

On the Usage of Sensor Pattern Noise for Picture-to-Identity Linking through Social Network Accounts

Riccardo Satta¹ and Pasquale Stirparo^{1,2}

¹*Institute for the Protection and Security of the Citizen
Joint Research Centre (JRC), European Commission, Ispra (VA), Italy*
²*Royal Institute of Technology (KTH), Stockholm, Sweden*

Keywords: Social Network, Account, Sensor Pattern Noise, Identity, Linking, Digital Image Forensics, Multimedia Forensics.

Abstract: Digital imaging devices have gained an important role in everyone's life, due to a continuously decreasing price, and of the growing interest on photo sharing through social networks. As a result of the above facts, everyone continuously leaves visual "traces" of his/her presence and life on the Internet, that can constitute precious data for forensic investigators. *Digital Image Forensics* is the task of analysing such digital images for collecting evidences. In this field, the recent introduction of techniques able to extract a unique "fingerprint" of the source camera of a picture, e.g. based on the Sensor Pattern Noise (SPN), has set the way for a series of useful tools for the forensic investigator. In this paper, we propose a novel usage of SPN, to find social network accounts belonging to a certain person of interest, who has shot a given photo. This task, that we name *Picture-to-Identity linking*, can be useful in a variety of forensic cases, e.g., finding stolen camera devices, cyber-bullying, or on-line child abuse. We experimentally test a method for Picture-to-Identity linking on a benchmark data set of publicly accessible social network accounts collected from the Internet. We report promising result, which show that such technique has a practical value for forensic practitioners.

1 INTRODUCTION

Nowadays, digital imaging devices have gained a prominent role in everyone's life. Mobile smart phones, tablets, digital cameras and camcorders have become progressively cheaper and affordable for every one; this goes hand in hand with the growing interest for sharing moments of our life using social networks (e.g., Facebook, Flickr) and Internet in general. As a result, everyone of us is continuously leaving visual "traces" of his/her presence and life on the Internet. Under the proper legal framework, law enforcers and forensic investigator can access this data (e.g., in undercover operations) in case it is relevant for investigations.

The task of analysing digital images for forensic purposes is usually referred to as *digital image forensics*. In this field, the recent introduction of techniques able to extract a unique "fingerprint" of the source camera of a picture (Dirik et al., 2008; Li, 2010; Lukas et al., 2006) has set the way for a series of useful tools for the forensic investigator. In particular, the Sensor Pattern Noise (SPN) left in the

image by the device sensor has been exploited in various forensic tasks like source device identification (Lukas et al., 2006), forgery detection (Li and Li, 2012), source device linking (Fridrich, 2009), or clustering of images with respect of the source camera (Li and Li, 2012).

In this paper, we present a novel usage of the SPN for digital image forensic purposes. We propose to exploit SPN fingerprints to find social network accounts belonging to a certain person of interest, who has shot a given, known photo. We name this task *Picture-to-Identity linking*. It can be useful in a variety of forensic cases, e.g., on-line child abuse, defamation, finding stolen camera devices. To the authors' best knowledge, this application of the SPN has never been proposed in the literature so far. We developed an implementation of Picture-to-Identity linking based on the SPN extraction method proposed by (Lukas et al., 2006), and tested it on a benchmark data set of social network accounts collected from the Internet. The reported results are promising and show evidence of a practical usefulness of such technique for forensic investigators.

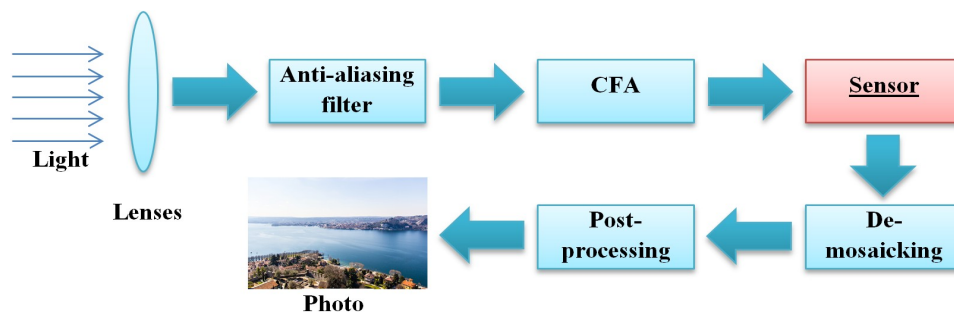


Figure 1: Image acquisition pipeline in typical camera devices.

The remainder of the paper is structured as follows. First, in Sect. 2 we review previous works on source camera identification, with particular regard to SPN. We then focus in greater detail on the usage of SPN for Picture-to-Identity linking in Sect. 3, providing an overview of the possible concrete applications and of the challenges that must be expected. Sect.4 describes and formalises a method for Picture-to-Identity linking, that is then experimentally tested on a benchmark data set in Sect. 5. Finally, Sect. 6 draws up conclusions and suggests possible directions for further research on this topic.

2 RELATED WORK

Digital images can be associated to various kinds of useful *metadata*. Examples are Exif data, image tags, or text associated to the image (e.g., contained in the same web page), etc. Exif metadata in particular has received much attention by forensic investigators, since it stores useful information about the device (e.g., camera model, serial number, etc.) that produced the content. However, from a forensic point of view, this information has to be taken into account with extreme care, as it is fairly easy to modify or remove it with image processing software (e.g., Photoshop) or with free tools available on the Internet (e.g., ExifTool). A robust cue that can be used in place of Exif data to identify the source camera of a picture is the noise pattern left by the sensor element of the camera (usually referred to as Sensor Pattern Noise or SPN) (Lukas et al., 2006). In fact, such noise pattern is univocal of a camera sensor, and can be seen as an unique fingerprint that identifies an individual device.

To proper understand how the SPN can be used as a fingerprint, it is worth to take a closer look at how a digital picture is typically produced by a camera (Li, 2010) (see a clarifying scheme in Fig. 1). The light coming from the scene arrives first to the camera lens. Then, it passes through an anti-aliasing filter, and reaches the Colour Filter Array (CFA), which

is placed just over the sensor and is used to capture colour information. The light finally reaches the sensor, a matrix of elementary sensitive elements each corresponding to a pixel, that converts light into a digital representation. The subsequent steps, de-mosaicking and post-processing, are respectively in charge of interpolating the missing two colours of each pixel, and of carrying out image processing operations (e.g., white balancing, de-noising, etc.) to increase the perceived image quality. Each step of this pipeline may leave artefacts on the image that can be used as a signature of the camera device.

Much research has been conducted in this direction, exploiting SPN (Kang et al., 2012; Li, 2010; Li and Li, 2012; Lukas et al., 2006), interpolation artefacts caused by de-mosaicking filter (Cao and Kot, 2009; Long and Huang, 2006; Popescu and Farid, 2005) and JPEG compression (Sorrell, 2009), traces of dust in the sensor (Dirik et al., 2008), and lens aberrations (Choi et al., 2006; Van et al., 2007), as possible fingerprints. Out of them, de-mosaicking and JPEG compression artefacts depend on the algorithms chosen by the manufacturer, which are usually specific of the model; therefore, they can be used only as a signature of the camera model (not of individual cameras). Dust traces affect mainly professional reflex cameras with interchangeable lens (dust may enter inside the camera when the photographer changes the lens) and are a fingerprint of the single device, that however exhibits a low stability over time (i.e. new dust particles may be deposited into the sensor). Regarding lens aberrations, their use as device fingerprint has been tested in a limited extent (Choi et al., 2006; Van et al., 2007) and its actual potential is still to be explored.

Differently from the above techniques, the Sensor Pattern Noise has the desired characteristics of uniqueness and stability, and has been studied and tested in various forensic tasks, e.g.: source device identification (Kang et al., 2012; Lukas et al., 2006; Li, 2010; Li and Li, 2012; Li and Satta, 2012), forgery detection (Li and Li, 2012), source device linking

(Fridrich, 2009), clustering of images with respect of the source camera (Li and Li, 2012).

Most techniques extract the SPN by exploiting the additive noise model first presented by (Lukas et al., 2006), which models SPN as an additive, high frequency noise. In the wavelet domain, this can be formulated as:

$$n_P = DWT(P) - F(DWT(P)) \quad (1)$$

where P is a picture, n_P is the SPN of P , $DWT(\cdot)$ is the Discrete Wavelet Transform, and F is a de-noising filter, which extracts the low-frequency (non-noise) components of P . The de-noising filter F used plays indeed a crucial role. The wavelet-based filter described in Appendix A of (Lukas et al., 2006) is an effective one and has been used in many other works (Li, 2010; Li and Li, 2012; Li and Satta, 2012).

In order to compare two SPNs n_{P1} and n_{P2} from two images $P1$ and $P2$, a common approach (Li, 2010; Li and Li, 2012; Li and Satta, 2012) is to evaluate the Normalised Cross-Correlation, which is defined as

$$\rho(n_{P1}, n_{P2}) = \frac{(n_{P1} - \bar{n}_{P1}) \cdot (n_{P2} - \bar{n}_{P2})}{\|n_{P1} - \bar{n}_{P1}\| \cdot \|n_{P2} - \bar{n}_{P2}\|} \quad (2)$$

where \bar{n}_{P1} and \bar{n}_{P2} are the means of n_{P1} and n_{P2} , respectively. The value of $\rho(n_{P1}, n_{P2})$ can be taken as the *matching score* between n_{P1} and n_{P2} .

3 PICTURE-TO-IDENTITY LINKING USING SPN

Most social networks (e.g., Facebook, Google Plus) offer users the possibility of uploading pictures to enrich their profile. Various social networks, like Flickr or Ipernity, are entirely devoted at sharing photos. Indeed, from a forensic viewpoint, these personal accounts may offer much useful information for investigations.

Finding social networks accounts that belong to a certain person of interest (e.g., who is relevant for a case) can be therefore very valuable. The task is a non-trivial one, even if one knows identity details. In fact, one should expect to find several homonyms, and the person of interest may use a pseudonym or a nick name. Nevertheless, from the perspective of digital forensics investigations, being able to go back to the person who shot a picture is important in several scenarios. Here we mention three of them, varying from lower to higher level of seriousness.

The first one is the case of a stolen smart-phone. After the theft, the thief will possibly start using the smart-phone, taking pictures and sharing them on-line

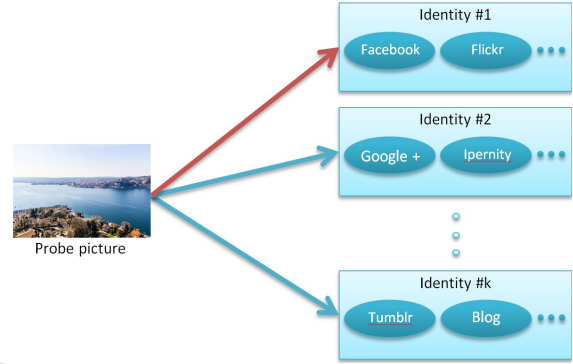


Figure 2: Picture-to-Identity Linking.

on social platforms. Having already pictures taken by the victim as sample, using the methodology we propose in this paper it would be possible to correlate the pictures in order to identify the thief. Another scenario is the case of defamation torts. Once they mainly used to happen verbally (e.g. during public speeches) or in a written form (e.g. releasing interviews to the newspaper), nowadays in the Internet era they are often perpetrated posting on websites/social networks pictures of which the targeted person may be ashamed of, maybe because taken without he/she being aware of it. A clear example of this is cyber-bullying. Finally, a serious crime which heavily involves digital forensics, is the on-line abuse of children, where the perpetrators often record moments of the ongoing crime by different means, taking video or photo to later share them over the Internet (On-line Child Abuse). In such cases, being able to identify the perpetrator as soon as possible is important due to the serial habit of this type of criminal in committing such crimes.

In this work, we propose to use the SPN signature as a mean to find social network accounts belonging to the person that has shot a given photo. In other words, we aim at answering the question: *Given a picture P^* ; how can I find an account belonging to the person who shot P^* ?* We name this task *Picture-to-Identity Linking* (see Fig. 2), and propose to address it by using the SPN. To the best of our knowledge, the task above has never been proposed in the literature so far. The rationale behind our proposal is that social network accounts might contain pictures taken with the digital imaging device(s) of the account owner. Given a *probe* picture P^* , one can extract its SPN and compare it against the SPNs of the pictures from a social network account. If a match is found, this means that the account contains a picture taken by the same camera of P^* ; in turn, the owner of the account is likely to be the person who shot P^* .

Picture-to-Identity linking via SPN poses several,

specific challenges. Among them we highlight:

- **Differences in Image Size.** The size of the probe image P^* and of the social network images can differ, e.g. because images uploaded to the social network are automatically resized (e.g., this is the case of Facebook). In turn, SPN sizes will differ.
- **SPN Misalignment.** The probe image P^* and/or the social network images might have been cropped by the author; consequently, the SPNs might be misaligned.
- **JPEG Compression.** Most images taken with digital imaging devices are compressed using JPEG. Some social networks might further compress them automatically after uploading. JPEG compression degrades the SPN as it attenuates or destroys high frequency image components.
- **Image Alteration.** Images can be altered for artistic purposes (e.g., contrast enhancement, HDR processing, or other visual effects). Any such image alteration in turn modifies the SPN.
- **Different ISO Settings.** Images shot in low light conditions will have a higher ISO, and thus show a stronger noise. On the contrary, images taken in good lighting conditions will have a low ISP and less noise.

Despite the challenges above, we will demonstrate that the Picture-to-Identity linking via SPN is feasible, and could be used as a tool for forensic investigations. A practical implementation of our proposal is presented in the next Section.

4 METHOD OVERVIEW AND IMPLEMENTATION

In this Section, we describe our method to perform Picture-to-Identity linking via SPN. We formalise the problem as follows. Let P^* be a probe (target) image, and let us define $\mathbf{I} = \{A_i\}$, $i = 1, m$ the set of m social network accounts A_i that belong to the person/identity \mathbf{I} . Let then $\mathbb{I} = \{\mathbf{I}_j\}$, $j = 1, k$ be a set of k candidate identities, each associated with her own set of accounts \mathbf{I}_j .

The problem of Picture-to-Identity linking can be formulated as finding the identity \mathbf{I}^* which owns the account containing the image with highest matching score to P^* (Fig. 2). Formally:

$$\mathbf{I}^* = \arg \max_{\mathbf{I}_j} S(P^*, \mathbf{I}_j), \mathbf{I}_j \in \mathbb{I} \quad (3)$$

where $S(P^*, \mathbf{I})$ is the *identity score*, i.e., the maximum SPN matching score between the SPN extracted

from P^* and each image of each account own by identity \mathbf{I} :

$$S(P^*, \mathbf{I}) = \max_{A \in \mathbf{I}} \max_{P \in A} \rho(n_{P^*}, n_P) \quad (4)$$

In practice, due to the problems stated in Sect. 3, the probability of wrong identification may be high (see next Sect. 5). Thus, in a practical implementation of the above technique, for a given probe P^* it is better to show an *ordered list* of the candidate identities $\{\mathbf{I}_1^*, \mathbf{I}_2^*, \dots, \mathbf{I}_k^*\}$, ranked with respect to their score $S(P^*, \mathbf{I}_j^*)$.

For the extraction of the SPN, we used the additive model of Eq. (1) with the Wavelet-based filter proposed by Lukas et al. (Lukas et al., 2006). The reliability of the SPN is usually better in the central portion of the image (Li and Satta, 2012; Li and Satta, 2011). Thus, before extracting the SPN, the image is cropped by taking a 256×256 pixels window in the image centre.

5 EXPERIMENTAL RESULTS

We evaluated the performance of the method above in the task of Picture-to-Identity Linking, on a data set of 1909 images, taken from social network accounts and/or personal blogs belonging to 10 different identities. For each identity, we found two social network accounts. All the accounts were publicly accessible (i.e., not restricted to friends only). The number of images per account, and the account type, are listed in Table 1. Images vary in size and are mostly small. In fact, in many social networks (e.g., Facebook) pictures are automatically scaled to a small size to fit better to the web page layout. Other social networks (e.g., Flickr) offer a low resolution preview of the image, and bigger versions of the same picture can be accessed by clicking on the preview. In these cases,

Table 1: Composition of the data set used for benchmark. For each social network account, the number of images is shown in brackets.

Identity	Accounts (Nr. of images)
#1	Flickr (148), Personal blog (93)
#2	Facebook (105), Flickr (118)
#3	Flickr (93), Google + (70)
#4	Facebook (3), Flickr (199)
#5	Flickr (143), Tumblr (26)
#6	Flickr (98), Personal blog (21)
#7	Flickr (112), Google + (99)
#8	Facebook (84), Flickr (109)
#9	Facebook (12), Flickr (192)
#10	Facebook (11), Flickr (173)

Table 2: Picture-to-Identity confusion matrix.

		Estimated identity									
		#1	#2	#3	#4	#5	#6	#7	#8	#9	#10
Real identity	#1	138	20	9	10	6	7	4	7	1	39
	#2	25	107	9	19	13	7	15	7	9	12
	#3	4	7	115	5	4	6	3	14	2	3
	#4	12	20	12	97	15	8	6	15	8	9
	#5	14	13	7	17	61	16	9	9	8	15
	#6	10	14	15	8	13	36	4	11	3	5
	#7	0	10	1	2	2	0	189	0	4	3
	#8	8	13	21	18	9	17	6	85	10	6
	#9	1	7	5	4	4	2	2	6	161	12
	#10	38	18	7	9	13	5	2	4	13	75

we took the preview image only. Before extracting the SPN, and in order to make them comparable, each image is (i) rotated so that it has a landscape orientation (if necessary) and (ii) resized so that the long side has a length of 4096 pixel, keeping the aspect-ratio.

Experiments have been carried out as follows. Each image of the data set was taken out from its source account and used as the query image P^* ; then, the identities were ranked with respect to the identity score, computed by means of Eq. (4).

In Table 2 we show the identification performance in terms of confusion matrix. The (i, j) -th element in the matrix is the number of images coming from an account belonging to the identity i that has been classified as belonging to the identity j . The ranking performance, instead, has been assessed by means of Cumulative Matching Characteristic (CMC), and of Synthetic Recognition Rate (SRR) curves. The former is the cumulative probability of finding the correct identity within the first n ranks; the latter is the probability of a correct recognition given m target identities. The SRR curve can be obtained from the CMC (hence the term *synthetic*) as $SRR(m) = CMC(k/m)$ where k is the number of identities (Gray et al., 2007). The resulting plots are shown in Fig. 3 and Fig. 4, respectively. The CMC and SRR curves corresponding to a random guess are shown as well for comparison.

The probability of correct recognition at the first rank is above 56%, which is far higher than the random guess, but obviously not enough for a precise identification. One could argue that these scoring results do not qualify data as evidence for its admissibility in court. It is correct to state that an evidence as such shall be 100% attributable, as level of accuracy, to a subject (e.g. the victim, the suspect, etc.). However, a ranked list of candidate identities can be valuable for the forensic practitioner in the investigation phase (Casey, 2009). In fact, there are circumstances where evidences lead to a pool of several people under suspicion. For example, if law enforcers are not able to identify and seize cameras belonging to each of the suspects, images taken from their social network

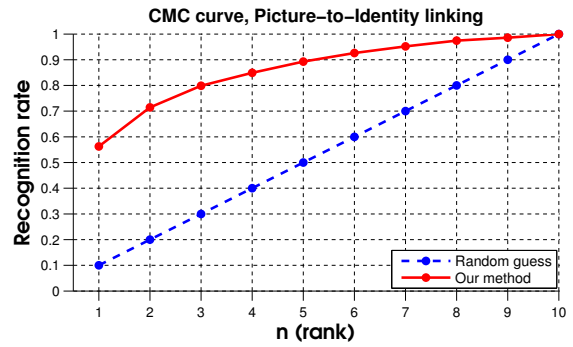


Figure 3: Picture-to-Identity linking performance, in terms of Cumulative Matching Characteristics curve, in a benchmark data set of 1909 images (see the text for details).

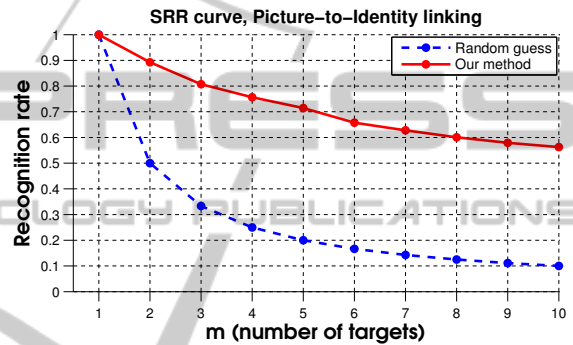


Figure 4: Picture-to-Identity linking performance, in terms of Synthetic Recognition Rate curve, in a benchmark data set of 1909 images (see the text for details).

accounts could be used as a set of candidates to test against the probe picture, using the ranking technique proposed in this paper. Another scenario is when the investigator knows, from other trails, that the person of interest has certain characteristics (e.g., age, nationality, or ethnic group); accordingly, he/she can exclude the non-relevant identities in the ranked list returned by the proposed technique. E.g., in all such cases the proposed method would help to skim off the number of candidate identities related to the case, which, in conjunction with other facts, may speed up the investigation and increase its accuracy.

For the sake of completeness, we also evaluated the CMC and SRR curves in the task of associating each image with the source account. In analogy with the terminology utilised so far in the paper, we refer to this task as *Picture-to-Account linking*. The corresponding plots are shown in Figs. 5-6. Note that, in this case, the number of targets is 20 (equal to the number of accounts), instead of 10 as in the previous experiment.

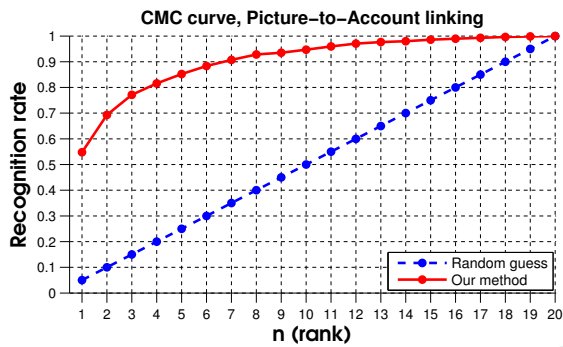


Figure 5: Picture-to-Identity linking performance, in terms of Cumulative Matching Characteristics curve, in a benchmark data set of 1909 images (see the text for details).

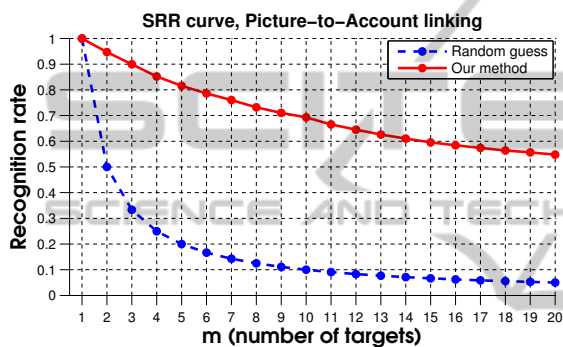


Figure 6: Picture-to-Identity linking performance, in terms of Synthetic Recognition Rate curve, in a benchmark data set of 1909 images (see the text for details).

6 CONCLUSIONS

In this paper, we presented a novel usage of SPN for digital image forensics, as a mean to find social network accounts belonging to a certain person, who has shot a given picture. This task, that we named Picture-to-Identity linking, has many potential applications in digital image forensics, e.g., identification of perpetrators of On-line Child Abuse, cyber-bullying, or thefts of smart-phones. Yet, it is a challenging task, due to differences in image size, misalignment of the extracted SPNs, compression artefacts, and other issues.

We proposed an implementation of Picture-to-Identity linking using the Wavelet-based filter of (Lukas et al., 2006) to extract the SPN, and Normalised Cross Correlation to compare noise signatures. We evaluated the performance on a benchmark data set built using publicly accessible social network accounts. The reported results are promising; furthermore, we see room for various improvements. A possible one, in order to tackle the problem of SPN misalignment that may happen e.g. if some images have

been cropped, is to use a sliding window approach to compare the SPN extracted from the probe image to the ones extracted from account images. Also, the performance could be improved by combining the SPN information with other metadata (e.g., Exif data, if available), or with other camera identification techniques (e.g., de-mosaicking artefacts, dust traces).

It is worth pointing out the limitations of the proposed evaluation, which is restricted to a reduced set of ten identities (and twenty social network accounts). Indeed, Picture-to-Identity linking makes much more sense with a large-scale data base of social network accounts. The experimental evaluation provided can be seen as a simulation of the case when a set of “candidate” social network accounts has already been found by other means. One of the directions of further research will therefore be to prepare a large-scale test bed for Picture-to-Identity linking, possibly improving the proposed technique as explained above.

REFERENCES

- Cao, H. and Kot, A. C. (2009). Accurate detection of demosaicing regularity for digital image forensics. *IEEE Transactions on Information Forensics and Security*, 4(4):899–910.
- Casey, E. (2009). *Handbook of Digital Forensics and Investigation*. Academic Press.
- Choi, K. S., Lam, E. Y., and Wong, K. K. Y. (2006). Automatic source camera identification using the intrinsic lens radial distortion. *Optics Express*, 14(24):11551–11565.
- Dirik, A. E., Sencar, H. T., and Memon, N. (2008). Digital single lens reflex camera identification from traces of sensor dust. *IEEE Transactions on Information Forensics and Security*, 3(3):539–552.
- Fridrich, J. (2009). Digital image forensic using sensor noise. *IEEE Signal Processing Magazine*, 26(2):26–37.
- Gray, D., Brennan, S., and Tao, H. (2007). Evaluating appearance models for recognition, reacquisition, and tracking. In *10th IEEE International Workshop on Performance Evaluation of Tracking and Surveillance (PETS)*.
- Kang, X., Li, Y., Qu, Z., and Huang, J. (2012). Enhancing source camera identification performance with a camera reference phase sensor pattern noise. *IEEE Transactions on Information Forensics and Security*, 7(2):393–402.
- Li, C.-T. (2010). Source camera identification using enhanced sensor pattern noise. *IEEE Transactions on Information Forensics and Security*, 5(2):280–287.
- Li, C.-T. and Li, Y. (2012). Color-decoupled photo response non-uniformity for digital image forensics. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(2):260–271.

- Li, C.-T. and Satta, R. (2011). On the location-dependent quality of the sensor pattern noise and its implication in multimedia forensics. In *Proceedings of the 4th International Conference on Imaging for Crime Detection and Prevention 2011 (ICDP 2011)*, page 37.
- Li, C.-T. and Satta, R. (2012). Empirical investigation into the correlation between vignetting effect and the quality of sensor pattern noise. *IET Computer Vision*, 6:560–566(6).
- Long, Y. and Huang, Y. (2006). Image based source camera identification using demosaicking. In *2006 IEEE 8th Workshop on Multimedia Signal Processing*, pages 419–424.
- Lukas, J., Fridrich, J., and Goljan, M. (2006). Digital camera identification from sensor pattern noise. *IEEE Transactions on Information Forensics and Security*, 1(2):205–214.
- Popescu, A. and Farid, H. (2005). Exposing digital forgeries in color filter array interpolated images. *IEEE Transactions on Signal Processing*, 53(10):3948–3959.
- Sorrell, M. J. (2009). Digital camera source identification through jpeg quantisation. In Li, C.-T., editor, *Multimedia forensics and security*. Information Science Reference.
- Van, L. T., Emmanuel, S., and Kankanhalli, M. (2007). Identifying source cell phone using chromatic aberration. In *2007 IEEE International Conference on Multimedia and Expo*, pages 883–886.