On the Reliability of Cell Phone Camera Fingerprint Recognition

Martin Steinebach[1], Mohamed El Ouariachi[1], Huajian Liu[1], and Stefan Katzenbeisser[2]

[1] Fraunhofer SIT, Darmstadt, Germany
www.sit.fraunhofer.de
[2] CASED, Darmstadt, Germany
www.cased.de

Abstract. Multiple multimedia forensic algorithms have been introduced allowing tracing back media copies back to its source by matching artifacts to fingerprint databases. While this offers new possibilities for investigating crimes, important questions arise: How reliable are these algorithms? Can a judge trust their results? How easy are they to manipulate? It has been shown that forensic fingerprints of digital cameras can be copied from one image to the next. Our aim is to develop new concepts for increasing the security of theses algorithms. In this work, we describe the state of our research work regarding attacks against forensics and provide an outlook on future approaches to increase their reliability.

Keywords: Camera forensics, fingerprints, copy attack, correlation modification.

1 Motivation

Multimedia forensic deals with the analysis of multimedia data to gather information on its origin and authenticity. One therefore needs to distinguish classical criminal forensics (which today also uses multimedia data as evidence) and multimedia forensics where the actual case is based on a media file. One area of multimedia forensics which has seen a lot of activity in recent years is the detection of physical devices used in the process of creating the multimedia data. This is also called multimedia ballistics (see e.g. [4]) in the style for classical ballistics for gun identification or source recognition.

But when using this technique in a truly forensic manner, the issue of reliability and security arises. So far most publications concentrate on the ability of identifying the source of a media data with a high probability. The only assumed attack often is degrading the quality of the media data, an approach similar to the evaluation of the robustness of watermarking algorithms. Still, first results like [5] by Gloe et al. show serious security risks, raising the question if the technology may be used in criminal investigations without the danger of facing simple but effective manipulations.

A simple scenario would be the following: Anne is a blog author, putting a lot of images made with her cell phone camera on the Internet. The images, and therefore the fingerprints of the camera, are available to anybody. And attacker Bert now wants to frame Anne. His plan is to accuse her of taking illegal photos with her cell phone.

S. Goel (Ed.): ICDF2C 2009, LNICST 31, pp. 69–76, 2010.

This can range from industrial espionage to child pornography. Bert downloads Anne's images and derives the fingerprint of Anne's camera. Then he gets hold of illegal photos, also via the Internet or by himself. This may be risky, but is often still rather easy. Now he copies the fingerprint of Anne's cell phone camera into the illegal photos. A ballistic detection would now positively respond to her camera. The only thing left to do is to distribute the images and accuse Anne of being the source. When Anne denies, a forensic analysis is suggested by Bert. The proof could be enough to ruin Anne. It becomes clear that the possibility of copying fingerprints can lead to a deadlock: Innocent persons can be framed; criminals can claim that proofs are the result of a forgery. At the end, the trust into the forensic approaches may be lost.

2 Camera Fingerprinting State-of-the-Art

In multimedia forensics, two areas have evolved. One research area deals with the origin of media data: It analyzes characteristics of the media data to find out which device created it; this includes approaches like camera, printer or scanner identification. The other research area deals with identifying content manipulations within the media data. This includes recognition of object removal or the combination of different images as well as utilitarian functions like the detection of scaling or multiple JPEG compression. In this work we focus on the first area, aiming to identify digital cameras. Similar to the goal of matching bullets and guns, it is therefore also called camera ballistics.

Source authentication tries to identify the source of a document, i.e. a specific camera, printer or scanner. For this purpose, artefacts of the hardware or the software of the device that created the media data can be utilized. In contrary to watermarking, these artefacts are not willingly embedded but are intrinsic to the creation process of the media data.

One example for utilizing software artefacts is the detection of JPEG quantization characteristics of different compression algorithms, as used in digital cameras [1]. In particular, quantization tables of the applied JPEG encoder can be identified, which are in 92% of all cases unique for one specific camera series. Some cameras choose from a set of potential tables depending on the nature of the image while others only use one table. Asymmetric tables are often applied, using different quantization values for the horizontal and the vertical direction. One important challenge in using the quantization table as a camera identifier is to perform identification even after the image was processed and re-compressed. The author of [1] shows that this is indeed possible.

Artefacts caused by characteristics of the hardware device used to create the media can also be used as a forensic hint. The most prominent example is digital camera recognition based on sensor noise matching as discussed by Fridrich et al. [2]. Digital cameras use CCD (charged coupled device) chips to convert light into a binary image during the A/D conversion process. Due to the production process these chips have some inherent characteristics and irregularities causing specific noise in the images produced by the camera. This noise, called pattern noise by Fridrich, can be used to match a camera model and in some cases even a specific camera. The approach requires a learning process where different images taken by one camera are used to

derive a statistical model of the characteristic pattern noise. Given this model, the presence of the noise can be verified for a given image, allowing to trace back an image to a camera it has been created with. An interesting fact in this type of forensics is that the reliability of the process is better with low-quality digital cameras as the cheaper production process causes more traceable irregularities in the chips. Similar approaches have been introduced to identify digital camcorders, printers and scanners. It has been shown by Fridrich et al. [3] that camera recognition is even possible after image printing.

Besides cameras, also algorithms for e.g. scanner [6] and laser printer [7], [8] identification have been introduced by authors, showing that at the end all devices creating data by sampling or reproducing data at the border between analogue and digital seem to be identifiable or traceable.

3 Cell Phone Image Fingerprint Implementation

As a starting point for our research, we implemented the approach introduced by Fridrich, Lucas and Golljan in [2]. We chose to focus on cell phone cameras, as they have two interesting characteristics: On the one hand they are widely available. This increases the probability that they will be used for evidence recordings. They are also common tools for a number of illegal photographic activities ranging from copyright violations to illegal erotic photography. On the other hand, the quality of the cameras is usually low compared to digital cameras. A low quality results in strong characteristic fingerprints due to low production standards.

Table 1. Cell phones and camera characteristics. C2 and C4 are two cell phones of identical type.

	Cell Phone	**Resolution**	**Flash**	**Autofocus**
C1	Nokia N95	2592x1944	Yes	Yes
C2	SonyEric. K550i	1632x1224	No	Yes
C3	Palm Treo 500v	1600x1200	No	No
C4	SonyEric. K550i	1632x1224	No	Yes
C5	Sony Eric K610i	1600x1200	No	No

To verify our implementation of [2] we used five cameras and took 300 photos with each of them creating a set of 1.500 photos, basically an array of the dimensions [300,5]. For each camera we then used the first 100 photos and calculated its average characteristic fingerprint. Then the correlation of this fingerprint was tested with the rest of the photos over all five cameras. The identification is successful and trustworthy if the correlation of the camera the photo was taken with and its average fingerprint is significantly higher then with the fingerprints of the rest of the cameras. Fig. 1 shows that this is the case for the vast number of photos. But one can observe also a number of samples where the correlation of the correct average fingerprint is similar or weaker than the fingerprint of the other cameras.

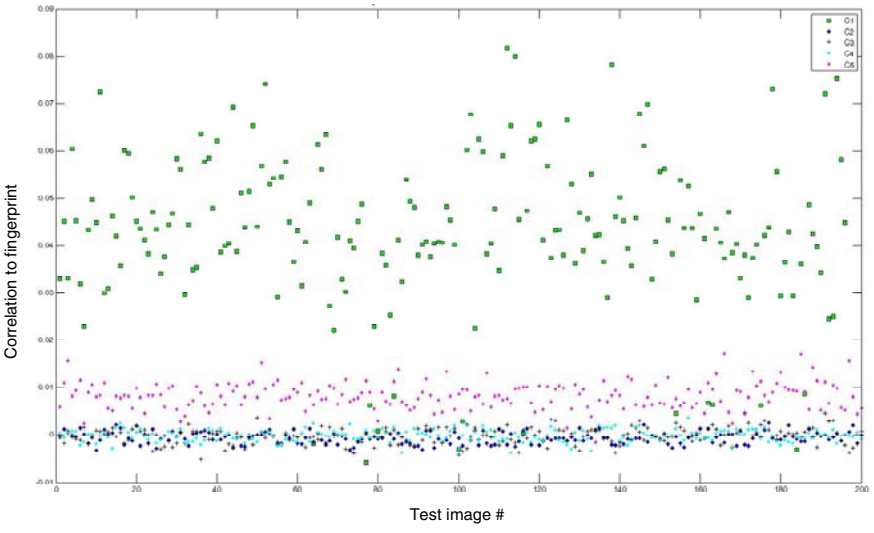

Fig. 1. Correlation results for 1000 images taken with C1 to C5 to the average fingerprint of C1

4 Attacking Cell Phone Image Fingerprints

To identify a camera, its average fingerprint is derived from a training set of images. The correlation of the noise within a given image to the average fingerprint is then used to identify the camera an image was taken with. This may enable a simple attack: One can calculate the fingerprint of one camera and then copy it into the image of another camera. Then the image could be said to be taken with the first camera

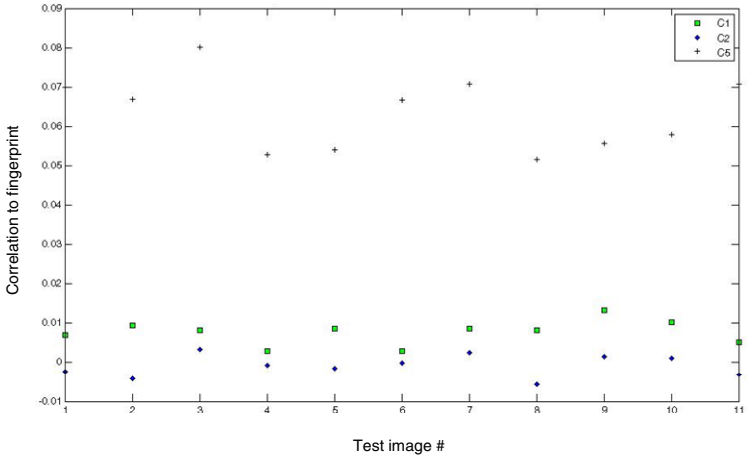

Fig. 2. Correlation without attack. All eleven test images taken with camera 5 show a significant peak at the correlation to the average fingerprint of camera 5

thereby potentially accusing its owner. This attack only requires access to a suitable number of photos from the first camera to calculate the average fingerprint. A first approach introducing the basic concept has been published by Gloe et al. in [7]. In our work, we verify their results, provide a large set of examples and describe a number of improvements masking the attack.

We analyze the behaviour of fingerprint correlations after the described attacks. For this, we use the set of five cameras from section 3 and 200 test images. The reference fingerprints of the cameras are derived from 100 photos taken with each camera. In the following figures we use 11 randomly selected images to illustrate our results.

Fig. 2 shows the state of the correlation without an attack. All 11 test images feature a high correlation to the average fingerprint of C5, the camera the images have been originally taken with.

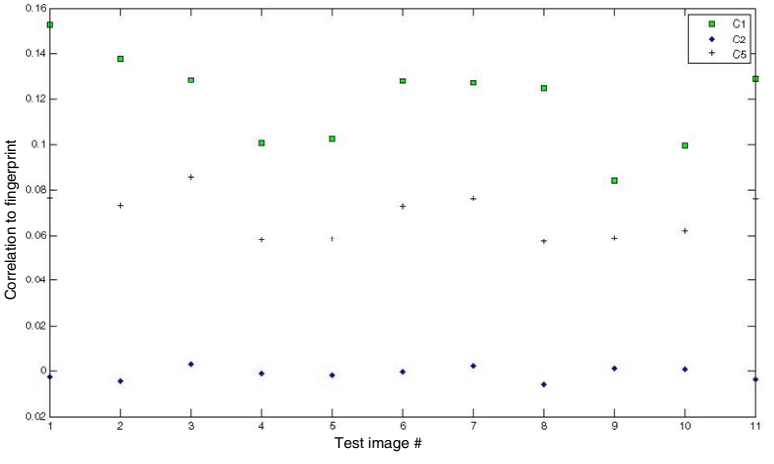

Fig. 3. Correlation after copy attack: When the average fingerprint of C1 is added to the original image, the strongest correlation is achieved with that fingerprint

The implementation of the attack is straight forward: In its most simple form, it takes the following steps:

1. Calculate average fingerprint F(K1) of camera K1 to be attacked
2. Take photo P with second camera K2
3. Copy F(K1) to photo P

Step 3 can be done with a simple addition of F(K1). If the dimensions of K(C1) and P do not match, tiling or cropping of F(K1) is necessary.

Fig. 3 shows the results of this attack. While the strength of the correlation between the images and the average fingerprint of C5 is not reduced, the correlation to the copied average fingerprint of C1 is about twice as strong as the original fingerprint. So while an attacker could successfully claim that a photo was taken with camera C1, a detailed analysis could show that there is also a weaker second fingerprint present within the image, giving a hint to the attack. If the attacker is the owner of the camera the photos of the attack have been taken with, the chance of confiscation and analysis of the camera's fingerprint exists. So the attacker will try to mask the usage of his camera.

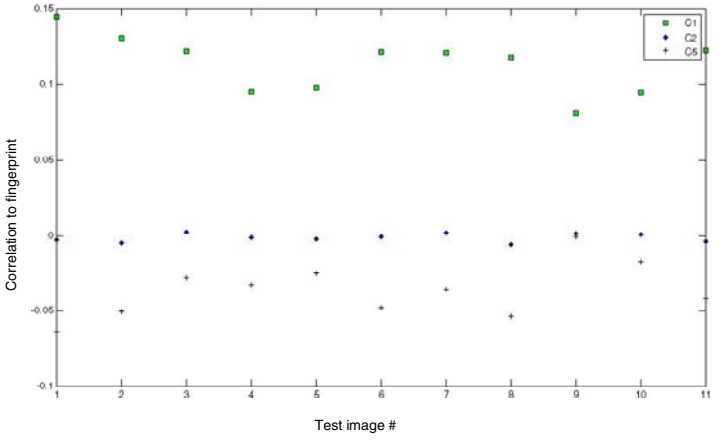

Fig. 4. Correlation after copy attack and suppression of original fingerprint. When the average fingerprint of the original camera C5 is subtracted from the image before adding the fingerprint of C1, its correlation becomes negative.

To achieve this, we present an advanced approach masking the usage of K2:

1. Calculate average fingerprint F(K1) of camera K1 to be attacked
2. Calculate average fingerprint F(K2) of camera K2
3. Take photo P with camera K2
4. Suppress F(K2) in P
5. Copy F(K1) to photo P

With suppressing F(K2) we ensure that there are not two correlation peaks when camera correlation is calculated. Therefore we first remove the original fingerprint and then add the new one. Fig. 4 shows the result of this attack. Removing the average fingerprint of C5 from images 1 to 11 ends up with a negative correlation, also a potential hint of the occurrence of an attack. The average strength of the negative correlation is about half the correlation of the original one.

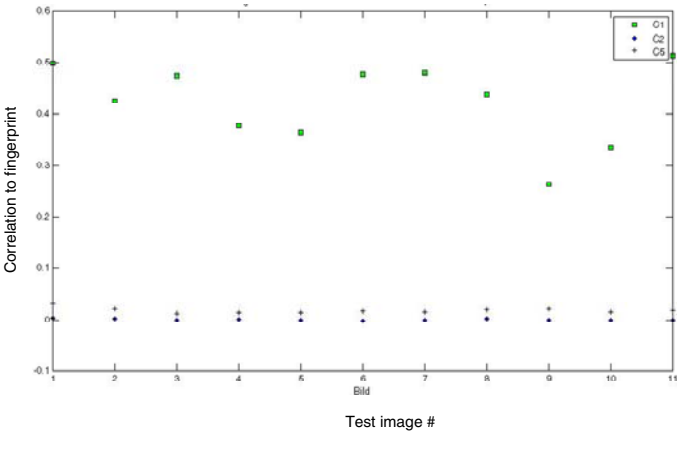

Fig. 5. Correlation after denoising and copy attack. The correlation of C5 is now successfully masked.

Our aim now was to improve the attack so that there is no hint of the camera used in the attack. Fig. 5 shows the results after applying a noise removal filter to P before suppressing F(K2) and copying the F(K1). Now the fingerprint of the camera the photo has been originally taken with is successfully masked, showing a correlation strength similar to the fingerprint of C2, a camera not participating in the attack.

To verify our test results, the same attack procedure took place with another set of cameras. Here C1 was the camera of the attacker and C2 was the camera to be attacked. The test results showed the same behaviour as above.

5 Outlook: Forensics beyond Ballistics

While the existing approach on camera ballistics seems to be vulnerable to attacks copying the fingerprint from one image to another, there may be a potential strategy for countering these attacks. Future research needs to identify if the fingerprints copied really fit to the images they are copied into.

The copied fingerprint is usually a result of averaging a large set of images. This means that the fingerprint represents the average noise caused by the individual CCD characteristics of the camera. It may be possible to distinguish average noise from noise belonging to a specific image. We assume that the noise is not independent from the image taken. Usually this noise will depend on shutter speed, lightness, temperature and so on. Now if it is possible to derive the influence of these characteristics on the fingerprint, one can compare the detected fingerprint with the fingerprint assumed to be typical for the given image. Only if both are similar, the fingerprint is really the one from the camera the photo has been taken with. Otherwise, a copy attack may have occurred.

Given this possibility of distinguishing original and copied fingerprints, attackers may start to set up illegal photos where the characteristics fit to those found within the photos from the person to be attacked. But this will at least significantly increase the efforts need to be taken to frame other persons by camera ballistics.

6 Summary

In this work we describe our first steps in improving the security of cell phone fingerprints. Cell phone cameras are applied as we see these as typical candidates for activities relevant for forensics. Due to their comparatively low quality they are also well suited for fingerprint calculation. We implement a widely known approach to derive fingerprints for photographs taken by digital cameras and provide test results showing the correct behavior of the implementation. Then we attack the fingerprints by suppressing the original fingerprint and copying another fingerprint into the photos. We achieve a state where the correlation of the copied fingerprint becomes significantly stronger than the original fingerprint.

This means that the current state of cell phone image forensics cannot be trusted and a significant amount of future research is necessary before this technology becomes trustworthy. Our direction from this point on is identifying the link between image and fingerprint. Similar to copy attacks in digital watermarking, a copied

fingerprint may show a high correlation to a characteristic average fingerprint, but on the other hand may feature a behavior in the given photo not fitting into the set of characteristics of that photo. We may end up with an additional mechanism calculating the probability of a fingerprint actually belonging to a given set of camera and photo.

References

1. Sorell, M.: Digital Camera Source Identification Through JPEG Quantisation. In: Li, C.-T. (ed.) Multimedia Forensics and Security, pp. 291–313. Idea Group Publishing (2008)
2. Fridrich, J., Lukáš, J., Goljan, H.: Digital Camera Identification from Sensor Noise. IEEE Transactions on Information Security and Forensics 1(2), 205–214 (2006)
3. Fridrich, J., Lukáš, J., Goljan, H.: Camera Identification from Printed Images. In: Electronic Imaging, Forensics, Security, Steganography, and Watermarking of Multimedia Contents X, San Jose, CA (2008)
4. Farid, H.: Digital Image Ballistics from JPEG Quantization. Technical Report TR2006-583 (2006), `http://www.cs.dartmouth.edu/farid/publications/tr06a.html`
5. Gloe, T., Kirchner, M., Winkler, A., Böhme, R.: Can We Trust Digital Image Forensics? In: Proceedings ACM Multimedia 2007, pp. 68–78 (2007)
6. Khanna, N., Chiu, G.T.C., Allebach, J.P., Delp, E.J.: Scanner Identification with Extension to Forgery Detection. In: Proceedings of the SPIE International Conference on Security, Steganography, and Watermarking of Multimedia Contents X, San Jose, CA (2008)
7. Suh, S., Allebach, J.P., Chiu, G.T.C., Delp, E.J.: Printer Mechanism-Level Information Embedding and Extraction for Halftone Documents: New Results. In: Proceedings of the IS&T's NIP 23: International Conference on Digital Printing Technologies, Anchorage, AK (2007)
8. Chiang, P., Mikkilineni, A., Delp, E.J., Allebach, J.P., Chiu, G.T.C.: Extrinsic Signatures Embedding and Detection in Electrophotographic Halftone Images through Laser Intensity Modulation. In: Proceedings of the IS&T's NIP 22: International Conference on Digital Printing Technologies, pp. 432–435 (2006)