

DIGITAL “BULLET SCRATCHES” FOR IMAGES

Jan Lukáš, Jessica Fridrich, and Miroslav Goljan

Department of Electrical and Computer Engineering
SUNY Binghamton, Binghamton, NY 13902-6000
{jan.lukas, fridrich, mgoljan}@binghamton.edu

ABSTRACT

The problem investigated in this paper is identification of sensor that was used to obtain a given digital image. We show that the high-medium frequency component of the sensor pattern noise is an equivalent of “bullet scratches” for digital images and can be used for reliable forensic identification. For each sensor, we first calculate its reference pattern (an estimate of the sensor pattern noise) by averaging the noise component from multiple images. This pattern serves as a unique identification fingerprint whose presence in a given image is established using a correlation detector. The proposed identification technique was tested on several thousand images obtained by nine digital cameras. In all cases, we were able to correctly identify the camera that took the image. We also show that it is possible to identify the camera from images subjected to combined processing, including lossy JPEG compression, gamma correction, recoloring, and resizing.

1. INTRODUCTION

With the rapid spread of digital imaging technology, the importance of reliable sensor identification from digital images will only increase. This is especially true for establishing the origin of images presented as evidence in the court. Also, the identification method could be used to prove that certain imagery has been obtained using a specific camera and is not a computer-generated image (e.g., in a child pornography case).

The simplest method for identification of digital images is the inspection of the image file itself (e.g., its EXIF header), image dimensions, or the JPEG quantization table. This metadata, 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 so easily replaced.

Another alternative is analysis of the color gamut or detecting artifacts of color interpolation algorithms and other camera processing. This information, however, is likely to be very fragile and may not survive common image processing. Moreover, it may not be possible to

distinguish between cameras that share the same processing algorithms.

A few manufacturers, such as Epson and Kodak, introduced cameras that embed an invisible fragile watermark in their images to protect their integrity. There have also been proposals to embed 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 is intriguing, unless all cameras do the same, it is only applicable to a closed environment, such as “secure cameras” used by forensic experts taking images at crime scenes.

It is possible to use defective pixels (hot or dead pixels) for reliable camera identification even from JPEG compressed images [2]. This approach fails, however, for cameras that do not have any defective pixels or cameras that eliminate defective pixels by processing their images on-board.

A qualitatively different approach to image identification that has recently proved to be quite successful is to extract a set of “features” from images and train a classifier to distinguish between images from different cameras [3]. It remains to be seen whether this method can distinguish between cameras with the same sensor or cameras that share the same on-board processing algorithms. Another concern is that the large number of images needed to train a classifier for each camera may not always be available.

In our previous work [4], we have shown that the high-medium frequency component of the pattern noise of CCD/CMOS arrays can be used for very reliable camera identification and can even distinguish between cameras of the exact same model. An important property of the pattern noise is that its high-medium frequency component is practically scene independent and relatively stable over the camera life span. The pattern noise is caused by pixel non-uniformity, dust specs on optics, interference in optical elements, dark currents [5][6], etc. Using the denoising filter described in [7], we extract the high-medium frequency component 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 a given image. An important advantage of this approach is its simplicity while avoiding the need to train classifiers on a large number of images.

In this paper, we subject this methodology to further testing, expanding the experiments on images subjected to a combination of common image processing operations. Section 2 briefly explains the basic signal processing inside a digital camera. In Section 3, we briefly describe the camera identification method based on detection of pattern noise and, in Section 4, we present the results of experiments. The paper is summarized in Section 5.

2. DIGITAL CAMERA SIGNAL PROCESSING

In a typical consumer-end 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 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 in most cameras (cameras with Foveon X3 sensor are exceptions) interpolated (demosaiced) using color interpolation algorithms. The colors are then processed using color correction and white balance adjustment. Further processing includes low-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, JPEG, or some other proprietary format).

3. NOISE PATTERN DETECTION

Assuming the high-frequency part of the pixel non-uniformity noise is an iid Gaussian signal $N(0, \sigma^2)$, the camera identification problem is detection of an iid Gaussian signal corrupted by noise – the image. Since there are no good statistical models of images in the spatial domain, we perform the signal-noise separation in the wavelet domain, where the image is modeled as an additive mixture of a non-stationary Gaussian signal (the denoised image) and a stationary Gaussian signal of a known variance (the pattern noise). Using the Gaussian denoising filter F_σ described in [7], we extract from the image an approximation to the high frequency part of the pattern noise. Denoting Y and $F_\sigma(Y)$ the spatial representation of the image and its denoised version, respectively, we take the difference signal $Y - F_\sigma(Y)$ as an approximation to the pattern noise.

The first step in camera identification is determining the reference camera pattern – the high-frequency component of its sensor pattern noise. The standard approach to obtain this pattern is using flat-fielding [6]. This needs to be done for the raw sensor data

before color interpolation and other on-board processing. Most consumer-end cameras, however, do not allow access to this data. Therefore, we opted for a different approach and extract the reference pattern by averaging the noise $Y - F_\sigma(Y)$ extracted from multiple images Y to eliminate the influence of the scene content on the output of the denoising filter and suppress other random noise. In most of our experiments, the reference pattern was obtained from about 300 images of natural indoor and outdoor scenes. Based on our previous work [4], we recommend at least 50 images for computation of the reference pattern. Fewer images might suffice if one can take images of uniformly lit scenes (e.g., blue sky shots).

Let P_C denote the reference noise from camera C . To decide whether image Y was taken by camera C , we calculate the correlation ρ_C between the image noise pattern $Y - F_\sigma(Y)$ and the reference 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 and $\|\cdot\|$ is the L_2 norm.

The detector performance is fairly insensitive to the filter parameter σ , as long as $\sigma > 1$. The value $\sigma = 5$ gave us the best overall performance.

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 by assigning the image to the camera whose reference pattern had the highest correlation with the image noise.

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_C 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.

4. EXPERIMENTS

In our experiments, we have used 9 cameras from different manufacturers with a variety of sensors and resolutions. They included Canon G2, Canon S40, Canon A10, Kodak DC290, Olympus C3030, Olympus C765 (two cameras of this same model), Sigma SD9 (with CMOS Foveon X3 sensor), and Nikon D100. With each camera, a little over 300 images in the RAW format (whenever possible) of natural indoor and outdoor scenes were taken, producing a database of more than 2700 images for our tests.

In the first set of tests, we have calculated the correlation between the noise from all images and all nine reference patterns. In all cases, the reference pattern from the camera that took the image produced the highest correlation value. As a representative example, we show in Figure 1 the correlation of the Olympus C765 images with reference patterns from all nine cameras. Note that it is possible to distinguish between two cameras of the same brand (two Olympus C765 cameras).




Figure 1: Correlation of noise from Olympus C765 images in TIFF format with 9 reference patterns.




Figure 2: The correlation (ρ_c) as a function of the JPEG quality factor.

In another experiment (performed only with 3 cameras – Canons G2 and S40 and Kodak DC290), we tested whether it is possible to perform reliable identification after lossy JPEG compression. Figure 2 shows the correlations between the DC290 reference pattern with the noise from JPEG images from all three cameras as a function of the JPEG quality factor. The crosses stand for the ρ_c between the reference pattern

from Kodak DC290 and the noise from Kodak DC290 JPEG images. Circles correspond to the correlation between the reference pattern from Kodak DC290 and the noise from Canon G2 and S40 JPEG images. For each quality factor, correlations were calculated for 100 randomly selected images from each camera. Squares denote the mean and plus signs mark the standard deviation. We can see that both the mean and the standard deviation of correlations of noise patterns with the correct reference pattern decrease with the decreasing quality factor while 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.

Next, we tested whether it is possible to identify images that were taken by a camera set to a lower resolution and saved as JPEG. For each camera in the test, whenever the dimensions of the noise pattern and the reference pattern did not match, the smaller of both was resized using bi-cubic interpolation to allow correlation computation.

Our experiments on 1600x1200 Canon G2 JPEG images (including those with JPEG quality factor around 72) revealed that reliable identification is still possible even for low quality low-resolution images.

Additionally, we have tested the possibility of reliable image identification from gamma corrected images. The pattern noise can be considered as a spread spectrum watermark [8] with most energy in the high-medium spatial frequencies. Since such watermarks are known to well survive point intensity transformations of the type of gamma correction, it is not a big surprise that correcting Olympus C765 TIFF images for gamma 1/1.4 decreased the separation displayed in the Figure 1 only negligibly. In our final test, we have gamma corrected low-resolution decompressed JPEG images.




Figure 3: Identification of low-resolution 1600x1200 Canon G2 gamma 1/0.7 corrected JPEG images.

In Figure 3, the noise from Canon G2 low-resolution 1600×1200 gamma 1/0.7 corrected JPEG images was 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. Again, in all cases, the camera that took the image produced the highest value of the correlation.

We conclude that the pattern noise survives well both linear and nonlinear point intensity transformations, such as histogram operations, brightness/contrast adjustment, or gamma correction.

It is also quite obvious that simultaneous application of several geometrical operations (e.g., cropping, resizing, and rotation) causes desynchronization and prevents easy detection, which now has to resort to expensive brute force searches.

In general, we believe that 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].

5. CONCLUSIONS

We present a new approach to the problem of camera identification from images based on pixel non-uniformity noise, which is a unique stochastic characteristic for both CCD and CMOS sensors. The presence of this noise is established using correlation as in detection of spread spectrum watermarks. Reliable identification is possible even from images processed using lossy JPEG compression, resizing, and point intensity transformations (e.g., gamma correction).

As the identifying pattern is a semi-robust spread-spectrum watermark, it should not be surprising that malicious processing becomes possible if an attacker has the camera in possession or has access to many images taken by the camera. In particular, we have shown [4] that it is possible to remove the pattern noise from an image beyond detection using (1) or plant a given reference pattern into another image.

An interesting question that we are currently investigating is whether it is possible to use the proposed technique on smaller blocks for identification of tampered areas (forgeries).

6. ACKNOWLEDGEMENT

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. We would like to thank Taras Holotyak for providing the code for the denoising filter, and to Paul Blythe, Peter Burns, James Adams, Chris Honsinger, John Hamilton, and George Normandin for many useful discussions.

7. REFERENCES

- [1] Blythe, P. and Fridrich, J.: “Secure Digital Camera”, *Digital Forensic Research Workshop*, Baltimore, Maryland, 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] Lukáš J., Fridrich J., and Goljan M.: “Determining Digital Image Origin Using Sensor Imperfections”, *Proc. SPIE Electronic Imaging, Image and Video Communication and Processing*, San Jose, California, pp. 249–260, January 16–20, 2005.
- [5] Holst, G. C.: *CCD Arrays, Cameras, and Displays*, 2nd edition, JCD Publishing & SPIE Pres, USA, 1998.
- [6] Janesick, J. R.: *Scientific Charge-Coupled Devices*, SPIE PRESS Monograph vol. PM83, SPIE—The International Society for Optical Engineering, 2001.
- [7] Mihcak M.K., Kozintsev, I., and Ramchandran, K.: “Spatially adaptive statistical modeling of wavelet image coefficients and its application to denoising,” in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing*, Phoenix, Arizona, vol. 6, pp. 3253–3256, March 1999.
- [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.