Determining Approximate Age of Digital Images Using Sensor Defects

Jessica Fridrich* and Miroslav Goljan

Department of Electrical and Computer Engineering, Binghamton University, State University of New York

ABSTRACT

The goal of temporal forensics is to establish temporal relationship among two or more pieces of evidence. In this paper, we focus on digital images and describe a method using which an analyst can estimate the acquisition time of an image given a set of other images from the same camera whose time ordering is known. This is achieved by first estimating the parameters of pixel defects, including their onsets, and then detecting their presence in the image under investigation. Both estimators are constructed using the maximum-likelihood principle. The accuracy and limitations of this approach are illustrated on experiments with three cameras. Forensic and law-enforcement analysts are expected to benefit from this technique in situations when the temporal data stored in the EXIF header is lost due to processing or editing images off-line or when the header cannot be trusted. Reliable methods for establishing temporal order between individual pieces of evidence can help reveal deception attempts of an adversary or a criminal. The causal relationship may also provide information about the whereabouts of the photographer.

Keywords: Temporal forensics, digital image forensics, imaging sensor defects, camera sensor aging, image origin.

1. INTRODUCTION

Establishing temporal order among individual pieces of evidence helps analysts deduce causal relationship of events in a court case. When the objects under investigation are represented in a digital form, we speak of temporal forensics of digital objects. This field is relatively new – the first works the authors are aware of include[14] and.[16] This paper deals with the problem of establishing an approximate acquisition time of digital images. It is assumed that an analyst is given a set of trusted digital images whose acquisition time is known – for example, it can be extracted from the EXIF header, the image content, or some other context. The analyst is simultaneously given one or more other images, whose acquisition time is unknown, and attempts to estimate it using the set of trusted images.

Recently, the following imperfections of digital imaging sensors were shown to be useful for forensic purposes: the photo-response non-uniformity,[4, 13] the dark current and defective pixels,[9] and sensor dust.[2] Not all such imperfections, however, exhibit changes over time and thus not all are suitable for determining the age of photographs. Reasonable candidates are the defective pixels and sensor dust because they accumulate over time. Both, however, have natural limitations. The sensor protective glass that collects the dust can be cleaned, which may mislead an analyst trying to extract temporal information by detecting traces of dust in images. Moreover, this method is applicable only to cameras with removable lens. Pixel defects seem more promising because they occur randomly in time and space on the sensor independently of each other[3, 10–12] and new defects appear suddenly and with a constant rate (It is a Poisson process.). Most importantly, once a defect occurs, it becomes a permanent part of the sensor and does not heal itself.

The main cause of new pixel defects is environmental stress, primarily due to impacting cosmic rays. In general, smaller pixels are more vulnerable to point defects than larger pixels. Sensors age both at high altitudes and at the sea level. They do age faster at high altitudes or during airplane trips where cosmic radiation is stronger. Consequently, the defect accumulation may not be linear in time.

* Jessica Fridrich: E-mail: fridrich@binghamton.edu, Telephone: +1 607 777 6177, Fax: +1 607 777 4464

The main technical problem of using pixel defects for temporal forensics is that they may not be easily detectable in individual images, depending on the image content, camera settings, exposure time, and processing. Also, the defects can be corrected for by upgrading the camera firmware or by post processing the images off-line. Despite these limitations, defective pixels may provide valuable forensic information in many practical scenarios as most photographers typically do not remove the defects from photographs. Consequently, the defects accumulate and thus provide basis for determining the acquisition time.

The main contribution of this paper is providing methodology for estimating pixel defects and their onsets from regular images of natural scenes and for estimating an approximate acquisition time for an unknown image. In Section 2, we describe the basic types of pixel defects and introduce a model for the output of a defective pixel. In Section 3, we derive a maximum-likelihood estimator for the model parameters and for the unknown acquisition time. Practical implementation issues are discussed in Section 4. The accuracy and limitations of the proposed method are the subject of the experimental Section 4. Finally, Section 6 summarizes the contribution and outlines future directions.

1.1 Notation

Everywhere in this chapter, boldface font will denote vectors (or matrices) of length specified in the text, e.g., $\mathbf{X}$ and $\mathbf{Y}$ can be vectors or matrices with $\mathbf{X}(i)$ denoting the ith component of $\mathbf{X}$. Unless mentioned otherwise, all operations among vectors or matrices, such as product, ratio, raising to a power, etc., are *elementwise*. For a logical statement P, we also make use of the Iverson bracket defined as $[P] = 1$ when P is true and $[P] = 0$ when P is false.

2. SENSOR DEFECTS

In the heart of every electronic device capable of taking digital pictures is an imaging sensor. There exist two types of sensors – CCD (Charge-Coupled Device) and CMOS (Complementary Metal-Oxide Semiconductor). Both sensors consist of a large number of photo detectors commonly called pixels. Pixels are made of silicon and capture light by converting photons into electrons using the photoelectric effect.[5, 7] The charge accumulated at every pixel is transferred out of the sensor, amplified, and then run through an AD converter that converts it to a digital signal. The digitized signal is further processed before the data is stored as an electronic file (JPEG, TIFF, etc.) on the camera storage device. The pixels are several microns across and have a rectangular shape. In theory, the amount of electrons (charge) outputted by a pixel should depend solely on the intensity of the incident light. In reality, however, there are many factors that introduce both systematic and random deviations. In this paper, we are interested in pixels that behave *systematically abnormally* to the point that the defect is detectable in a single image. Such abnormally-behaving pixels are extreme cases of photo-response non-uniformity and dark current. Both concepts are explained below.

2.1 Photo-response non-uniformity

The charge generated in a pixel depends on the physical dimensions of the pixel photosensitive area and on the homogeneity of silicon. The pixels' physical dimensions slightly vary due to imperfections in the manufacturing process. Also, the inhomogeneity naturally present in silicon contributes to variations in quantum efficiency among pixels (the ability to convert photons to electrons). The variations in quantum efficiency among pixels can be captured with a matrix $\mathbf{K} \in \mathbb{R}^{m \times n}$ of the same dimensions as the sensor. When an imaging sensor is illuminated with light intensity $\mathbf{I} \in \mathbb{R}^{m \times n}$, in the absence of other noise sources or imperfections, the sensor would register a noisy scene $\mathbf{I} + \mathbf{IK}$ instead. (We remind that the product $\mathbf{IK}$ is an element-wise product of matrices) The term $\mathbf{IK}$ is usually referred to as the photo-response non-uniformity or PRNU. A large value of $\mathbf{K}$ leads to a point defect called "pixel with abnormal sensitivity."

2.2 Dark current

Even when a pixel is not exposed to light during picture taking, it contains a small number of free electrons due to thermal effects. Their number increases with temperature and exposure.[5-7] It is also affected by ISO setting. In the absence of all other defects, the pixel's output is $\mathbf{I} + \tau \mathbf{D} + \mathbf{c}$, where $\tau \mathbf{D}$ is called the dark current and $\mathbf{c}$ the offset. Here, $\tau \geq 0$ is a multiplicative factor whose value is determined by the temperature, exposure, and ISO

Figure 1. Top: A stuck pixel in an image and its close up. The pixel happens to be red because it has a red color filter in front of it. Bottom: An example of a dark frame. Notice that there appears to be a degree of "hotness" among the pixels captured by the matrix $\mathbf{D}$.

(higher ISO leads to a larger value of τ), and $\mathbf{D}, \mathbf{c} \in \mathbb{R}^{m \times n}$ are matrices. In images taken with a short exposure (e.g., 1/60th of a second or shorter), the dark current is usually very weak. However, it may start dominating the sensor output for dark scenes ($\mathbf{I} \approx 0$) when τ becomes large (large ISO and/or temperature and/or long exposure).

An extremely high value of $\mathbf{D}$ produces the most common defect called a hot pixel. A high value of the offset $\mathbf{c}$ leads to another defect type commonly recognized as a stuck pixel. Both defects were proposed for forensic tasks in 1999 by Kurosawa,[9] who demonstrated that as long as a video clip contained some dark frames, hot/stuck pixels can be used to uniquely identify digital video cameras. In this paper, we propose to use these defects for determining an approximate image acquisition time.

2.3 Pixel output model

Considering the impact of all the defects discussed so far, we arrive at the following model for the raw output of a sensor:

$$\mathbf{Y} = \mathbf{I} + \mathbf{IK} + \tau \mathbf{D} + \mathbf{c} + \mathbf{\Theta}. \tag{1}$$

We remind that τ is a scalar multiplicative factor whose value is determined by exposure, temperature, and ISO settings. The matrix $\mathbf{c}$ is the matrix of offsets and the dark current factor $\mathbf{D}$ is a noise-like signal due to leakage of electrons into pixels' electron wells (Fig. 1 right). Finally, the reader recognizes $\mathbf{K}$ as the PRNU factor. The modeling noise $\mathbf{\Theta}$ is a collection of all other noise sources, which are mostly random in nature and thus difficult to use for forensic purposes (readout noise, shot noise, also known as the photonic noise, quantization noise, etc.).

For spike pixel defects, such as hot or stuck pixels, at least one of the three parameters, $\mathbf{K}$, $\mathbf{D}$, and $\mathbf{c}$, becomes large. Even though the parameters could be accurately estimated by taking special test pictures, such as pictures of dark scenes with long exposure, in forensic setting the camera may not be available and the challenge becomes to estimate these three matrices from the limited set of images available to the analyst. We also need to design an estimator for the onset of a defect and an estimator for the acquisition time of an unknown image by detecting in it the presence of defects. This is the subject of the next section.

3. DERIVING THE ESTIMATORS

Even though sensor defects can be easily estimated in a laboratory environment by taking test images under controlled conditions, a forensic analyst must work with a given set of images taken with camera settings that may be quite unfavorable for estimating certain defects. For example, the dark current is difficult to estimate reliably from images of bright scenes taken with a short exposure time and low ISO.

To improve the signal to noise ratio between the signal of interest (defect) and the observable $\mathbf{Y}$, we will work with the noise residual $\mathbf{W} = \mathbf{Y} - F(\mathbf{Y})$, obtained using a denoising filter F. Since hot pixels (and stuck pixels with a large offset) are spiky in nature, denoising filters of the type[15] that extract additive white Gaussian noise are likely a poor choice for estimating point defects. Instead, non-linear filters, such as the median filter, are more likely to extract the spike correctly. In this paper we use the median filter.

The model of the residual becomes:

$$\begin{aligned}
\mathbf{W} &= \mathbf{Y} - F(\mathbf{Y}) \\
&= \mathbf{IK} + \tau\mathbf{D} + \mathbf{c} + \mathbf{I} - F(\mathbf{Y}) + \mathbf{\Theta} \\
&= \mathbf{IK} + \tau\mathbf{D} + \mathbf{c} + \mathbf{\Xi},
\end{aligned} \tag{2}$$

where $\mathbf{\Xi}$ stands for the sum of the modeling noise and the remnant of the content $\mathbf{I} - F(\mathbf{Y})$ due to the inability of the denoising filter to separate content from noise. The term $\mathbf{I} - F(\mathbf{Y})$ is especially large in textured regions and around edges.

Let $\mathbf{Y}_k$, $k = 1, \ldots, d$, be d images of regular scenes with their noise residuals $\mathbf{W}_k$ taken at known times $t_1, \ldots, t_d$:

$$\mathbf{W}_k(i) = \mathbf{I}_k(i)\mathbf{K}(i) + \tau_k\mathbf{D}(i) + \mathbf{c}(i) + \mathbf{\Xi}_k(i), \tag{3}$$

where k and i are the image and pixel indices, respectively. We have experimented with two models for the noise residual − the Gaussian model and the Laplacian model. Since the experiments with the much more complex Laplacian noise model lead to very similar results, we next describe the estimators for the Gaussian model only. Thus, we model $\mathbf{\Xi}_k(i)$, $k = 1, \ldots, d$, as an i.i.d. Gaussian sequence with zero mean and variance $\sigma^2(i)$. The known quantities in the model are $\mathbf{I}_k \approx F(\mathbf{Y}_k)$ and τ_k (the exposure time). Note that $\mathbf{W}_k$ are observables.

From now on, all derivations will be carried out for a fixed pixel i. This will allow us to drop the pixel index and make the expressions more compact. The unknown vector parameter $\boldsymbol{\theta} = (\mathbf{K}, \mathbf{D}, \mathbf{c}, \sigma)$ (for a fixed pixel, $\boldsymbol{\theta} \in \mathbb{R}^4$) can be estimated using the maximum likelihood principle:

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} L(\mathbf{W}_1, \ldots, \mathbf{W}_d | \boldsymbol{\theta}), \tag{4}$$

where L is the likelihood function

$$L(\mathbf{W}_1, \ldots, \mathbf{W}_d | \boldsymbol{\theta}) = (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{k=1}^{d} (\mathbf{W}_k - \mathbf{I}_k\mathbf{K} - \tau_k\mathbf{D} - \mathbf{c})^2\right). \tag{5}$$

Because the modeling noise $\mathbf{\Xi}_k$ is Gaussian, the ML estimator becomes the least squares estimator.[8] Additionally, the model linearity allows us to find the maximum in (9) in the following well-known form:

$$(\hat{\mathbf{K}}, \hat{\mathbf{D}}, \hat{\mathbf{c}}) = (\mathbf{HH}')^{-1}\mathbf{H}'\mathbf{W}, \tag{6}$$

$$\hat{\sigma}^2 = \frac{1}{d} \sum_{k=1}^{d} \left(\mathbf{W}_k - \mathbf{I}_k\hat{\mathbf{K}} - \tau_k\hat{\mathbf{D}} - \hat{\mathbf{c}}\right)^2, \tag{7}$$

where

$$\mathbf{H} = \begin{pmatrix} \mathbf{I}_1(i) & \tau_1 & 1 \\ \mathbf{I}_2(i) & \tau_2 & 1 \\ \cdots & \cdots & \cdots \\ \mathbf{I}_d(i) & \tau_d & 1 \end{pmatrix} \tag{8}$$

is a matrix of known quantities and $\mathbf{W} = (\mathbf{W}_1, \ldots, \mathbf{W}_d)'$ the vector of observations (noise residuals).

3.1 Determining defect onset

We now extend the estimator derived above to the case when we have observations (pixel intensities) across some time span during which the pixel becomes defective. Our goal is to estimate $\boldsymbol{\theta}$ before and after the onset, $\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}$, and the onset time j. The estimator derived in the previous section easily extends to this case

$$(\hat{\boldsymbol{\theta}}^{(0)}, \hat{\boldsymbol{\theta}}^{(1)}, \hat{j}) = \arg \max_{(\boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}, j)} L_j(\mathbf{W}_1, \ldots, \mathbf{W}_d | \boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}), \tag{9}$$

where

$$L_j(\mathbf{W}_1, \ldots, \mathbf{W}_d | \boldsymbol{\theta}^{(0)}, \boldsymbol{\theta}^{(1)}) = L(\mathbf{W}_1, \ldots, \mathbf{W}_j | \boldsymbol{\theta}^{(0)}) \times L(\mathbf{W}_{j+1}, \ldots, \mathbf{W}_d | \boldsymbol{\theta}^{(1)}) \tag{10}$$

is the likelihood function written in terms of (5). Because of the form of (10), the maximization in (9) factorizes:

$$(\hat{\boldsymbol{\theta}}^{(0)}, \hat{\boldsymbol{\theta}}^{(1)}, \hat{j}) = \arg \max_j \left\{ \arg \max_{\boldsymbol{\theta}^{(0)}} L(\mathbf{W}_1, \ldots, \mathbf{W}_j | \boldsymbol{\theta}^{(0)}) \times \arg \max_{\boldsymbol{\theta}^{(1)}} L(\mathbf{W}_{j+1}, \ldots, \mathbf{W}_d | \boldsymbol{\theta}^{(1)}) \right\}, \tag{11}$$

which converts the problem of onset estimation to the problem of estimating the pixel defect parameters addressed in the previous section.

3.2 Determining approximate acquisition time

We are now ready to develop the algorithm for placing a given image $\mathbf{I}$ under investigation among other d images, $\mathbf{I}_1, \ldots, \mathbf{I}_d$, whose time of acquisition is *known*, monotone increasing, and whose pixel defects are known, including the onset time and the parameters before and after the onset. This problem is again addressed using the ML approach. This time, only the time index j of $\mathbf{I}$ is the unknown as the parameters of all defective pixels are already known. Denoting the set of all defective pixels $\mathcal{D}$, the estimator becomes:

$$\hat{j} = \arg \max_j \prod_{i \in \mathcal{D}} L(\mathbf{W}_\mathbf{I}(i) | \boldsymbol{\theta}^{([j > j(i)])}(i)) \tag{12}$$

written in terms of (5). The superscript of $\boldsymbol{\theta}$ is the Iverson bracket.

3.3 Confidence intervals

In forensic setting, the analyst will likely be interested in statements of the type "the probability that image $\mathbf{I}$ was taken in time interval $[t, t']$ is at least p." The approach outlined above allows the analyst to quantify the results in this way because the conditional probabilities $\Pr(\mathbf{W}_\mathbf{I} | j) = \prod_{i \in \mathcal{D}} L(\mathbf{W}_\mathbf{I}(i) | \boldsymbol{\theta}^{[j > j(i)]}(i))$ are known for each j. From the Bayes formula

$$\Pr(j | \mathbf{W}_\mathbf{I}) = \frac{\Pr(\mathbf{W}_\mathbf{I} | j) \Pr(j)}{\Pr(\mathbf{W}_\mathbf{I})}. \tag{13}$$

Thus, the probability that $\mathbf{I}$ was taken in time interval $[t, t']$ is

$$\Pr(j \in [t, t'] | \mathbf{W}_\mathbf{I}) = \frac{\sum_{k \in [t, t']} \Pr(\mathbf{W}_\mathbf{I} | k) \Pr(k)}{\sum_k \Pr(\mathbf{W}_\mathbf{I} | k) \Pr(k)}. \tag{14}$$

Depending on the situation at hand, the prior probabilities $\Pr(k)$ may be known from other forensic evidence or may be completely unknown. In the absence of any information about the priors, one could resort to the principle of maximum uncertainty and assume the least informative prior distribution – that the owner of the camera was taking images at a uniform rate, leading to $\Pr(k) = 1/(t_k - t_{k-1})$ for all k, where t_k is the time when the kth trusted image was taken.

4. PRACTICAL IMPLEMENTATION ISSUES

In this section, we discuss some implementation issues when applying the above framework in practice. The next section reports the results of specific experiments.
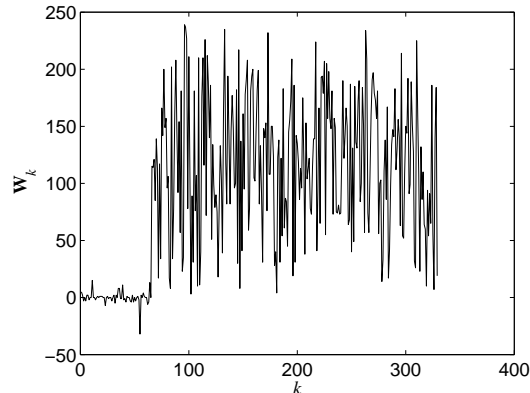
Figure 2. Example of a noise residual for a defective pixel with onset at $k \approx 64$. Note that, in agreement with the model (3), the residual of a correctly-functioning pixel (before $k = 64$) is much smaller than after the defect onset.

4.1 Selecting candidate defective pixels

The first issue concerns the estimation of the model parameters $\hat{\boldsymbol{\theta}}$. Running the estimation procedure of Section 3 for all pixels would not be feasible due to excessive computational demands. Thus, prior to estimation, we select a much smaller set of "candidate defects" that contains with high probability all defective pixels.

Fig. 2 shows an example of the noise residual $\mathbf{W}_k(i)$ for a defective pixel with defect onset at around $k = 64$. One can clearly see that the noise residual of defective pixels fluctuates much more than for non-defective pixels. Thus, we identify the set of candidate defects simply by thresholding the sample variance of $\mathbf{W}_k(i)$, $k = 1, \ldots, d$. Note that any pixel that appeared to be located at the image 1-pixel border was not considered since the median filter has a small support there.

4.2 Verifying the model

Fig. 3 shows the noise residual $\mathbf{W}_k(i)$ for $k = 1, \ldots, d$ and a fixed i versus the model (3) obtained using the method described in Section 3. The left figure shows the case of a defective pixel that follows the model rather well, while the right figure is for a defect whose behavior is compatible with the model to a lesser degree. We hypothesize that the misfit is probably caused by an unknown parameter in the model (3), which is the ISO setting. Some cameras set the ISO automatically, based on the scene, but do not store this information in the EXIF header. Since we determine the factor τ from the exposure only, it does not accommodate the potentially varying ISO. Thus, hot pixels with a large of $\mathbf{D}$ (right figure) are expected to be modeled less accurately than pixels with abnormal sensitivity that have a large $\mathbf{K}$ (left figure) and small $\mathbf{D}$. A potential remedy is to use a simpler model for the defective pixel, the so-called Kc model, in which $\mathbf{D}$ is set to 0, and the dark current is included into the modeling noise. The parameter estimation simplifies as only two parameters now need to be estimated for each defect − $\mathbf{K}$ and $\mathbf{c}$.

4.3 Other issues

The likelihood function maximized in (12) returns an index $\hat{j}$ − the approximate acquisition time for the tested image. However, the likelihood is constant between two consecutive defect onsets and thus it is not clear which time should be returned as the estimated age of the tested image (follow Fig. 4). In the absence of any other information, it is tempting to return the average time of all images $\mathbf{I}_j$ with the same maximal value of the likelihood. However, this may lead to a bias in the estimate because it is far more likely that a defect is not detected than detecting a false defect. Therefore, the estimated time tends to be smaller than the true date.

5. EXPERIMENTS

5.1 Tests on Canon PS SD400 #1

The first set of images we experimented with came from a Canon PS SD 400 #1 with native resolution of 2592×1944 pixels. As all cameras included in this experimental section, this camera has a 1/2.5" CCD sensor.
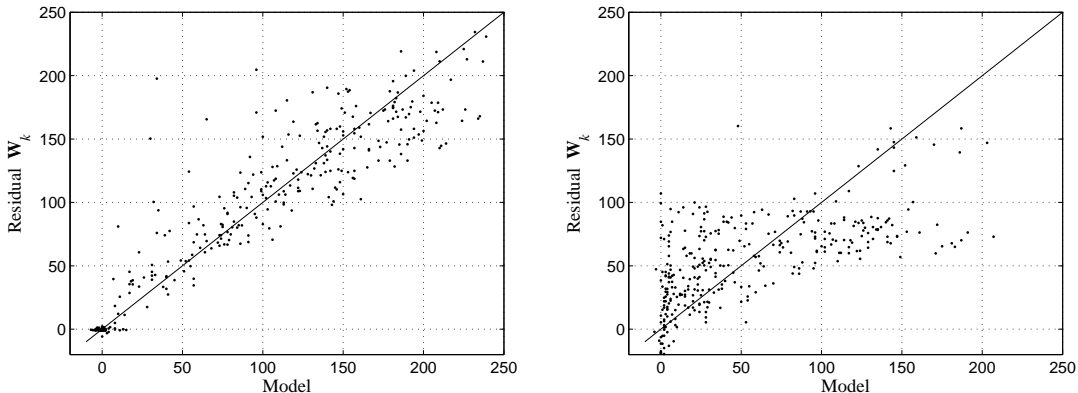
Figure 3. The noise residual $\mathbf{W}_k(i)$ versus the model (3) for two defective pixels i. Left: Estimated defect parameters: $\mathbf{K} = -0.895, \mathbf{D} = 63.9, \mathbf{c} = 219.7$. Right: Estimated defect parameters: $\mathbf{K} = -0.333, \mathbf{D} = 953.0, \mathbf{c} = 61.34$.
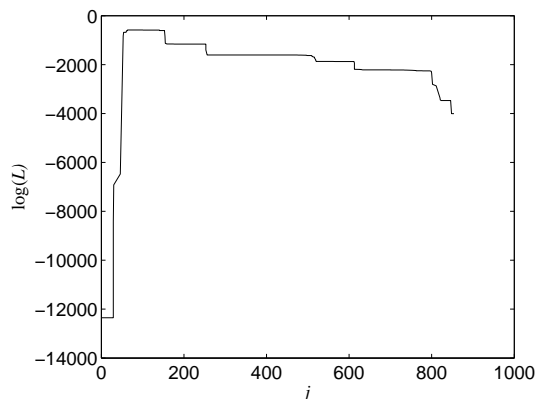

Figure 4. Example of the log-likelihood function with maximum average at $\hat{j} = 85$ while the true date is at $j = 141$.

The available set of images was split to $d = 329$ images of a known acquisition time, from which the parameters of defective pixels were obtained, and a set of 155 test images that would be tested by the proposed method. Before the experiment, all d images were converted to grayscale and 200 candidate defective pixels were located by the largest sample variance of the noise residual $\mathbf{W}_k(i)$, $k = 1, \ldots, d$, obtained using a 3×3 median filter. An initial visual inspection of the plots of $\mathbf{W}_k(i)$ lead to a set of defective pixels $\mathcal{D}_0$ containing 110 pixels. The pixels formed 12 clusters on the camera sensor while two defective pixels appeared isolated (see Fig. 5).

We tested the proposed framework with several different models for the noise residual: the KDc model (3) containing three unknown parameters, $\mathbf{K}, \mathbf{D}, \mathbf{c}$, and the simpler Kc model with $\mathbf{D} \equiv 0$. This simplified version of the noise residual was included because the factor τ in the model depends on an unknown ISO setting, as discussed in Section 4.2. Thus, it might be beneficial to simply include this uncertainty into the noise component instead of attempting to model it. We also tested the Laplacian noise model for $\mathbf{\Xi}$. The results of all experiments were evaluated using the Mean Absolute Error (MAE) expressed in the number of days (see Table 1).

Overall, the Laplacian noise model did not lead to more accurate estimates when compared to the Gaussian model. And since it is significantly more complicated for implementation, the Gaussian model is the preferred choice. The best average error of about 77 days was obtained using the Gaussian KDc model. This result is quite reasonable considering the fact that the 329 trusted images spanned the total of 802 days. The accuracy with which one can expect to estimate the acquisition time obviously depends on the number of new defects during the time span at hand. Thus, in addition to MAE we report the relative time estimation error as the ratio between the MAE and the average time between onsets of new defects. Both MAE and the relative error are shown in Table 1. We remind that there were total 19 new defects for this camera. In five cases, two defects appeared at the same time (both were counted as separate events).
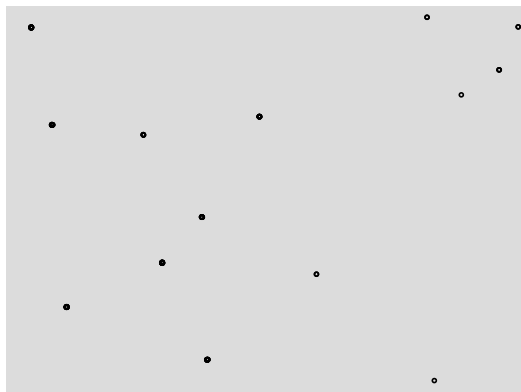
Figure 5. Defective pixel locations for CanonSD400 #1 detected by visual inspection of noise residuals with large variance (enhanced by small circles).

	Kc model Gauss	KDc model Gauss	Kc model Lapl.	KDc model Lapl.
MAE of dating / relative error	78.30 / 1.74	76.97 / 1.71	85.99 / 1.91	81.06 / 1.80

Table 1. MAE and the relative error for 155 images from Canon PS SD400 #1 for four different estimators differing by the number of estimated parameters and the modeling noise distribution.

Fig. 6 shows the true date versus the estimated date in days for all 155 tested images obtained from the Gaussian KDc model. The crosses on the diagonal correspond to the set of 329 trusted images with known dates. They show the temporal distribution of the trusted images. Note that no images appear between time 400 and 500, limiting thus our ability to date images within this time interval.

5.2 Tests on Canon PS SD400 #2

To obtain a better insight about the reliability of the proposed method, we repeated the same experiments as in the previous section for another camera of the same model – Canon PS SD400 #2. We worked with 300 images spanning 649 days for establishing pixel models and with 200 test images. This camera had substantially fewer new defects that were clustered in five groups (see Fig. 7 left), which yielded a larger MAE than for the first camera. The results appear in Table 2. The MAE was approximately 111 days. Even though this error may seem rather large, it may be sufficient to reliably answer the question whether a given image was taken before or after a certain event.

The detected onset of each of the 58 defective pixels plotted in Fig. 7 right is in agreement with five physical defects. (shown in Table 2).

5.3 Tests on Canon PS S2

In the last experiment with the third camera, 237 images obtained in a time span of more than 1000 days were divided into a set of 150 trusted images and 87 test images. Total of 42 defective pixels were identified from the trusted set. Grouping these pixels by their location revealed one very large defect (consisting of 23 pixels) that appeared between May 21 and May 29, 2006. There were three other smaller defects and a few individual defective pixels that lead to the total of 6 defect onsets during the span of 1411 days. The onsets are separated by dashed lines in Fig. 8. The low number of new defects limits the accuracy of date estimation for images from this camera.

The MAE and the relative error appears in Table 3.

	Kc model Gauss	KDc model Gauss	Kc model Lapl.	KDc model Lapl.
MAE of dating/ relative error	113.3 / 0.87	113.3/ 0.87	111.6 / 0.86	111.2 / 0.86

Table 2. MAE and the relative error for 200 images from Canon PS SD400 #2 for four different estimators.
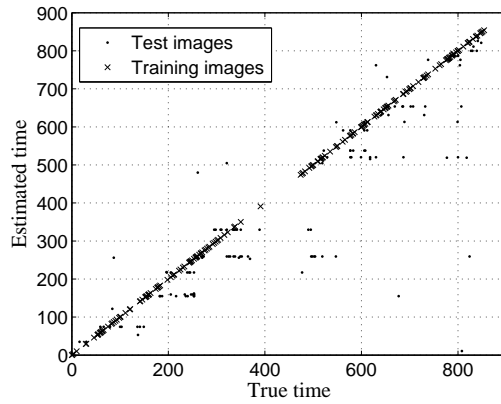
Figure 6. True acquisition date (in days) versus date estimated using (12) based on detecting pixel defects.
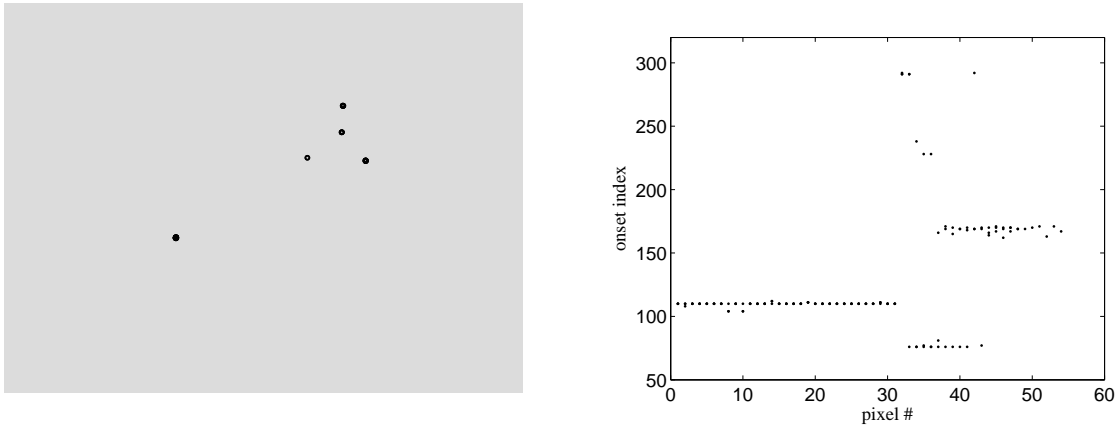

Figure 7. Left: Defective pixels for camera CanonSD400 #2 form five clusters. Right: The occurrence of defects in time.
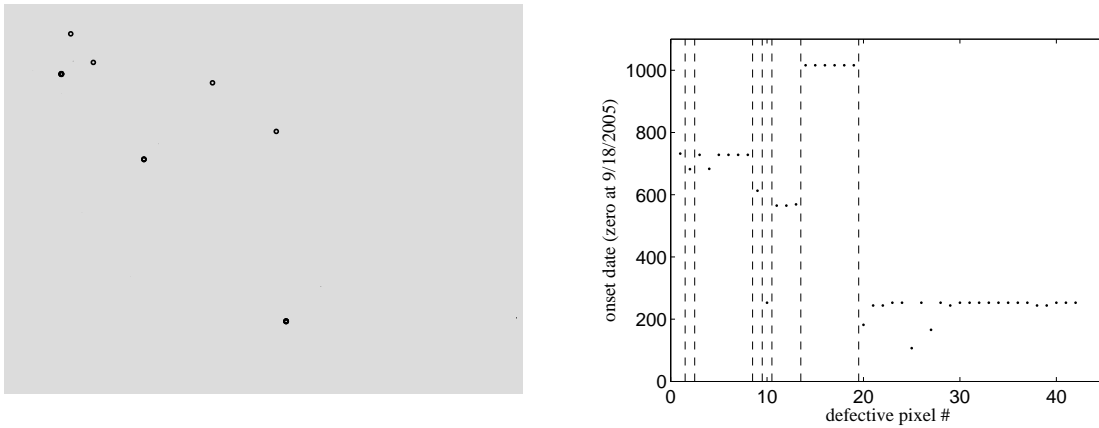

Figure 8. Canon PS S2: Left: Defective pixels location. Right: Defect occurrence in time grouped by pixel's location.

	Kc model Gauss	KDc model Gauss	Kc model Lapl.	KDc model Lapl.
MAE of dating / relative error	150.25 / 0.64	140.49 / 0.60	138.63 / 0.59	141.27 / 0.60

Table 3. MAE and the relative error for 87 images from Canon PS SD400 #2 for four different defect estimators.

6. CONCLUSIONS

Causal forensic of digital objects is a relatively new field. The first works the authors are aware of are[14] and.[16] The former proposed the mutual information between sensor fingerprint estimates as the quantity that can establish causal relation among digital objects via the information inequality.[1] The authors applied this framework to digital images using estimates of sensor fingerprints[13] as the forensic quantity. Assuming the individual elements of a sensor fingerprint are realizations of an iid Gaussian signal, the mutual information is the normalized correlation between fingerprints. Even though the authors reported that normalized correlation can be used to establish temporal information among digital objects, we were unable to repeat the experiments reported in.[14] In fact, it appears that the experiments were negatively affected using an unintentional programming error.

We determined experimentally that the only temporal changes in the fingerprint are due to onset of new defective pixels. This motivated our effort presented in this paper. The methodology is applicable whenever the analyst has a set of trusted images with known acquisition times and wishes to determine an approximate acquisition time for an unknown image. It works by first estimating the model parameters of all defective pixels, including their onset using a maximum-likelihood estimator. Then, the unknown time of acquisition is established by detecting the presence of defective pixels in the image under investigation. Again, a maximum-likelihood approach is used. The accuracy of the proposed method depends on how often new defects appear on the sensor. For one camera, we were able to determine the age with an average accuracy of about 80 days while new defects appeared on average every 45 days. For other cameras with a lower frequency of occurrence of new defects, the method was less accurate (average accuracy of about 111 and 140 days). Even this lower accuracy may be useful in court in situations when only a rough (binary) time granularity is needed, such as when the question is whether an image was taken before or after a certain event.

7. ACKNOWLEDGMENTS

REFERENCES

1. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. New York: John Wiley & Sons, Inc., 2006.
2. A. E. Dirik, H. T. Sencar, and N. Memon. Digital single lens reflex camera identification from traces of sensor dust. *IEEE Transactions on Information Forensics and Security*, 3(3):539–552, September 2008.
3. J. Dudas, L. M. Wu, C. Jung, G. H. Chapman, Z. Koren, and I. Koren. In R. A. Martin, J. M. DiCarlo, and N. Sampat, editors, *Identification of in-field defect development in digital image sensors*, volume 6502, page 65020Y. SPIE, 2007.
4. J. Fridrich. Digital image forensic using sensor noise. *IEEE Signal Processing Magazine*, 26(2):26–37, 2009.
5. G. E. Healey and R. Kondepudy. Radiometric CCD camera calibration and noise estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3):267–276, March 1994.
6. G. C. Holst. *CCD Arrays, Cameras, and Displays*. JCD Publishing & SPIE Press, USA, 2nd edition, 1998.
7. J.R. Janesick. *Scientific Charge-Coupled Devices*, volume PM83. SPIE Press Monograph, 2001.
8. S. M. Kay. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, volume II. Upper Saddle River, NJ: Prentice Hall, 1998.
9. K. Kurosawa, K. Kuroki, and N. Saitoh. CCD fingerprint method – identification of a video camera from videotaped images. In *Proc. IEEE International Conference on Image Processing (ICIP)*, pages 537–540, October 1999.
10. J. Leung, G. H. Chapman, I. Koren, and Z. Koren. Automatic detection of in-field defect growth in image sensors. In *DFT '08: Proceedings of the 2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems*, pages 305–313, Washington, DC, 2008. IEEE Computer Society.

11. J. Leung, G. H. Chapman, I. Koren, and Z. Koren. Characterization of gain enhanced in-field defects in digital imagers. *Defect and Fault-Tolerance in VLSI Systems, IEEE International Symposium on*, 0:155–163, 2009.

12. J. Leung, J. Dudas, G. H. Chapman, I. Koren, and Z. Koren. Quantitative analysis of in-field defects in image sensor arrays. In *Defect and Fault-Tolerance in VLSI Systems, 2007. DFT '07. 22nd IEEE International Symposium on*, pages 526–534, September 2007.

13. J. Lukáš, J. Fridrich, and M. Goljan. Digital camera identification from sensor pattern noise. *IEEE Transactions on Information Forensics and Security*, 1(2):205–214, June 2006.

14. J. Mao, O. Bulan, G. Sharma, and S. Datta. Device temporal forensics: An information theoretic approach. In *Proc. IEEE International Conference on Image Processing (ICIP)*, volume 1, November 2009.

15. M. K. Mihcak, I. Kozintsev, K. Ramchandran, and P. Moulin. Low-complexity image denoising based on statistical modeling of wavelet coefficients. *IEEE Signal Processing Letters*, 6(12):300–303, December 1999.

16. A. D. Rosa, F. Uccheddu, A. Costanzo, A. Piva, and M. Barni. Exploring image dependencies: a new challange in image forensics. In N. Memon, J. Dittmann, and A. Alattar, editors, *Proceeedings of SPIE Electronic media Forensics and Security XII*, volume 7541, pages 0X–1–0X–12, January 2010.