

Erwin J. Alles,^{1,†} M.Sc.; Zeno J. M. H. Geradts,¹ Ph.D.; and Cor J. Veenman,^{1,2} Ph.D.

Source Camera Identification for Heavily JPEG Compressed Low Resolution Still Images*

ABSTRACT: In this research, we examined whether fixed pattern noise or more specifically Photo Response Non-Uniformity (PRNU) can be used to identify the source camera of heavily JPEG compressed digital photographs of resolution 640 × 480 pixels. We extracted PRNU patterns from both reference and questioned images using a two-dimensional Gaussian filter and compared these patterns by calculating the correlation coefficient between them. Both the closed and open-set problems were addressed, leading the problems in the closed set to high accuracies for 83% for single images and 100% for around 20 simultaneously identified questioned images. The correct source camera was chosen from a set of 38 cameras of four different types. For the open-set problem, decision levels were obtained for several numbers of simultaneously identified questioned images. The corresponding false rejection rates were unsatisfactory for single images but improved for simultaneous identification of multiple images.

KEYWORDS: forensic science, digital source camera identification, fixed pattern noise, JPEG compression, digital fingerprint, still images

In an ever increasing number of criminal cases, digital photographs or video footage is an important or even crucial part of the incriminating evidence. In a number of such cases, the origin of the footage is questioned and thus has to be determined. For instance, in child pornography cases source camera identification can conclude whether a suspect merely owned or actually produced the evidential images.

When the origin of images has to be determined, two scenarios exist. Either two (sets of) photographs are present, of which one is known to originate from a suspect and it has to be determined whether the other set is obtained with the same camera, or one (set of) photograph(s) and a suspected source camera are present, in which case it has to be determined whether the images originated from that camera. The latter case can present much stronger evidence to whether or not the suspect produced the incriminating images as camera properties can be studied more thoroughly. In this research we will mainly focus on the latter situation, assuming the suspected source camera to be available and in working order.

To determine the origin of a given digital image, several techniques have been developed. For instance, the aspect ratio of the photograph, color quantization tables, and effects caused by color interpolation schemes (1) can be used to determine the camera model. These methods, however, only discriminate between camera models and thus cannot distinguish between the suspected source camera and a different camera of the exact same model.

To overcome this limitation, use has to be made of unique features, i.e., features that differ from camera to camera. Such unique features may be found in the extended file information

(EXIF) header. An EXIF header is additional information embedded in digital images in which the camera manufacturer, model, and serial number can be found (2). If the serial number extracted from questioned images matches that of the suspected source camera, then this is the strong evidence that the suspected camera is indeed the source camera. However, EXIF headers can easily be modified or removed and in these cases this feature cannot be used.

A different approach is to study the traces left in the images by the imaging sensor. One strong type of such a trace is pixel defects. By studying the type, locations, and numbers of pixel defects in the questioned images and the suspected source camera and similar cameras, it can be concluded whether the suspected camera is likely to be the source camera (3). However, pixel defects could be absent or not visible in the photograph under study, and in the case of lossy compressed images, in slightly different locations.

So instead of the limited number of pixel defects, traces introduced by the imaging sensor affecting all pixels can be used. One such trace is the so-called fixed pattern noise (4). Within a digital camera, several mechanisms introduce stochastic noise causing pixel intensities to deviate randomly from the value that is expected based on the photographed scene. However, noise in the production process caused by fluctuations in manufacturing conditions results in static differences in the response of the pixels. This fixed pattern noise can thus be seen as a kind of fingerprint within digital images and can be used to identify the source camera.

Part of the fixed pattern noise, the so-called Photo Response Non-Uniformity (PRNU) (5), has been successfully used in source camera identification in (5). In this work, the source cameras of around 3000 supposedly questioned images were selected without error from a group of nine digital cameras of resolution 1280 × 960 and higher. Both uncompressed and mildly to moderately JPEG-compressed images (JPEG quality factors between 100 and around 70) were attempted. Out of these nine cameras, only two were of the same model. Even though distinction could be made between these two cameras, more research on the unicity of the PRNU patterns of different cameras of the same model is required.

¹Department of Digital Technology and Biometry, Netherlands Forensic Institute, The Hague, Holland.

²Intelligent Systems Lab, University of Amsterdam, Amsterdam, Holland.

*A portion of this work was presented at the International Conference on Computational Science and its Applications (ICCSA) in Perugia, Italy from June 30 to July 3, 2008.

[†]Present Address: Faculty of Applied Sciences, Department Imaging Science and Technology, Laboratory of Acoustical Imaging and Sound Control, Delft University of Technology, Delft, Holland.

Received 11 Feb. 2008; and in revised form 17 July 2008; accepted 20 July 2008.

Photo Response Non-Uniformity has also been used to identify the origin of video footage (6). For three camcorders of two different models and various compression techniques and qualities, the correct source camera could be pointed out for clips of resolution 536×720 pixels. For high-quality clips, a duration of 40 sec was required, for strongly compressed footage durations up to 10 min were necessary. These durations correspond to as many as 600 and 9000 images, respectively.

In this research, we will focus on source camera identification based on PRNU for full color still images of a low resolution of 640×480 pixels, acquired with webcams and phone cameras. To save on bandwidth, these types of cameras use heavy compression with JPEG quality factors as low as 30, comparable with the strongly compressed footage (6).

As we focus on still images, requiring 600 images or more for reliable results is not in general realistic. The performance of the identification scheme (6) thus has to be improved significantly. Therefore, we will propose new techniques to address problems introduced by lossy JPEG compression and additional information being extracted along with the PRNU pattern.

For the experiments, a total of 38 cameras of four different models were used, with a minimum of eight cameras of the same model to test the unicity of the patterns. For proper comparison of the different cameras, only cameras with the same native resolution were used. In this research only unaltered photographs, i.e., uncropped, unscaled, etc., were experimented. This is not a serious restriction as for the selected cameras used in this research, i.e., webcams and phone cameras, photographs are usually not altered. As EXIF headers are well-known identifiers and trivially removed from images, these headers are considered unavailable.

In the remainder of this article, we will first state the problem that has to be solved, followed by our proposed scheme for the extraction of PRNU from both reference and questioned images. Then we will propose techniques to address difficulties in the extraction of the reference and questioned patterns, followed by the performed experiments and the resulting performance of the identification scheme. Finally some pointers to future work are given.

Problem Statement

Source camera identification is based on the comparison of discriminating features extracted from both the questioned images and the suspected camera. To denote the similarity between the reference and questioned features, a similarity measure s_i is required:

$$s_i = f(P_i, P_q) \quad (1)$$

where P_i is the extracted reference feature, P_q is the feature extracted from the questioned photographs and f an arbitrary function.

Source camera identification appears in two scenarios. The first scenario is to select, from a set of cameras, the camera that is most likely to be the source camera of a collection of photographs. This problem will be referred to as the “closed-set” problem and is solved by selecting camera j , the camera corresponding to the maximum value for similarity s_i :

$$j = \arg \max_i s_i. \quad (2)$$

Consequently, in the closed-set problem, in all cases one camera will be identified as being the source camera. Naturally to reduce the number of false identifications a decision threshold can be

introduced. However, as the closed set acts merely as a proof of concept, the closed-set threshold is not further researched.

The second and harder scenario is to establish that a certain camera is the source camera of a given (collection of) photograph(s). In principle, this so-called “open-set” problem requires knowledge of the characteristics of all digital cameras currently in use.

The open-set problem is solved by establishing a suitable decision threshold level d on s_i : if $s_i > d$, the corresponding camera is decided to be the source camera of the photographs under study. This approach can point out multiple cameras as the source camera or none at all if all values $s_i < d$.

Both the open and closed-set problems will be addressed in this paper using a selection of cameras that is assumed to be representative for all cameras of the same resolution. The closed-set problem can be seen as a proof of concept; its evidential value is limited as it can only exclude cameras from a group of suspected source cameras when its similarity is much lower than the maximum similarity. The open-set problem is much more valuable and might answer the question whether a (collection of) photograph(s) is taken with the suspected, available source camera.

PRNU

In this section, we will elaborate on the previously mentioned fixed pattern noise. Fixed pattern noise is introduced by small differences in production conditions for each separate sensor and results in small variations in pixel size and performance (4) which form a fixed pattern within the sensor. As all sensors of the same type are produced in the same process, it is expected that the resulting fixed pattern is partly unique and partly the same for all sensors of the same type. This statement is supported by Fig. 1. In this figure the average of all reference PRNU patterns extracted from one camera model is shown. Very clear features remain after averaging, indicating that those features are present in all individual patterns. The common features were not further researched. It thus has to be determined whether the unique part is strong enough for reliable source camera identification.




FIG. 1—The average of the 10 reference PRNU patterns of all Motorola V360 phone cameras used in this research. Clearly visible are the horizontal and vertical lines present in this averaged pattern. These lines are also visible in all individual PRNU patterns (not shown). PRNU, Photo Response Non-Uniformity.

Due to the large amount of pixels in present day digital consumer cameras, typically between 75,000 and 12 million, many different patterns are possible. As the patterns are for a large part random in nature, they are likely to be unique. Accordingly, fixed pattern noise can be considered as the fingerprint of a digital camera.

Fixed pattern noise consists of two parts: dark current and PRNU. The former is caused by thermally generated free charge within a pixel leading to additional intensity being registered, and a fixed pattern emerges due to slight inhomogeneities in material properties introduced during manufacturing causing some pixels to generate more dark current than others.

The latter contribution is also caused by slight differences between the pixels in material and construction. Pixels that are slightly smaller or pixels of slightly less pure material composition are somewhat less sensitive than average. Being caused by sensitivity differences, PRNU is a multiplicative signal and its effect depends on the image content.

Therefore, we model pixel (i,j) of image I being stored to file, consistent with image degradation models in image restoration (7)

$$I(i,j) = F(i,j) \cdot O(i,j) + D(i,j) + N(i,j) \quad (3)$$

where O is the ideal, noise-free value; F , the PRNU; D , the dark current contribution; and N , the stochastic noise of pixel (i,j) . The latter two contributions are additive signals; they are independent of the pixel value O .

Dark current has been successfully used in source camera identification (8). However, the resulting pattern is relatively weak and only detectable in dark scenes. Therefore from here on dark current will be neglected, and thereby Eq. (3) simplifies to

$$I(i,j) \approx F(i,j) \cdot O(i,j) + N(i,j). \quad (4)$$

As PRNU is a multiplicative signal, it will be invisible in very dark or saturated scenes.

PRNU Extraction

Being caused by per-pixel variations in construction and material, PRNU results in a per-pixel and thus high-frequency pattern. To extract F' , an estimate of the high-frequent PRNU pattern F together with stochastic noise, from an image I , a high-pass filter implemented as

$$F' = I - G * I \quad (5)$$

is applied, where G is a low-pass filter and $*$ denotes convolution.

In this research we apply a two-dimensional Gaussian filter of variance σ^2 in the spatial domain. This filter yielded in experiments, not reported here, results comparable with the wavelet-domain-based method used in Eq. (5), but requires less computation time. The influence of the variance parameter will be studied in the Experiments section.

Information Extracted Along with PRNU

Scene Content

A consequence of applying a high-pass filter is that besides high-frequency noise and PRNU, also high-frequency scene content is affected. Accordingly, part of the scene content would be considered PRNU. As is visible in Fig. 2, this extra content is significantly stronger than the actual PRNU.




FIG. 2—Example of additional information being extracted along with noise and PRNU. The original image (a) can easily be recognized in the extracted noise, (b), and the noise and PRNU are barely visible. The high-frequency content in (b) is obtained using a Gaussian filter of variance 0.6 pixels, and contrast scaled to improve visibility. PRNU, Photo Response Non-Uniformity.

Reference Pattern—As in our case the suspected source camera is assumed available and in working condition, influence of scene content in extracted PRNU patterns can be avoided for the reference patterns. One possibility is to use flat field photographs. Ideally, a flat field image is the result of a uniformly illuminated sensor and therefore no detail is present. Any information present in a flat field image is thus a result of PRNU or stochastic noise. In the notation of Eq. (4), the intensity I of pixel (i,j) after illuminating the sensor uniformly is

$$I(i,j) = F(i,j) \cdot C + N(i,j) \quad (6)$$

in which C is a constant intensity due to the uniform illumination.

However, perfectly uniform illumination is not easily achieved as the optics and sensor give rise to vignetting, a decrease of intensity towards the edges of the image. Even with the optics removed, flat fielding is a cumbersome task as this requires a parallel beam of light of uniform intensity. Correcting for vignetting effects is not straightforward as simply raising or lowering the intensity of certain areas affects the PRNU and noise levels as well.

As the PRNU pattern of interest is of high frequency, the Gaussian filter used to extract this pattern will have a high cut-off frequency. Therefore, reference images are not required to be perfect flat field images as long as the contained detail is of lower frequency than that of the PRNU. Thus, using a high-pass filter to extract PRNU from reference images allows for simple setups; we photographed a white sheet of paper in fluorescent light under varying angles to suppress any possible constant contribution due to scene content. Where possible the cameras were defocused to further avoid any detail.

Questioned Images—Upon obtaining the reference pattern, one can acquire any required image to extract the PRNU pattern as efficient as possible. However, scene content which is extracted along with PRNU is unavoidable in questioned images. This additional scene content has to be suppressed in F'_q , the high-frequency content extracted from the questioned image, otherwise identification will be (partly) based on this scene content.

As the PRNU pattern has certain maximum amplitude, any pixel in PRNU patterns extracted from questioned images of higher absolute value than this amplitude cannot be part of the PRNU. Therefore, we propose to suppress the scene content by applying a threshold; only pixels in the questioned pattern with absolute value smaller than a certain threshold are considered reliable pattern pixels. The other pixels are considered unreliable and masked out.

The threshold levels for the various camera types used are determined in the Experiments section.

Stochastic Noise

Extracting the high frequencies from an image will not only yield an approximation of the PRNU pattern, but also a contribution of the stochastic noise. Assuming the noise contributions $N(i,j)$ to be zero-mean, i.i.d. variables, the stochastic noise can be removed by averaging the extracted high-frequency contents $F'_n(i,j)$ from multiple images:

$$F'_{\text{ref}}(i,j) = \frac{1}{N} \sum_{n=1}^N F'_n(i,j) \quad (7)$$

where $F'_{\text{ref}}(i,j)$ is the average of pixel (i,j) of all the contents extracted from reference images $n = 1, 2, \dots, N$. Averaging multiple high-frequency contents also averages out any possible scene content present in the reference image and thus in the reference pattern, like for instance fine structure within the paper sheet and edges introduced by shadows.

Above a certain number of reference images N , the stochastic noise is averaged out and what remains is a stable PRNU pattern. As a consequence of convergence of all pixels, the sum of the absolute values of all pixel values will converge to some value as well. The quadratic sum Q of all pixel values is thus a good measure of the convergence of the pattern:

$$Q = \sum_{\text{all } (i,j)} F'_{\text{ref}}(i,j)^2 \quad (8)$$

where the summation runs over all pixels and $F'_{\text{ref}}(i,j)$ is pixel (i,j) of the averaged reference pattern as defined in Eq. (7). A stable reference pattern is reached if Q converges to a constant value.

In Fig. 3 the value of Q is plotted against the number of images acquired with a Motorola V360 phone camera included in the averaged reference pattern. Judging from this figure, a stable reference pattern is obtained after averaging the high-frequency contents extracted from around 300 reference images, though Q remains roughly the same after 50 images. The same number of images were required for the reference patterns (Eq. [5]). Different camera

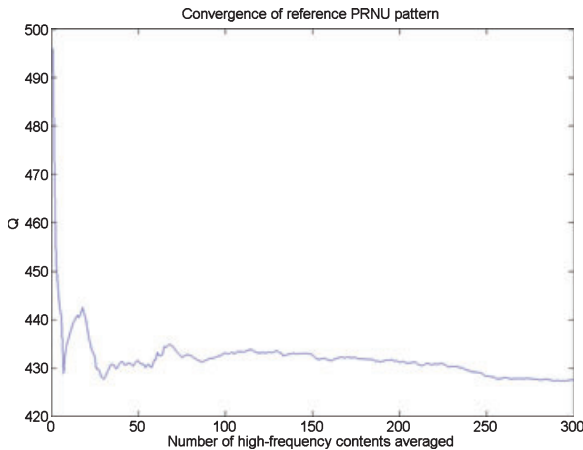


FIG. 3—Convergence of Q versus the number of reference images included in the averaged reference pattern. From around 300 images Q becomes stable. Images originated from a Motorola V360 phone camera. PRNU, Photo Response Non-Uniformity.

models yielded similar plots, and as the camera is assumed available, we used 300 reference images to obtain a reference pattern.

JPEG Compression

All cameras used in this research applied JPEG compression to their images to save on bandwidth. Upon JPEG compression, the discrete cosine transform (DCT) of groups of 8×8 pixels is calculated and the resulting coefficients are stored. The reduction of required storage space is reached by suppressing or removing some of the higher frequencies within each such DCT block.

However, the content of the surrounding DCT blocks is not taken into account at all and therefore continuity between neighboring DCT blocks is not guaranteed. This discontinuity results in a clear pattern present in the extracted high-frequency content. As visible in Fig. 4, this pattern coincides with the edges of the DCT blocks and is similar in all JPEG compressed images of the same resolution. The resulting pattern is even for reference images, mainly caused by scene content and stronger than the reference PRNU pattern.

A pattern that is present in all photographs and reference patterns and stronger than the pattern caused by PRNU will severely complicate source camera identification. Therefore, this blocking effect should be suppressed to obtain a camera fingerprint that is mainly based on the actual image content.

Noting the periodicity of the pattern caused by JPEG compression, a straightforward approach would be to suppress the corresponding frequency components within the image. This approach is used in the study by Chen et al. (6). In this work, the edge artifacts are modeled as a spike-train. However, as can be seen in Fig. 4, this pattern is only approximately periodic: the edges of the DCT blocks are present in differing strengths and colors. Furthermore, both the left- and right-hand side of the DCT blocks show edge artifacts. Thus, the assumption of a spike-train is not optimal.

In the context of image enhancement, in the study by Singh et al. (9), several other techniques to suppress these edge artifacts are treated. However, all of these techniques have only limited applicability as they aim at removing the pattern to yield a more pleasing image. For that purpose, the artifacts are mostly smoothed instead of removed.

Instead of using the periodicity of the edges, we propose an approach that utilizes the fixed locations of the JPEG edge artifacts. In this approach, multiple pixels are per color channel, averaged into one “macro element.” This way the effect of the DCT block




FIG. 4—Example of a part of a reference pattern obtained from 300 reference images. The edges of the discrete cosine transform blocks are clearly present and dominant over the other content within the pattern.

edges is averaged out over these multiple pixels and thereby strongly suppressed. However, by averaging multiple pixels into one macro element, the resolution of the images, and therefore of the reference and questioned estimates of the PRNU patterns is decreased.

As each DCT block consists of $8 \times 8 = 64$ pixels, there are several ways in which the edge effects can be suppressed. As the DCT is calculated using only information within the DCT block, only pixels within one block should be averaged to suppress the edges in the most effective way. Naturally as much information as possible should be maintained and therefore all pixels within each block should be included upon averaging. This leaves four possibilities: averaging groups of 8×8 , 4×4 , and 2×2 into one effective pixel or no pixel averaging at all.

The choice as which of these options is to be used is governed by the trade-off between the effective resolution of the resulting pattern, and thus the number of possible unique PRNU patterns and the remaining strength of the pattern introduced by JPEG compression. This trade-off will be studied extensively in the Experiments section. In Fig. 5 the effect of averaging groups of 8×8 pixels into one macro element on an extracted reference PRNU pattern is demonstrated.

PRNU Detection

If the reference pattern obtained by filtering reference images is strongly present in PRNU extracted from questioned images, it is very likely that the questioned images originated from the camera corresponding to that reference pattern. To measure this, we need a similarity function s_i as in the Eq. (1). As PRNU is a multiplicative signal, the resulting pattern will differ from image to image in intensity and contrast, and a comparison method should be insensitive to both intensity and contrast.

Therefore, an appropriate measure for the presence of reference pattern F'_{ref} in the questioned pattern F'_q is the correlation coefficient ρ between the two patterns as used in Eq. (5):

$$s_i = \rho(F'_{ref}, F'_q) = \frac{\sum_{\text{all}(i,j)} (F'_{ref}(i,j) - \overline{F'_{ref}}) \cdot (F'_q(i,j) - \overline{F'_q})}{\sqrt{\sum_{\text{all}(i,j)} (F'_{ref}(i,j) - \overline{F'_{ref}})^2} \cdot \sqrt{\sum_{\text{all}(i,j)} (F'_q(i,j) - \overline{F'_q})^2}} \quad (9)$$

where $\overline{F'_q}$ is the average of all pixel values within natural PRNU pattern F'_q . The higher the value for ρ , the more the two patterns are alike.




FIG. 5—The effect of pixel averaging on extracted PRNU patterns. On the left is the original extracted pattern, on the right is the pattern after averaging 8×8 pixels into one macro element. Note that even though much information is discarded, similar features are still present. Only the red channel is selected for these images. PRNU, Photo Response Non-Uniformity.

Any PRNU pattern whether extracted from reference or questioned images will be of the same dimensions as the source image: it consists of three two-dimensional patterns, one for each of the three color channels. Due to the color interpolation applied to each pixel (10), the three color channels are, however, not independent. To preserve this dependence, we compare the resulting three color channels separately using Eq. (9) and the three correlation coefficients are averaged. Upon comparison, only the pixels in F'_q that are considered reliable after thresholding are compared with the corresponding pixels in F'_{ref} .

Experiments

In this section, all the above is put together and the accuracy of the resulting algorithm is tested. The questioned images in this research were shot at a wide variety of settings to obtain the most general results. The 38 cameras used are listed in Table 1. Assuming them to be most commonly used, all cameras were set to the following settings: illumination and white balance are set to automatic, zoom is not used, images are taken at native resolution, and JPEG compression is enabled. No control over the compression ratio was possible. Four examples of questioned images are shown in Fig. 6.

Questioned images can be images of any scene that is of relevance in criminal investigations and therefore can be taken under a wide variety of conditions. To obtain the most general results images used in this research as questioned images are also acquired under various conditions. Indoor and outdoor, motion blurred, defocused, high-detail, over- and underexposed scenes, and scenes virtually free from detail are used.

Parameter Estimation

To summarize our method, we extract PRNU from both reference and questioned images by applying a Gaussian filter, and we suppress additional scene content extracted along with the PRNU patterns from questioned images by masking out all pixels with intensity higher than a certain threshold. To suppress JPEG edge artifacts, we average several pixels into one macro element. These three techniques all have one parameter that has to be optimized, and the three parameters are estimated below. Experiments not reported here showed that the order in which these three parameters are optimized is irrelevant, and thus that the three parameters are virtually independent. The optimal macro element size is researched in section Camera Identification Performance.

Variance of the Gaussian Filter—The Gaussian filter has one free parameter; σ^2 , the variance of the kernel. This variance determines the cut-off frequency of the filter and thus what part of a photograph is considered to be scene content or stochastic noise and PRNU. To determine the optimal value for σ , the following test is performed for each camera model.

TABLE 1—List of used cameras.

Model	Quantity	Camera Numbering
Motorola V360	10	7.1–7.10
Vodafone 710 (internal camera)	10	9.1–9.10
Creative Live! Cam Video IM	8	11.1–11.8
Logitech QuickCam Communicate STX	10	12.1–12.10



FIG. 6—Some examples of questioned images used in this research.

The reference pattern of one camera per model is determined from 300 reference images using a Gaussian filter of varying variance σ . To remove DCT-block artifacts, groups of 4×4 pixels are averaged into one macro element as this will turn out to be the best of the four JPEG edge artifact suppression options previously discussed. For each value of σ the high-frequency content is extracted using the same Gaussian filter from 106 questioned photographs of the same camera (match) and from 106 images of a different camera (mismatch) of the same model.

Using Eq. (9), 106 correlation coefficients between each questioned image and the reference pattern are calculated for both the match and mismatch situation. The mean of these correlation coefficients together with twice the standard deviation are plotted against filter variance σ for two Creative Live! Cam Video IM webcams in Fig. 7.

The plots in Fig. 7 can roughly be divided into three regions. In region I, $\sigma < 0.6$, the Gaussian kernel does not extend past the

central pixel, i.e., the cut-off frequency is higher than the highest frequency present in the image. Therefore, no information is extracted from the images.

In region III, when $\sigma > 1.4$, the Gaussian kernel is too wide; the corresponding cut-off frequency is too low, leading to too much non-PRNU information being extracted from the images. Therefore, additional patterns are present in the extracted information and correspondingly a higher spread in the correlation coefficients is found. Also, in this region there is no separation between match and mismatch correlation coefficients; all mismatch correlation coefficients found are within the match correlation coefficient distribution.

Only in region II, for $0.6 < \sigma < 1.4$, separation between match and mismatch correlation coefficients is found; the lower boundary for mismatch correlations is lower than that of the match correlations. Therefore, choosing σ in this region will result in the best performance. As can be seen in the enlargement of region II in

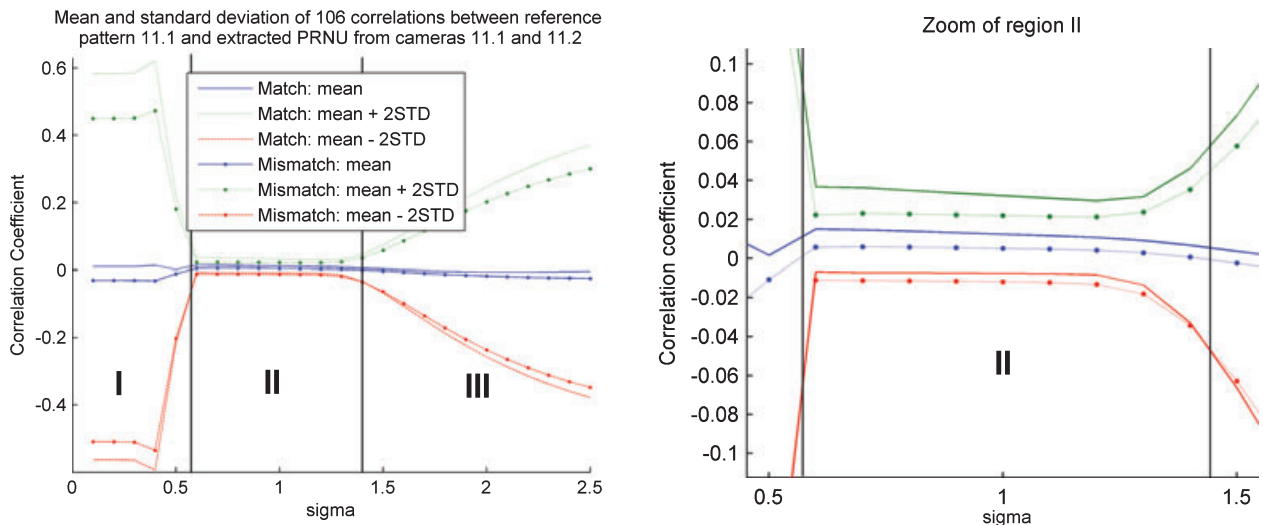


FIG. 7—Plot of correlation distribution between one reference pattern and 106 questioned images from the same (match) or a different (mismatch) Creative Live! Cam Video IM camera against filter variance σ . In the enlarged plot of section II it can be seen that only for $0.6 \leq \sigma \leq 1.4$, the match correlation lower boundary is higher than the mismatch lower boundary. PRNU, Photo Response Non-Uniformity.

Fig. 7, the separation is largest for $\sigma = 0.6$, and therefore a Gaussian filter of variance $\sigma = 0.6$ will be used. The same conclusion is reached for the other three camera models used in this research and is easily explained using $\sigma = 0.6$, the Gaussian kernel extends just to the surrounding pixels and therefore the correct cut-off frequency is achieved. This value is thus a universal value.

Threshold Levels to Suppress Scene Content—In the subsection Questioned Images, a method to suppress scene content from extracted content was proposed; only pixels from the extracted patterns with absolute value lower than a certain threshold value based on the PRNU amplitude are considered reliable. To determine the appropriate threshold level for a certain camera model, we performed the following experiment.

For one camera, 300 reference images are acquired. From these photographs, due to the limited amount of data present, only 200 images are used to calculate a reference pattern in the way described above as the same set of photographs is used. The PRNU and noise of the remaining 100 images are extracted using a Gaussian filter of variance $\sigma = 0.6$ and thresholding is applied for various t ; all pixels with absolute value higher than t are masked out.

For various t the 100 correlation coefficients between reference pattern and extracted contents from a Creative Live! Cam Video IM webcam are calculated and the mean and standard deviation of these 100 correlations are plotted against t in Fig. 8. Note that the mean match correlation is much larger than that found in Fig. 7. This is easily explained as the data in Fig. 8 was obtained by comparing reference images to a reference pattern, whereas in Fig. 7 questioned images were compared with a reference pattern, thereby introducing spurious effects caused by high-frequency image content.

In Fig. 8 it can be observed that for thresholds $t > 4$ the distribution of correlation coefficients remains the same, whereas for lower values where more and more of the pattern is suppressed, the correlation drops rapidly. This implies that for this camera, below a threshold level of $t = 4$, the PRNU itself is suppressed significantly. Note that the threshold should be as low as possible to remove as much scene content as possible from questioned images without suppressing the PRNU itself.

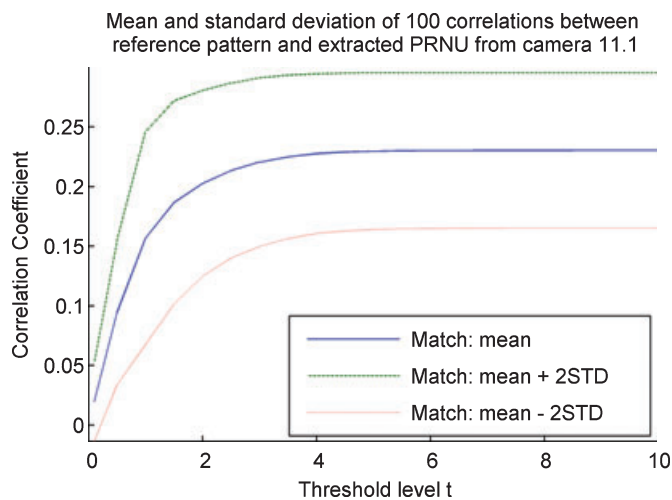


FIG. 8—Plot showing the distribution of the 100 correlation coefficients between PRNU patterns estimated from single reference images and the pattern obtained from 200 reference patterns. Below $t = 4$ the correlation decreases rapidly, indicating that $t = 4$ is the effective amplitude of the PRNU for this camera, a Creative Live! Cam Video IM. PRNU, Photo Response Non-Uniformity.

For this camera, the effective amplitude of the PRNU is thus found to be a pixel value of 4. In total three Creative Live! Cam Video IM webcams were tested and all three resulted in the same threshold level $t = 4$. Using the same method, we obtained threshold levels for the remaining three camera models. For these camera models, different cameras of the same model yielded the same threshold level as well.

Camera Identification Performance

Closed-Set Problem—The closed-set problem concerns selecting from a fixed group of cameras the one that is most likely to be the source camera. Therefore, its evidential strength is lower than in the open-set problem as not all possible cameras can be included. However, the closed-set problem is still valuable as cameras can be determined not to be the source camera.

To determine the performance of the closed-set source camera identification scheme, the following experiment has been performed. First, the reference patterns of all 38 cameras in Table 1 are calculated. With each of the 38 cameras, 106 supposedly questioned images are acquired. From all questioned images the high-frequency content is extracted and using thresholding, scene content within the estimated pattern is suppressed.

For each camera, 100 random selections were made from the 106 questioned images. Each of these 100 images was matched to one of the 38 cameras based on a maximum correlation between reference and estimated questioned patterns.

In Table 2 the results of the above experiment are shown in the form of a confusion matrix. On each row, the 100 randomly selected images originating from the camera corresponding to that row are based on a maximum correlation coefficient as in Eq. (2), matched to one of the 38 cameras.

Each time a certain camera was concluded to be the source camera, its corresponding entry was raised by one. From the 100 questioned images from camera 7.1 for instance, a Motorola V360, 83 were concluded to originate from the correct camera, whereas two questioned photographs were incorrectly matched to camera 7.5, etc.

Table 2 shows that the diagonal elements, i.e., correct identifications, have high values. This suggests that in the closed-set problem, source cameras can quite accurately be pointed out based on a maximum correlation coefficient. To quantify the performance of this source camera identification, the accuracy a is calculated from the confusion matrix:

$$a = \frac{\text{Number of correct identifications}}{\text{Total number of identifications}} \quad (10)$$

For the confusion matrix of identification of single questioned images in Table 2, this accuracy is $a = 83.7\%$.

The above accuracy of $a = 83.7\%$ was reached by averaging groups of 4×4 pixels into one macro element in order to suppress the DCT-block edge artifacts. When instead of pixel averaging, frequency suppression is applied as in Eq. (6), the accuracy for this experiment is only 41.2%. It can thus be concluded that our proposed technique yields significantly better results.

The experiment was repeated for simultaneous identification of $n = 2, 5, 10, 17,$ and 25 questioned images as well, i.e., the high-frequency contents of several randomly selected questioned images of the exact same camera were averaged, and the resulting pattern was then compared to all reference patterns. Also the size of the pixel group being averaged into one macro element to suppress DCT-block edge artifacts was varied to determine the optimal method. All resulting accuracies are shown in Table 3.

TABLE 3—Accuracies of the closed-set source camera identification problem.

<i>n</i>	Size of Averaged Pixel Group			
	1 × 1	2 × 2	4 × 4	8 × 8
1	47.6%	80.4%	83.7%	54.7%
2	59.7%	87.9%	93.1%	70.1%
5	74.2%	94.7%	99.0%	89.7%
10	81.8%	97.5%	99.9%	96.1%
17	85.9%	98.0%	100%	98.3%
25	87.1%	99.0%	100%	99.3%

Table 3 shows that the accuracy of the source camera identification increases when more questioned images are identified simultaneously. This is easily explained as averaging multiple extractions to one questioned PRNU pattern will suppress both stochastic noise and scene content.

Based on Table 3, the source camera, in a closed set of 38 cameras of four different models, can be quite reliably identified for simultaneous identification of $n = 5$ questioned images. For compressed footage of comparable resolution and compression factors, for a closed set of three cameras, duration of 40 sec or more was required for a reliable identification, corresponding to more than 600 images (6). It can thus again be concluded that pixel averaging is much more effective in removing JPEG edge artifacts than frequency suppression. Requiring as few as five frames for a reliable identification is reasonable for source camera identification of still images.

As a final remark, most of the images of which the source camera was incorrectly identified either contained lots of high-frequency detail or contained saturated or very dark regions. In images with high-frequency details, separating PRNU from this detail is complicated, and in saturated images the effect of the PRNU is lost as all pixels will have the same maximum value.

Optimal Macro Element Size—Table 3 also shows that the best performance is reached when groups of 4×4 pixels are averaged to suppress JPEG edge artifacts. Even though in this case only four macro elements remain per DCT-block and the number of possible unique patterns is decreased by a factor of 16, the edge artifacts are strongly suppressed resulting in an improvement of the performance of the identification scheme.

Open-Set Problem—The closed-set problem always results in one camera to be chosen as most likely source camera as this decision is based on determining the maximum correlation between reference and questioned PRNU patterns. This approach is, however, not suitable to the open set as this would require all cameras in existence to be compared.

Instead of determining the most likely source camera, a decision level d for the correlation coefficient has to be obtained above which it can, with reasonable certainty, be concluded that the camera corresponding to the reference pattern yielding this correlation is indeed the source camera.

To derive this decision level d , knowledge of the distribution of the correlation coefficients for a representative group of cameras is required. Once the distribution is known, it is relatively straightforward to impose an appropriate false acceptance rate (FAR), i.e., the rate at which the wrong camera is misidentified as being the source camera, and to determine the corresponding decision level d .

The correlation coefficient distribution was obtained using the same experiment as for the closed-set problem; for each of the 38

cameras, 100 random selections of one or multiple questioned images from the same camera were made. The PRNU patterns extracted from these selections of questioned images were compared with all reference patterns, including the reference pattern corresponding to the actual source camera.

The correlation coefficients between matching, i.e., reference and questioned patterns from the same camera, and mismatching questioned/reference patterns are shown in Fig. 9. Again, $\sigma = 0.6$ and macro elements of size 4×4 pixels were used.

From the histograms in Fig. 9 it can be observed that the separation between match and mismatch correlation is larger when more questioned images are simultaneously identified. This is to be expected as the closed-set problem showed higher accuracies for higher numbers of simultaneously identified questioned images. In Fig. 9 it is also clear that especially for simultaneous identification of multiple questioned images, the mismatch distribution is bimodal, which causes a large overlap between matching and mismatching distributions. The cause of this overlapping lump is unknown; however, an experiment not reported here showed that it was not caused by similarities between reference patterns of cameras of the same model.

From the histograms in Fig. 9, d is determined. First, a rather arbitrary FAR of 1‰ is chosen, implying that one in a thousand correlation coefficients between reference and nonmatching questioned patterns is greater than d . Thus, once in every thousand attempts the wrong camera is identified as being the source camera. This FAR certainly may not be any higher for reliable results and preferably should be lower still.

Requiring a FAR of 1‰ implies selecting a decision level d for which

$$\text{FAR} = \frac{\text{Number of } \rho_{\text{mismatch}} > d}{\text{Total number of } \rho_{\text{mismatch}}} = \frac{1}{1000}. \quad (11)$$

As the match and mismatch histograms exhibit some overlap, imposing a FAR will lead to some correct identifications to be discarded as well. The rate at which this occurs is called the false rejection rate (FRR)

$$\text{FRR} = \frac{\text{Number of } \rho_{\text{match}} < d}{\text{Total number of } \rho_{\text{match}}}. \quad (12)$$

Furthermore, each d will lead to different FAR and FRR values. For a certain value of d , the FAR and FRR are equal and this rate is called the equal error rate (EER). The EER is a good measure of the performance of the source camera identification scheme as a low value indicates low error rates and thus few incorrect identifications and few rejections of correct identifications. For a FAR of 1‰, the corresponding decision level d and FRR are given in Table 4 for all n . In that same table, the EERs are given for all n .

From Table 4 it appears that requiring FAR = 1‰ leads to high values of FRR for all n . Therefore, a large part of the correct identifications will be rejected. If more than one questioned image is simultaneously identified, the FRR drops as does the EER. Again, the performance is shown to improve upon identifying multiple images originating from the same camera simultaneously.

Discussion

For 38 cameras of four different models at a low resolution of 640×480 pixels, closed-set problems could be solved with accuracies of 83% for single up to 100% for selections of 25 questioned

Histograms of 3800 match and 140600 mismatch correlation coefficients

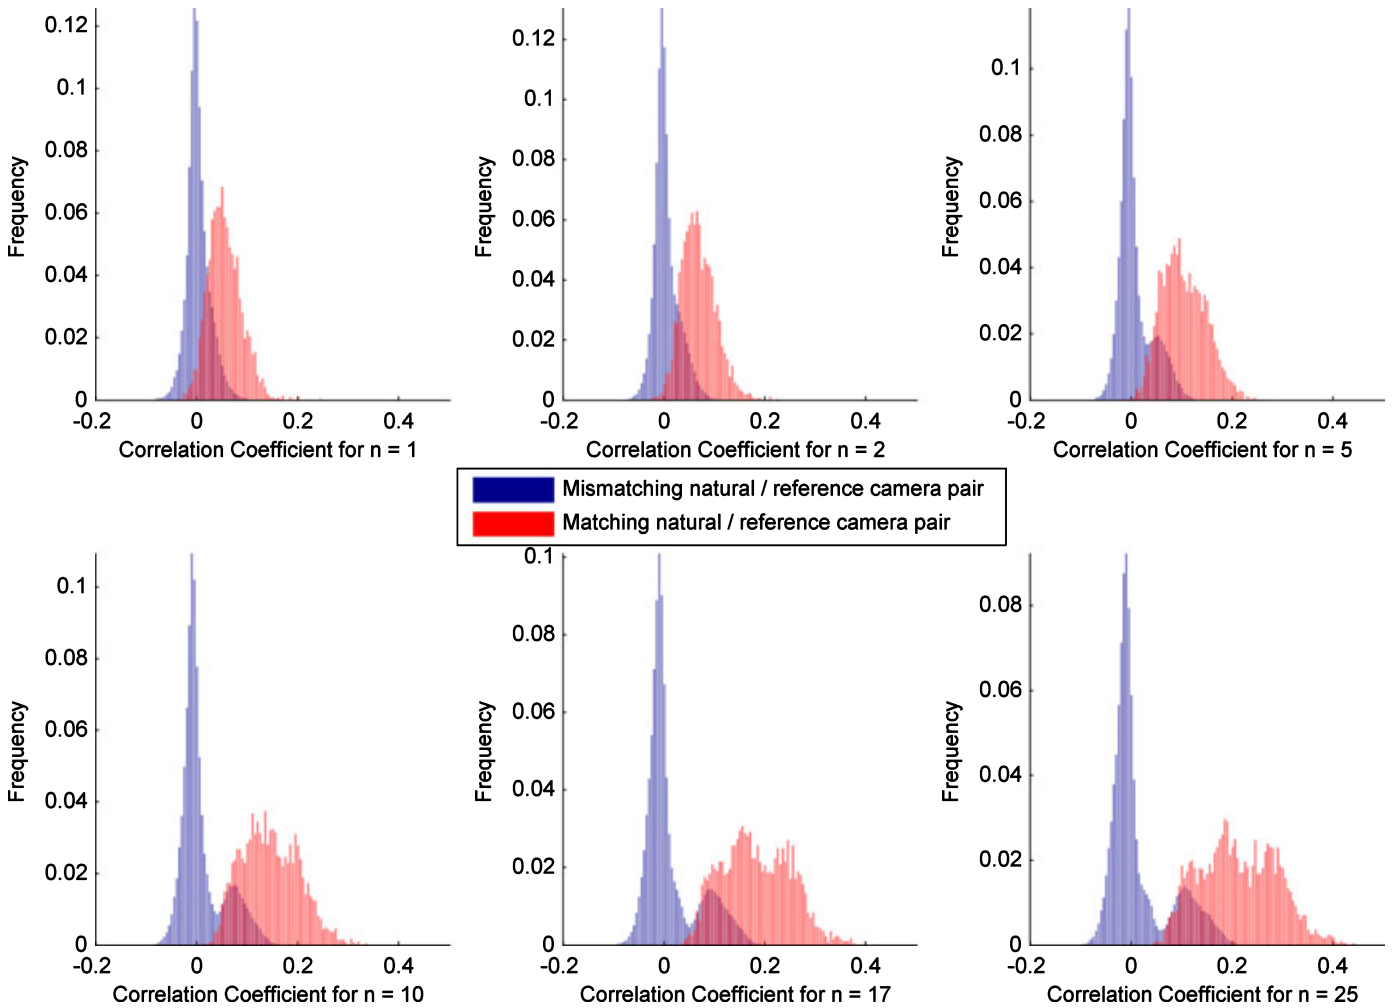


FIG. 9—The normalized histograms of correlation coefficients between matching and mismatching camera pairs for various numbers (n) of simultaneously identified images. The larger the number of simultaneously identified questioned images, the greater the separation between match and mismatch distributions and thus the better the performance of the identification scheme. In all six cases, some amount of overlap persists.

TABLE 4—Decision levels plus corresponding FRR for a FAR of 1‰, and EERs.

	FAR = 1‰		EER
	D	FRR	
n = 1	0.101	89.4%	0.1605
2	0.102	83.0%	0.1320
5	0.126	67.2%	0.1186
10	0.159	60.0%	0.1154
17	0.186	53.3%	0.1146
25	0.206	48.8%	0.1118

EER, equal error rate; FAR, false acceptance rate; FRR, false rejection rate.

images, implying that of at least 83% of all (selections of) questioned images, the correct source camera was identified from this group of cameras. Thus we have shown that PRNU is even unique among cameras of the exact same type.

We focused on the situation where both questioned photograph(s) and suspected source cameras are present and the camera is in working order. This enabled the acquisition of flat field images so that the reference patterns could easily be extracted. The

research can, however, still be carried out by extracting the reference pattern from a set of questioned photographs of which the source camera is known to be the suspected camera as shown in the study by Lukáš et al. (5). However, the only conclusion that can be drawn is that two sets of images originate from the same camera. The origin of the reference set is to be determined separately.

Using the same set of cameras, we obtained decision levels to solve the open-set problem, i.e., determining how likely it is for a given camera to be the source camera without directly comparing it with other cameras. When the reference pattern of a certain camera has a correlation coefficient with the PRNU extracted from the questioned image(s) that is higher than this decision level, this camera is concluded to be the source camera. This decision level is found by requiring a false acceptance ratio of 1‰, and the corresponding FRRs were found to be rather high. This means that in most cases, a single image will not suffice and thus that multiple images are required.

The results show that source camera identification based on PRNU is possible despite heavy JPEG compression which suppresses high-frequency signals such as PRNU. To remove DCT-block edge artifacts introduced by JPEG compression, we have proposed the simple method of averaging multiple pixels into

one macro element. This method proved significantly more effective than previously reported methods even though the effective resolution of the images, and thus the number of possible unique patterns to choose from decreases. In addition, the techniques proposed in this research can readily be applied to video footage.

In this research we showed that simple and computationally efficient techniques enable effective source camera identification, especially for the closed-set problem, i.e., PRNU extraction using a two-dimensional Gaussian filter, detection by calculating a correlation coefficient, JPEG edge artifact suppression by pixel averaging, and scene content suppression by thresholding enabled accurate identification in closed-set problems.

However, the open-set problem performance has to be improved significantly. If it can be determined what causes the bimodal behavior in the histograms in Fig. 9, the resulting overlap between matching and mismatching distributions may be addressed. If this overlap can be reduced, the false rejection and EERs will drop significantly.

Future Work

Future work will include both improving the performance of the identification scheme in this research and extending this research to different situations and applications. The greatest improvement is expected to be achieved by researching more elaborate methods to suppress scene content being extracted along with PRNU and stochastic noise. One of the possibilities is to apply edge detection schemes (11) to determine the affected pixels.

Extensions can be made by changing either the source of the images, the type of problem, or the type of images. For instance, digital cameras with different resolutions, cameras taking uncompressed images, or camcorders resulting in video footage can be used as carried out by Lukáš et al. (5) and Singh et al. (9), respectively. In addition, in this research we only considered images that came directly from the camera, i.e., images were not cropped, scaled, recompressed, etc. The effects of the various alterations were not researched, but it is expected that certain actions, for instance rotation or scaling, will alter the PRNU pattern significantly (12).

Finally, the PRNU might yield information of at least the brand or model of an image. During this research we have observed that for certain camera models, the calculated reference patterns of different cameras of the same model showed similarities. However,

this has to be studied in more detail before conclusions can be drawn.

References

1. Bayram S, Sencar H, Memon N, Avcibas I. Source camera identification based on CFA interpolation, Proceedings of International Conference on Image Processing, Vol. 3, Genova, Italy, September 11–14. Los Alamitos, CA: IEEE, 2006;69–72.
2. Alvarez P. Using extended file information (EXIF) file headers in digital evidence analysis. *Int J Digit Evid* 2004;2(3).
3. Geradts ZJM, Bijhold J, Kieft M, Kurosawa K, Kuroki K, Saitoh N. Digital camera identification. *J Forensic Identif* 2002;52(5):621–32.
4. Hytti HT. Characterization of digital image noise properties based on RAW data. In: Cui LC, Miyake Y, editors. Image quality and system performance III. Proceedings of SPIE-IS&T Electronic Imaging, San Jose, CA, January 17. Bellingham, WA: SPIE, 2006;1–12.
5. Lukáš J, Fridrich J, Goljan M. Determining digital image origin using sensor imperfections. In: Said A, Apostolopoulos JG, editors. Proceedings of SPIE Imaging and Video Communications and Processing, San Jose, CA, January 18. Bellingham, WA: SPIE, 2005;249–60.
6. Chen M, Fridrich J, Goljan M, Lukáš J. Source digital camcorder identification using sensor photo response non-uniformity. In: Delp EJ, Wong PW, editors. Proceedings of SPIE Security, Steganography, and Watermarking of Multimedia Contents IX; San Jose, CA, January 29. Bellingham, WA: SPIE, 2007;65051G.
7. Banham MR, Katsaggelos AK. Digital image restoration. *IEEE Signal Process Mag* 1997;14(2):24–41.
8. Kurosawa K, Kuroki K, Saitoh N. CCD fingerprint method—identification of a video camera from videotaped images. Proceedings on International Conference on Image Processing, Kobe, Japan, Oct 24–28. Los Alamitos, CA: IEEE, 1999;537–40.
9. Singh S, Kumar V, Verma HK. Reduction of blocking artifacts in JPEG compressed images. *Digit Signal Process* 2007;17(1):225–43.
10. Ramanath R, Snyder WE, Bilbro GL, Sander WA III. Demosaicking methods for Bayer color arrays. *J Electron Imaging* 2002;11(3):306–15.
11. Canny J. A computational approach to edge detection. *IEEE Trans Pattern Anal Mach Intell* 1986;8(6):679–98.
12. Lukáš J, Fridrich J, Goljan M. Detecting digital image forgeries using sensor pattern noise. In: Delp EJ, Wong PW, editors. Proceedings of SPIE Security, Steganography, and Watermarking of Multimedia Contents VIII, San Jose, CA, January 16. Bellingham, WA: SPIE, 2006;362–72.

Additional information and reprint requests:

Cor J. Veenman, Ph.D.

Department of Digital Technology and Biometry

Netherlands Forensic Institute

The Hague, Holland

Email: c.j.veenman@uva.nl