doi: 10.5072/zenodo.12761

Heidelberg Collaboratory

for Image Processing

Compression by Noise Equalization

HCI Technical Report No. 1

B. Jähne

Heidelberg Collaboratory for Image Processing (HCI) am
Interdisziplinären Zentrum für Wissenschaftliches Rechnen (IWR)
Universität Heidelberg

Version A, November 29, 2010
Version B, November 30, 2010
Version C, June 25, 2013 :
Corrections and practical computation of LUT considering offsets and variable compression quality
Version D, June 30, 2013 : Addition of resolution in Section 3.3, cumulative histograms in Section 4
Version D, July 25, 2013 : Test of simplified LUT, Version 1, implemented for pco.edge

Abstract

For any relation of the noise variance with the gray value, $\sigma^2(g)$, a nonlinear transform $h(g)$ can be applied so that the variance of the transformed signal h is constant. The number of bits required to represent the noise-equalized signal is in good approximation equal to the maximum signal/noise ratio (SNR_{max}). Thus the noise-equalized signal of any imaging sensor with a full-well capacity of less than 2^{16} can be represented by only 8 bit or less with only a slight increase in the overall noise level caused by additional quantization noise. The procedures are illustrated using the pco.edge camera.

Contents

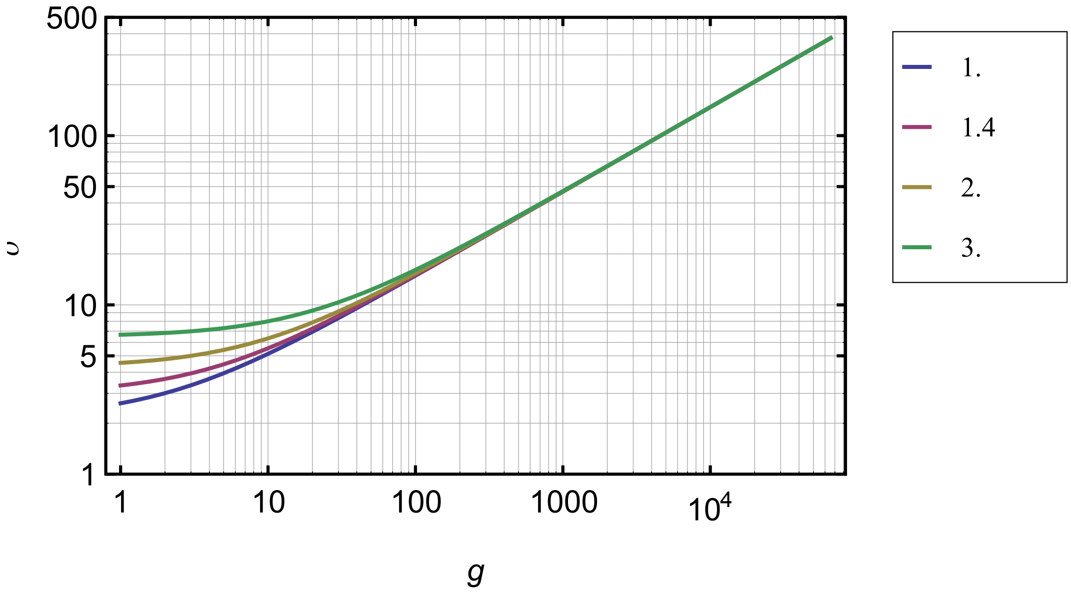

Figure 1: *Standard deviation σ as a function of the mean gray value in a double logarithmic graph. The values for the pco.edge with a perfect linear characteristic curve are taken: $\sigma^2 = \sigma_0 + Kg$, with $K = 2.17\,DN/e^-$ and $(\sigma_0 = \{1.0, 1.4, 2.0, 3.0\}e^- K$ (as indicated in the figure legend) and 16 bit quantization.*

1 Introduction

Modern high-end cameras show a low dark noise and therefore the variance of the noise increases strongly with the mean digital value. Any equidistant quantization is a clear mismatch to this situation. The quantization must be fine enough to resolve the small random fluctuations at low signal levels but is therefore much too fine at high signal levels. As an illustration the standard deviation of the noise of the pco.edge is shown in Fig. 1. The standard deviation covers more than two orders of magnitude. It is a little more than two for the dark image but more than 300 close to saturation.

Thus the question arises, whether it is possible to apply a nonlinear transform in such a way that the standard deviation in the independent of the gray value. The interesting question is: Given the noise variance / gray value relation, what is the minimum number of bits with which a signal can be presented without any significant loss of information? It would be especially useful, if an 8 bit quantization could be achieved. Transforming a 16-bit image to an 8-bit image constitutes already a compression factor of 2.

PCO applied already this idea by using a combined linear / square root gray value transform [1, 4]. This solution is suboptimal in two respects. Firstly, there is a jump in the slope of the nonlinear transform at the transition from the linear to the square root part of the curve [1, Fig. 1]. Secondly, the standard deviation is *not* constant over the whole gray value range.

Here a continuous analytic transformation curve is derived for a given noise variance / gray value relation In the transformed signal, the standard deviation is constant from the dark signal to saturation. In this way, an optimal compression of the signal can be reached.

In Section 2 the necessary theoretical background about quantization of noisy signals is addressed. In particular, an answer is given to the question about the required ratio between the standard deviation of the noise to the quantization levels. Then in Section 3, the nonlinear transformation for noise equalization is discussed and it is derived how many bits are required to represent a camera signal in a noise-equalized signal without significant signal degradation. Finally, Section 4 gives some practical examples, including the pco.edge sCMOS camera.

2 Quantization of Noisy Signals

Quantization limits the resolution of output values. As simple as it appears at first glance, it is not easy to describe the effects of quantization theoretically. This is due to the fact that quantization is a nonlinear function

$$q = \text{floor}\left(\frac{g + 0.5}{\Delta g}\right), \tag{1}$$

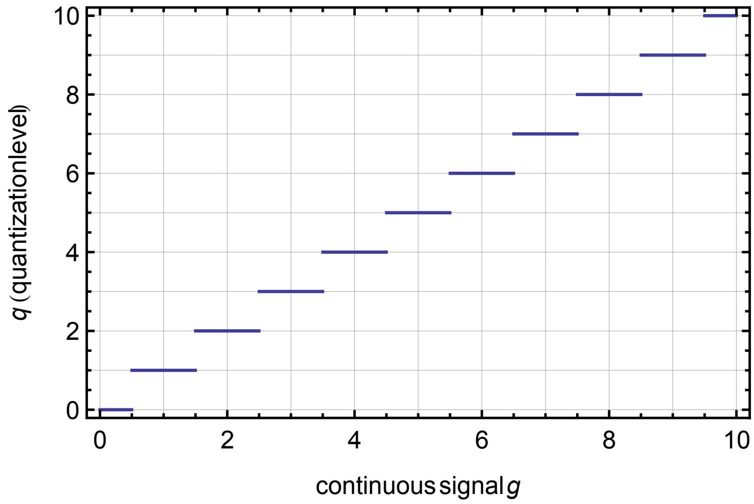

Figure 2: *Quantization is a nonlinear staircase function, see Eq. (1).*

mapping real numbers g to an integer numbers q. Δg is the distance between the equidistant quantization intervals. Quantization results in a staircase function (Fig. 2). Every value g within the the interval $[q\Delta g - 1/2, q\Delta g + 1/2[$ results in the quantized value q. Only one thing is easy to state, the maximum error caused by quantization of a signal. It is half of the quantization interval Δg.

But it is impossible to reconstruct the original values from the quantized values. Interestingly, this is in contrast to sampling. If a continuous signal is sampled and the conditions of the *sampling theorem* are met (in simple terms, every periodic component contained in the signal is sampled at least twice per period), then it is possible to reconstruct the continuous signal *exactly* from the sampled signal.

Because of the nonlinear nature of the quantization function in Eq. (1), there is no theorem for quantization equivalent to the sampling theorem. It is also not necessary, because any real signal is uncertain in any way by its random nature. Thus the question is rather how fine the quantization must be so that the statistical properties of the signal are not disturbed. Or in other words, a proper theory of quantization must answer the question: can we reconstruct the probability density function (PDF) of the continuous signal from the PDF of the quantized signal? If this is the case, then we can estimate, e. g., the mean and the variance of the random signal from the quantized signal without *any* error.

This problem has been studied in detail by Widrow and Kollar [5] and resulted in quantization theorems that are very similar to the sampling theorem if we replace the continuous signal g by its PDF. There are actually several *quantization theorems* [5, Section 4.3]. The first states the condition that the continuous PDF can be *exactly* reconstructed from the PDF of the quantized signal, the second states the condition that only the moments of the continuous PDF can be *exactly* reconstructed from the PDF of the quantized signal. The latter conditions are less stringent and are of more importance, because the PDF normally is known and all what we need to know is the mean and the variance of the signal. The conditions for exact reconstruction require that the PDF is bandlimited, for details see Widrow and Kollar [5, Section 4.3]. If these conditions are meet the mean of the quantized signal is *exactly* the mean of the original signal and the quantization adds pseudo noise (PQN model, [5, Section 4.2]) to the quantized signal with a uniform distribution and a variance

$$\sigma_q^2 = \frac{1}{12}(\Delta g)^2. \tag{2}$$

The variance of the original signal, σ_g^2, is then given from the measured variance of the quantized signal, σ_q^2, by

$$\sigma_g^2 = \sigma_q^2 - \frac{1}{12}(\Delta g)^2. \tag{3}$$

The big difference to sampling is that this condition is, unfortunately, never met *exactly*, neither by the normal distribution nor by the Poison distribution. Therefore Eq. (3) will only be an approximation. The question is just, how fine the quantization must be so that the error is negligible. This is discussed in detail by Widrow and Kollar [5, Chapter 5]. The answer is surprising: Widrow and Kollar [5, Section 5.5] come to the conclusion that the errors both in the mean and the variance computed by Eq. (3) are negligible when

$$\Delta g \le 2\sigma_g \quad \text{or} \quad \sigma_g \ge \Delta g/2. \tag{4}$$

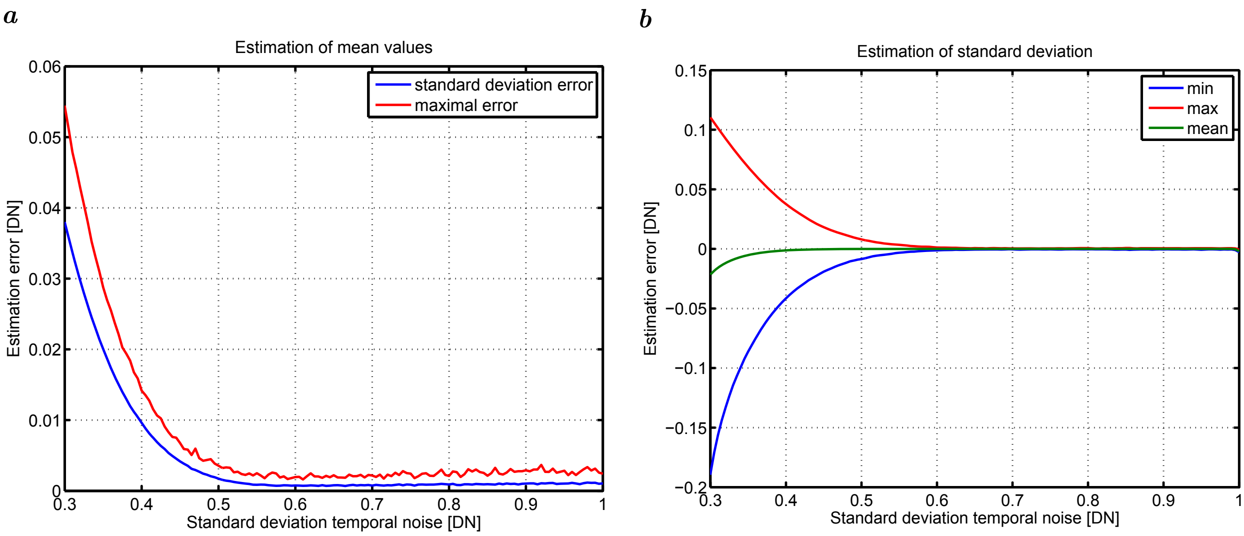

Figure 3: *Results from Monte Carlo simulations for the estimation of **a** the mean value and **b** the standard deviation of the digital gray value as a function of the standard deviation in units of digital numbers [DN] in the range between 0 and 1.*

This means that the standard deviation of the noise σ_g can be as low as half of the quantization resolution Δg.

We verified these theoretical results by independent numerical Monte-Carlo simulations. For the simulations 201 mean gray values equally distributed between 0 and 1 were taken and zero-mean normally distributed noise was added to the values. The estimated mean and variances were averaged over 900 000 realizations of each value. Finally, the deviations in the estimations were averaged over all 201 values.

The results are shown in the range for $\sigma_g/\Delta g = [0.3, 1]$ in Fig. 3. The mean gray value can be estimated with a maximum error of less than 0.014 DN even for standard deviations as low as 0.4 DN (Fig. 3b). The maximum error of the estimate of the standard deviation remains below 0.04 even for standard deviations as low as 0.4.

Therefore the Monte Carlo simulations fully agree with the theoretical results from Widrow and Kollar [5]. For the nonlinear transform this means that we can apply it in such a way that the (constant) standard deviation in the nonlinear signal h is 0.5 DN.

3 Noise Variance Equalization

3.1 General Solution for Noise Variance Equalization

As already stated in Section 1, the variance $\sigma_g^2(g)$ is a function of the mean gray value. We start with any arbitrary function to show that it is possible for almost any function $\sigma_g^2(g)$ to device a nonlinear gray value transfer so that the standard deviation is constant in the transformed signal h. This procedure has first been proposed by [2]. It is derived in the following.

By the laws of error propagation (see, e. g., [3, Section 3.2], the variance of $h(g)$ is given in first order by

$$\sigma_h^2 \approx \left(\frac{\mathrm{d}h}{\mathrm{d}g}\right)^2 \sigma_g^2(g) \tag{5}$$

If we set σ_h^2 to be constant, we can rearrange Eq. (5) to

$$\mathrm{d}h = \frac{\sigma_h}{\sqrt{\sigma^2(g)}}\mathrm{d}g.$$

Integration yields

$$h(g) = \sigma_h \int\limits_0^g \frac{\mathrm{d}g'}{\sqrt{\sigma^2(g')}}. \tag{6}$$

The integration constant is chosen in such a way that $h(0) = 0$. Equation (6) clearly says that an analytical solution exists for any function $\sigma_g^2(g)$ for which the integral can be expressed by an analytic function. If this is not the case, we can still solve Eq. (6) numerically.

3.2 Specific Solution for Linear Camera Model

No we can derive the specific solution for the linear sensor model. Then the variance increases linearly with the mean gray value:

$$\sigma_g^2(g) = \sigma_0^2 + Kg \tag{7}$$

The following equations become simpler, if we introduce the new variable

$$\tilde{g} = \frac{g}{g_{\max}} \quad \text{and} \quad \tilde{h} = \frac{h}{h_{\max}} \tag{8}$$

as the fraction of saturation of gray values in the range $[0, g_{\max}]$ and $[0, h_{\max}]$, respectively, and the maximum variance as

$$\sigma_{\max}^2 = \sigma_g^2(g_{\max}) \tag{9}$$

Then Eq. (7) can be written as

$$\sigma_g^2(g) = \sigma_0^2 + (\sigma_{\max}^2 - \sigma_0^2)\tilde{g} \tag{10}$$

With the linear variance function Eq. (10), the integral in Eq. (6) yields

$$h(g) = \frac{2\sigma_h}{K}\left(\sqrt{\sigma_0^2 + Kg} - \sigma_0\right). \tag{11}$$

We use the free parameters σ_h to map the values of h into the interval $[0, h_{\max}]$. This implies the conditions $h(g_{\max}) = h_{\max}$ and we obtain

$$\tilde{h} = \frac{\sqrt{\sigma_0^2 + (\sigma_{\max}^2 - \sigma_0^2)\tilde{g}} - \sigma_0}{\sigma_{\max} - \sigma_0}, \quad \sigma_h = \frac{h_{\max}}{2} \cdot \frac{\sigma_{\max} + \sigma_0}{g_{\max}}. \tag{12}$$

Now we can use this equation to answer the important question how many bits are required to quantize the equalized signal h. This is given by the value of $h_{\max}$, which can be expressed by

$$h_{\max} = 2\sigma_h \cdot \frac{g_{\max}}{\sigma_{\max} + \sigma_0} = 2\sigma_h \cdot \mathrm{SNR}_{\max} \bigg/ \left(1 + \frac{\sigma_0}{\sigma_{\max}}\right). \tag{13}$$

This equation gives an important result. Because the dark noise σ_0 is much smaller than the maximum noise $\sigma_{\max}$ for high-end cameras, the correction by σ_0 is very small. It makes the number of required bits, $h_{\max}$, anyway only smaller. Using the threshold for the standard deviation of the noise, $\sigma_h = 0.5$, derived at the end of Section 2, yields

$$h_{\max} \le \mathrm{SNR}_{\max}. \tag{14}$$

This means that the number of required bits depends in first order only on the maximum SNR. For any camera with a full-well capacity of less than $2^{16} = 65\,536$ electrons, the $\mathrm{SNR}_{\max}$ is smaller than 2^8 and thus only 8 bits are required to quantize the noise-equalized signal without any significant loss of signal quality.

Replacing $g_{\max}$ in Eq. (13) using Eq. (7), $h_{\max}$ can also be related to the system gain K:

$$h_{\max} = 2\sigma_h \cdot \frac{\sigma_{\max} - \sigma_0}{K}. \tag{15}$$

3.3 Inverse Relation, Interval Widths, and Resolution

In order to convert the noise-equalized signal back to the original signal, it is required to know the inverse relation to Eq. (12). This is given by

$$\tilde{g} = \tilde{h} \cdot \frac{\tilde{h}(\sigma_{\max} - \sigma_0) + 2\sigma_0}{\sigma_{\max} + \sigma_0}. \tag{16}$$

Another important question is the width of the intervals in the original signal, Δg, as a relation to the constant unit interval width of h. This relation can be derived by differentiating Eq. (16) and using Eq. (12), resulting in

$$\Delta g = \frac{\sigma_0 + \tilde{h}(\sigma_{\max} - \sigma_0)}{\sigma_h} = \frac{\sqrt{\sigma_0^2 + (\sigma_{\max}^2 - \sigma_0^2)\tilde{g}}}{\sigma_h}, \quad \Delta g \in \left[\frac{\sigma_0}{\sigma_h}, \frac{\sigma_{\max}}{\sigma_h}\right]. \tag{17}$$

The intervals Δg in the original signal space g are well adapted to the standard deviation σ_g. Inserting Eq. (10) into Eq. (17) yields

$$\Delta g = \frac{\sigma_g}{\sigma_h} \quad \text{or} \quad \sigma_g = \sigma_h \Delta g. \tag{18}$$

In the limiting case of $\sigma_h = 0.5$, the standard deviation is half the interval width and a variation of the signal in a range of $6\sigma_g$ is just 3 bins wide. This is just enough to sample the distribution into three to for bins from which the standard deviation and mean values can be estimated with sufficient accuracy as described in Section 2.

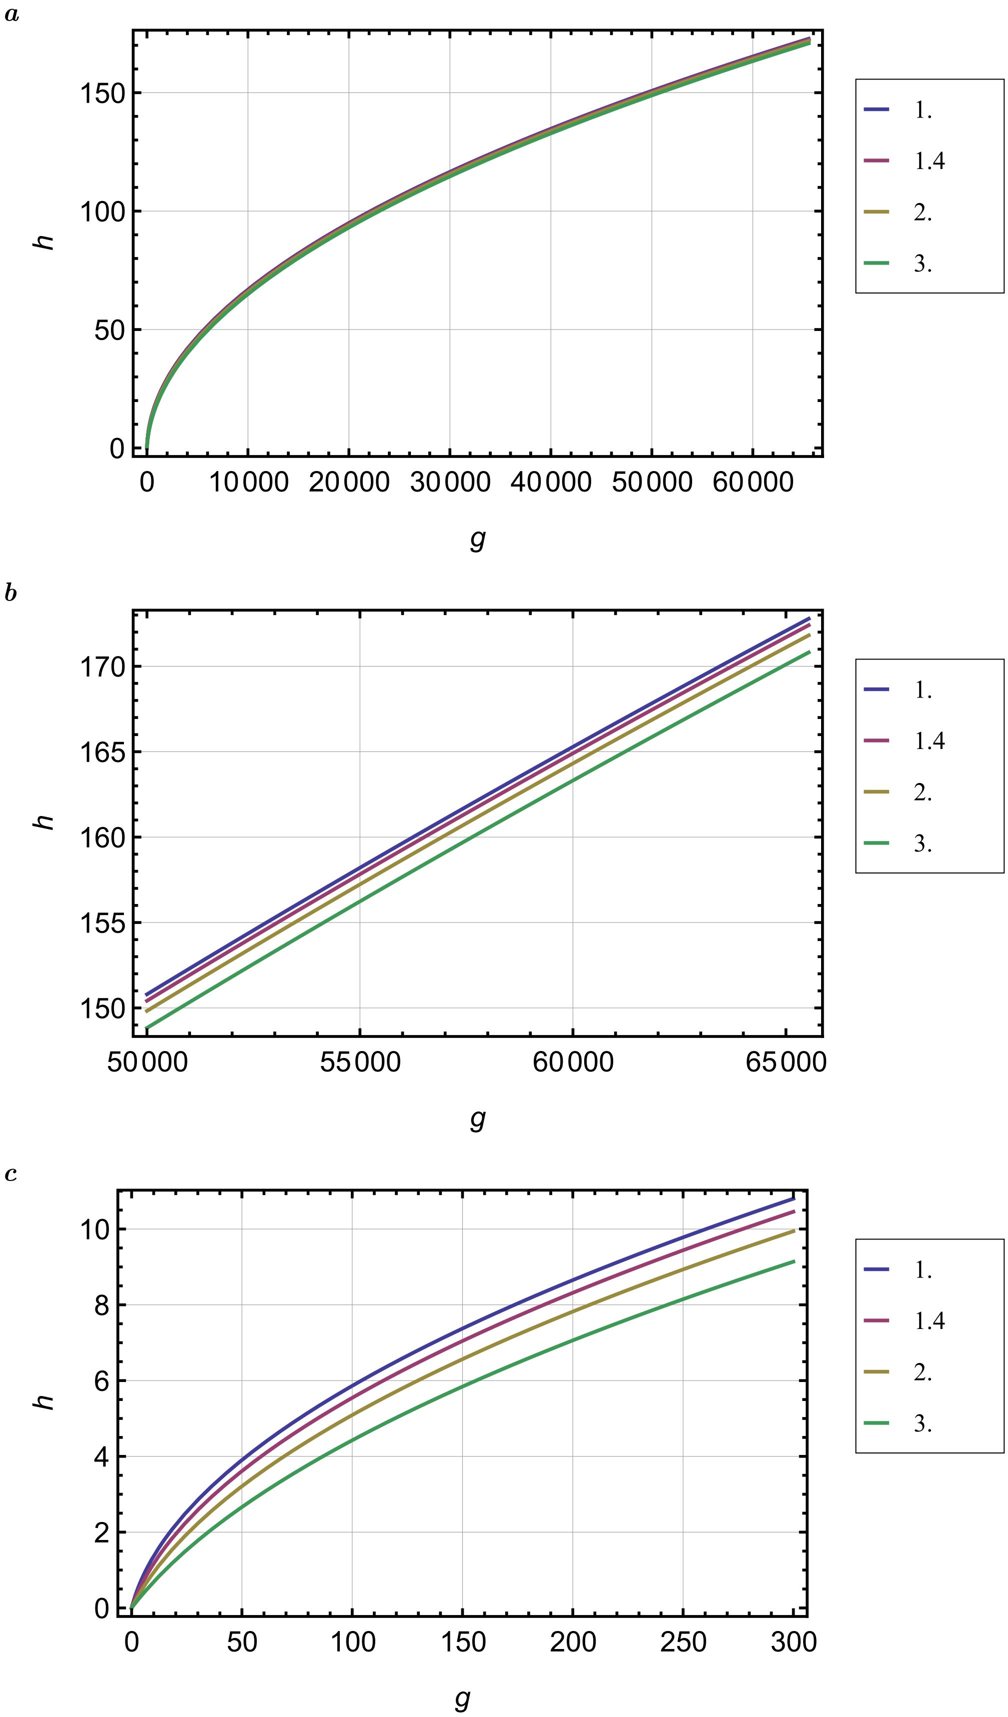

Figure 4: *Nonlinear noise-equalizing forward transform for the pco.edge camera with the same parameters as in Fig. 1: **a** whole gray value range, **b** small section close to saturation, and **c** 5% section close to the dark value*

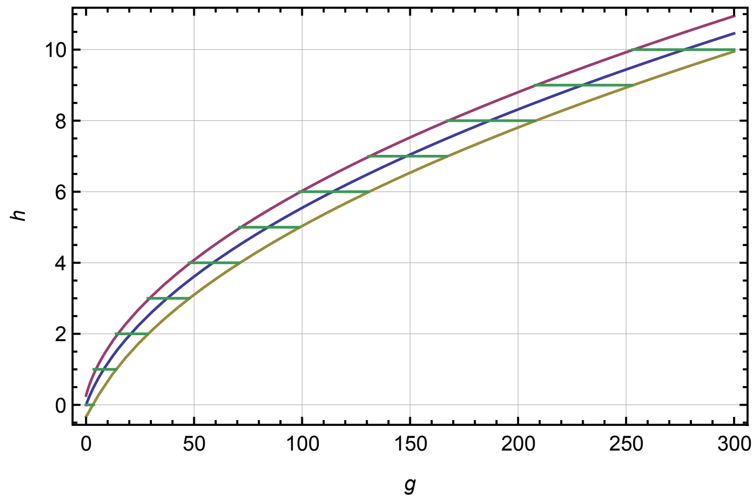

Figure 5: *Illustration of interval widths with nonlinear transformations: shown is the forward transform (blue curve) with a ± one sigma margin σ_g according to Eq. (7). Because $\sigma_h = 0.5$ was chosen, the values between the lower margin (pink curve) and the upper margin (cyan curve) just fall into one h integer bin (green horizontal lines).*

3.4 Computation of LUTs

The computation of the lookup tables for the forward and backward transform can directly be derived from Eqs. (12) and (16):

$$\begin{aligned} &\text{Forward} \quad h = \frac{h_{\max}}{\sigma_{\max} - \sigma_0}\left(\sqrt{\sigma_0^2 + (\sigma_{\max}^2 - \sigma_0^2)g/g_{\max}} - \sigma_0\right) = 2\sigma_h \cdot \frac{\sqrt{\sigma_0^2 + Kg} - \sigma_0}{K}, \\ &\text{Backward} \quad g = \frac{h}{h_{\max}} \cdot \left(2\sigma_0 + (\sigma_{\max} - \sigma_0)\frac{h}{h_{\max}}\right) / (\sigma_{\max} + \sigma_0) = \frac{h}{\sigma_h} \cdot \left(\sigma_0 + \frac{Kh}{4\sigma_h}\right). \end{aligned} \quad (19)$$

In the rightmost version of the equations in Eq. (12), only two camera parameters are required to compute the transforms:

- standard deviations of the dark noise σ_0
- camera system gain K (DN/e$^-$)

The parameter σ_h determines the compression ratio. For optimal compression $\sigma_h = 0.5$ should be used as discussed in Section 2. For an adjustable σ_h, its value can be computed using Eq. (13):

$$\sigma_h = \frac{h_{\max}}{2\,\mathrm{SNR}_{\max}}\left(1 + \frac{\sigma_0}{\sigma_{\max}}\right) = \frac{1}{2} \cdot \frac{h_{\max}K}{\sqrt{\sigma_0^2 + Kg_{\max}} - \sigma_0}. \quad (20)$$

The correct implementation of Eq. (19) can easily be checked with the values $h(0) = 0$, $h(g_{\max}) = h_{\max}$, $g(0) = 0$, $g(h_{\max}) = g_{\max}$ For an unbiased computation, it is required to round correctly of the floating point values x from Eq. (19) to integer values, i. e., the integer value q is given by $q = \mathrm{floor}(x + 0.5)$.

Equation (13) does not consider values smaller than the mean dark value. These values occur however due to noise. If we assume to take values down to m times the standard deviation σ_h, also the h values get an offset. For values smaller than the dark value g_0, a linear relation is proposed, which has the same slope as Eq. (19) for $g = 0$. Then we end up with the modifed LUT

$$h = m\sigma_h + \begin{cases} 2\sigma_h \dfrac{\sqrt{\sigma_0^2 + K(g_- g_0)} - \sigma_0}{K} & g \geq g_0 \\ \dfrac{\sigma_h}{\sigma_0}(g - g_0) & g < g_0 \\ 0 & h < 0 \end{cases} \quad (21)$$

The inverse LUT is given by

$$g = g_0 + \begin{cases} 2\sigma_h \dfrac{h - m\sigma_h}{\sigma_h}\left(\sigma_0 + \dfrac{K(h - m\sigma_h)}{4\sigma_h}\right) & h \geq m\sigma_h \\ \dfrac{\sigma_0}{\sigma_h}(h - m\sigma_h) & h < m\sigma_h \end{cases} \quad (22)$$

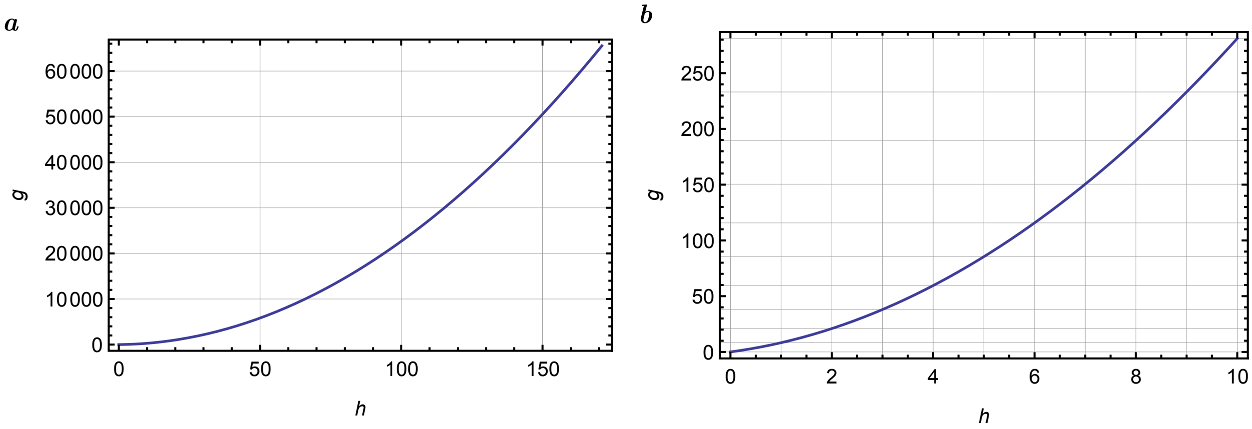

Figure 6: *Nonlinear backward transformations for the pco.edge with a dark noise of 1.4 e⁻: **a** Whole range, **b** small sector at dark value.*

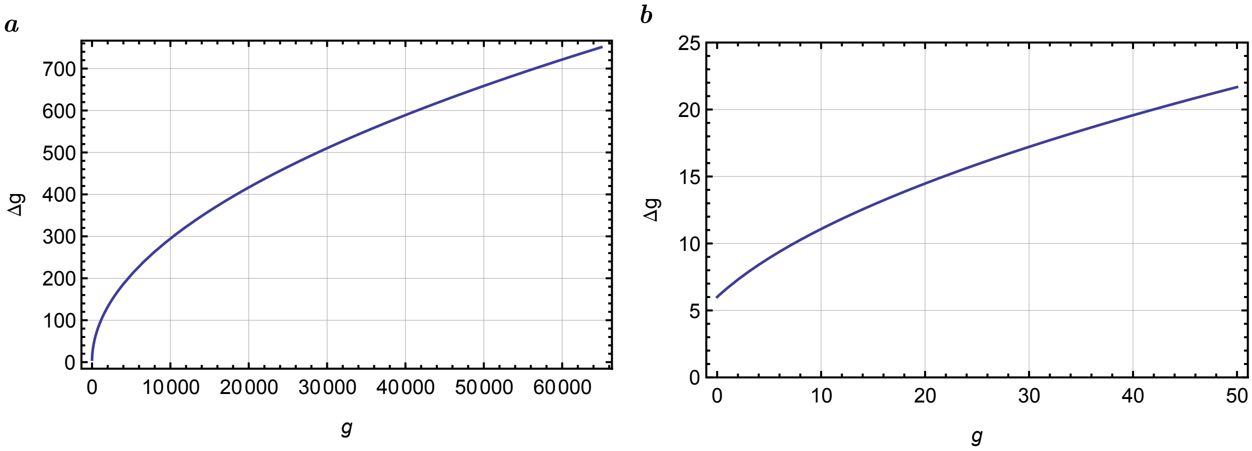

Figure 7: *Interval width of original signal g mapped to one interval of the noise-equalized signal h: **a** Whole range, **b** small sector at dark value.*

4 Examples

4.1 PCO.edge

As an illustrative example, we take the pco.edge camera with the parameters as shown in Fig. 1. Its original linear output signal has 16 bits ($g_{\max} = 2^{16}$) and with a full well capacity of 30 000 (almost 2^{15}, the noise-equalized signal can be expressed according to Eq. (14) with a maximum digital value of $h_{\max} = 171 \approx 2^{7.5}$ using $\sigma_h = 0.5$. Therefore only 7.5 bits are required to represent the noise-equalized signal. The nonlinear noise-equalizing transform curves are shown in Fig. 4, the backward transform in Fig. 6.

Figure 5 illustrates which intervals of g-values close to the dark vale are mapped into one unit interval of the noise equalized signal h and how the interval width is related to the standard deviation σ_g of the original singal g.

Figure 7 directly shows the interval widths of g that are mapped to one value in the noise equalized signal h. It can be seen that at least 6 values are mapped to one h value. Therefore no significant rounding errors should occur. This could be verified by computing histograms of the noise-equalized signal h.

4.2 Test with EMVA 1288 standard

Because it appears hard to believe that the signal of a camera with a high dynamic range can just be represented by 8 bits, a practical test was performed. A pco.edge camera (SN 1013) was measured according to the EMVA 1288 standard. The results are contained in the first row of Table 1. These values where taken to compute a compresion LUT with 2^{16} entries and 8 bit values according to Eq. (21) with $\sigma_h = 0.67$ and $m = 6$. The value 0.67 was chosen for σ_h so that the complete range of 16 bit for the values of g can be covered, see Eq. (20).

After applying the non-linear compression LUT, almost constant, signal independent noise is achieved

Table 1: *Verification of the compression by EMVA 1288 measurements. The upper row*

	η	σ_0 (DN)	σ_0 (e$^-$)	K (DN/e$^-$)	g_0 (DN)	$\mathrm{SNR}_{\mathrm{max}}$	DR
Direct (m0387)	0.454	3.91	1.97	1.975	96.32	185	12570
$\sigma_h = 0.67$ DN (m0389)	0.385	4.36	1.86	2.334	96.15	172	11530
corrected	0.455	4.36	2.21	1.975	96.15	172	11530

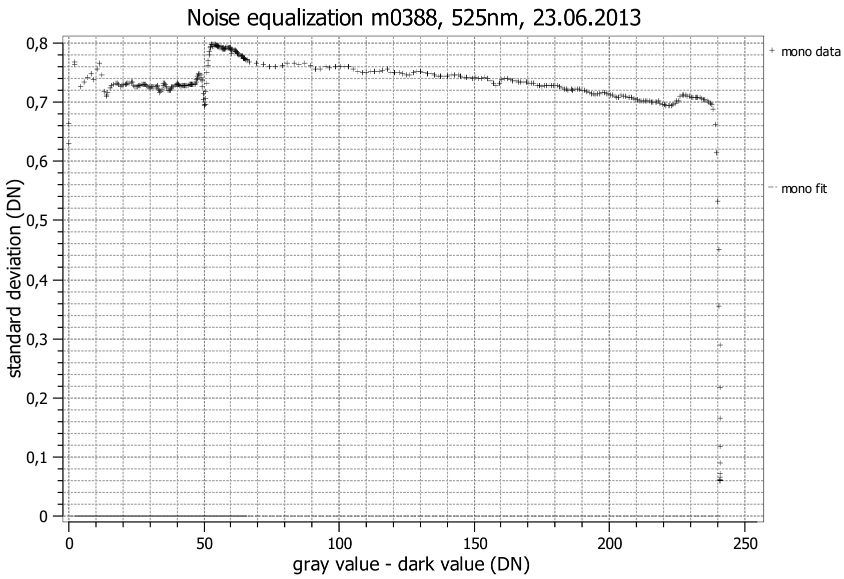

Figure 8: *Measured temporal noise in the compressed signal h.*

Fig. 8. The measured standard deviation is slightly higher than 0.67, because of the additional quantization noise with a variance of 1/12. Addition of the variances yields a resulting standard deviation of 0.73 DN, which is 8.9% higher and in good agreement with the measured values, shown in Fig. 8.

The influence of the compressing LUT to the EMVA 1288 parameters was tested in the following way. Firstly, the nonlinear compression LUT is applied using Eq. (21), then the inverse LUT after Eq. (22). These two processing steps result again in 16 bit values, but now at most 256 values are occupied in the 16-bit space, as illustrated in Fig. 9. With images processed in this way, a second EMVA 1288 measurements is performed. The results are shown in the second row of Table 1.

The influence of the additional quantization noise results in a higher gain value K. Because k scales with the variance, it should be 18.6% higher. The measurements show that it is 18.1% higher, which is an excellent agreement. If the correct gain is used, see third row in Table 1, the same quantum efficiency is computed as with the unprocessed images (row one in Table 1). The values of dark noise σ_0 will remain lower, the maximum SNR and dynamic range DR lower by the same factor, because of the additional noise caused by quantization. These values should also be 8.9% higher, because they scale with the standard deviation. The measured values are 11.5%, 7.6%, and 9.0%, respectively. Given the deviations from linear camera response and a linear photon transfer curve, the agreement is very good.

4.3 Test of simplified LUT Version 1, implemented in pco.edge

In the first version of a new firmware, PCO implemented a simplified version of the nonlinear LUT with fixed gain and no initial linear part:

$$h = \sqrt{g - g_{\mathrm{offs}}}. \tag{23}$$

A comparison to Eq. (21) gives

$$\sigma_h = \frac{\sqrt{K}}{2} \quad \text{and} \quad g_0 = g_{\mathrm{offs}} + \frac{\sqrt{\sigma_0^2}}{K}. \tag{24}$$

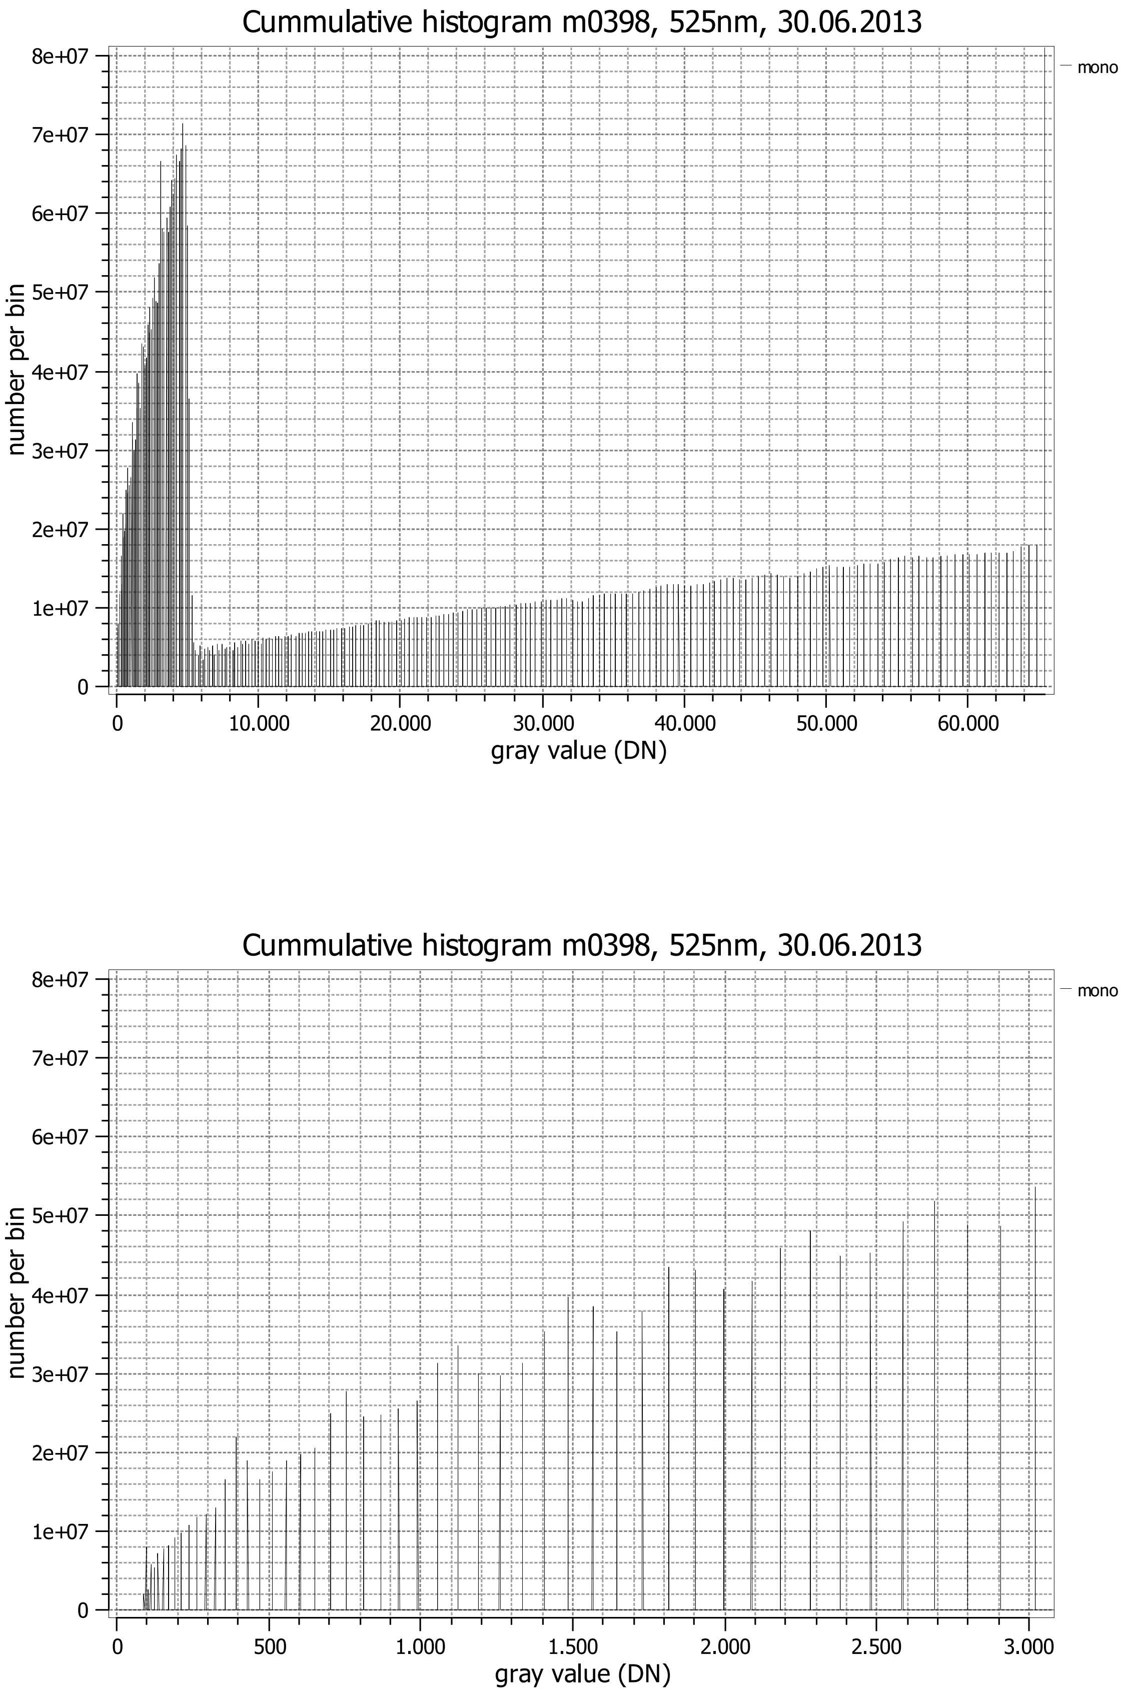

Figure 9: *Cumulative histogram of an EMVA 1288 measurement using 200 low-irradiation steps and 200 high-irradiation steps. Only 256 values in the 16-bit space are occupied after decompression:* ***a*** *full range and* ***b*** *small section with low irradiation.*

a

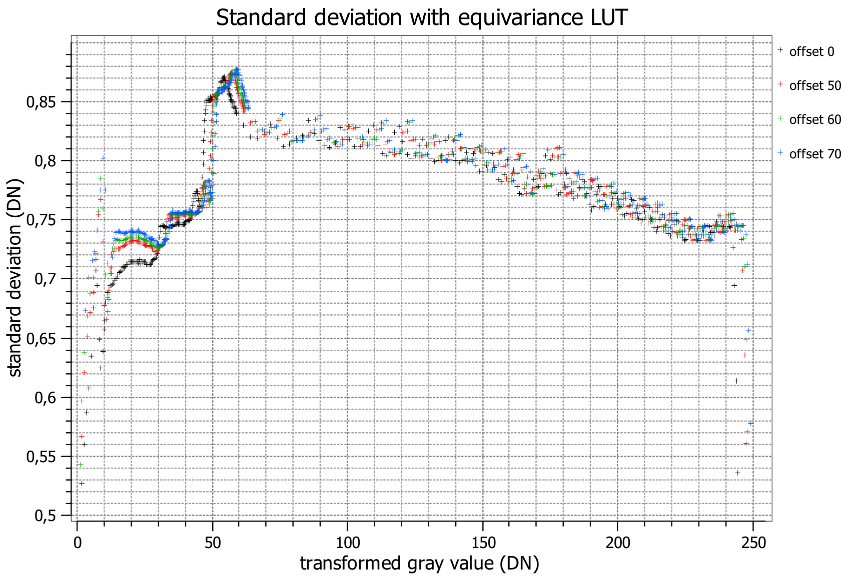

b

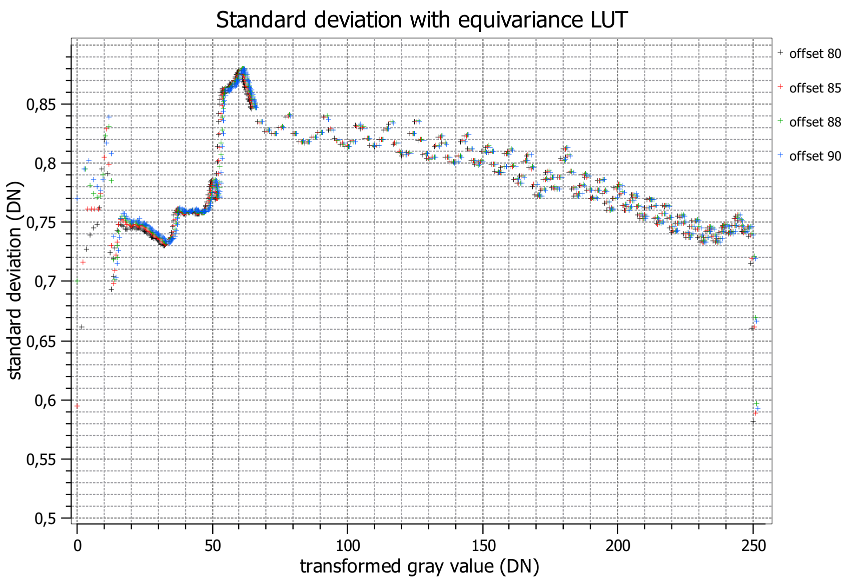

Figure 10: *Temporal noise in the compressed signal h with the simplified LUT and offsets as indicated.*

This means that the simplified LUT correctly maps the exact equation with the exception that the gain is fixed. The question remains, however, whether the missing linear part of Eq. (21), which was introduced to handle values below the mean dark values, remains a problem. From Eq. (24) it is obvious that no values lower than σ_0^2/K below the mean dark value can be handled. With a typical $K = 2$ and $\sigma_0^2 = 4$, this is not sufficient. Thus the only way to artifically lower the offset value.

For this purpose, a series of EMVA 1288 measurements applying the nonlinear LUT according to Eq. (23) with different offsets. As expected, the dependence of the estimated standard deviation is limited to low mean gray values (Fig. 10). If no offset is used, the standard deviation drops to values of 0.5 and lower, which is not acceptable. But for offsets in the range between 60 to 90, the dependency is really weak.

More significant effects can be seen for the DSNU histograms in Fig. 11. All distributions become oblique and the computed spatial DSNU standard deviations strongly depend on the chosen offset. As compared to the ideal LUT (Fig. 11a), offset values between 80 and 85, i.e., about 10 to 15 gray values below the mean dark value seem to be best.

In summary, the simplified LUT seems to be a good compromise between functionality and complexity. The only missing feature is a variable gain for a better sampling with lower additional quantization noise if for an application the full range of gray values is not required or used. Therefore the following modification is suggested:

$$h = a\sqrt{g - g_{\text{offs}}}, \tag{25}$$

where a is a gain factor with values between one and a given maximum value. An efficient FPGA implementation could be

$$h' = \sqrt{g - g_{\text{offs}}} \quad , \quad h = h' + [(bh') >> 3] \quad \text{with} \quad 0 \leq b < 256. \tag{26}$$

Thus $b = 8(a - 1)$ would be an 8-bit multiplication factor for a maximum gain of close to 33. The gain $a = 1 + b/8$.

5 Conclusions and Outlook

Noise equalization proofs to be a valuable processing step to represent images with fewer bits even for cameras with a high dynamic range. Even images from sCMOS cameras such as the pco.edge can be compressed to 8 bits. The price to be paid is only that dark noise, maximum SNR and dynamic range are at most 10% higher. If a smaller gray value range is required for low-light imaging applications, a higher value of σ_h can be chosen and these effects become much lower. If, for instance, only 14 bit are required (half the maximum SNR), $\sigma_h = 1.34$ can be chosen, and the increase in the standard deviation of dark noise and the decrease in the maximum SNR and dynamic range will be reduced to a negligible 2.3%.

Therefore, it is proposed to alter the firmware of the pco.edge in such a way that the 16-bit LUT with 8-bit values can be loaded into the edge. The following steps are required to compute this LUT:

1. Determine the gray value range, $g_{\max}$, required by the application.
2. Perform an EMVA 1288 measurement of the target camera to determine the dark noise value σ_0 and the gain K in the operation mode, the camera should be used.
3. Choose a value for m, the range of values below the mean dark value, measured in units of σ_h. Typically, a value of $m = 6$ should be appropriate.
4. Use Eq. (20) to compute σ_h.
5. Use Eq. (21) to compute the LUT.
6. If the inverse LUT is required, use Eq. (22) to compute it.

6 Bibliography

[1] Franke, R. Datenkompression mit Hilfe der Wurzelfunktion. Technical report, PCO AG, Kelheim, 2010. 2

[2] Förstner, W. Image preprocessing for feature extraction in digital intensity, color and range images. In Dermanis, A., A. Grün, and F. Sanso, eds., *Geomatic Methods for the Analysis of Data in the Earth Sciences*, vol. 95 of *Lecture Notes in Earth Sciences*. Springer, Berlin, 2000. 4

[3] Jähne, B. *Digitale Bildverarbeitung*. Springer, Berlin, 6th edn., 2005. 4

[4] PCO AG. 12 zu 8 Bit Komprimierung. Technical report, PCO AG Kelheim, 2006. 2

[5] Widrow, B. and I. Kollar. *Quantization Noise — Roundoff Error in Digital Computation, Signal Processing, Control, and Communications*. Cambridge Univ Press, 2008. URL www.cambridge.org/9780521886710. 3, 4

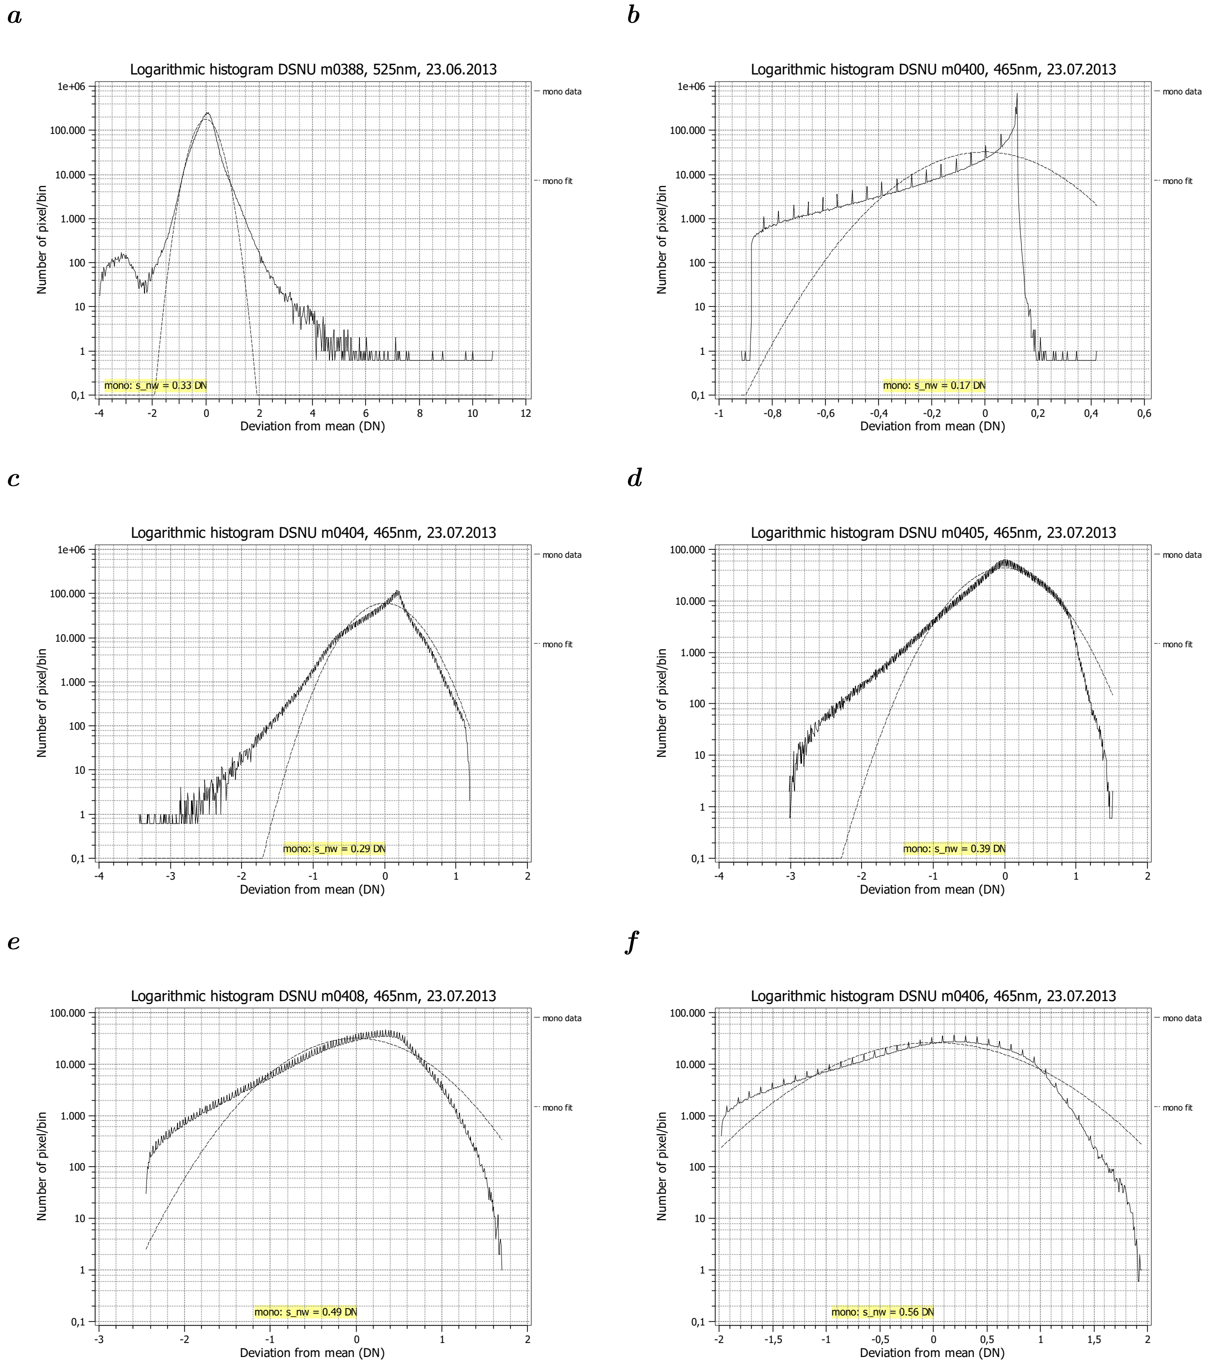

Figure 11: *Histogram of the dark signal nonuniformity with **a** the LUT according to Eq. (21) and the LUT according to Eq. (23) with offsets **b** 0, **c** 80, **d** 85, **e** 88, and **f** 90. The mean dark value is about 96.3.*