https://doi.org/10.2352/ISSN.2470-1173.2020.4.MWSF-215

Score-based likelihood ratios for camera device identification

***Stephanie Reinders*[+], *Li Lin*[+], *Wenhao Chen*[†], *Yong Guan*[†], *Jennifer Newman*[+]**
[+]*Department of Mathematics,* [†]*Department of ECPE, Iowa State University, Ames, Iowa, USA,* {*srein, llin, wenhaoc, guan, jlnewman*}*@iastate.edu*

Abstract

Many areas of forensics are moving away from the notion of classifying evidence simply as a match or non-match. Instead, some use score-based likelihood ratios (SLR) to quantify the similarity between two pieces of evidence, such as a fingerprint obtained from a crime scene and a fingerprint obtained from a suspect. We apply trace-anchored score-based likelihood ratios to the camera device identification problem. We use photo-response non-uniformity (PRNU) as a camera fingerprint and one minus the normalized correlation as a similarity score. We calculate trace-anchored SLRs for 10,000 images from seven camera devices from the BOSSbase image dataset. We include a comparison between our results the universal detector method.

Introduction

Digital image forensics consists of two main types problems: source identification problems and forgery and tampering detection problems. Source identification problems can be further divided into two key problems: *camera model identification* and *camera device identification*. The goal of camera model identification is to identify the particular camera make and model used to capture a digital image. The camera device identification problem goes a step further and aims to identify not only the make and model but the specific camera device used to capture an image. This paper applies score-based likelihood ratios (SLR) to the camera device identification problem. We consider the scenario where, given a digital image from an unknown camera device and a specific known camera device, we want to determine whether the image originated from that specific device. Not only do we want to decide whether the image originated from that device, we also want to determine the strength or weakness of the evidence in favor of that decision. SLRs afford us a means of quantifying the weight of the evidence.

The vast majority of work in camera device identification employs a *universal detector* method [1, 2, 3, 4, 5, 6, 7, 8]. Photo-response non-uniformity (PRNU) [9] serves as a camera fingerprint for the specific camera device under consideration. The noise residual is estimated from the image in question and compared to the camera fingerprint using a similarity score, such as peak-to-correlation energy (PCE). A set of *matching* similarity scores – scores between noise residuals of images and camera fingerprints that originate from the same camera device – is generated from a database of training images. Likewise, a set of *non-matching* scores – scores between noise residuals of images and camera fingerprints that do not originate from the same camera device – is generated. The sets of matching and non-matching scores are used to select a decision threshold t. The score between the noise residual of the image in question and the camera fingerprint of the specific device is compared to the threshold t and a binary decision of match or non-match is made. The motivation behind using a universal detector is that once a reference dataset of matching and non-matching scores is created, the decision threshold does not need to be readjusted if the input image or camera fingerprint changes.

Score-based likelihood ratios (SLR) quantify the weight of evidence, providing more information than a binary decision of match or non-match. SLRs have been applied to various fields of forensics, such as shoeprint [10], glass [10], and handwriting [11] evidence. Hepler et. al. [11] showed that the end result of the SLR for handwriting evidence depends heavily on the method used to define non-matches and in some cases different definitions lead to different conclusions on the same evidence. They present three methods for defining non-matches – source-anchored, trace-anchored, and general match. We adapt the three definitions for non-matches presented by Hepler et. al. to fit the device identification problem. The source-anchored method considers non-matching scores between a camera fingerprint from the specific camera device and noise residuals of images from other camera devices. The trace-anchored method considers non-matching scores between the noise residual of the image in question and camera fingerprints from devices other than the specific device. Lastly, the general match method considers non-matching scores between the noise residual of an image from a device other than the specific device and a camera fingerprint from a second device other than the specific device.

SLRs for camera device identification use a similarity score and PRNU as the camera fingerprint. They create sets of matching and non-matching scores like the universal detector method. Instead of choosing a decision threshold, SLRs fit probability density functions to both sets of scores. Both pdfs are evaluated at the score between the noise residual of the image in question and the camera fingerprint, and the SLR is the ratio of the results. SLRs were first applied to the camera device identification problem by Nordgaard and Hoglund in 2011 [12]. Van Houten, Alberink, and Geradts wrote a follow-up paper later that year [13]. Both papers only consider the source-anchored definition of non-matches. To the best of our knowledge, SLRs have not appeared in the camera device identification literature since. The main contribution of this paper is to introduce the trace-anchored definition of SLRs to camera device identification.

Likelihood ratios provide another method for quantifying the weight of evidence. The key difference between score-based likelihood ratios (SLR) and likelihood ratios (LR) is that SLRs model similarity scores applied to features while LRs directly model features. Initially, LRs were applied to the camera model identification problem for RAW images [14] and JPEG images [15]. Later they were adapted to be suitable for the camera device identification problem for RAW images [16] and JPEG images [17]. The

universal detector method, which relies on similarity scores and PRNU, has been tested against cropping [3], scaling [3], gamma correction [18], denoising [18] and compression [7, 18]. The universal detector method has also been extended to account for lens distortion [19]. Because of its relative newness in the literature, the LR model has not yet been tested, to our knowledge, against such transformations. For this reason, we chose to expand the use of SLRs in device identification, instead of focusing on LRs.

We present a framework for trace-anchored score-based likelihood ratios for the camera device identification problem in the Score-Based Likelihood Ratios section. We summarize relevant work in the Related Work section. Our methodology and experiments are detailed in the Experiments section.

Score-Based Likelihood Ratios

In this section, we first discuss photo-response non-uniformity and similarity scores. Then we introduce our framework for trace-anchored score-based likelihood ratios.

Photo Response Non-Uniformity for Device Identification

Lukas et. al. [9] introduce the method of using photo-response non-uniformity (PRNU) as a *camera fingerprint* to address the device identification problem. Each camera imprints its unique camera fingerprint in every image that it takes, but the fingerprint is very faint and exists at the noise level of the image. Lukas, Fridrich and Goljan introduce methods of extracting the PRNU from a camera and the noise residual from individual images [9].

We use the PRNU extraction code [20] in our experiments. Here we summarize the PRNU framework used in [5]. Let I be the image output from a camera that includes noise, and let I_0 be the perfect image absent any noise. Denote the true PRNU, or *camera fingerprint*, as K, and denote all other noise components from the image as Θ. Then the image output I can be modeled as

$$I = I_0 + I_0 K + \Theta \tag{1}$$

where multiplication is performed element-wise. In practice, we are unable to acquire the true camera fingerprint K, so instead we estimate it with a maximum likelihood estimator $\hat{K}$. We obtain $\hat{K}$ by first acquiring N images $I^{(1)}, I^{(2)}, ...I^{(N)}$ from the camera in question. Using a denoising filter F we estimate the noise residuals from each image $W^{(i)} = I^{(i)} - F\left(I^{(i)}\right)$ for $i = 1, 2, ...N$. Then the PRNU estimate is

$$\hat{K} = \frac{\sum_{i=1}^{N} W^{(i)} I^{(i)}}{\sum_{i=1}^{N} \left(I^{(i)}\right)^2} \tag{2}$$

Similarity Score

Normalized correlation was used as the similarity score in early camera device identification works [1, 9, 18, 21]. The peak-to-correlation energy (PCE) score later replaced normalized correlation as the similarity score [3, 4, 5]. Goljan showed that if a periodic signal such as linear pattern is present, the decision threshold needs to be adjusted if normalized correlation is the score, but does not need to be adjusted if PCE is the score [4]. We used PCE as the similarity score for our first score-based likelihood ratio experiments, but the large variance of PCE scores (on the order of 10^6) on our image data resulted in a large number of inconclusive SLR values. (Inconclusive results occur when the SLR defined in equation 5 has zero in the denominator and a very small number in the numerator.) Normalized correlation has much smaller variance on our image data (on the order of 10^{-3}) than PCE so it produces many fewer inconclusive SLR results. The primary benefit of the PCE is that the decision threshold is more stable than with normalized correlation. Because we compute a decision threshold every time we compute a new SLR, we can safely use normalized correlation as our similarity score. More precisely, we use one minus the normalized correlation as our similarity score because it can be implemented as a built-in distance function in Matlab.

SLR Framework

In this section we will introduce a framework for a score-based likelihood ratio for the camera device identification problem. This framework applies to the situation where forensic analysts have two pieces of evidence: (1) a digital image e_u from an unknown camera device that was involved in a crime; and (2) a camera fingerprint e_s from a suspect's camera device c_s. We make the assumption that the analysts have access to the device c_s and are able to take images on it for analysis. We refer to the device c_s as the *specific known device*. The analysts want to determine whether the questioned image e_u and the camera fingerprint e_s originate from the same camera device, namely the suspect's device c_s. In order to compute an SLR the analysts would also identify a set c_a of relevant *alternative camera devices*. These alternative devices might be cameras obtained from other suspects for example.

The prosecution's hypothesis is that questioned image e_u and a camera fingerprint e_s originated from the same camera device, while the defense's hypothesis is that image e_u and camera fingerprint e_s did not originate from the same camera device:

$$H_p : e_u \text{ and } e_s \text{ originated from the same camera device} \tag{3}$$

$$H_d : e_u \text{ and } e_s \text{ didn't originate from the same camera device} \tag{4}$$

Because it is hard to compare e_u and e_s directly, we will instead compare the noise residual x_u of image e_u with the camera fingerprint e_s. We will use the similarity score Δ to compare x_u and e_s. We choose Δ to be one minus the sample correlation (implemented using the Matlab function pdist2 with the correlation distance metric).

In order to determine the range of scores that could be expected if the prosecution's hypothesis is true we need to generate a set of *matching* scores D_s for the specific device c_s. We generate this set of scores by randomly selecting 50 images from c_s and computing a PRNU reference image e_t. Note that the camera fingerprint e_t is not necessarily the same as the camera fingerprint e_s although they are both from the same camera. We randomly select another 50 images from c_s not used to create e_t. For each of these 50 images we compute the score between e_t and the noise residual of the image. We repeat this process five times, each time randomly selecting 50 images for the PNRU reference image and a disjoint set of 50 images for comparison.

To determine the range of scores that could be expected if questioned image e_u originated from a camera device other than

c_s, we generate a set of *trace-anchored non-matching* scores D_u conditioning on the questioned image e_u. For each device in the set of alternative camera devices c_a we randomly select 50 images and create a PRNU reference image and calculate the score between the noise residual x_u of image e_u and the PRNU reference image. We repeat this 5 times for each device in c_a.

We fit probability density functions f_m to the set of *matching* scores D_m and f_n to the set of *non-matching* scores D_n with the Matlab function `fitdist` using a kernel distribution object and the normal kernel function. Similar methods are used in [22, 23]. The *trace-anchored score-based likelihood ratio* is defined

$$SLR_{trace} = \frac{f_m(\Delta(x_u, e_s) = \delta \mid e_s, H_p)}{f_n(\Delta(x_u, e_s) = \delta \mid x_u, H_d)}. \quad (5)$$

The numerator of the SLR is the pdf f_s evaluated at the score $\delta = \Delta(x_u, e_s)$, while the denominator is the pdf f_u evaluated at δ. If the SLR is greater than 1 then the value of the numerator is larger than the value of the denominator. This means that the score δ is more like the *matching* scores than the *non-matching* scores. The larger the SLR is than 1, the stronger the evidence is in favor of the prosecution's hypothesis over the defense's hypothesis. Conversely, an SLR less than 1 supports the defense's hypothesis over the prosecution's, and the closer the SLR is to zero the stronger that support is.

Related Work

In this section we summarize the previous work that has been done with source-anchored SLRs and universal detectors in camera device identification.

Source-anchored SLR for camera device identification

Score-based likelihood ratios were first applied to the device identification problem by Nordgaard and Hoglund [12], although the authors use the term "approximate likelihood ratio." They use the source-anchored method to define non-matches and perform simulations on a set of two camera devices. Van Houten, Alberink and Geradts expanded Nordgaard and Hoglund's experiments to include ten devices [13], also using the source-anchored SLR. We build upon this previous work by considering the trace-anchored SLR, which has not been presented in device identification before.

Universal Detector Methods

Ideally, an analyst would simply be able to compare the noise residual of a questioned image with the PRNU of the suspect's camera device and determine whether they *match*. Unfortunately, no such direct method exists. Instead, the analyst looks at many examples of *matches*, comparing the noise residuals of images with the PRNU of their source cameras, and *non-matches*, comparing the noise residuals of images with the PRNU of other cameras. More specifically, a similarity score is used to measure the similarity between a noise residual and a PRNU estimate. The analyst considers a set of *matching* scores and *non-matching* scores and determines whether the particular score δ between the questioned image and the PRNU of the suspect's camera looks more like a *matching* score or a *non-matching* score. The traditional approach chooses the threshold t that produces a specified false acceptance rate, and then compares δ to the threshold t to decide *match* or *non-match*.

Experiments

Methodology

We use $10,000$ images from the BOSSbase image database [24]. The images come from seven digital still cameras: Canon EOS 400D, Canon EOS 40D, Canon EOS 7D, Canon Rebel XSi, Pentex K20 D, Nikon D70 and Leica M9. We start with the images in their native camera format and save them as TIFF files in Photoshop. To avoid the complexity introduced by different sizes of images, we center-crop the images to 512×512 and save them in the png file format in Matlab. We use the code referenced in [5] and made available on their website [20] to estimate PRNU reference images and extract noise residuals from individual images. We use 50 images to create PRNU reference images in all of our experiments. We use one minus the normalized correlation as the similarity score. We use kernel density estimation with Gaussian kernels to fit the probability density functions following work done in [22, 23].

Results

We compute trace-anchored SLRs for each image in the BOSSbase dataset as follows: we in turn treat each image as the questioned image. For each questioned image e_u, we in turn treat each BOSSbase device as the specific known device c_s. We estimate the noise residual x_u from e_u and the camera fingerprint e_s from 50 images from c_s and calculate the trace-anchored SLR for e_u given e_s as described in the Score-Based Likelihood Ratios section.

Figure 1 shows boxplots of the SLRs calculated for images from the Canon EOS 7D device when each BOSSbase device is in turn treated as the specific known source device. The first boxplot on the left in Figure 1 represents the case where we have true matches. I.e. we compute the SLR between images from the Canon EOS 7D and camera fingerprints from the Canon EOS 7D. The other six boxplots in Figure 1 show the cases where we have non-matches. I.e. we calculate the SLR between images from the Canon EOS 7D and camera fingerprints from other devices. As desired, the log(SLR) values are greater than zero for the matching cases, and mostly less than zero for the non-matching cases. We see similarly good results for the Canon Rebel XSi, Leica M9, Nikon D70, and Pentex K20D in Figures 2, 3, 4, and 5 respectively.

The results are not as good for images from the Canon EOS 40D and the Canon EOS 400D and are displayed in Figures 7 and 6. For the Canon EOS 40D, we see that the log(SLR) values for true matches are greater than zero. However, we see that the log(SLR) values are incorrectly greater than zero when the Canon EOS 400D is the specific source device. In this case the SLR results in misleading evidence in support of the prosecution hypothesis H_p. We see a similar phenomenon for images from the Canon EOS 400D when the Canon EOS 40D is the specific known source device.

Figures 1-7 exhibit two types of misleading evidence. If the log(SLR) value for a true non-match is positive, it is misleading evidence in favor of the prosecution hypothesis H_p, while If the log(SLR) value for a true match is negative, it is misleading evidence in favor of the defense hypothesis H_d. Table 1 shows the rates of misleading evidence in support of the prosecution hypothesis H_p. Each row corresponds to the device whose images are each treated as the questioned image. The rate for questioned im-

Figure 1. *Boxplots of $\log_{10}(SLR)$ values for images from device Canon EOS 7D by specific known source device*

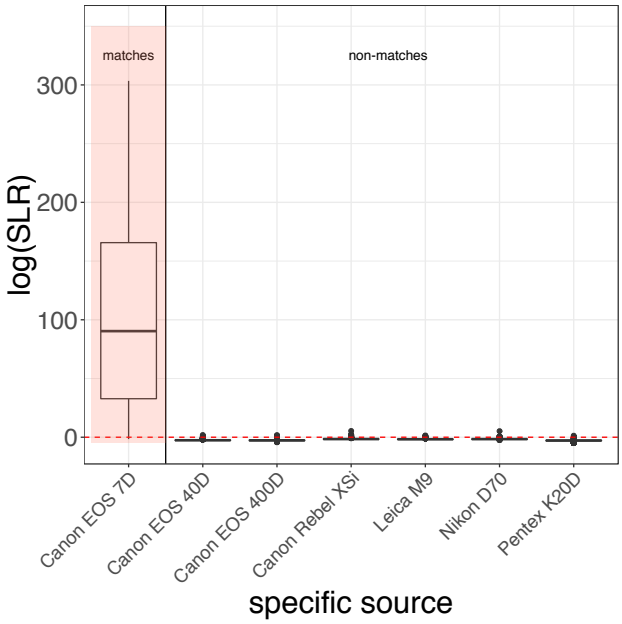

Figure 2. *Boxplots of $\log_{10}(SLR)$ values for imag es from device Canon Rebel XSi by specific known source device*

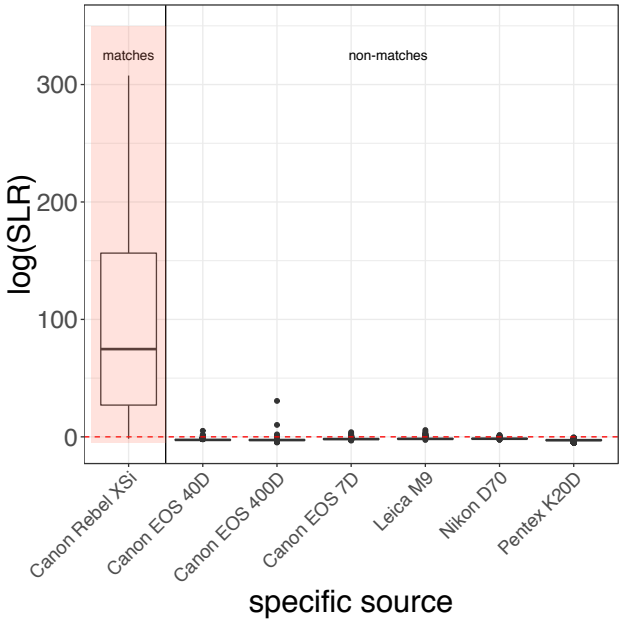

Figure 3. *Boxplots of $\log_{10}(SLR)$ values for images from device Leica M9 by specific known source device*

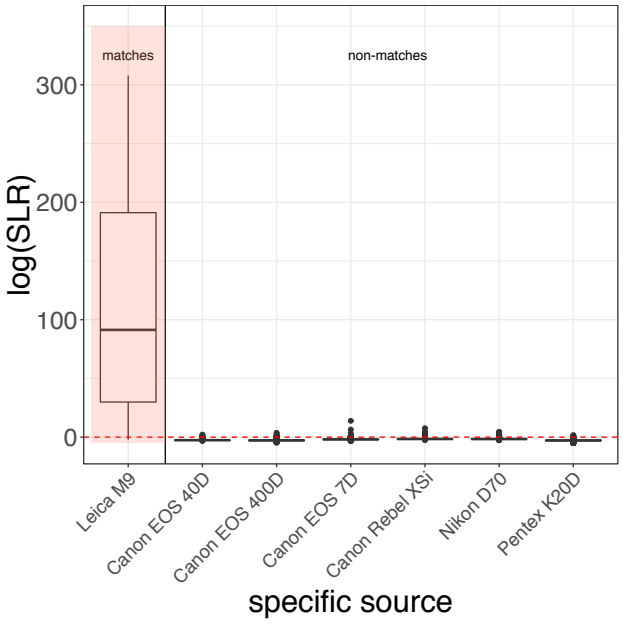

Figure 4. *Boxplots of $\log_{10}(SLR)$ values for images from device Nikon D70 by specific known source device*

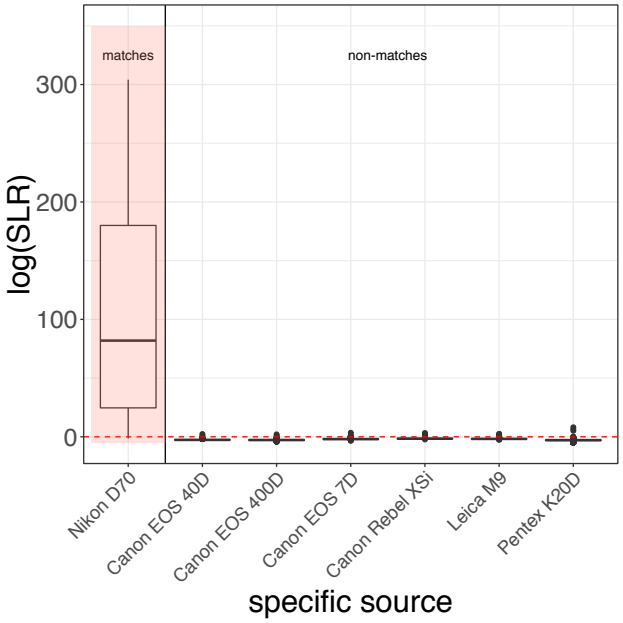

Figure 5. *Boxplots of $\log_{10}(SLR)$ values for images from device Pentex K20D by specific known source device*

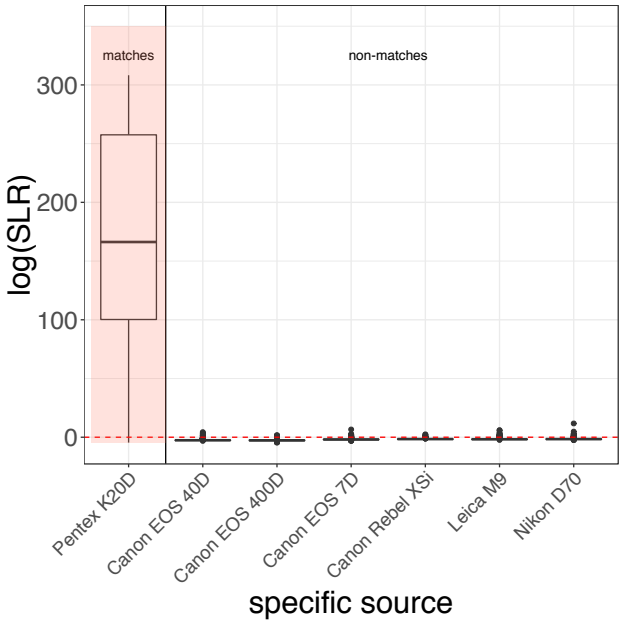

Figure 6. *Boxplots of $\log_{10}(SLR)$ values for images from device Canon EOS 400D by specific known source device*

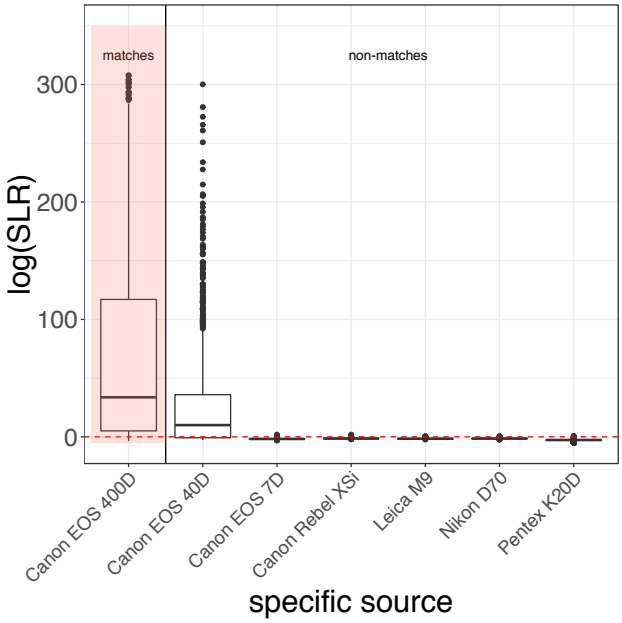

Figure 7. *Boxplots of $\log_{10}(SLR)$ values for images from device Canon EOS 40D by specific known source device*

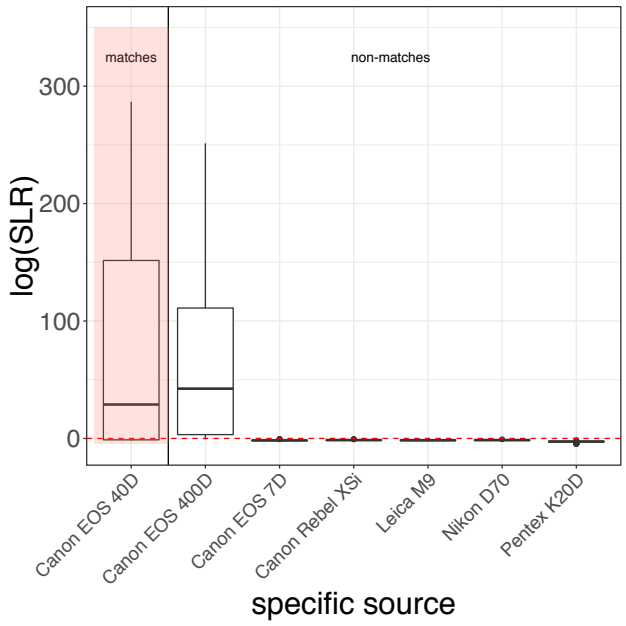

age device Canon EOS 400D and specific known device Canon Rebel XSi is 0.0060 and is calculated as the number of images from the Canon EOS 400D whose log(SLR) for specific device Canon Rebel XSi is greater than zero divided by the total number of images from the Canon EOS 400D. As expected from the boxplots, we see that the rate of misleading evidence in support of H_p is quite high when the questioned image device is the Canon EOS 400D and the specific known device is the Canon EOS 40D and vice versa. Despite the few outliers that appear in the boxplots, we see from table 1 that the rates of misleading evidence in support of H_p for all other cases is quite low. Table 4 shows the rates of misleading evidence in support of the defense hypothesis H_d. We see that 24.59% of the questioned images from the Canon EOS 40D were mistakenly identified as non-matches with that device. The rates of misleading evidence in support of H_d is less than 0.0628 for the other six devices.

In order to better understand the cause of the high rates of misleading evidence for Canon EOS 40D and 400D, we implement the universal detector method where we select a decision threshold to minimize the average of the false acceptance rate and false rejection rate and the SLR method for 50 randomly selected questioned images from the Canon EOS 400D. Table 1 shows the confusion matrix for the class predictions from the universal detector method and Table 2 shows the confusion matrix of class predictions from the SLR method. We see that the universal detector and SLR methods both do well in the case of true matches and their prediction results are similar. The SLR method misclassifies 37 out of 50 true non-matches as matches, which is roughly comparable to the false acceptance rate we observe in Figure 6. Table 1 shows that the universal detector method also misclassifies some true non-matches as matches, although not as many as the SLR method. Because both methods rely on the same PRNU estimates and the same similarity score, if the PRNUs and similarity scores used in the universal detector method are unable to

Table 3. **Rates of misleading evidence in support of H_p**

questioned image device	specific known device						
	Canon EOS 40D	Canon EOS 400D	Canon EOS 7D	Canon Rebel XSi	Leica M9	Nikon D70	Pentex K20D
Canon EOS 40D		0.9016	0	0	0	0	0
Canon EOS 400D	0.6987		0.0030	0.0060	0.0037	0.0015	0.0007
Canon EOS 7D	0.0030	0.0037		0.014	0.0052	0.0089	0.0015
Canon Rebel XSi	0.0029	0.0059	0.0069		0.0127	0.0093	
Leica M9	0.0018	0.0040	0.0083	0.0156		0.0091	0.0018
Nikon D70	0.0019	0.0019	0.0087	0.0165	0.0097		0.0039
Pentex K20D	0.0057	0.00423	0.0150	0.0114	0.0129	0.0165	

Table 4. **Rates of misleading evidence in support of H_d**

questioned image device	specific known device						
	Canon EOS 40D	Canon EOS 400D	Canon EOS 7D	Canon Rebel XSi	Leica M9	Nikon D70	Pentex K20D
Canon EOS 40D	0.2459						
Canon EOS 400D		0.0628					
Canon EOS 7D			0.0037				
Canon Rebel XSi				0.0044			
Leica M9					0.0062		
Nikon D70						0.0058	
Pentex K20D							0.0043

Table 1. **Confusion matrix of class predictions of universal detector for 50 randomly selected Canon EOS 400D images and specific known devices Canon EOS 400D and Canon EOS 40D**

		predicted class	
		match	non-match
True Class	match	48	2
	non-match	7	43

Table 2. **Confusion matrix of class predictions of SLR for 50 randomly selected Canon EOS 400D images and specific known devices Canon EOS 400D and Canon EOS 40D**

		predicted class	
		match	non-match
True Class	match	47	3
	non-match	37	13

perfectly distinguish between the two classes, we don't expect the SLR method to be able perfectly distinguish between the classes either. Thus we can attribute some of the false acceptance observed in the SLR method to the underlying inseparability of the classes. However, this does not account for all of the observed false acceptance rate, and we plan to study this further in future work.

Conclusions and Future Work

We develop the framework for trace-anchored score-based likelihood ratios (SLR) as a method to quantify the weight of evidence in the camera device identification problem. We calculate the trace-anchored SLRs for 10,000 png images from seven devices from the BOSSbase image dataset. For five of the BOSSbase devices, we achieve good results where the log(SLR) values for true matches are greater than zero and the log(SLR) values for most non-matches are less than zero. Images from two of the devices have positive log(SLR) values for a non-matching device, which is considered misleading evidence.

We plan to extend the experiments in this paper to larger images datasets with more devices. We also plan to implement and compare the source-anchored and general match SLRs with the trace-anchored SLR.

Acknowledgements

This work was partially funded by the Center for Statistics and Applications in Forensic Evidence (CSAFE) through Cooperative Agreement #70NANB15H176 between NIST and Iowa State University, which includes activities carried out at Carnegie Mellon University, University of California Irvine, and University of Virginia.

We would also like to thank Danica Ommen for her guidance and suggestions.

References

[1] J. Lukas, J. Fridrich, and M. Goljan, "Digital camera identification from sensor pattern noise," *IEEE Transactions on Information Forensics and Security*, vol. 1, no. 2, pp. 205–214, 2006.

[2] M. Goljan, M. Chen, and J. Fridrich, "Identifying common source

digital camera from image pairs," in *2007 IEEE International Conference on Image Processing*, vol. 6. IEEE, 2007, pp. VI–125.

[3] M. Goljan and J. Fridrich, "Camera identification from cropped and scaled images," in *Security, Forensics, Steganography, and Watermarking of Multimedia Contents X*, vol. 6819. International Society for Optics and Photonics, 2008, p. 68190E.

[4] M. Goljan, "Digital camera identification from images–estimating false acceptance probability," in *International workshop on digital watermarking*. Springer, Berlin, Heidelberg, 2008, pp. 454–468.

[5] M. Goljan, J. Fridrich, and T. Filler, "Large scale test of sensor fingerprint camera identification," in *Media forensics and security*, vol. 7254. International Society for Optics and Photonics, 2009, p. 72540I.

[6] H. B. Costa, R. F. Zampolo, D. M. Carmo, A. R. Castro, and E. P. Santos, "On the practical aspects of applying the prnu approach to device identification tasks," in *International Conference on Multimedia Forensics, Surveillance and Security*, 2012.

[7] M. Goljan, M. Chen, P. Comesaña, and J. Fridrich, "Effect of compression on sensor-fingerprint based camera identification," *Electronic Imaging*, vol. 2016, no. 8, pp. 1–10, 2016.

[8] S. Samaras, V. Mygdalis, and I. Pitas, "Robustness in blind camera identification," in *23rd International Conference on Pattern Recognition (ICPR)*. IEEE, 2016, pp. 3874–3879.

[9] J. Lukas, J. Fridrich, and M. Goljan, "Determining digital image origin using sensor imperfections," in *Image and Video Communications and Processing 2005*, vol. 5685. International Society for Optics and Photonics, 2005, pp. 249–260.

[10] S. Park, "Learning algorithms for forensic science applications," Ph.D. dissertation, Iowa State University, 2018.

[11] A. B. Hepler, C. P. Saunders, L. J. Davis, and J. Buscaglia, "Score-based likelihood ratios for handwriting evidence," *Forensic science international*, vol. 219, no. 1-3, pp. 129–140, 2012.

[12] A. Nordgaard and T. Höglund, "Assessment of approximate likelihood ratios from continuous distributions: a case study of digital camera identification," *Journal of forensic sciences*, vol. 56, no. 2, pp. 390–402, 2011.

[13] W. van Houten, I. Alberink, and Z. Geradts, "Implementation of the likelihood ratio framework for camera identification based on sensor noise patterns," *Law, Probability and Risk*, vol. 10, no. 2, pp. 149–159, 2011.

[14] T. Thai, R. Cogranne, and F. Retraint, "Camera model identification based on the heteroscedastic noise model," *IEEE Trans. Image Proces*, vol. 23, no. 1, pp. 250–263, 2014.

[15] R. C. T.H. Thai, F. Retraint, "Camera model identification based on dct coefficient statistics," *Digit. Signal Process*, vol. 40, pp. 88–100, 2015.

[16] T. Qiao, "Statistical detection for digital image forensics," Ph.D. dissertation, Troyes, 2016.

[17] T. Qiao, F. Retraint, R. Cogranne, and T. H. Thai, "Individual camera device identification from jpeg images," *Signal Processing: Image Communication*, vol. 52, pp. 74–86, 2017.

[18] M. Chen, J. Fridrich, M. Goljan, and J. Lukás, "Determining image origin and integrity using sensor noise," *IEEE Transactions on information forensics and security*, vol. 3, no. 1, pp. 74–90, 2008.

[19] M. Goljan and J. Fridrich, "Sensor-fingerprint based identification of images corrected for lens distortion," in *Media Watermarking, Security, and Forensics 2012*, vol. 8303. International Society for Optics and Photonics, 2012, p. 83030H.

[20] M. Goljan, J. Fridrich, M. Chen, and T. Filler. Camera fingerprint - matlab implementation. [Online]. Available: http://dde.binghamton.edu/download/camera_fingerprint/

[21] M. Chen, J. Fridrich, and M. Goljan, "Digital imaging sensor identification (further study)," in *In Security, steganography, and watermarking of multimedia contents IX*, vol. 6505. International Society for Optics and Photonics, 2007, p. 65050P.

[22] C. G. Aitken and D. Lucy, "Evaluation of trace evidence in the form of multivariate data," *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, vol. 53, no. 1, pp. 109–122, 2004.

[23] A. van Es, W. Wiarda, M. Hordijk, I. Alberink, and P. Vergeer, "Implementation and assessment of a likelihood ratio approach for the evaluation of la-icp-ms evidence in forensic glass analysis," *Science & Justice*, vol. 57, no. 3, pp. 181–192, 2017.

[24] P. Bas, T. Filler, and T. Pevný, "Break our steganographic system: The ins and outs of organizing boss," in *Information Hiding*. Springer, 2011, pp. 59–70.

Author Biography

Stephanie Reinders received her B.A. in Journalism and Asian Languages and Literatures from the University of Minnesota (2005). After working for several non-profit organizations as an administrative assistant, she returned to school to earn a graduate degree in mathematics. She received a post-baccalaureate certificate in Mathematics from Smith College (2013) and currently is pursuing a Ph.D. in Applied Mathematics and Computer Engineering at Iowa State University.

Li Lin received his B.S. degree in Mathematics from Capital Normal University, Beijing, China. He is currently pursuing the Ph.D degree in Applied Mathematics at Iowa State University, Ames, Iowa. His research interests include statistical image forensics, steganalysis, and statistical learning.

Wenhao Chen received his B.E. and M.S. degree in Computer Science from Xi?An Jiaotong University. He is currently a Ph.D. candidate in Computer Engineering at Iowa State University. His research interests include program analysis, malware detection, and steganalysis.

Dr. Yong Guan is a Professor of Electrical and Computer Engineering, the Associate Director for Research of Information Assurance Center, and the cyber forensics coordinator for NIST-CSAFE at Iowa State University. He received his Ph.D. degree in Computer Science from Texas A&M University. Supported by NSF, NIST, IARPA, ARO and Boeing, his research focuses on security and privacy issues, including digital forensics, network security, and privacy-enhancing technologies for the Internet.

Dr. Jennifer Newman received her BA in Physics from Mount Holyoke College and her PhD in Mathematics from the University of Gainesville, FL. She is an Associate Professor of Mathematics at Iowa State University in the Department of Mathematics, her research focusing on image processing, stochastic modeling, steganalysis and image forensics. She is a member of SIAM and IS&T.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

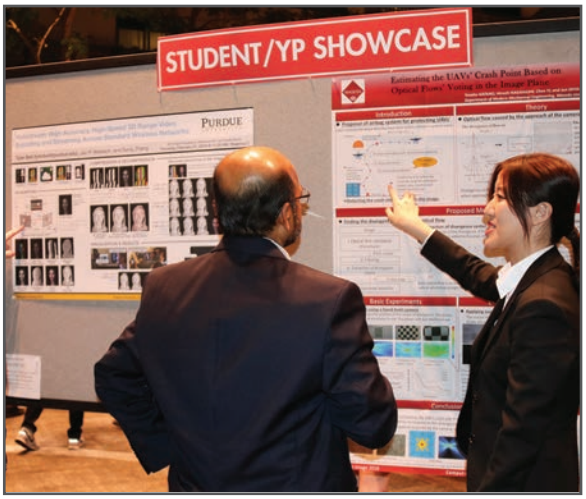

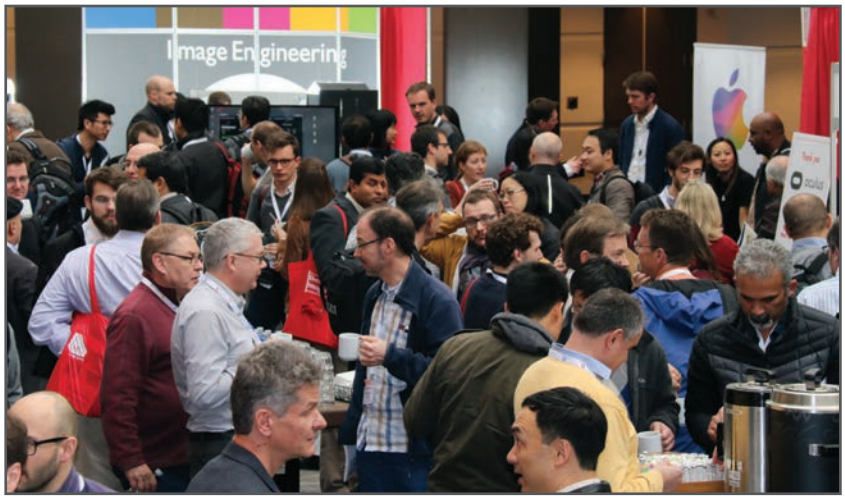

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org