Noise measurement for raw-data of digital imaging sensors by automatic segmentation of non-uniform targets

Alessandro Foi, Sakari Alenius, Vladimir Katkovnik, and Karen Egiazarian, *Senior member*, *IEEE*

***Abstract*— In this paper we present a new method for measuring the temporal noise in the raw-data of digital imaging sensors (e.g., CMOS, CCD). The method is specially designed to estimate the variance function which describes the signal-dependent noise found in raw-data. It gives the standard-deviation of the noise as a function of the expectation of the pixel raw-data output value.**

In contrast with established methods (such as the ISO 15739), our method does not require the use of a specific target or a particular calibration. This is possible due to an automatic segmentation embedded in the data analysis.

We show experimental results for the raw-data from two different CMOS sensors of commercial cameraphones.

***Index Terms*— CMOS, CCD, digital imaging sensors, noise measurement, photon-limited, raw-data, noise modeling.**

I. INTRODUCTION

Noise characteristics are one of the critical elements in the evaluation and choice of digital imaging sensors. However, noise characteristics are not only used in the comparison between different sensors (e.g. to determine which sensor is better), but are also an important parameter in the digital image processing chain (e.g., digital gain, denoising, interpolation, color processing, compression) in which the raw-data goes before the final color image is obtained. Here, by raw-data we mean the unprocessed digital output of the sensor; roughly speaking, this output is obtained after the photon-to-electron, electron-to-voltage, and voltage-to-digit conversions which are performed within the monolithic sensor. By applying modern adaptive filtering algorithms on the raw-data produced by a cheap and noisy sensor, it is possible to obtain an image of quality comparable to that expected from a more expensive sensor. Nevertheless, to fully exploit the potential of these modern algorithms it is necessary that an accurate noise model is used when processing the raw-data. Noise modeling for digital sensors and their raw-data has been a vivid research area in the recent years (e.g., [1], [2], [3], [9], [10]).

A. Foi, V. Katkovnik, and K. Egiazarian are with the Institute of Signal Processing, Tampere University of Technology, P.O. Box 553, 33101 Tampere, Finland. E-mail: firstname.lastname@tut.fi

S. Alenius is with the Multimedia Technologies Laboratory, Nokia Research Center, Tampere, Finland, E-mail: firstname.lastname@nokia.com

This paper includes material from the authors' pending U.S. patent "Apparatus, method, mobile station and computer program product for noise estimation, modeling and filtering of a digital image" (Application No. 11/519,722, September 2006, and No. 11/426,128, June 2006).

The authors would like to thank Adrian Burian (Nokia Research Center) for the development of the cameraphone software and Heli Hytti (Tampere University of Technology) for providing the facilities and equipment for the experiments.

This work was supported by the Finnish Funding Agency for Technology and Innovation (Tekes), AVIPA project, and by the Academy of Finland, project No. 213462 (Finnish Centre of Excellence program 2006 - 2011).

The above issues become more and more relevant with the introduction of digital imaging sensors (CMOS, CCD) having a dramatically increased resolution. This is mainly achieved by an increase of the pixel density. With the size of each pixel becoming smaller and smaller, the sensor output signal's susceptibility to photon noise becomes greater and greater.

Because of the inherent "photon-counting" process, the noise in the raw-data is predominantly signal-dependent, with brighter parts of the image having a larger noise standard-deviation. Hence, when measuring this standard-deviation from the recorded data it is necessary to either have uniform data, or take into account for the data non-uniformity.

In practice it is very difficult (sometimes even impossible) to guarantee uniform recordings. Even with perfectly uniform targets and illumination, the lens of the camera introduce a systematic "vignetting" effect where the center of the image is much brighter (and hence noisier) than the peripheral area. As a result, the measurements taken under these assumptions are inherently biased. Unfortunately, the current international standards ISO 15739 and ISO 14524 [8], [7] and proposed approaches (e.g., [1], [2], [6]) for measuring the noise assume known uniform targets and thus provide results that are of a global nature (i.e., "average" values which are meant to be valid for the whole sensor). Such global noise estimates are rough when applied to an individual pixel of the sensor. In particular they are inadequate for high-quality image processing steps which can take place further in the imaging chain, e.g. denoising or deblurring. These techniques commonly aim at solving ill-posed inverse problems and therefore require accurate pointwise (i.e., pixelwise) knowledge of the noise characteristics, in order to properly restore the image details.

In this paper we present a new approach for measuring the temporal noise in the raw-data of digital imaging sensors (e.g., CMOS, CCD). The method is specially designed to estimate the curve which describes the standard-deviation of the noise as a function of the expectation of the pixel raw-data output. Based on an automatic segmentation of the recorded images, we separate samples with different expected output and calculate their standard-deviations. Thus, while other techniques require a uniform target, in our approach we benefit from the target non-uniformity by simultaneously estimating the variance function over a large range of output values.

Because of the automatic segmentation embedded in the procedure, our approach has a number of advantages over current noise measurement standards:

- Our approach does not require a specific target (i.e. a test chart). In fact, any fixed target or scene can be used for performing the noise measurements.
- The target and illumination do not need to be known in advance, i.e. no calibration is required before performing the measurements.
- The method is not influenced by the focusing or by the presence of the lens of the camera.
- Illumination does not need to be uniform in space. It is sufficient that the illumination is constant in time.
- The method is applicable without modifications for gray as well as for color sensors (Bayer pattern).
- Fixed-pattern noise does not influence the measured temporal noise.
- With a single experiment we measure a standard-deviation vs. expectation curve, whereas previous techniques aimed at estimating only a single standard-deviation value for a given uniform intensity.

Overall the method is simple, easy to implement, and allows for accurate measurement with much simplified laboratory equipment. It is thus a cost-effective alternative to other noise measurement techniques.

In this paper we concentrate on the measurement methodology for a given sensor with fixed parameters. We present experimental results obtained for two different CMOS sensors used in Nokia cameraphones and show how a parametric model can be fit to the measured curve. We refer the reader to our publication [5], where a parametric Poissonian-Gaussian noise model for the raw-data is analytically derived and fitted to the measurements (from single or multiple images), with the model parameters related to the sensors characteristics (e.g., analog gain, pedestal, quantum efficiency). Application of the proposed model to the imaging chain has been shown, e.g., in [4].

II. Observation model

We consider an observation model of the form

$$z(x) = y(x) + \sigma(y(x))\,\xi(x), \qquad x \in X, \tag{1}$$

where X is the set of the sensor's active pixel positions, z is the actual raw-data output, y is the ideal output, ξ is zero-mean random noise with standard deviation equal to 1, and σ is a function y, modulating the standard-deviation of the overall noise component. The function $\sigma(y)$ is called *standard-deviation function* or *standard-deviation curve*. The function $\sigma^2(y)$ is called *variance function*. Since $E\{\xi(x)\} = 0$ we have $E\{z(x)\} = y(x)$ and $\operatorname{std}\{z(x)\} = \sigma(E\{z(x)\})$. There are no additional restrictions on the distribution of $\xi(x)$, and different points may have different distributions.

In practice, $z(x)$ is the recorded value of the raw-data at the pixel x, and $y(x)$ is the ideal value to be recorded if no quantization or noise would be present. The (signal-dependent) signal-to-noise ratio (SNR) of the imaging sensor can be expressed as $\text{SNR}(y) = \frac{y}{\sigma(y)}$.

A good estimate of $y(x)$ can be obtained as the pointwise average of a large enough number N of observations $z_n(x)$, $n = 1, \ldots, N$, of the form (1):

$$\frac{1}{N}\sum\nolimits_{n=1}^{N} z_n(x), \qquad z_n(x) = y(x) + \sigma(y(x))\,\xi_n(x). \tag{2}$$

III. The method

The experimental realization of the approach described by (2) requires that the deterministic terms of the equation are truly invariant with respect to the replication index n. That is, the underlying true signal $y(x)$ must not change over time. In practice this means that during the acquisition process the camera must not move, the acquisition parameters (e.g., exposure, aperture, gain) must be fixed and the illumination is constant in time.

A. Acquisition and averaging

Under these conditions, we record a number N of images in raw-data format. These shots are averaged, to obtain an approximation $\bar{z}$ of the noise-free y,

$$\bar{z}(x) = \frac{1}{N}\sum\nolimits_{n=1}^{N} z_n(x) = y(x) + \frac{\sigma(y(x))}{\sqrt{N}}\tilde{\xi}(x), \quad x \in X. \tag{3}$$

Here $\tilde{\xi}(x)$ is again some zero-mean noise with unitary variance. It is recommended to take a large number of images, so that the factor $1/\sqrt{N}$ in (3) is small. In what follows we assume that N is large enough and we consider $\bar{z}(x) = E\{z(x)\} = y(x)$, for all x.

B. Segmentation

The average image $\bar{z}$ is segmented into a number of uniform regions $\{S\}$, or segments. Ideally, within these regions the value of $\bar{z}(x)$ should be constant:

$$S(y) = \{x : \bar{z}(x) = y\}.$$

However, this may lead to uncertain results as there may be too few (or maybe none) samples, i.e. pixels, that satisfy the equality $\bar{z}(x) = y$. Pragmatically, it is convenient to consider a larger estimation set of the form

$$S_\Delta(y) = \{x : \bar{z}(x) \in [y - \Delta/2, y + \Delta/2)\},$$

where $\Delta > 0$.

Let y and Δ be fixed and denote by x_m, $m = 1, \ldots, M$, the coordinates of the M pixels that constitute the segment $S_\Delta(y)$, $\{x_m\}_{m=1}^{M} = S_\Delta(y)$. On this segment we have that the observations (2) satisfy

$$\begin{aligned} z_n(x_m) &= y(x_m) + \sigma(y(x_m))\,\xi_n(x_m) = \\ &= y + \Delta d_{n,m} + \sigma(y(x_m))\,\xi_n(x_m), \end{aligned} \tag{4}$$

where $d_{n,m} \in [-1/2, 1/2)$. It is reasonable to assume that $d_{n,m}$ are uniformly distributed. Hence, Equation (4) can be rewritten as

$$z_n(x_m) = y + \sqrt{\frac{\Delta^2}{12} + \sigma^2(y(x_m))}\xi'_n(x_m), \tag{5}$$

where ξ'_n is again some zero-mean random noise with standard deviation equal to 1.

C. Measurement of the standard deviation

For any fixed y we proceed as follows in order to compute the estimate $\hat{\sigma}(y)$ of $\sigma(y)$.

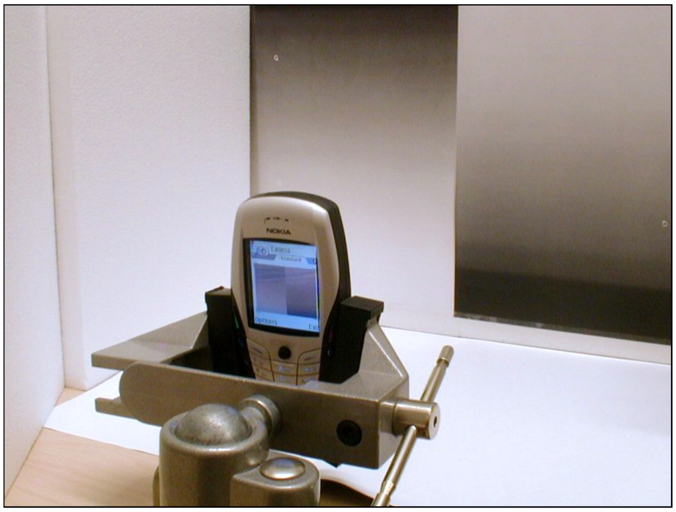

Fig. 1. Experimental setup.

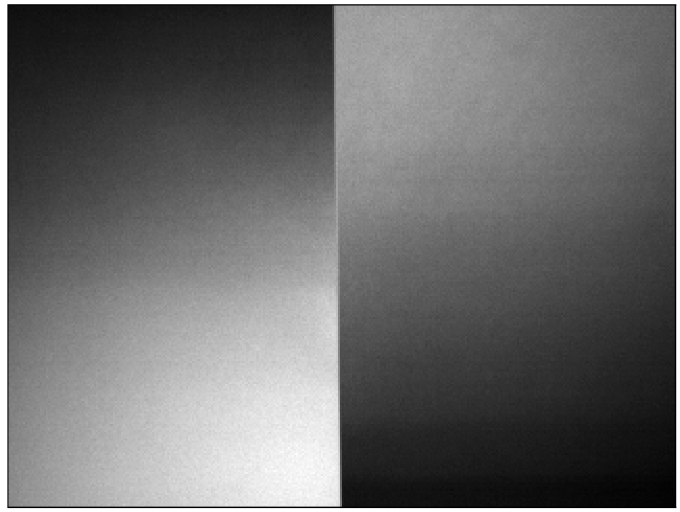

Fig. 2. One of the acquired shots z_n.

1) Variance for a single image and output value: To contain memory requirements, the variances are first computed independently for each shot z_n. The variance $\mathrm{svar}_n(y)$ corresponding to expected output y and shot n is calculated as the unbiased sample variance estimator on the segment,

$$\mathrm{svar}_n(y) = \frac{\sum_{m=1}^{M} \left(z_n\left(x_m\right) - \tilde{z}_n\left(y\right)\right)^2}{M-1}, \tag{6}$$

where

$$\{x_m\}_{m=1}^{M} = S_\Delta\left(y\right) = \{x : \bar{z}\left(x\right) \in [y - \Delta/2, y + \Delta/2)\}$$

and $\tilde{z}_n\left(y\right)$ is the mean value of z_n over $S_\Delta\left(y\right)$, $\tilde{z}_n\left(y\right) = \frac{1}{M}\sum_{m=1}^{M} z_n\left(x_m\right)$.

By (5), and assuming that within the segment $\sigma\left(y\left(x_m\right)\right)$ can be well approximated by $\sigma\left(y\right)$, we obtain

$$E\left\{\mathrm{svar}_n(y)\right\} = \frac{\Delta^2}{12} + \sigma^2\left(y\right). \tag{7}$$

2) Averaging over the shots: The estimate $\hat{\sigma}\left(y\right)$ of the standard deviation $\sigma\left(y\right)$ is given according to (7) by the average of the above N estimates:

$$\hat{\sigma}\left(y\right) = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\mathrm{svar}_n(y) - \frac{\Delta^2}{12}}. \tag{8}$$

Here we assume that the expectation under the square-root is non-negative. In practice, one can neglect the extra term due to Δ in (7) and (8) by choosing Δ significantly smaller than $\sigma\left(y\right)$. In this case we simply have

$$\hat{\sigma}\left(y\right) = \sqrt{\frac{1}{N}\sum_{n=1}^{N}\mathrm{svar}_n(y)}.$$

We remark that while the recorded raw-data has a certain fixed precision (e.g. 10 bits), the values attained by $\bar{z}$ are much denser because of the averaging. Therefore Δ can be taken much smaller than the quantization step of the raw-data. In Section IV we show results obtained for much different choices of Δ which demonstrate that the proposed technique is very stable with respect to this parameter.

3) Standard-deviation curve (variance function): The procedure described by the previous two steps is repeated for different values of y. In this way the standard-deviation curve is found. Typically, one would use discrete values $y \in \{\Delta i, i \in \mathbb{N}\} \cap [\min\{\bar{z}\}, \max\{\bar{z}\}]$, which ensures that the segments are nonoverlapping and hence that the measurements for different values of y are independent. The estimate of the variance function is simply $\hat{\sigma}^2\left(y\right)$, $y \in [\min\{\bar{z}\}, \max\{\bar{z}\}]$.

IV. Experimental results

The standard-deviation curves $\sigma\left(y\right)$, as from Equation (1), have been measured from the raw-data of two different CMOS sensors used in Nokia cameraphones. We denote them here as sensor/cameraphone “U” and “V”. They are, respectively, an older 660×492 (VGA) 1/4" sensor (5.4μm pixel pitch) with global shutter, and a newer 1296×1040 (1.3 Mpixel) 1/3.3" sensor (3.3μm pixel pitch) with rolling shutter. Both sensors have a Bayer pattern color filter array (CFA) with red, green, and blue filters (R, G_1, G_2, and B).

A. Setup

The sensors had not been separated from the phones, which were held clamped in a vice. We used a non-uniform target composed by grayscale vertical ramps going from white to black[1]. This simple setup is shown in Figure 1. To ensure constant-in-time (flicker-free) illumination, measurements where taken in a darkroom where the only source of light was an array of white LED lights powered by stabilized DC power supplies.

The devices were configured to take multiple shots automatically and without user intervention (as this would introduce mechanical vibration). A total of 50 shots were taken for each experiment. In Figure 2 we show an example of the shots which were taken. Enlarged details of this shot and of the average image $\bar{z}$ (3) are shown in Figure 3. The raw-data had 10-bit precision, which we normalize on the range $[0, 1]$ dividing by 2^{10}=1024. The exposure time, gain, and illumination were fixed in such a way that the recordings were not saturated (i.e., clipped) and thus our raw-data output is

[1]The target was printed on normal paper using a common office printer and then enlarged to A3-size using a copy machine.

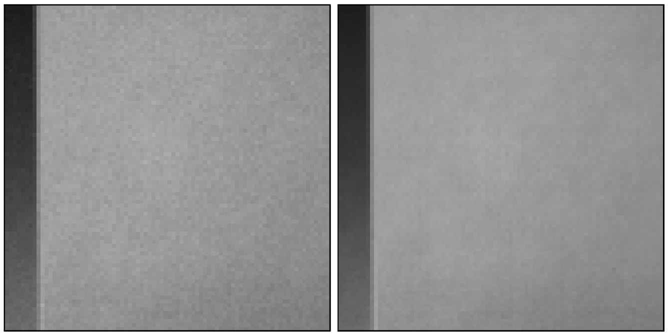

Fig. 3. Enlarged detail of the acquired shot z_n shown in Figure 2 and the corresponding fragment of the image $\bar{z}$, which is obtained by averaging N such z_ns.

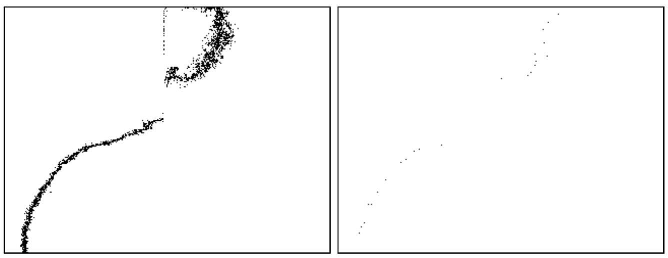

Fig. 4. Two segments $S_\Delta(y)$ obtained for Δ=0.01 (left) and Δ=0.0001 (right). The value of y is the same for both segments.

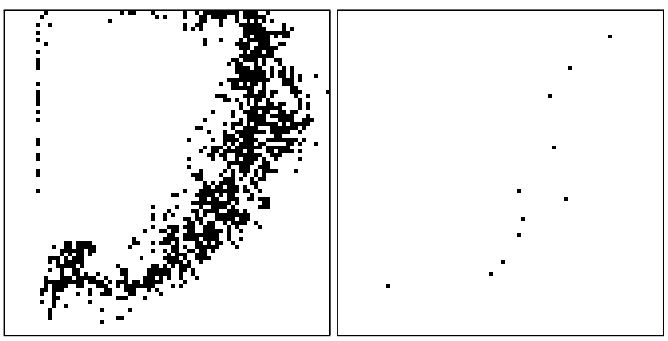

Fig. 5. Enlarged details of the segments $S_\Delta(y)$ obtained for Δ=0.01 and Δ=0.0001, shown in Figure 4. These details correspond to the same enlarged location shown in Figure 3.

typically concentrated inside the range $[0, 0.5]$. Furthermore, because of pedestal offset, the output is always larger than a certain minimum value, which for the two considered devices is about 0.05.

B. Segmentation

In Figures 4 and 5 we show examples of a segment $S_\Delta(y)$ obtained for $\Delta = 0.01$ and $\Delta = 0.0001$. These values are respectively about ten times and one tenth of the quantization step 1/1024. Observe that the segment corresponding to the smaller Δ is a subset of the other one. It is interesting to observe that despite the target was composed by vertical ramps, because of the nonuniform illumination and the vignetting, the segments do not have a particular horizontal/vertical "shape". Especially for small values of Δ the segments are usually composed of separated (i.e., disconnected) pixels.

C. Standard-deviation curves

In Figure 6, we show the standard-deviation curves obtained for $\Delta = 0.005, 0.0005, 0.00005$. It can be seen that, although the curve becomes noisier for very small Δ, the three plots are essentially the same. This demonstrates the accuracy of (7-8) and the stability of our procedure with respect to the choice of Δ. We remark also that the noisiness of the plot obtained for $\Delta = 0.00005$ is well compensated by its higher "sampling" density (there are about 10000 samples in the plot) which allows for very accurate smoothing or parametric fitting.

Due to the automatic segmentation, the standard-deviation vs. expectation curve can been measured using the whole sensor at once as well as using each color channel separately. The two approaches are in general equivalent. In Figure 7 we show the plots obtained separately for each one of the four color channels (R, G_1, G_2, and B), whereas the plots shown in Figure 6 are due to measurements using the whole sensor. The four color channels exhibit the same behavior shown in Figure 6.

We note that the sharp vertical drops in the estimated standard-deviation $\hat{\sigma}$ visible in the plots (e.g., for the blue and for the red channels in Figure 7(left)) do not correspond to a real drop in the standard deviation of the noise. Instead, they are only due to the segments $S_\Delta(y)$ becoming singletons or empty sets when y approaches or exceeds the bounds of the interval $[\min\{\bar{z}\}, \max\{\bar{z}\}]$. For empty or singleton segments $S_\Delta(y)$, the Equation (6) loses its meaning. Such drops in the estimated standard-deviation are therefore to be ignored when analyzing the estimated curve $\hat{\sigma}(y)$.

D. Fitting a parametric model

The following parametric model based on the Poissonian (photon-limited) nature of the sensor [5] achieves a near-perfect fit to the experimentally measured data:

$$\sigma(y) = q\sqrt{y-p}, \qquad \begin{array}{ll} q_{\mathrm{U}} = 0.0060, & p_{\mathrm{U}} = 0.050, \\ q_{\mathrm{V}} = 0.0092, & p_{\mathrm{V}} = 0.021, \end{array} \tag{9}$$

The plots of these functions are shown in Figure 7. We remark again that the same curve $\sigma(y)$ fits equally well all four color channels. It can be seen that the newer sensor "V" presents a lower SNR. This can be justified by the increased pixel density.

V. CONCLUSIONS

In this paper we have presented a novel method for the accurate measurement of the standard-deviation of the temporal noise in the raw-data of digital imaging sensors. The method is based on an automatic segmentation of the recorded images. It allows to use non-uniform targets and illumination, thus can be implemented easily without special calibrated test charts and lighting. While previous techniques estimate only single value for the standard deviation for a given uniform raw-data, the proposed method can measure a complete standard-deviation vs. expectation curve in a single experiment. The experimental results obtained by the proposed measurement method show a near-perfect agreement with an analytically derived noise model for the raw-data.

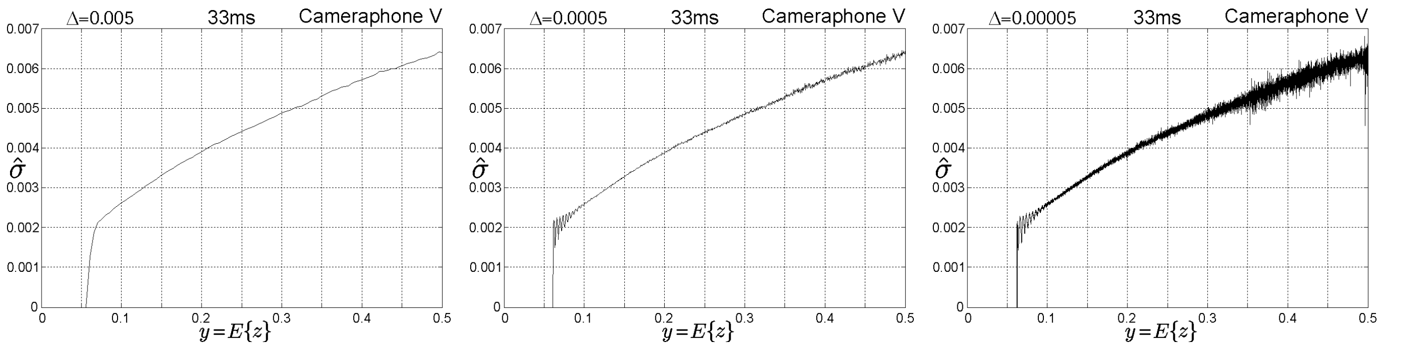

Fig. 6. Standard-deviation curves for the raw-data from CMOS sensor "V" measured for different values of the segmentation parameter Δ=0.005, 0.0005, 0.00005. The curves are calculated using Equation (8) and give an estimate $\hat{\sigma}$ of the standard-deviation of the temporal noise in the raw-data as function of the expected (i.e., noise-free) output raw-data value $y = E\{z\}$.

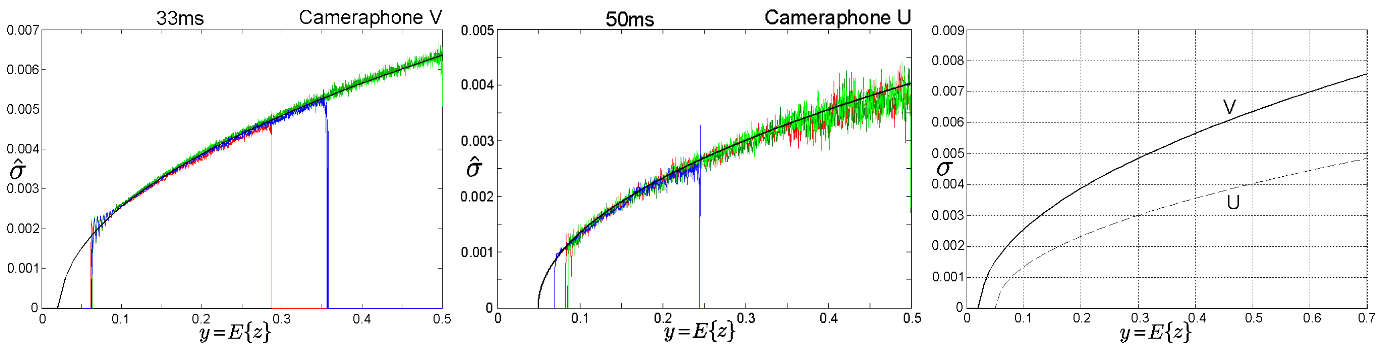

Fig. 7. Measured standard-deviation curves (and corresponding fitted model) for the raw-data from CMOS sensor "V" (left) and "U" (center). Plots for the different color channels (R, G_1, G_2, B) are drawn in their corresponding colors. The model fits equally well all color channels. A direct comparison between the two models (right), shows that for the same expected raw-data values the newer sensor is noisier than the older one.

References

[1] Blanksby, A.J., M.J. Loinaz, D.A. Inglis, and B.D. Ackland, "Noise performance of a color CMOS photogate image sensor", *IEEE Int. Electron Devices Meeting 97 Tech. Dig.*, pp. 205-208, 1997.

[2] Blanksby, A.J., and M.J. Loinaz, "Performance analysis of a color CMOS photogate sensor", *IEEE Trans. Electron Devices*, vol. 47, no. 1, pp. 55-64, January 2000.

[3] Costantini, R., and S. Süsstrunk, "Virtual Sensor Design", *Proc. IS&T/SPIE Electronic Imaging 2004: Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications V*, vol. 5301, pp. 408-419, 2004.

[4] Foi, A., R. Bilcu, V. Katkovnik, and K. Egiazarian, "Anisotropic local approximations for pointwise adaptive signal-dependent noise removal", *Proc. XIII European Signal Proc. Conf., EUSIPCO 2005*, September 2005.

[5] Foi, A., M. Trimeche, V. Katkovnik, and K. Egiazarian, "Practical Poissonian-Gaussian noise modeling and fitting for single image raw-data", submitted for publication to *IEEE Trans. Image Processing*, 2007 (in review).

[6] Hytti, H., "Characterization of digital image noise properties based on RAW data" *Proc. SPIE-IS&T Electronic Imaging, Image Quality and System Performance III*, vol. 6059, 60590A, Jan. 2006.

[7] "Photography - Electronic still-picture cameras - Methods for measuring opto-electronic conversion functions (OECFs)", ISO 14524:1999, Dec., 1999.

[8] "Photography - Electronic still-picture imaging - Noise measurements", ISO 15739:2003, May, 2003.

[9] Tian, H., B. Fowler, and A. El Gamal, "Analysis of Temporal Noise in CMOS Photodiode Active Pixel Sensor", *IEEE Journal of Solid-State Circuits*, vol. 36, no. 1, pp. 92–101, Jan. 2001.

[10] Wach, H., and E.R. Dowski, Jr., "Noise modeling for design and simulation of computational imaging systems", *Proc. SPIE, Visual Information Processing XIII*, vol. 5438, pp. 159-170, July 2004.

PLACE PHOTO HERE

A lessandro Foi has received the M.Sc. and the Ph.D. degrees in mathematics from Università degli Studi di Milano (Italy) in 2001 and from Politecnico di Milano in 2005, respectively. His research interests include mathematical and statistical methods for signal processing, functional analysis, and harmonic analysis. Currently, he is a researcher at the Institute of Signal Processing, Tampere University of Technology (Finland). His work focuses on spatially adaptive algorithms for denoising and deblurring of digital images and on noise modeling for digital imaging sensors.

PLACE PHOTO HERE

S akari Alenius is a senior research scientist at Nokia Research Center since 2003. His Dr.Tech. thesis (1999, Tampere University of Technology) dealt with image reconstruction in Positron Emission Tomography. In 2007-2008, he will conduct research at TUT in "Computational imaging concepts and algorithms"-project funded by the Academy of Finland.

PLACE PHOTO HERE

V ladimir Katkovnik received the M.Sc., Ph.D., and D.Sc. degrees in technical cybernetics from the Leningrad Polytechnic Institute, Leningrad, Russia, in 1960, 1964, and 1974, respectively. From 1964 to 1991, he held the positions of Associate Professor and Professor at the Department of Mechanics and Control Processes, Leningrad Polytechnic Institute. From 1991 to 1999, he was a Professor of statistics with the Department of the University of South Africa, Pretoria. From 2001 to 2003, he was a Professor of mechatronics with the Kwangju Institute of Science and Technology, Korea. From 2000 to 2001 and since 2003 he is a Research Professor with the Institute of Signal Processing, Tampere University of Technology, Tampere, Finland. He has published seven books and more than 200 papers. His research interests include stochastic signal processing, linear and nonlinear filtering, nonparametric estimation, imaging, nonstationary systems, and time-frequency analysis.

PLACE PHOTO HERE

K aren Egiazarian (SM'96) was born in Yerevan, Armenia, in 1959. He received the M.Sc. degree in mathematics from Yerevan State University in 1981, the Ph.D. degree in physics and mathematics from Moscow State University, Moscow, Russia, in 1986, and the D.Tech. degree from the Tampere University of Technology (TUT), Tampere, Finland, in 1994. He has been Senior Researcher with the Department of Digital Signal Processing, Institute of Information Problems and Automation, National Academy of Sciences of Armenia. Since 1996, he has been an Assistant Professor with the Institute of Signal Processing, TUT, where he is currently a Professor, leading the Transforms and Spectral Methods Group. His research interests are in the areas of applied mathematics, signal processing, and digital logic.