Novel Security and Privacy Perspectives of Camera Fingerprints (Transcript of Discussion)

Jeff Yan(✉)

University of Lancaster, Lancaster, UK
jeff.yan@lancaster.ac.uk

I will talk about three very simple ideas about camera fingerprints. I didn't have time to put all the details in my slides, so please feel free to ask for any clarification or any question anytime. I didn't realize until it was too late that my laptop charger didn't work, and Ross came to my rescue just a while ago. I borrowed his charger and did some quick-hack slides.

First of all, some quick backgrounds about camera fingerprints. On the left-hand side, this is human fingerprints which is used for many security applications. On the right-hand side, this is a camera fingerprint which can be used to identify individual digital cameras.

The first question: Where does this camera fingerprint come from? We know that CCD or CMOS sensors are a digital camera's heart. Most cameras use CCD sensors, but some use CMOS sensors. All these sensors are made in semiconductor foundries. The sensor manufacturing process is critical, but never perfectly controlled.

Therefore, each sensor in a CCD or CMOS chip responds to light differently. The difference is not huge, and each pixel responds to light in a slightly different way. The effect collectively is that all pixels in the camera will leave systematic artifacts in each image created by this particular camera or video camera.

A way of modelling or capturing this slight variation under illumination is called PRNU or Photo Response Non-Uniformity. Basically, PRNU captures variations among each individual pixels in their capability of converting photons to electrons. That's actually the physical nature of camera fingerprints.

A high-level concept is like this illustration. The first image is an image produced by a camera, and we can consider that this image is a combination of a perfect image, or a pure image, and a camera fingerprint. This camera fingerprint captures the variation of pixel responsivity under illumination for all pixels. That's why the dimension of a camera fingerprint is large, the same size as the image.

If we want to get the camera fingerprint, we apply signal processing methods to the original image to get a pure (sort of) image without noise. When we remove this pure image from the original image, the remaining image is mostly of noise signals, which are the camera's fingerprint. That's the high-level concept.

This line of research was started in the electronic engineering community and there are several known applications. For example, source camera identification.

© Springer International Publishing AG 2017
J. Anderson et al. (Eds.): Security Protocols XXIV, LNCS 10368, pp. 96–102, 2017.
DOI: 10.1007/978-3-319-62033-6_12

The question to answer is, if the police have confiscated some cameras, they'd be interested to know: which camera was used to produce a known porn image?

You have an offensive image, and you have a set of suspicious cameras. Then via camera fingerprint you can trace back the source camera which was used to produce the image. The police have particular interest in this application of source camera identification.

Another application that the police are interested is device linkage. You might have physical access to cameras or might not, but you are interested to answer this question: whether two images are produced by the same camera.

The third application is detecting digital image forgery. If you modified some images, by examining camera fingerprints in the images, you can detect this kind of forgery. All these known applications are for forensics purposes. Good algorithms for extracting camera fingerprints are important for these applications.

When I became interested in camera fingerprint, I had a look at the literature, which unfortunately was very chaotic, to say the least. Many fingerprint extraction algorithms had been published by then, but it's funny enough, nearly everybody claims that, "My method is the best." When you look at a single paper, by looking at its experiment results and curves and figures in the paper, the numbers in the paper do suggest that their method looks like the best. But when you put together all these papers, you don't know which one you're going to trust.

Another problem is, all the studies are hardly comparable because their experiment configurations were never consistent. Images used are different. Extraction methods are different. Comparison metrics, different too. It's very surprising and chaotic.

I decided to do a controlled experiment, which is expensive and meticulous, but I wanted to understand whether this line of work is indeed useful or not. We spent quite some time, but I think we got some good results with this experiment. At least we understand a lot of misconceptions in the literature and we also understand how to best extract camera fingerprints.

We wrote up a paper, which on one hand is well-received. It's highly commended by the PC committees and the leading researcher of camera fingerprints. They all say that this is a very good piece of research. On the other hand, this paper was hated by some people because we basically put together all those competing methods under the same testing condition to see which is indeed good. Of course, our results contradict many research teams' claims.

We had our paper published last year at ACM Information Hiding and MMSec, which was started by Ross 20 years ago as Information Hiding Workshop, but has evolved into a premium venue in the fields of information hiding, and multimedia security and forensics. I think the most important output from our experiment is that we've gained confidence that indeed, camera fingerprint is not snake oil.

The reason that I decided to do this very expensive experiment (in terms of human resources) was simple: I had three simple ideas which I wanted to build on the experiment.

Jonathan Anderson: Did the fingerprints survive transformations like RAW-to-JPEG and re-scaling and things like that?

Reply: Yes, the fingerprints survive all those image processing procedures.

Frank Stajano: Frank Stajano. I have a maybe related question. In moderately high-end cameras like prosumer and that thing, there are facilities to eliminate noise on the sensors, as you have hair or dust, by taking picture of a blank screen and then, if we are presenting, the camera then subtracts that. Would that remove also the camera fingerprint?

Reply: No. Those kind of operations will not remove the fingerprints.

Frank Stajano: Why is that?

Reply: Because camera fingerprints are unlike noise caused by hair or dust. As PRNU captures slight variations among individual pixels in their capacity of converting photon to electrons, camera fingerprints come from such a low-level effect. That's, the thing is deep at the bottom of an image, and manifests intrinsically in each pixel.

Frank Stajano: If each picture taken by the camera is then corrected by this presumably static bias that you examine by taking natural picture, how would this not also eliminate these differences?

Reply: The reason is that you need the specific algorithms to extract this fingerprint. You can decrease the quality of the fingerprint by doing some random processing, but the fingerprint will not be destroyed unless you know exactly their extracting algorithms. If you know this algorithm, indeed you can remove the fingerprints.

To give a history background, many years ago, actually before electronic engineers started this topic, this phenomenon of camera fingerprint was known in astronomy. Why? Because they had huge cameras. They wanted to capture images of outer space. In these areas, their image signals are very weak. If you allow weak noise signals like PRNU, then the pictures taken by those huge cameras would be useless for astronomy researchers.

Thirty years back, those people actually had to spend a lot of money to suppress noise signals like PRNU. For hugely expensive cameras, they can claim camera fingerprints do not exist because they're suppressed, but for consumer cameras, PRNU fingerprints are there.

Frank Stajano: You test after the dust removals?

Reply: Yeah.

We tested 50 different models of cameras.

The first idea we want to explore is to build camera fingerprint into security protocols. For example, an obvious application would be authentication. We've had in mind a new authentication scheme, 'any photo you take are you'.

The motivation is simple because now actually a lot of research efforts aim to get people to use hardware tokens for authentication in their daily lives. But deployment is a challenging issue. In our case, it's good for us because camera phones are everywhere. If our authentication scheme works well, it's easy to have a large-scale deployment.

To explain how our idea works, let's look at a simplified basic protocol. There is first an enrolment process, where each camera is enroled to our system. Each camera will take many photos, and each of these photos will be used to extract a fingerprint for the same camera. Each extracted fingerprint will be averaged to produce a high-quality fingerprint. We want keep high-quality fingerprints as reference fingerprints.

Then, when the user comes for authentication, she takes a photo of anything, and send the photo to an authentication server. The server will extract a camera fingerprint from the photo, and the fingerprint will be compared with a reference fingerprint stored.

This comparison can be complex. It's not a yes or no binary decision, but calculating a pixel-wise correlation, typically called a normalized cross correlation, which involves with some matrix manipulations.

Of course, there are many issues with this basic protocol. For example, fingerprint leakage would be a serious problem, because a lot of people use their cameras, use their phones to take photos and post them online, but these images will contain camera fingerprints. Adversaries can easily do a fingerprint replay attack. Initially, I thought a challenge–response mechanism might work against such replay attacks, but it doesn't.

The main lesson is the following. The common techniques for verifying message freshness in security protocols don't guarantee a camera fingerprint's liveness.

One possible solution is not to allow a user to enrol an old camera. If we enforce a policy that a user can enrol his camera only when it's fresh, i.e. a new camera, then we would not worry about images leaked before enrolment. Then afterwards, we'll enforce that fingerprints will not be leaked from the system.

This means, for authentication purposes, the system will keep a camera fingerprint intact in each image and send the image to the authentication server, but otherwise, if a user wants to post a photo online, then this image should be processed first – effectively, the fingerprint should be removed before the image leaves the camera.

The advantage of this solution is that it doesn't require hardware modification. It doesn't require a modified operating system, either. But there could be some serious security usability issues because users might forget to remove fingerprints. Although we can have a software program to alert each user that she should remove the camera fingerprint from each image that is to be publicly shared, she might not actually comply with the policy or she might forget it.

To address this problem, a solution is to introduce system level controls deciding for which images their camera fingerprints should be removed and for which the fingerprints will stay. Basically, this intends not to involve with any users in this decision-making process. The good thing, compared to the previous solution, is that we will have good usability and security, but the disadvantage is that we might have to modify a phone's operating system.

Of course, there is space to design a lot of fancier authentication protocols. For example, we shouldn't allow servers to store camera fingerprints in plain text because they are very sensitive materials, and we want to prevent the servers from leaking those camera fingerprints. We also should prevent users from leaking their camera fingerprints to a phishing site. It's important to protect a users privacy as well as the confidentiality of reference fingerprints. A common line of research on privacy-preserving authentication protocols is highly relevant here.

So far it's all about our first idea.

The second idea is about privacy. The motivation is very simple. A lot of digital images and video clips are available online, and most people do not even realize the existence of camera fingerprints in these images. We can conceive some privacy intrusion studies.

For example, we can use camera fingerprints to reveal people who post photos anonymously and we can link people who have multiple digital personas, for example, multiple accounts on social networking sites. We can link them together by exploring camera fingerprints.

In the security literature, there are a lot of papers talking about writing style analysis (or stylometry) for privacy intrusion, but there is not a single paper looking into privacy invasion based on camera fingerprints. Camera fingerprints can also complement stylometry for cybercrime investigation.

Frank Stajano: Frank Stajano. I guess, unless people go to extreme lengths to scrub their photos before posting them online, in fact, the metadata that's embedded in the JPEGs or whatever already has the serial number of the camera, the model and blah blah blah, so you don't have to go to very fancy things for de-anonymizing 99% of photos.

Reply: Some sites do remove those metadata, but they do not remove camera fingerprints.

Frank Stajano: Which sites remove metadata?

Jonathan Anderson: I think Flickr is an option.

Reply: Flickr, Yes. I'm not sure of Facebook because I do not use Facebook.

Frank Stajano: I thought that Flickr made a big deal of showing you all the metadata on photos. [inaudible] Interesting.

Reply: What I'm interested is, what are other interesting threat models in this study of privacy invasion. The result could be alarming if we have a number of interesting threat models and if we show this privacy intrusion method is very effective.

Bill Roscoe: It seems to me that anyone with whom I authenticate myself using this method will know my fingerprint and therefore potentially be able to steal my identity, so there's not really a strong form of signature, if you like, as with many cryptographic primitives.

Reply: Yes. You're right. Camera fingerprint can be used to trace you, and identify you.

Bill Roscoe: In other words, if anybody comes to verify, if I reveal my signature, he will know my signature. Therefore, potentially, forge photos with my signature.

Reply: That's true. If we use camera fingerprint for authentication, your camera becomes very sensitive, like biometrics. Therefore, privacy-preserving protocols are important for protecting the fingerprint, as I mentioned earlier.

Jiří Kůr: I was wondering if a server could use the fingerprint for remote attestation of an application platform, and the server connects to the application, takes a picture, and okay, then the fingerprint [inaudible]. The server can authenticate the remote device by taking a picture and verifying if the fingerprint is there.

Reply: Yeah, that's actually the first idea I was talking about, using the fingerprint as an authentication mechanism.

Jiří Kůr: I was talking about maybe a different context, that basically the server authenticates the device. Yeah. Sorry, probably you're right. I'm not sure.

Reply: Never mind. We can actually discuss that offline.

Jiří Kůr: Okay.

Reply: The last slide.

My third idea is using the camera fingerprint to fight against revenge porn. What's revenge porn? It's a relatively new socio-technical problem, where an ex-partner posts online a victim's sexually explicit photos or video clips. There are many incidents reported in the news.

Several countries now outlaw revenge porn, but many do not. There's no technical solution to this problem. A possible defense is to combine porn detection, face detection and face recognition, but they all have false positives and false negatives, and will not offer a good solution to the detection of revenge porn.

I think camera fingerprint looks like a good and simple solution, because if you have a concerned camera or you have access to other images produced by the camera, then you can extract their fingerprints from those devices or images. Then you do online search. Sort of like camera fingerprints provide a simple 'side channel' approach to an otherwise complicated problem.

Hugo Jonker: How would you technically identify that this is a malicious posting?

Reply: I think that's not my concern, because I assume the victim or anybody who worry about the sexually explicit photos or videos posted, will come to ask for help, and I will ask her for other images taken by the same camera, the concerned camera, or I will ask for her camera so that I can establish a camera fingerprint from there. Then I use this to ...

Hugo Jonker: I understand that part, but I'm just wondering, if you find a bunch of pictures but you wouldn't know if those were consensual or non-consensual.

Reply: The victim would know.

Hugo Jonker: The victim would know, and you wouldn't. So you're not proposing a technical method to automatically determine whether or not it's ...

Reply: No. You are right.

Hugo Jonker: So this is essentially a search service... ?

Reply: You are right. This can be a search service that help the victims.

Hugo Jonker: So this is essentially a search function ... find porn starring you online?

Reply: Yes.

Frank Stajano: You'd mostly find all the other photos that this camera has taken of deserts and flowers and so on that were ever posted on Flickr, right?

Reply: Yeah.

Audience: Would the platforms that are used for this revenge porn be interested in collaboration? Why would they care? As a way to fight this problem as a whole, why would the platforms care ...

Reply: Victims care.

Audience: But if they make money on high volumes of people viewing their content, then they don't really care. They are used to ...

Reply: Victims care, and this is also actually a crime in some countries. The law enforcement there would care, too.

Audience: Yes, but the offender, if they want to use this kind of revenge on somebody, they may use the platform that's not interested in collaborating, in a country where jurisdiction doesn't really prosecute this kind of activity.

Reply: You're right, but that still can be tackled with our technical approach, right? We can identify revenge porn for the victims and then take the proper measures, for example, taking down those images.

Hugo Jonker: I think this will provide evidence in civil court cases.

Reply: Yeah. Exactly.

Hugo Jonker: You can show it's your camera, it was coming from your camera, and now the photos are there.

Reply: Yeah. For countries where revenge porn is outlawed, of course this service will produce evidence, so you can sue the bad guy.