

UC Santa Cruz

UC Santa Cruz Previously Published Works

Title

Empirical Characterization of Camera Noise

Permalink

<https://escholarship.org/uc/item/76n7n6ck>

Authors

Baumgartner, Jeremy

Hinsche, Markus

Manduchi, Roberto

Publication Date

2013-06-01

Empirical Characterization of Camera Noise

Jeremy Baumgartner¹, Markus Hinsche², and Roberto Manduchi¹

¹University of California, Santa Cruz, CA
{jbaumgar,manduchi}@ucsc.edu

²Hasso Plattner Institute, Potsdam, Germany
markus.hinsche@student.hpi.uni-potsdam.de

Abstract. Noise characterization is important for several image processing operations such as denoising, thresholding, and HDR. This contribution describes a simple procedure to estimate the noise at an image for a particular camera as a function of exposure parameters (shutter time, gain) and of the irradiance at the pixel. Results are presented for a Pointgrey Firefly camera and are compared with a standard theoretical model of noise variance. Although the general characteristic of the noise reflects what predicted by the theoretical model, a number of discrepancies are found that deserve further investigation.

1 Introduction

The quantitative estimation of image noise is critical for basic operations such as denoising [5], thresholding [6], and HDR [2]. The simplifying assumption of “uniformly distributed Gaussian noise” in images is well known to be incorrect: for the same camera, the statistical characteristics of noise depend on the exposure parameters as well as on the irradiance received by the pixel under consideration. Theoretical noise models and procedures for noise parameter estimation have been described by several authors [3, 7, 4, 8]. In general, previous work either assumes access to raw data from the sensor, or tries to “reverse engineer” the image signal processor (ISP) that performs operations such as gamma correction, gamut mapping, and white point calibration, in order to estimate the “true” irradiance at a pixel and the noise characteristics of the acquisition process. Published work ranges from methods that assume a well-controlled illumination and reflection surface [3], to approaches that attempt to estimate relevant parameters from a single image [5, 1]. In this work we take an intermediate stance: we assume that a number of pictures of a stationary backdrop are taken with a number of different exposure settings, but make no particular assumption about the illumination and reflectance characteristics of the scene, except that they should remain constant during data acquisition. This can be easily achieved in a standard lab environment. For each exposure setting, pixels with similar mean value of brightness are pooled together for noise variance estimation. This procedure produces a characterization of camera noise as a function of the mean brightness value and of exposure parameters of interest (shutter time and gain). The results can be used to validate theoretical models

and compute relevant parameters, or could be stored as a look-up table to be queried when needed. This may be especially useful when using image data that has been already processed by the ISP with unknown non-linear characteristics.

This contribution is organized as follows. We first summarize in Sec. 2 the standard noise model for photodiode-based sensors. We then describe our simple experimental setup in Sec. 3 and our procedure for estimation of noise variance in Sec. 4. This procedure, which pools together pixels with similar mean value of brightness, is an original contribution of this work. We then present an analysis of the results vis-a-vis the theoretical model in Sec. 5. Sec. 6 has the conclusions.

2 Noise Model

In this section we summarize the standard camera noise model [3]. Let P be the power of light impinging on a pixel (equal to the irradiance at the pixel multiplied by the pixel area times the fill factor). Light generates a photocurrent I , equal to P times the quantum efficiency of the photodiode. Due to photon noise, I should be regarded as a Poisson process with rate parameter \bar{I} . Dark current I_{DC} also contributes to the charge accumulated at the pixel. This can also be modeled a Poisson process with rate parameter \bar{I}_{DC} . We will assume the I and I_{DC} are independent.

The charge accumulated at a pixel is equal to $Q = (I + I_{DC}) \cdot T$, where T is the shutter time. Since Q is the sum of two independent Poisson variables, it is itself distributed as a Poisson variable with mean $(\bar{I} + \bar{I}_{DC}) \cdot T$ and variance equal to its mean. This charge is then transferred to the amplifier. The main purpose of the amplifier is to increase the range of the signal that is passed to the quantizer, thus increasing the quantization SNR. Read-out (thermal) noise n_{RO} is generated by the amplifier; this noise is usually modeled as a zero-mean Gaussian variable with variance σ_{RO}^2 . If G is the amplifier gain, the amplified signal is thus equal to $((I + I_{DC}) \cdot T + n_{RO}) \cdot G$. The amplified signal is then quantized with N bits. The quantization process can, in first approximation, be modeled by the addition of a uniform random variable with zero mean and variance equal to $\sigma_{qt}^2 = \Delta^2/12$, where Δ is the quantization interval. Note that, referring to the measured values represented as integers, the quantization interval is equal to 1, and thus the quantization noise variance of the measured values (“brightness”, denoted by B) is equal to $1/12$. Summarizing, the brightness value B measured at a pixel is a random variable with mean \bar{B} equal to:

$$\bar{B} = (\bar{I} + \bar{I}_{DC}) \cdot TG \quad (1)$$

and variance σ_B^2 equal to:

$$\sigma_B^2 = (\bar{I} + \bar{I}_{DC}) \cdot TG^2 + \sigma_{RO}^2 \cdot G^2 + \frac{1}{12} = \bar{B} \cdot G + \sigma_{RO}^2 \cdot G^2 + \frac{1}{12} \quad (2)$$

where we have assumed that the sources of noise are uncorrelated with each other. Note that the read-out noise is expressed in units of measured values. It is noteworthy that the noise on the measured values is independent of the shutter time.




Fig. 1. Our experimental setup.

3 Experimental Setup

We devised a data collection procedure that can be easily replicated in standard lab settings, without the need for expensive equipment or for equipment that requires accurate calibration. The goal is to collect a large number of images of a stationary background with different values of shutter time and gain. We do not attempt to measure the irradiance at each pixel directly, nor do we expect that, for a given picture, the irradiance is uniform across the image. In fact, due to non-uniform illumination, shading, off-axis attenuation and other factors, a highly non-uniform distribution of irradiance values should be expected for a given picture. The main requirement for our data collection is that the illumination of the (stationary) background being imaged is constant with time. In practice, this means using an artificial light source, and avoiding other sources of possibly changing light as well as shadows due to moving people in the room. For our experiments, we placed the camera inside a cardboard box as shown in Fig. 1. A lamp illuminated the inside of the box through a hole in one side of the box. The lamp was left on for a certain amount of time before data collection to ensure stable illumination.

We used a Firefly MV 0.3 MP Color USB 2.0 camera with a Fujinon 2.8mm lens. This camera uses an Aptina MT9V022 CMOS sensor. Data from the sensor is first adjusted for gain and brightness (bias) and then quantized at 10 bits, after which gamma correction is applied followed by pixel correction, demosaicing and white balance, before quantization to 8 bits. We should note that the Firefly camera give the option of accessing raw data; we decided to use the 8-bit, demosaiced output as this is the format most commonly used in practical applications. However, we disabled the gamma correction and white balance in our tests in order to obtain a linear and scene-independent response. Color data was transformed to greyscale (luminosity) using the standard conversion formula.

We collected 50 images of the scene for several combinations of shutter time T and gain G . We considered values of T from 3 ms to 13 ms in steps of 2 ms, and values of G from 0 dB to 12 dB in steps of 2 dB. Overall, 2100 pictures were collected.

4 Variance Estimation Procedure

The goal of this work is to estimate the variance of the measured value B as a function of its mean value \bar{B} , shutter time T , and gain G . Inspired by [8], we use a “pooling” approach, whereby pixels that are assumed to receive the same irradiance are grouped together to compute useful statistics. Unlike [8], though, we don’t pool together pixels in an image with similar values of B . Instead, for each choice of T and G , we first estimate the mean value \bar{B} at each pixel by averaging over the 50 images taken for that (T, G) pair. The mean value \bar{B} is again quantized to 8 bits. For each one of the 256 possible values of \bar{B} , we compute the location of pixels that take on that particular value of \bar{B} . This set of pixels are assumed to receive the same irradiance, and thus to be samples of the same random variable. The sample variance computed from this pool of pixels over the 50 images represents our estimation of the variance σ_B^2 . In order to reduce the effect of estimation variance, only values of \bar{B} for which the pool contains at least 10 pixels (500 values overall) are considered for this computation.

5 Analysis of Results




Fig. 2. The average value \bar{B} at a pixel as a function of shutter time (left) and gain (right). Note that the gain is expressed on a linear scale. The dots represent measured values, while each line is a least square linear fit. Each line in the left (right) plot corresponds to one specific value of gain (shutter time).

5.1 Validation of Linearity Assumption

According to Eq. (1), the mean measured value \bar{B} at a pixel should be proportional to both shutter time T and gain G . In order to validate this assumption, we plotted in Fig. 2 the mean value \bar{B} of one pixel in the image (averaged over all 50 frames taken for each exposure setting) as a function of T and G . (Note that, for the pixel selected for these plots, the value \bar{B} at the maximum gain and

shutter time was equal to 95.) For each plot (representing a fixed value of G or T), we also computed and plotted the least squares linear fit.

From Fig. 2 it is clear that, for fixed G , the function $\bar{B}(T, G)$ is linear with non-null intercept. Least-squares fitting with a common intercept for all values of G results in an intercept of $T_0 = 2.95$ ms with reasonably small residual. Thus, the linear form in Eq. (1) represents an appropriate model if one substitutes T with $T - T_0$. The characteristic $\bar{B}(T, G)$ for fixed T is also linear with non-null intercept, but least-squares estimation of a common intercept produces fairly large residuals.

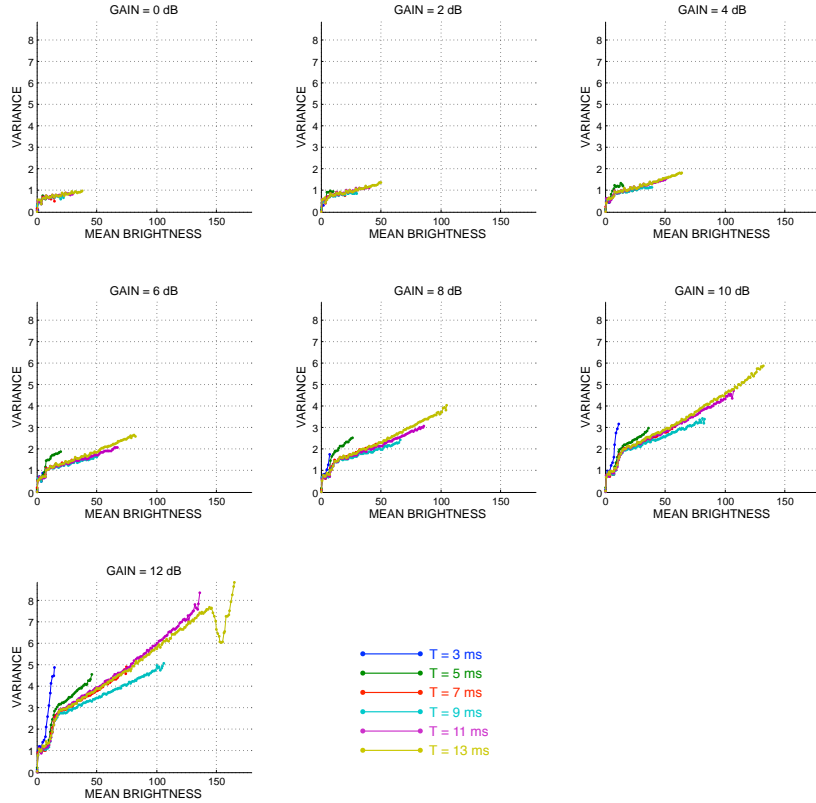


Fig. 3. The variance of B as a function of the mean brightness \bar{B} . Each subfigure represents a different value of the gain G ; each curve corresponds to a different value of the shutter time T .

5.2 Noise Variance as a Function of \bar{B}

According to Eq. (2), for a given value of the gain G , the variance of the measured values B should be a linear function of \bar{B} with intercept on the Y axis at $\sigma_{RO}^2 G^2 +$

1/12. In Fig. 3 we plotted the measured values of σ_B^2 as a function of \bar{B} for all values of G and T . (Note that, when the gain and/or shutter time are small, then \bar{B} can only take small values.) The plots show that, for each value of the gain G , the characteristic of σ_B^2 as a function of \bar{B} is indeed approximately linear with intercept for \bar{B} larger than a certain value. However, when \bar{B} is smaller than this value, the variance drops. This phenomenon could be explained in part by the “floor bias”: since B cannot take on negative values, the noise distribution is skewed for small \bar{B} . In addition, the linear slope seems to have an unexpected dependency on the shutter time. In particular, for small values of T (3-5 ms), the slope is quite large; for $T = 7$ ms, 11 ms and 13 ms the slope is constant, while for $T = 9$ ms the slope is noticeably smaller. Another strange inconsistency can be noticed for values of \bar{B} around 150 with $T=13$ ms and $G=12$ dB, where the curve of σ_B^2 has an unexpected dip.

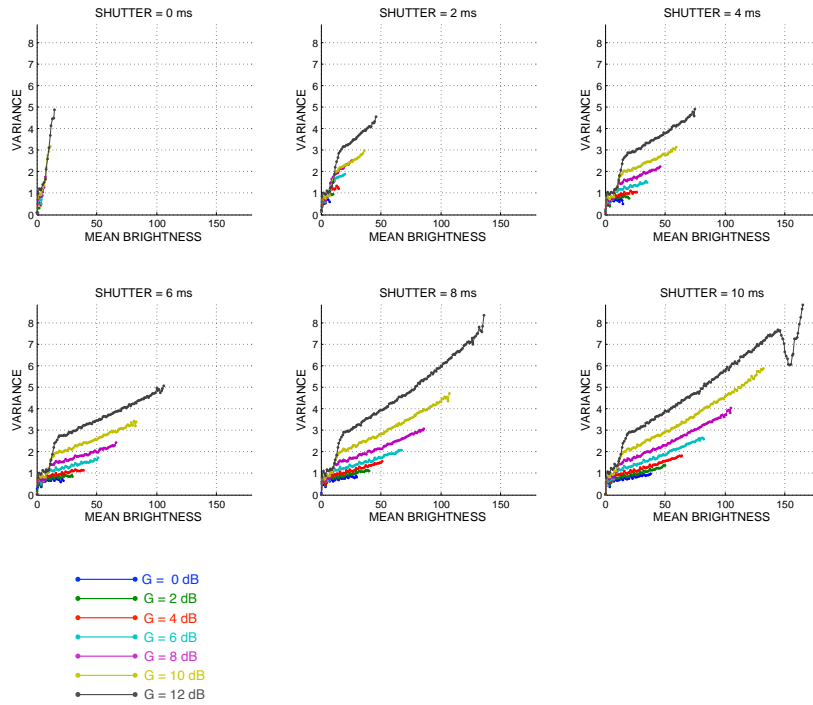


Fig. 4. The variance of B as a function of the mean brightness \bar{B} . Each subfigure represents a different value of the shutter time T ; each curve corresponds to a different value of the gain G .

Fig. 4 plots the same data but grouping together curves with the same shutter time. As expected, the slope and intercept of the linear part of each plot depend on the gain G .

5.3 Noise Variance as a Function of T

As noted earlier, the shutter time T does not appear in the expression of the variance of B (2). Hence, for fixed \bar{B} and G , the characteristic of σ_B^2 as a function of T is expected to be flat. Fig. 5 plots σ_B^2 against T for different values of \bar{B} and G . It can be seen that the characteristic is only approximately flat, with noticeable variations around the mean value.

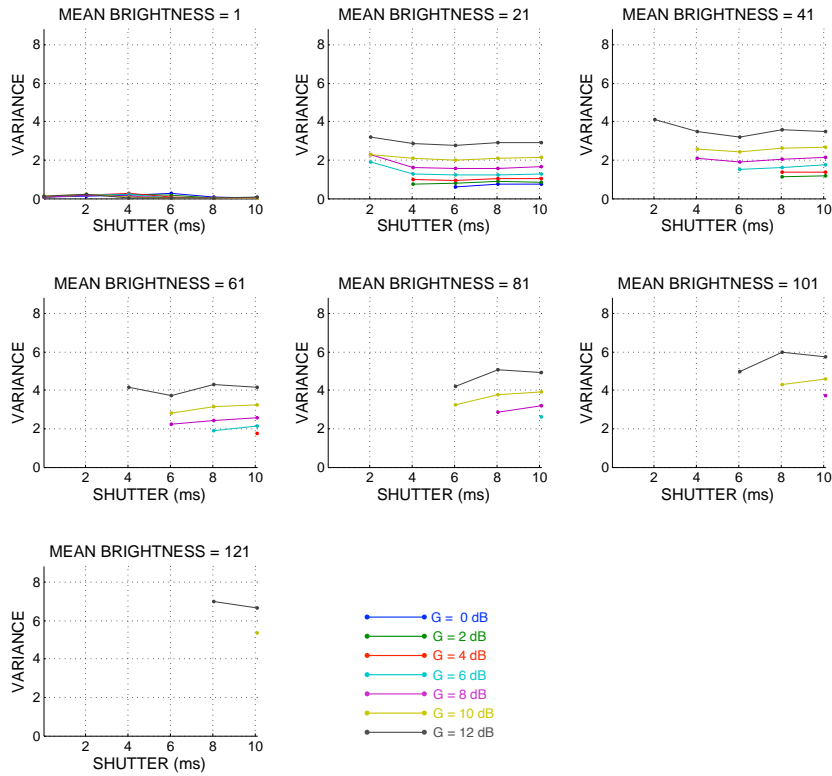


Fig. 5. The variance B as a function of the shutter time T . Each subfigure represents a different value of the mean brightness \bar{B} ; each curve corresponds to a different value of the gain G .

5.4 Noise Variance as a Function of G

Each subfigure in Fig. 6 shows the variance of B as a function of the gain G for different values T and for a fixed value of \bar{B} . According to Eq. (2), these plots should overlap with each other (since they are independent of T) and should have a parabolic characteristic. It is seen that the plots do tend to overlap, except for $T=5$ ms and $T=9$ ms. A light concavity can be observed, consistent with the theory, although the linear term $\bar{B}G$ in Eq. (2) dominates in this interval of values. Note that in both Fig. 5 and 6, the variance σ_B^2 takes on very small values for small B ; this phenomenon was already observed and discussed earlier in relation to Figs. 3 and 4.

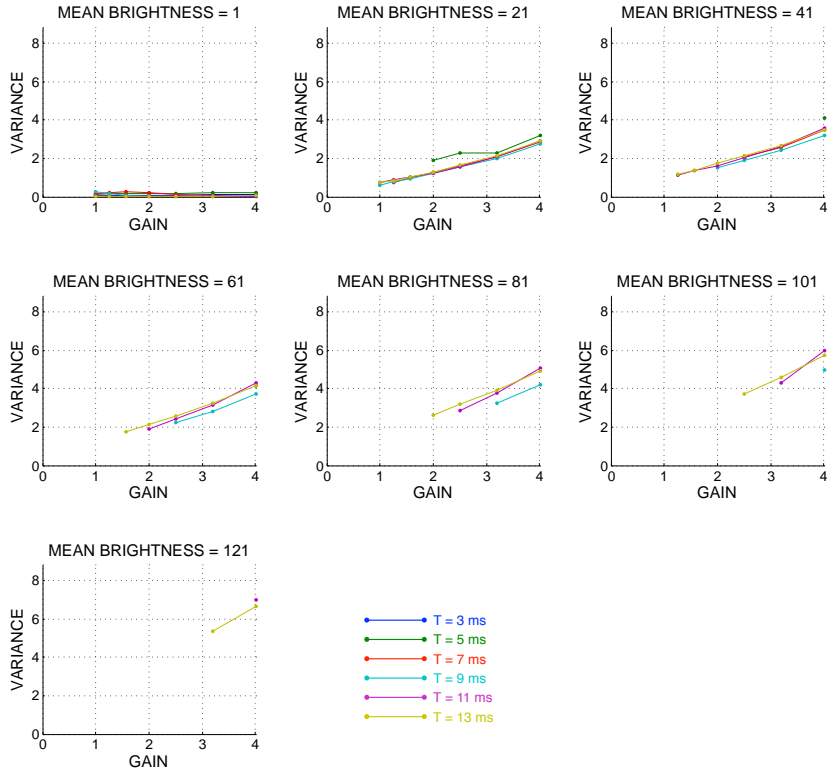


Fig. 6. The variance B as a function of the the gain G (shown in linear scale). Each subfigure represents a different value of the mean brightness \bar{B} ; each curve corresponds to a different value of the shutter time T .

6 Conclusions

We have described a simple approach to computing the noise characteristics at a pixel as a function of the mean brightness value at the pixel and of the exposure parameters. This procedure can be easily reproduced without the need for calibrated illuminators or uniformly reflective material, and thus can be used to quickly characterize different camera models. We compared the characteristics of the mean brightness value and of the noise variance as a function of exposure parameters against the theoretical model. Our analysis brought to light a number of discrepancies that deserve further study, thus confirming the importance of direct empirical measurements for correct camera modeling.

Acknowledgment

This material is based upon work supported in part by the National Science Foundation under Grant No. IIS-083564. The authors would like to thank Dr. Orazio Gallo for useful advice.

References

1. A. Foi, M. Trimeche, V. Katkovnik, and K. Egiazarian. Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing*, 17(10):1737–1754, 2008.
2. M. Granados, B. Ajdin, M. Wand, C. Theobalt, H.-P. Seidel, and H. Lensch. Optimal HDR reconstruction with linear digital cameras. In *Proc. CVPR*, 2010.
3. G. Healey and R. Kondepudy. Radiometric CCD camera calibration and noise estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(3):267–276, 1994.
4. J. Janesick. *Scientific charge-coupled devices*. SPIE Press, 2001.
5. C. Liu, R. Szeliski, S. Kang, C. Zitnick, and W. Freeman. Automatic estimation and removal of noise from a single image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 30(2):299–314, 2008.
6. P. Rosin. Thresholding for change detection. In *International Conference on Computer Vision*, pages 274–279, 1998.
7. H. Tian. *Noise Analysis in CMOS Image Sensors*. PhD thesis, Stanford University, 2000.
8. Y. Tsin, V. Ramesh, and T. Kanade. Statistical calibration of CCD imaging process. In *IEEE International Conference on Computer Vision (ICCV'01)*, July 2001.