrich, J., Goljan, M., and Lukáš, J.: “Determining Image Origin and Integrity Using Sensor Noise,” *IEEE Transactions on Information Security and Forensics*, vol. 1(1), pp. 74–90, March 2008.
- [3] Khanna, N., Mikkilineni, A.K., Chiu, G.T.C., Allebach, J.P., and Delp, E.J.: “Scanner Identification Using Sensor Pattern Noise,” *Proc. SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX*, vol. 6505, San Jose, CA, pp. 1K–1L, 2007.
- [4] Gloe, T., Franz, E., and Winkler, A.: “Forensics for Flatbed Scanners,” *Proc. SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents IX*, vol. 6505, San Jose, CA, pp. 1I–1J, 2007.
- [5] Filler, T., Fridrich, J., and Goljan, M.: “Using Sensor Pattern Noise for Camera Model Identification,” *Proc. IEEE, ICIP 08*, San Diego, CA, September 2008.
- [6] Goljan, M.: “Digital Camera Identification from Images – Estimating False Acceptance Probability,” *Proc. IWDW 2008*, Busan, South Korea, November 10–12, 2008.
- [7] Janesick, J. R.: *Scientific Charge-Coupled Devices*, SPIE PRESS Monograph, vol. PM83, SPIE–The International Society for Optical Engineering, January, 2001.
- [8] Holst, G. C.: *CCD Arrays, Cameras, and Displays*, 2nd edition, JCD Publishing & SPIE Pres, USA, 1998.
- [9] Healey, G. and Kondepudy, R.: “Radiometric CCD Camera Calibration and Noise Estimation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 16(3), pp. 267–276, March, 1994.
- [10] Holt, C.R.: “Two-Channel Detectors for Arbitrary Linear Channel Distortion,” *IEEE Trans. on Acoustics, Speech, and Sig. Proc.*, vol. ASSP-35(3), pp. 267–273, March 1987.
- [11] Fridrich J., Goljan M., and Chen, M.: “Identifying Common Source Digital Camera from Image Pairs,” *Proc. ICIP’ 07*, San Antonio, Texas, September 14–19, 2007.
- [12] Kay, S.M., *Fundamentals of Statistical Signal Processing*, Volume II, Detection theory, Prentice Hall, 1998.
- [13] Kumar, B.V.K.V., and Hassebrook, L., “Performance measures for correlation filters,” *Applied Optics*, vol. 29(20), pp. 2997–3006, 1990.
- [14] Goljan, M. and Fridrich, J.: “Camera Identification from Cropped and Scaled Images,” *Proc. SPIE, Electronic Imaging, Forensics, Security, Steganography, and Watermarking of Multimedia Contents X*, vol. 6819, San Jose, California, January 28–30, pp. 0E-1–0E-13, 2008.

APPENDIX

Table. List of all 150 camera makes and models taken from the EXIF header (converted to lower case) together with the total number of images for each model.

Camera Model	Dim	$ \mathcal{D}(c,u) $	Camera Model	Dim	$ \mathcal{D}(c,u) $	Camera Model	Dim	$ \mathcal{D}(c,u) $
canon eos 10d	3072×2048	1198	fujifilm FP 2600 z	1600×1200	633	olympus c150	1600×1200	385
canon eos Rebel	3072×2048	10599	fujifilm FP 3800	2048×1536	333	olympus c300z	1984×1488	333
canon PS a40	1600×1200	667	fujifilm FP a340	2272×1704	400	olympus c310z	2048×1536	600
canon PS a400	2048×1536	1017	fujifilm FP a345	2304×1728	16590	olympus c350z	2048×1536	13597
canon PS a430	2272×1704	861	fujifilm FP e550	2848×2136	1164	olympus c4000z	2288×1712	1295
canon PS a510	2048×1536	3047	fujifilm FP f10	2848×2136	897	olympus c460z	2288×1712	400
canon PS a520	2272×1704	121498	fujifilm FP f30	2848×2136	345	olympus c50z	2560×1920	200
canon PS a530	2592×1456	83	fujifilm FP s5000	2816×2120	9327	olympus c740uz	2048×1536	1582
canon PS a540	2816×2112	3880	fujifilm FP s5100	2272×1704	1183	olympus c750uz	2288×1712	1085
canon PS a60	1600×1200	3466	fujifilm FP s7000	2848×2136	3560	olympus c765uz	2288×1712	800
canon PS a610	2592×1944	2440	fujifilm FP s9000	3488×2616	7372	olympus s300	2048×1536	39944
canon PS a70	2048×1536	3353	hp photosmart 735	2048×1536	400	olympus s400	2272×1704	680
canon PS a75	2048×1536	2791	hp photosmart r707	2592×1952	144	olympus s410	2272×1704	29496
canon PS a80	2272×1704	3507	kodak c330 z	2304×1728	1228	olympus s600	2816×2112	1012
canon PS a85	2272×1704	2326	kodak cx6330 z	2032×1524	500	olympus sv	2272×1704	206
canon PS a95	2592×1944	3491	kodak cx7300	2080×1544	18086	panasonic dmc-fx01	2816×2112	30451
canon PS g2	2272×1704	625	kodak cx7330 z	2032×1524	749	panasonic dmc-fx7	2560×1920	14946
canon PS g3	2272×1704	620	kodak cx7430 z	2304×1728	388	panasonic dmc-fx9	2816×2112	487
canon PS g5	2592×1944	1747	kodak dx4330	2160×1440	335	panasonic dmc-fz20	2560×1920	596
canon PS g6	2048×1536	487	kodak dx4530 z	2580×1932	1099	panasonic dmc-fz30	2048×1536	597
canon PS s110	1600×1200	1204	kodak dx6490 z	2304×1728	735	panasonic dmc-fz5	2560×1920	1208
canon PS s1 is	2048×1536	2472	kodak dx7440 z	2304×1728	586	panasonic dmc-fz7	2816×2112	29997
canon PS s200	1600×1200	1424	kodak dx7590 z	2576×1932	978	panasonic dmc-tz1	2560×1920	1055
canon PS s230	2048×1536	1288	kodak dx7630 z	2856×2142	399	pentax optio s4	2304×1728	621
canon PS s2 is	2592×1944	4548	kodak z740 z	2576×1932	31739	sony dsc-f828	2048×1536	399
canon PS s30	2048×1536	1333	minolta dimage x50	2560×1920	200	sony dsc-h1	2592×1944	923
canon PS s3 is	2816×2112	69683	minolta dimage xt	2048×1536	13832	sony dsc-h2	2816×2112	200
canon PS s400	2272×1704	1505	minolta dimage z1	2048×1536	17290	sony dsc-p10	2592×1944	479
canon PS s410	2272×1704	1034	nikon coolpix 2100	1600×1200	1029	sony dsc-p100	2048×1536	752
canon PS s45	2272×1704	713	nikon coolpix 3100	2048×1536	1708	sony dsc-p200	3072×2304	38840
canon PS s50	2592×1944	937	nikon coolpix 3200	2048×1536	43945	sony dsc-p41	2304×1728	283
canon PS s500	2592×1944	1924	nikon coolpix 4100	2288×1712	504	sony dsc-p72	2048×1536	1165
canon PS s60	2592×1944	374	nikon coolpix 4300	2272×1704	31483	sony dsc-p73	2304×1728	1116
canon PS sd10	2272×1704	200	nikon coolpix 4600	2288×1712	52513	sony dsc-p8	2048×1536	663
canon PS sd100	2048×1536	1531	nikon coolpix 5200	2592×1944	1338	sony dsc-p92	2592×1944	711
canon PS sd110	2048×1536	679	nikon coolpix 5600	2592×1944	988	sony dsc-s40	2304×1728	169
canon PS sd200	2048×1536	1775	nikon coolpix 5700	2560×1920	918	sony dsc-s500	2816×2112	200
canon PS sd30	2592×1944	773	nikon coolpix 5900	2592×1944	488	sony dsc-s600	2816×2112	242
canon PS sd300	2272×1704	2259	nikon coolpix 775	1600×1200	898	sony dsc-t3	2592×1944	489
canon PS sd400	2592×1944	89844	nikon coolpix l1	2816×2112	577	sony dsc-t5	2592×1944	363
canon PS sd450	2592×1944	2211	nikon coolpix l3	2592×1944	1240	sony dsc-t7	2592×1944	1022
canon PS sd600	2816×2112	29813	nikon coolpix l4	2272×1704	1240	sony dsc-t9	2816×2112	306
canon PS sd630	2816×1584	197	nikon coolpix s1	2048×1536	282	sony dsc-v1	2592×1944	1112
canon PS sd700 is	2816×2112	2544	nikon d100	3008×2000	724	sony dsc-w1	2592×1944	915
canon PS sd750	3072×2304	107	nikon d40	3008×2000	118778	sony dsc-w30	2816×2112	1215
casio ex-s500	2560×1920	132	nikon d50	3008×2000	4432	sony dsc-w5	2592×1944	326
casio ex-s600	2816×2112	1178	nikon d70	3008×2000	4032	sony dsc-w50	2816×2112	34895
casio ex-z50	2560×1920	600	nikon d70s	3008×2000	1800	sonyericsson k750i	1632×1224	958
casio ex-z60	2816×2112	752	nokia n70	1600×1200	405	sonyericsson k800i	2048×1536	1709
casio ex-z750	3072×2048	154	nokia n73	2048×1536	1241	sonyericsson w810i	1632×1224	609

Abbreviations: 'PS' = powershot, 'Rebel' = digital rebel, 'FP' = finepix, 'z' = zoom.