

Determining Digital Image Origin Using Sensor Imperfections

Jan Lukáš, Jessica Fridrich, and Miroslav Goljan
Department of Electrical and Computer Engineering
SUNY Binghamton, Binghamton, NY 13902-6000

ABSTRACT

In this paper, we demonstrate that it is possible to use the sensor's pattern noise for digital camera identification from images. The pattern noise is extracted from the images using a wavelet-based denoising filter. For each camera under investigation, we first determine its reference pattern, which serves as a unique identification fingerprint. This could be done using the process of flat-fielding, if we have the camera in possession, or by averaging the noise obtained from multiple images, which is the option taken in this paper. To identify the camera from a given image, we consider the reference pattern noise as a high-frequency spread spectrum watermark, whose presence in the image is established using a correlation detector. Using this approach, we were able to identify the correct camera out of 9 cameras without a single misclassification for several thousand images. Furthermore, it is possible to perform reliable identification even from images that underwent subsequent JPEG compression and/or resizing. These claims are supported by experiments on 9 different cameras including two cameras of exactly the same model.

Keywords: Fixed pattern noise, digital camera identification, forensic, pattern noise, pixel non-uniformity

1. INTRODUCTION

In this paper, we ask the following questions: Is it possible to find an equivalent of gun identification from bullet scratches for identification of digital cameras from images? How reliably can we distinguish between images obtained using different sensors or cameras? Is reliable identification possible from processed images?

In classical film photography, there are methods for camera identification that are commonly used in forensic science. Some of these methods use camera imperfections, such as scratches on the negative caused by the film transport mechanism. As digital images and video continue to replace their analog counterparts, reliable, inexpensive, and fast identification of digital image origin increases on importance. Reliable digital camera identification would especially prove useful in the court. For example, the identification could be used for establishing the origin of images presented as evidence, or, in a child pornography case, one could prove that certain imagery has been obtained using a specific camera and is not a computer-generated image.

The process of image identification can be approached from different directions. On the most obvious and simplest level, one could inspect the electronic file itself and look for clues in headers or any other attached or associated information. For example, the EXIF header contains a plethora of direct information about the digital camera type and the conditions under which the image was taken (e.g., exposure, time, etc.). Additional information could be obtained from the quantization table in the JPEG header (some cameras use customized quantization matrices). This header data, however, may not be available if the image is resaved in a different format or recompressed. Another problem is the credibility of information that can be easily replaced. Recognizing this problem, there has been some effort in the digital watermarking community to embed an invisible fragile watermark (Epson PhotoPC 700/750Z, 800/800Z, 3000Z), robust, or visible watermark (Kodak DC 290) in the image that would carry information about the digital camera, a time stamp, or even biometric of the person taking the image [1]. While the idea to insert the "bullet scratches" in the form of a watermark directly into each image the camera takes is an elegant and empowering solution to the camera identification problem, its application is limited to a closed environment, such as "secure cameras" used by forensic experts taking images at crime scenes. Under these controlled conditions, the secure cameras can provide a solution to the problem of

evidence integrity and origin. This watermark solution, however, cannot solve our problem in its entirety unless all cameras insert watermarks in their images.

Another approach to camera identification is analysis of pixel defects. In [2], the authors point out that defective pixels, such as hot pixels or dead pixels, could be used for reliable camera identification even from lossy compressed images. This approach fails for cameras that do not contain any defective pixels or cameras that eliminate defective pixels by post-processing their images on-board. Also, the defective pixels may not be obvious in every scene and thus, in order to identify the defective pixels, one either needs to have access to the camera or have sufficiently many images from which the defective pixels can be determined.

Kharrazi et al. [3] proposed a different idea in which a vector of numerical features is extracted from the image and then presented to a classifier built from a training set of features obtained from images taken by different cameras. The feature vector is constructed from average pixel values, correlation of RGB pairs, center of mass of neighbor distribution, RGB pairs energy ratio, and it also exploits some small scale and large scale dependencies in the image expressed numerically using a wavelet decomposition previously used for image steganalysis [4]. This “blind identification” appears to work relatively reliably. Tests on a larger number of cameras could bring a more decisive answer if this approach has a good distinguishing power with respect to different cameras. One of the concerns is whether this method is capable of distinguishing between similar camera models (e.g., models with the same sensor type) or between cameras of the exactly same model. Also, the large number of images needed to train a classifier for each camera may not always be available.

In this paper, we propose another approach to digital camera identification that uses the pattern noise of CCD arrays. The pattern noise is caused by several different factors, such as pixel non-uniformity, dust specs on optics, interference in optical elements, dark currents, etc. Using the denoising filter described in [5], we extract the high frequency part of the pattern noise and then use correlation (as in robust watermark detection using spread spectrum) to evaluate the presence of the pattern noise in the image. We show that this approach is an effective and reliable way to identify the camera from its images even after lossy compression or resizing. This approach is computationally simple and there is no need to train classifiers on a large number of images. Also, this approach is able to distinguish between cameras of the exact same model.

2. DIGITAL CAMERAS

We first briefly describe the operation of a digital camera as well as particular cameras used in our experiments.

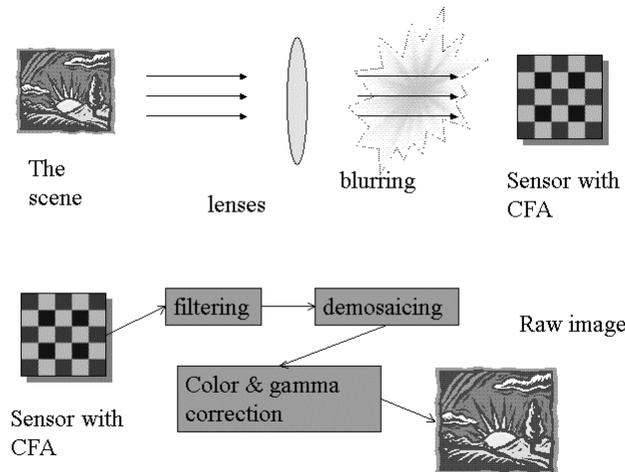


Figure 1: Image acquisition using a typical digital camera.

2.1 Signal processing inside a digital camera

In the classical film camera, the light from the scene passes through the lenses and interacts with a photo responsive film. Similarly, in a typical consumer digital camera the light from the photographed scene passes through the camera lenses, but before reaching a photo responsive sensor, the light goes through an antialiasing (blurring) filter and then through a color filter array (CFA). The photon counts are converted to voltages, which are subsequently quantized in an A/D converter. This digital signal is interpolated (demosaiced) using color interpolation algorithms. The colors are then processed using color correction and white balance adjustment. Further processing includes high-pass filtering and gamma correction to adjust for the linear response of the imaging sensor. Finally, the raw image is written to the camera memory device in a user-selected image format (e.g., TIFF or JPEG). The picture-capturing process is displayed schematically in Figure 1.

The imaging sensor is a device that converts light into electrical signal. Currently, the most frequently used sensors are charge-coupled devices (CCD), but CMOS sensors are quickly emerging as another alternative. There are also cameras on the market that use JFET sensor (Nikon D2H digital SLR) or the Foveon X3 sensor (Sigma SD9 and SD10 digital SLRs); the latter is also based on CMOS technology. In order, to process the picture, the sensor is divided into very small minimal addressable picture elements (pixels) that collect photons and transfer them into voltage. The Foveon X3 sensor is the only sensor that is able to capture all three basic colors at every pixel. All remaining sensors can capture only one particular basic color at any single pixel and the remaining colors must be interpolated. The CFA assigns each pixel its appropriate color it is supposed to capture. There are many types of CFAs; one example is shown in Figure 1.

R	G	R	G	R
G	B	G	B	G
R	G	R	G	R
G	B	G	B	G

Figure 1: Bayer color filter array.

2.2 Cameras and test images

Camera brand	Sensor	Maximal resolution	Image format
Canon PowerShot A10	1/2.7-inch CCD	1280×960	JPEG
Canon PowerShot G2	1/1.8-inch CCD	2272×1704	CRW, JPEG
Canon PowerShot S40	1/1.8-inch CCD	2272×1704	CRW, JPEG
Kodak DC290		1792×1200	TIFF, JPEG
Olympus Camedia C765 UZ - 1	1/2.5-inch CCD	2288×1712	TIFF, JPEG
Olympus Camedia C765 UZ - 2	1/2.5-inch CCD	2288×1712	TIFF, JPEG
Nikon D100	23.7×15.5 mm Nikon DX CCD	3008×2000	NEF-RAW, TIFF, JPEG
Sigma SD9	20.7×13.8 mm CMOS-Foveon X3	2268×1512	X3F-RAW
Olympus Camedia C3030	1/1.8-inch CCD	2048×1536	TIFF, JPEG

Table 1: Cameras used in experiments and their properties.

For our work, we prepared an image database containing 320 images from each camera listed in Table 1 with a large variety of outdoor and indoor scenes, including close-ups and landscapes.

The Canon PowerShot G2 and Canon PowerShot S40 both have the same CCD sensor type. Both cameras have similar features but different optics. The images were stored in the proprietary Canon raw format (*.CRW) and then converted using Canon Utilities RAW image converter version 1.2.1 to the 24-bit true color TIFF format.

All images taken with the Kodak DC290 camera were taken in the uncompressed 24-bit true color TIFF format.

The Nikon D100 images were taken in the Nikon proprietary NEF-RAW format. All images were converted by Nikon Capture 4.0 into the 8-bit TIFF format.

Although the Olympus Camedia C3030 camera supports TIFF images, the maximum size SmartMedia™ memory card it supports could only store 5 TIFF images. This is the reason why we collected the database images in the JPEG format.

We have purchased two cameras of the *exact same model* – the Olympus Camedia C765 Ultrazoom. Images taken with these cameras were stored in the uncompressed 24-bit true color TIFF format.

The Sigma SD9 is a semiprofessional digital SLR camera with the CMOS-Foveon X3 sensor with 10 million photo detectors. It only outputs images in its proprietary X3F-RAW format. All images were converted by the supplied software Sigma PhotoPro 2.1 using automatic processing. For almost half of these images, however, the processing had to be manually adjusted to produce pleasingly looking images. The results were saved in the 8-bit TIFF format.

3. PATTERN NOISE

There are many sources of noise in images obtained using CCD arrays, such as dark current, shot noise, circuit noise, fixed pattern noise, etc. [6],[7]. In this paper, we are only interested in the systematic part of the noise that does not change from image to image and is relatively stable over the camera life span and a reasonable range of conditions (temperature). The only noise components that are not reduced by frame averaging [6] are fixed pattern noise and photoresponse non-uniformity noise, together referred to as pattern noise, pixel noise, or pixel non-uniformity [6].

Even an absolutely evenly lit CCD element will exhibit small changes in charge collected by individual pixels. This is partly because of the shot noise, which is a random component, and partly because of pattern noise – a component that is approximately the same for each image. Pattern noise is introduced during the manufacturing process. It is also influenced by the clock bias [7]. The pattern noise is typically of the order of several percent measured by photon transfer and it varies for different CCDs.

The fixed pattern noise (FPN) is one part of the pattern noise caused by dark currents. It primarily refers to the pixel-to-pixel differences when the CCD array is in the dark. The FPN is an additive noise and thus can be corrected for by subtraction of dark frame [7], which is a reference pattern usually obtained by averaging dark images (covered lenses) obtained with a similar exposition as the scene image. Denoting the raw scene image as X and the dark frame D , the FPN correction is performed as $X \leftarrow X - D$. Some middle to high-end consumer cameras automatically subtract the dark frame every time the camera is turned on. Cheaper cameras may not perform this step.

The photoresponse non-uniformity noise (PRNU) is usually the dominant part of the pattern noise. It is caused mainly by pixel non-uniformity, which is a signal primarily concentrated in high spatial frequencies. Light refraction on dust particles, on optical surfaces, and the CCD element itself are low frequency signals that also contribute to the photoresponse non-uniformity noise. Because in our work, we use a denoising filter to extract the noise, we essentially work with the portion of the PRNU noise caused by pixel non-uniformity.

The linear response of CCDs enables a simple correction of PRNU using a process called flat fielding [6], [7], in which the image is divided by a normalized reference pattern extracted from a uniformly lit scene. It is noted in [6] that simple images of uniform blue sky should be sufficient for this purpose. Denoting the flat-fielding frame F , the noise correction process (including the dark frame subtraction) is expressed as $X \leftarrow \frac{X - D}{F}$. This processing must be done before any further nonlinear image processing is performed. Most consumer cameras do not flat-field their images because it is not simple to achieve a uniform sensor illumination inside the camera.

Other imaging sensors, CMOS, JFET, or CMOS-Foveon X3, are also built from semiconductors using a process similar to manufacturing CCDs. Therefore, it seems reasonable to expect that the pattern noise in these sensors will exhibit similar properties. Our experiments with the CMOS-Foveon X3 based Sigma SD9 confirm the presence of pattern noise that survives frame averaging.

4. DETECTION BY CORRELATION

As explained in the introduction, assuming the high-frequency part of pixel non-uniformity noise is an iid Gaussian signal, we formulate the camera identification problem as detection of high-frequency low-amplitude spread spectrum Gaussian noise watermark – the pattern noise. To improve the low watermark-to-image SNR, prior to detection we extract the noise component from the image using a denoising filter F_σ [5]. This filter extracts from the image a Gaussian noise with variance σ^2 based on the assumption that in the wavelet domain, the image and the noise form an additive mixture of a non-stationary Gaussian signal and a stationary Gaussian signal with a known variance. This filter is one of the best filters currently available for image denoising.

Let Y and $F_\sigma(Y)$ denote the spatial representation of the image and its denoised version, respectively. The difference signal $Y - F_\sigma(Y)$ is an approximation to the pattern noise.

For each camera C , we first determine its reference pattern noise P_C . To decide whether a specific image Y was taken by camera C , we calculate the correlation ρ_C between its noise pattern $Y - F_\sigma(Y)$ with the reference noise pattern P_C

$$\rho_C(Y) = \text{corr}(Y - F_\sigma(Y), P_C) = \frac{(Y - F_\sigma(Y) - E[Y - F_\sigma(Y)]) \cdot (P_C - E[P_C])}{\|Y - F_\sigma(Y) - E[Y - F_\sigma(Y)]\| \|P_C - E[P_C]\|}, \quad (1)$$

where $E[\]$ stands for the mean value.

The performance of the detector is not very sensitive to the filter parameter σ , as long as $\sigma > 1$. After experimenting with different values of the parameter σ , we finally settled at $\sigma=5$ as the value that gave us the best performance across all cameras.

There are two types of identification problems one can encounter in practice. The first, and easier, problem is to determine from several cameras the camera that most likely took a given image. This can be achieved simply by assigning the image to the camera whose reference pattern has the highest correlation with the noise from the image. In our experiments with 9 cameras, we have not observed any misclassification in tests performed on more than 3000 images including low-resolution low quality JPEGs.

The second, harder, problem is to evaluate the evidence that a given image was taken by a specific camera. In this case, it is necessary to compare the value of ρ for images produced by the camera and other cameras and to determine an appropriate measure (e.g., threshold) for ρ_C to reach a conclusion about the origin of the image. The proper measures and thresholds are discussed in Section 6.

4.1 Reference camera pattern

The reference camera pattern can be obtained in several different ways depending on whether or not we have the camera in possession. We obtained the best results by averaging the noise extracted from multiple images to eliminate the influence of the scenes on the output of the denoising filter. The more images we use, the more accurate approximation to the pattern we obtain and the more reliable our identification becomes. In most of our experiments, the reference pattern was obtained from approximately 300 images of natural scenes. In Figure 2, we show the average correlation ρ between the pattern noise from 20 test images with the reference pattern from Olympus C765 obtained by averaging pattern noise from N_p images. One can see that the correlation falls pretty quickly once the number of images for calculating the reference pattern falls below roughly 50, which we recommend as the minimal number of images for calculating the reference pattern.

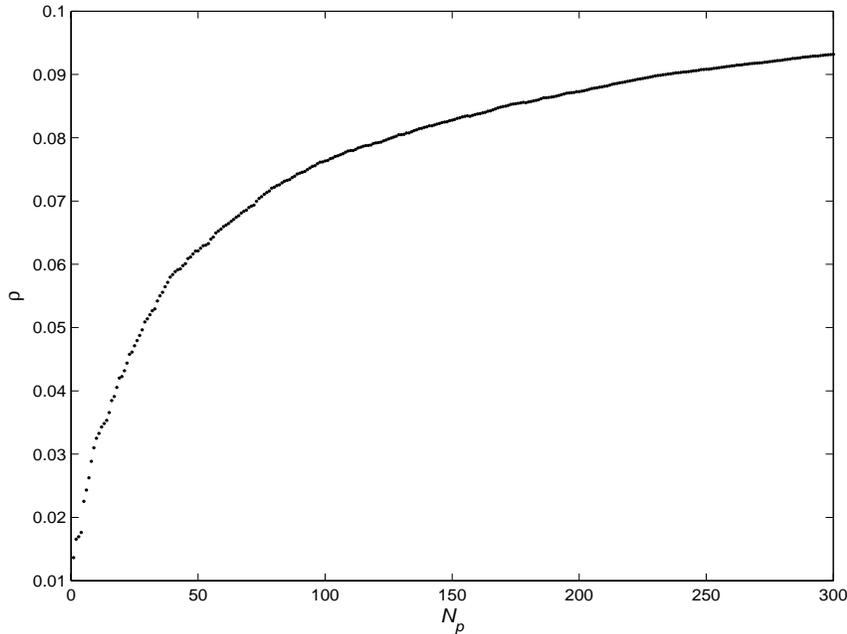


Figure 2: Average correlation ρ as a function of the number of pictures N_p used for calculating the reference pattern for Olympus C765 TIFF images.

We have also experimented with other methods for calculating the reference pattern, such as denoising using a median filter or using the process of flat-fielding. By comparing the results for both experiments it appears that the best choice for calculating the reference pattern is the one based on averaging the noise component from multiple images of natural scenes. This method is also advantageous because it does not require us to have the camera in possession. One problem with extraction of the reference pattern using flat-fielding is that it should be done for the raw sensor data before demosaicing and other processing. Most consumer-end cameras, however, do not allow access to this data.

4.2 Geometrical transformations

Since our camera identification is essentially based on detection of a high frequency watermark using correlation, geometrical transformations of images, such as rotation, cropping, resizing, or fish-eye lens processing cause desynchronization and must be corrected for prior to calculating the correlation. One very common geometrical transformation is rotation by 90 degrees. Since we cannot say which way the user held the camera, we rotated each image both ± 90 degrees, extracted the noise (we can rotate the extracted noise only), computed 2 correlations, and then used the higher value. Originally, we experienced problems with images that were automatically rotated in the Kodak DC290 camera. A simple experiment revealed that the in-camera rotation introduces a shift by one pixel.

5. EXPERIMENTS

5.1 Image identification (raw format)

Figure 3 shows the correlation of the noise from approximately 300 Olympus C765-1 JPEG (left) and 300 Olympus C3030 TIFF (right) images of natural scenes with the reference pattern from all 9 cameras (for each camera, the reference pattern was obtained by averaging the noise component from 300 images). The 20 images on the right were not used for the computation of reference pattern and serve as a validation set. The correct reference pattern always produces the highest correlation value.

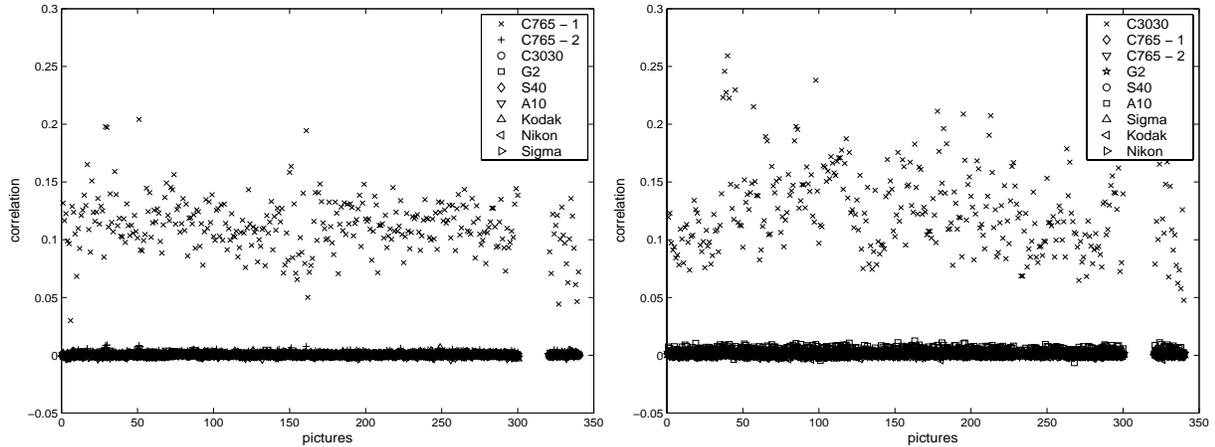


Figure 3: Correlation of noise from Olympus C765 (left) TIFF and Olympus C3030 (right) JPEG images with 9 reference patterns.

Note that it is possible to distinguish between two cameras of the same brand (two Olympus C765 cameras). This is not surprising as the dominant component of the pattern noise is the pixel non-uniformity, which is a stochastic phenomenon introduced in the manufacturing process and thus different from camera to camera. The outlier image (from Olympus C765) with a low correlation of approximately 0.03 on the left is a completely dark scene. This low correlation is to be expected because the dominant component of the reference noise pattern is suppressed in dark images. When computing the correlation between an image noise and a reference pattern of different sizes, the larger of the two has always been cropped to the smaller size before calculating the correlation. Another option here is to resize one of the noise patterns and this possibility is investigated in Section 5.3.

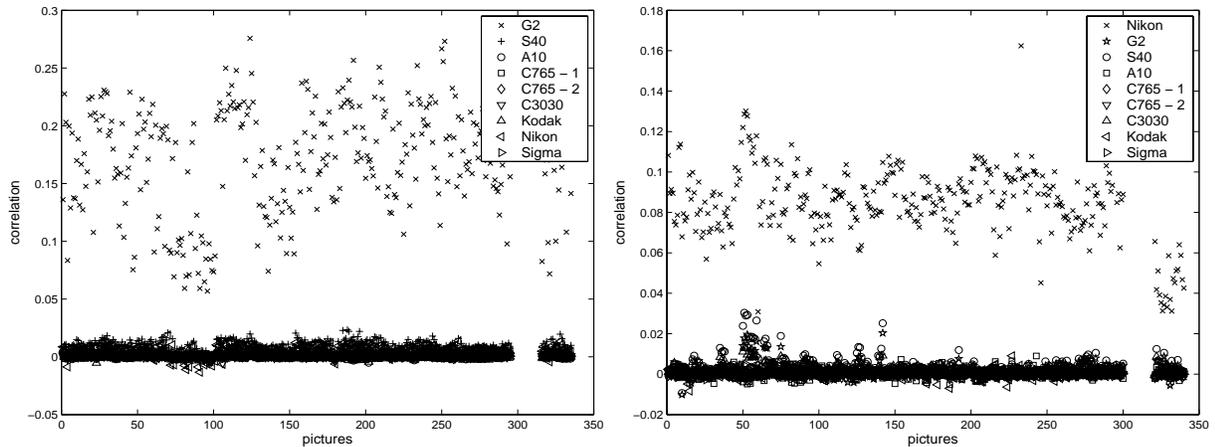


Figure 4: Correlation of noise from Canon G2 (left) and Nikon D100 (right) TIFF images with 9 reference patterns.

The outlier image from Nikon D100 with index 60 (right) is again a completely dark image of a night sky with the Moon in the middle. Also, we have observed a small positive correlation between the noise from the Canon G2, S40, and Nikon D100, which can be seen in Figure 4. We hypothesize that this is caused by similarities in post processing in these cameras, which leads to some residual dependence between corresponding reference patterns. This issue is investigated in more detail in Section 6.

5.2 Image identification after JPEG compression

We now investigate the influence of JPEG compression on camera identification. Since most of the energy of the pixel non-uniformity pattern is in high spatial frequencies, we expect that ρ_C will decrease with JPEG

compression. While this is, indeed, true, at the same time, the variance of ρ_C for an incorrect reference pattern also decreases.

In Figure 5, the circles stand for the mean of ρ_C between the noise from Kodak DC290 images with the reference pattern from Kodak DC290. Diamonds correspond to the mean of the correlation between the noise from Canons G2 & S40 images and the reference pattern from Kodak DC290. The figure shows how the mean and the standard deviation of correlations decrease with decreasing JPEG quality factor. For this experiment, uncompressed TIFF images were converted to grayscale (to reduce computation time and complexity of the computer experiments).

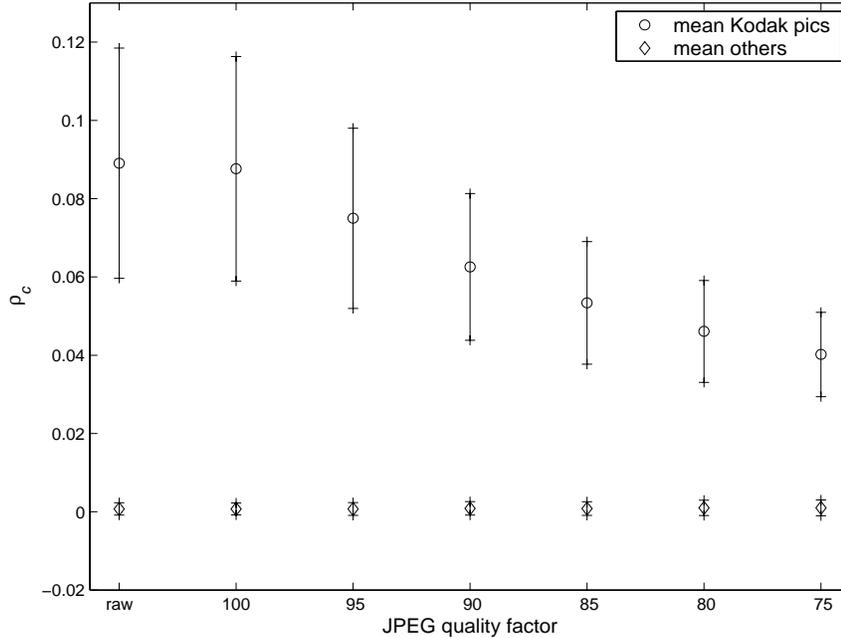


Figure 5: The mean and standard deviation of ρ_C as a function of the JPEG quality factor.

Continuing the description of the JPEG experiment, out of 320 images for each camera, we have randomly selected 100 images from each camera and computed the reference pattern from them. The reference pattern was computed only once from uncompressed images. Then we took another set of randomly selected 100 images, JPEG compressed them with different quality factors and computed correlations with the reference patterns. We can see that both the mean and the standard deviation of correlations of noise patterns with the correct reference pattern decrease with JPEG compression but the variance of correlations with incorrect reference patterns remains almost constant. We conclude that it is possible to obtain reliable camera identification even after subsequent JPEG compression.

5.3 Image identification after resampling and JPEG compression

In this section, we investigate whether it is possible to identify images obtained at a lower resolution than the maximal resolution. This problem is more involved than it might seem at the first sight. Let us take a look at, for instance, 1600×1200 JPEG images with the task to determine the camera that captured them, given a set of cameras. First of all, we should note that Nikon D100, Sigma SD9, and Kodak DC290 produce images with an aspect ratio 3:2 instead of 4:3, as for the rest of our cameras, and thus cannot take 1600×1200 images. Also, the Canon PowerShot A10 has the maximum supported resolution smaller than 1600×1200. The Olympus C765 camera presents another challenge. The ratio of its largest supported image size is not exactly 4:3. Although, this camera supports 1600×1200 images simply resizing them into maximum resolution supported by this camera would slightly distort the aspect ratio.

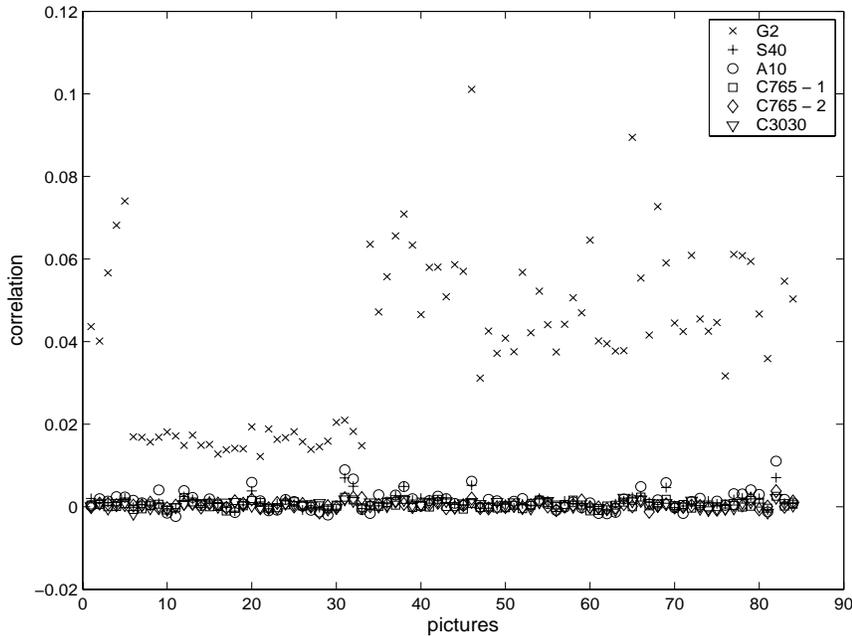


Figure 6: Identification of low-resolution 1600×1200 Canon G2 JPEG images.

We repeat that our problem is to find for a given 1600×1200 image, which camera from our set took the image. Let us also assume that we have a prior knowledge that images were taken at lower resolution or were rescaled in the computer, but have not been cropped. For Canon G2, S40, A10, and Olympus C3030, we simply resample the image to the maximum resolution that each camera supports and then we compute the correlation with its reference pattern (which means downgrading the image for comparison with Canon A10 reference pattern). Bicubic resampling worked well in our experiments. We are aware that Canon A10 normally does not produce images in such resolution, however because it has the same aspect ratio, we have included it in the experiment. For Olympus C765, we also resample the image into the maximum resolution 2288×1712. Although, it slightly distorts the aspect ratio, our investigation revealed, that the Olympus camera does produce lower resolution images this way. Nikon D100, Sigma SD9, and Kodak DC290 do not support 1600×1200 images and are therefore not included in this experiment.

Figure 6 shows the results of such experiments. Noises from Canon G2 low-resolution 1600×1200 JPEG images were correlated (after resizing) with the high-resolution reference patterns obtained from 6 different cameras. Images No. 6–33 were compressed with a very low JPEG quality factor (around 72) while the rest of the images were compressed using JPEG with an average quality factor of around 98. This test indicates that it is possible to correctly identify the camera even from JPEG images taken at lower resolution.

5.4 Forgeries and malicious processing

Since we expect that the identification techniques are likely to be used in the court, we need to address issues, such as removing the pattern noise from an image to prevent identification as well as extracting the noise and copying to another image to make it appear as though the image was taken with a particular camera. Although this problem is currently under investigation, in this section we provide some preliminary results.

Addressing the issue of intentionally removing the pattern noise, we have tried to extract the noise from a denoised image. Although the correlations with the correct pattern were about two times smaller, the property that the highest correlation occurs for the reference pattern from the correct camera was always fulfilled (see Figure 7). The 20 extra images on the right were not used for reference pattern computation and serve as validation set. This experiment also indicates that the denoising filter does a sub-optimal job in extracting the pattern noise.

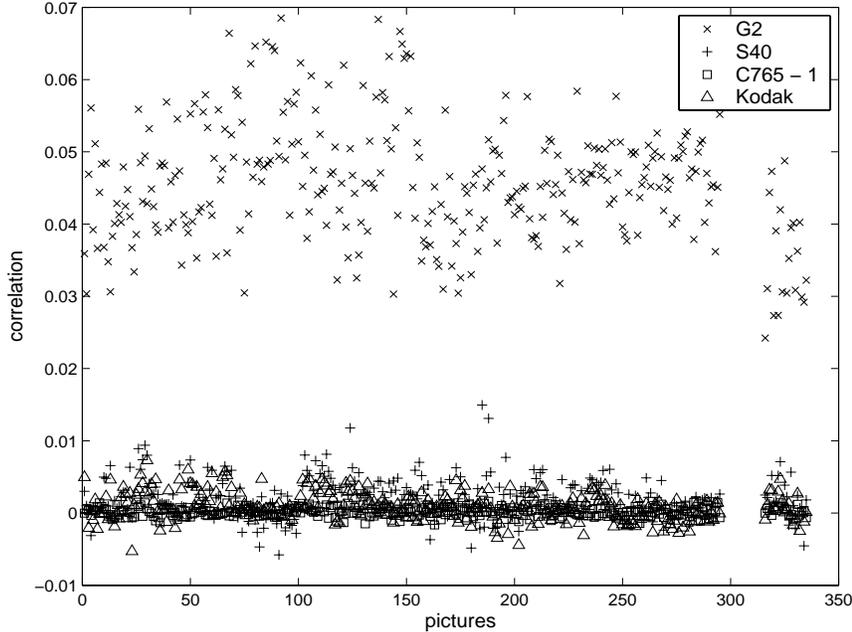


Figure 7: Correlation of denoised Canon G2 images with reference patterns from four cameras.

As noted in Section 3, the pattern noise could be theoretically removed by flat-fielding and dark frame subtraction. However, since the images were processed in the camera using non-linear operations and operations that combine neighboring pixel values (e.g., demosaicing, color correction, or high-pass filtering), it is in general not possible to perform the flat fielding correctly without accessing the raw CCD data before demosaicing.

The pattern noise is a low-amplitude high-frequency “natural watermark” and as such, it is well known that the easiest way to prevent its detection is desynchronization, such as slight rotation, possibly combined with other processing that might include resizing, cropping, and JPEG compression. We also point out that it is known from digital watermarking [8] that anybody with an access to the reference pattern P_C can arrange for $\rho_C = 0$ for any image Y taken with C by solving the equation $\rho_C(Y + \alpha P_C) = 0$ with respect to α and taking $Y + \alpha P_C$ as the forged image.

The second problem we now investigate is whether it is possible to make an arbitrary image look as if it was taken by a specific camera. Again, having access to the reference patterns or the cameras makes this indeed possible. We denoised 20 Canon G2 pictures and added to them the reference pattern from Canon S40. We increased the amplitude of the added noise, till we reached a correlation that was higher than the image previously had with Canon G2 reference pattern. The peak signal to noise ratio (PSNR) for the forged images was above 37.5dB and the images were visually indistinguishable from the originals. The forgeries did have slightly higher correlations with Canon G2 reference pattern than expected from different camera images, but this could be eliminated using some of the techniques mentioned above.

6. THEORY AND MODELS

We may attempt to model the extracted noise from an image as well as the reference pattern as iid Gaussian signals. Also, if two signals are coming from two different cameras, they should be statistically independent. In practice, however, these assumptions are not satisfied. It can be shown that the correlation of two iid independent Gaussian signals of length N is approximately Gaussian $N(0, 1/N)$ for large N . For example, the standard deviation for the correlations between two independent Gaussian signals of the length corresponding to Canon G2 and Olympus C765 reference patterns should be approximately 3×10^{-4} . Thus, the probability of obtaining the

measured value of 2.1×10^{-3} (see Table 2) is, in fact, less than 10^{-12} . Therefore, we must conclude that most values in the table are highly improbable to occur given our assumptions about the pattern noise.

	Nikon	C765-1	C765-2	G2	S40	Sigma	C3030	Kodak	A10
Nikon	1	0.0017	-0.0001	0.0335	0.0497	0.0082	0.0198	0.0030	0.0034
C765-1	0.0017	1	0.0215	0.0009	0.0034	0.0017	0.0018	0.0036	0.0032
C765-2	-0.0001	0.0215	1	0.0021	0.0025	-0.0006	0.0002	0.0050	0.0014
G2	0.0335	0.0009	0.0021	1	0.0579	0.0051	0.0072	0.0047	0.0060
S40	0.0497	0.0034	0.0025	0.0579	1	0.0060	0.0104	0.0064	0.0086
Sigma	0.0082	0.0017	-0.0006	0.0051	0.0060	1	0.0044	0.0055	0.0064
C3030	0.0198	0.0018	0.0002	0.0072	0.0104	0.0044	1	0.0019	0.0452
Kodak	0.0030	0.0036	0.0050	0.0047	0.0064	0.0055	0.0019	1	0.0052
A10	0.0034	0.0032	0.0014	0.0060	0.0086	0.0064	0.0452	0.0052	1

Table 2: Mutual correlations of 9 reference patterns.

There are multiple reasons for this discrepancy. First, the wavelet denoising filter F_σ for example assumes that the image in the wavelet domain is a non-stationary Gaussian signal and the pattern noise is a stationary Gaussian signal. Since these assumptions are satisfied only approximately, the pattern noise extracted using the denoising filter is not Gaussian, either. Another problem is that the filter is applied to the image on slightly overlapping blocks and it also pads image borders with zeros. This leads to a small residual dependence between all extracted noises. Furthermore, we point out, that even camera reference patterns are often slightly correlated due to similar or even the same image processing algorithms inside the cameras.

Thus, coming back to the problem (Section 4) of evaluating the absolute (rather than relative) evidence that a given image was taken by a specific camera, the decision thresholds should be obtained from measurements rather than models. In particular, we should estimate the distribution of correlation of pattern noise with incorrect reference patterns from the data. For instance, let us look at Canon G2 RAW images. The correlations with non G2 reference patterns (see Figure 5 left) have mean 0.0026 and standard deviation 0.0042. We can estimate proper thresholds for evaluating the evidence for the G2 camera by modeling the correlations with incorrect reference patterns as a Gaussian random variable $N(0.026, 0.0042)$. However, in our 2360 correlations with non G2 reference patterns, there appear to be several outliers that suggest that the Gaussian assumption is not valid and the true distribution of correlations with incorrect reference patterns is non-Gaussian with larger tails.

It seems that the only feasible solution would be to determine the decision thresholds experimentally. This means, we have to use a larger number of images from many cameras and divide them into groups. One group should be used for reference pattern computations, the second for threshold selection, and the third group for performance testing. Also, it appears that different thresholds would be needed for JPEG images with different quality factors, lower resolution images, etc.

7. CONCLUSIONS

We present a new approach to the problem of camera identification from images. The identification is based on pixel non-uniformity noise, which is a unique stochastic characteristic for both CCD and CMOS based cameras. The presence of this noise is established using correlation as in detection of spread spectrum watermarks. Reliable identification is possible even from images that were resampled and JPEG compressed. By testing the approach on 9 different digital cameras, we were able to correctly match several thousand images to the correct camera. We were also 100% successful in distinguishing between images taken by two cameras of the same model.

While the proposed identification technique can be used for reliable camera identification from original images or images processed using JPEG compression or resizing, simultaneous application of other geometrical operations (e.g., cropping resizing, rotation) causes desynchronization and thus increases the computational complexity of pattern detection because the detection will likely have to resort to brute force searches. At this point, we would like to point out that reliable camera identification should be approached from multiple directions, combining the

evidence from other methods, such as the feature-based identification [3], which is less likely to be influenced by geometrical transformations. We also may be able to retrieve information about geometrical operations using the technique described in [9].

Our future research will include further investigation of the effect of malicious tampering as discussed in Section 5. Also, we plan to investigate if it is possible to use our technique to identify tampered areas (forgeries). Finally, we intend to extend this technique to the problem of scanner identification.

ACKNOWLEDGEMENTS

The work on this paper was supported by Air Force Research Laboratory, Air Force Material Command, USAF, under a research grant number F30602-02-2-0093. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation there on. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, or the U. S. Government. Special thanks belong to Taras Holotyak for providing us with Matlab code for the denoising filter. We would also like to thank to Peter Burns, James Adams, Chris Honsinger, John Hamilton, and George Normandin for many useful discussions.

REFERENCES

- [1] Blythe, P. and Fridrich, J.: "Secure Digital Camera", *Digital Forensic Research Workshop*, Baltimore, August 11–13, 2004.
- [2] Geradts, Z., Bijhold, J., Kieft, M., Kurosawa, K., Kuroki, K., and Saitoh, N.: "Methods for Identification of Images Acquired with Digital Cameras", *Proc. of SPIE, Enabling Technologies for Law Enforcement and Security*, vol. 4232, pp. 505–512, February 2001.
- [3] Kharrazi, M., Sencar, H. T., and Memon, N.: "Blind Source Camera Identification", *Proc. ICIP' 04*, Singapore, October 24–27, 2004.
- [4] Farid, H. and Lyu, S.: "Detecting Hidden Messages Using Higher-Order Statistics and Support Vector Machines", in F.A.P. Petitcolas (ed.): *5th International Workshop on Information Hiding, LNCS vol. 2578*, Springer-Verlag, Berlin-Heidelberg, New York, pp. 340–354, 2002.
- [5] Holotyak, T., Fridrich, J., and Goljan, M.: "Estimation of Message Length Embedded Using ± 1 Embedding", *Proc. SPIE Electronic Imaging, Steganography, Security, and Watermarking of Multimedia Contents VII*, San Jose, California, January 16–20, 2005.
- [6] Holst, G. C.: *CCD Arrays, Cameras, and Displays*, 2nd edition, JCD Publishing & SPIE Pres, USA, 1998.
- [7] Janesick, J. R.: *Scientific Charge-Coupled Devices*, SPIE PRESS Monograph vol. PM83, SPIE–The International Society for Optical Engineering, January, 2001.
- [8] Cox, I., Miller, M.L., and Bloom, J.A.: *Digital Watermarking*, Morgan Kaufmann, San Francisco, 2001.
- [9] Popescu, A.C. and Farid H.: "Statistical Tools for Digital Forensic", in J. Fridrich (ed.): *6th International Workshop on Information Hiding, LNCS vol. 3200*, Springer-Verlag, Berlin-Heidelberg, New York, pp. 128–147, 2004.
- [10] Adams J., Parulski K., and Spaulding K.: "Color Processing in Digital Cameras", *Micro, IEEE*, vol. 18, No. 6, pp. 20–30, November-December 1998.
- [11] Geradts, Z. and Bijhold, J.: "Pattern Recognition and Image Processing in Forensic Science", <http://geradts.com/html/Documents/IMVIP00/IMVIP.htm>
- [12] Geradts, Z., Bijhold, J., and Kieft, M.: "Methods for Identification of Images Acquired with Digital Cameras", *AAFS*, Reno, February 2000, <http://forensic.to/camera/cameras.htm>