

QUANTILE ANALYSIS OF IMAGE SENSOR NOISE DISTRIBUTION

Jiachao Zhang, Keigo Hirakawa*

Xiaodan Jin

University of Dayton
Dayton, OH 45469 USA
{zhangj33, khirakawa1}@udayton.edu

BOE Technology Group Co.
Beijing, 100176 China
jinxiaodan@boe.com.cn

ABSTRACT

This paper describes a study aimed at comparing the *real* image sensor noise distribution to the models of noise often assumed in image denoising designs. Quantile analysis in pixel, wavelet, and variance stabilization domains reveal that the tails of Poisson, signal-dependent Gaussian, and Poisson-Gaussian models are too short to capture real sensor noise behavior. Noise model mismatch would likely result in image denoising that undersmooths real sensor data.

Index Terms—image denoising, image sensor, Poisson

1. INTRODUCTION

Image sensor noise is present in all commercial, professional, and scientific cameras. Unlike fixed pattern and banding noises that are largely predictable and therefore correctable [2,3], random phenomena such as photon emission, photon transfer/recapture, dark current, thermal noise, reset noise, quantization, etc. introduce sensor measurement uncertainties [4]. Although an analytical form of noise distribution is unknown, noise distribution *model* plays a critical part of image denoising, aimed at computationally reversing the effects of the degradation caused by random noise [5–16].

We provide a comprehensive study of the discrepancy between the *model* of noise commonly used in image denoising algorithms and the distribution of real sensor noise acquired by real sensor hardware. There are major differences between our noise model validation study and the many existing empirical studies [17–23]. First, existing studies overwhelmingly focused on the relationship between the pixel intensity and noise variance, but there has been little emphasis on the tail behavior of the noise distribution that greatly influences denoising performance. Second, most modern image denoising techniques incorporate linear transformations that give rise to sparse signal representation and nonlinear transforms that decouple noise and signal. As such, we developed an analytical technique aimed at scrutinizing the noise models in the transform domains, rather than focusing exclusively on the pixel noise distribution model as the previous investigations have done. Our investigation provides an insight to improving the image denoising algorithms.

2. EXPERIMENTAL SETUP

2.1. Data Acquisition

We obtained samples of noisy sensor data by capturing X-Rite ColorChecker (Figure 1(a)) at 3.4 lux using Nikon D90, Canon 550D,

*Based upon work supported in part by the National Science Foundation under Grant No. CCF-1422104. The work was completed while X. Jin was a research fellow at Harvard University, Cambridge, MA 02138. This paper extends our preliminary results reported in [1].

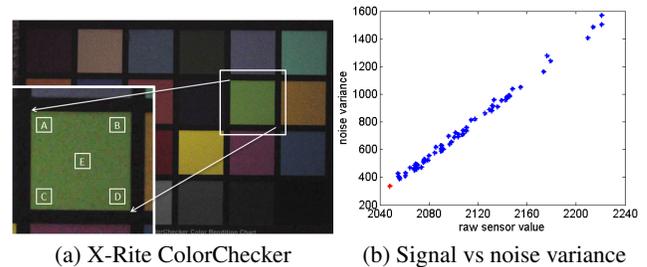


Fig. 1. (a) ANOVA over regions A-E ensure uniformity of each ColorChecker patch. (b) Noise variance scales linearly with signal. Red dot shows electronic noise (see text).

and Fuji Pro 1 in raw sensor format with 1/200 second exposure (1/250 for Fuji), ISO 1600, and f/4.0. Color filter array sampled data is partitioned into red, blue, and green measurements, which are treated separately (i.e. $3 \times 24 = 72$ ColorChecker patches). Under an ideal scenario, measured pixel component values from the same ColorChecker patch are drawn from the same probability distributions.

Despite our best efforts, however, uneven lighting, vignetting, and camera angle introduce additional variabilities. For this reason, we detect non-uniformity of ColorChecker patches by the *analysis of variance* (ANOVA) over five 10×10 regions (labeled A-E in Figure 1(a)) cropped from each ColorChecker patch. Any ColorChecker patch that rejects the null hypothesis (i.e. means of A-E are equal) at the 99% confidence level is removed from the experiment. We also removed green pixels from blue-green rows because color crosstalk contaminations affect green pixels in red-green and the blue-green rows differently [24]. Each accepted ColorChecker patch has over 25,000 samples.

We also acquired another image (with the same camera settings) with a cap placed over the lens. By blocking the incoming light, this “blank” image offers an indication for the circuit noise that is independent of the signal strength.

2.2. Heteroscedastic Pixel Noise Model

Empirical mean and variance for each ColorChecker patch are shown in Figure 1(b). Assuming that fixed pattern and banding noises are insignificant, this figure presents a convincing evidence that noise variance scales linearly with signal strength. The linearity is usually attributed to the Poisson process of the photon emission in the literature [9–19, 25]. Under this scenario, a sensor observation h is modeled as h_P , an affine transformed Poisson count data g_P :

$$h_P := \alpha \cdot g_P + \beta, \quad g_P \sim \mathcal{P}(f_P), \quad (1)$$

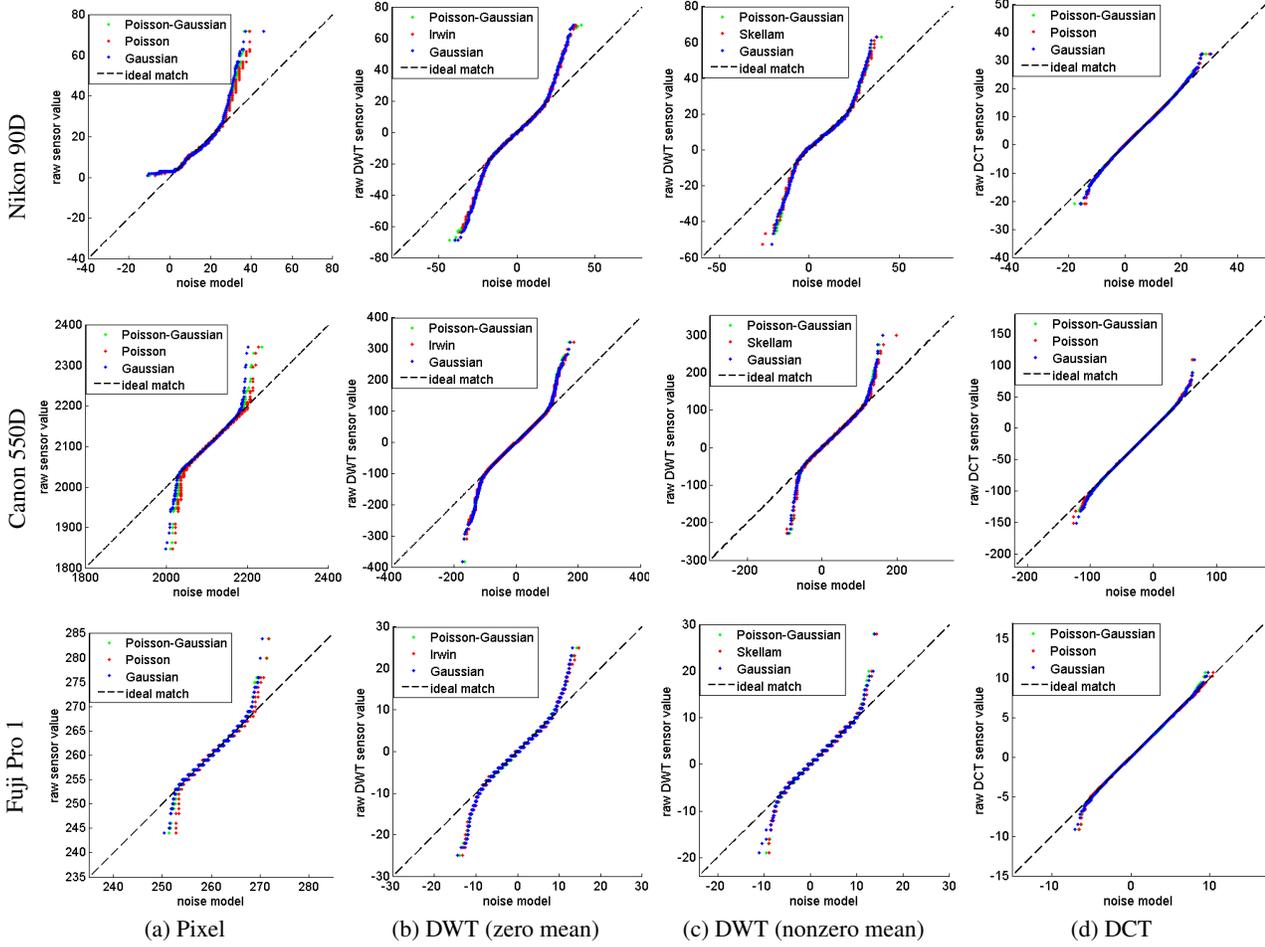


Fig. 2. QQ plot comparing the distribution of the sensor measurements to the noise models of (1-3).

where f_P is the latent intensity, and the subscript P denotes Poisson-based model. The parameters α and β are learned from regressing signal strength and variance in Figure 1(b).

Another way to capture the signal dependence of noise is to couple the variance of a normal random variable [23]:

$$h_G \sim \mathcal{N}(\alpha f_G + \beta, \alpha^2 f_G). \quad (2)$$

where the subscript G denotes signal-dependent Gaussian noise model ((2) is the normal approximation of h_P).

Poisson-Gaussian hybrid model treats signal-dependent and signal-independent noises separately [5, 6, 21, 26]. The observation $h_H = \alpha g_S + g_C$ is a combination of signal g_S and circuit g_C noise:

$$g_S \sim \mathcal{P}(f_S), \quad g_C \sim \mathcal{N}(\mu_C, \sigma_C^2), \quad (3)$$

where (μ_C, σ_C^2) is the signal-variance pair of lens cap image (red dot in Figure 1(b)). Subscript H denotes hybrid model.

3. NOISE UNDER LINEAR TRANSFORMATION

3.1. Pixel Noise

The noise models above are heuristic approximations at best, and model discrepancies deteriorate the image denoising performance. We employ quantile analysis to robustly compare the distribution of

sensor data and the noise model. Consider a parametric curve of the form [27]:

$$x(t) = F_{data}^{-1}(t), \quad y(t) = F_{model}^{-1}(t) \quad (4)$$

where $F_{data}, F_{model} : \mathbb{R} \rightarrow [0, 1]$ are the cumulative distribution functions of data and model distribution, respectively. This so-called quantile-quantile plot (QQ plot) lies on the 45° line if the empirical data and noise model distributions are well-matched. If the data variables are found to be an affine transformation of the model variables, then the QQ plot forms an affine line as well (but not necessarily on 45° line with zero intercept). QQ plot is useful for detecting the deviations of the model from data, particularly in the tails of distributions where the samples are sparse.

The QQ plots shown in Figure 2(a) compare the distribution of the measured pixel data within a ColorChecker patch against the models in (1-3). For $F_{data}(h)$, the empirical histogram of the pixels measured within each ColorChecker patch was used. For $F_{model}(h)$, we derived the model parameters for each ColorChecker patch by:

$$f_P = f_G = \frac{\mu_H - \beta}{\alpha}, \quad f_S = \frac{\mu_H - \mu_C}{\alpha}, \quad (5)$$

where μ_H is the sample mean of each ColorChecker patch, and α, β and μ_C are as described in Section 2.2.

Each QQ plot describes the variation of real and model noise within *one ColorChecker patch* only.¹ As evidenced by the 45° line formed by a portion of the QQ plot, the noise model is accurate near the median. However, sensor measurements are clearly more heavy tailed than the model (Nikon D90's short negative tail is likely due to saturation).

3.2. Discrete Wavelet Transform

Though we are interested in the impact of linear transforms on the noise, quantile analysis in the transform domain is challenging. With the precise form of real noise distribution unknown, an analytical form of noise distribution in discrete wavelet (DWT) and cosine (DCT) transform domains cannot be derived. Yet, unlike the pixel domains, it is difficult to obtain a large number of *real* noisy DWT/DCT coefficients drawn from the same distribution. We rejected the idea to take the raw sensor data from a video sequence of a scene—although applying DWT/DCT to each frame yields a large number of noisy coefficients, noise is unnaturally coupled with the temporal hysteresis of reset noise.

We developed a new strategy to obtain a large number of noisy coefficients from the ColorChecker image. Consider Haar wavelet transform (HWT)—finest level noisy wavelet ($w(n)$) and scaling ($s(n)$) coefficients at location n are

$$\begin{aligned} w(n) &= h(2n) - h(2n + 1) \\ s(n) &= h(2n) + h(2n + 1) \end{aligned} \quad (6)$$

where $h(2n)$ and $h(2n + 1)$ are neighboring pixels. Thanks to sparsity, the majority of DWT coefficients have zero mean:

$$\mathbb{E}w(n) = 0 \quad \Leftrightarrow \quad \mathbb{E}h(2n) = \mathbb{E}h(2n + 1). \quad (7)$$

Hence $h(2n)$ and $h(2n + 1)$ are assumed to be drawn from the same distribution. We obtain a large number of $w(n)$ samples corresponding to a *mean zero* DWT coefficient by taking a difference between two observed samples *drawn at random* from the *same* ColorChecker patch. By contrast, coefficients with nontrivial mean have the property:

$$\mathbb{E}w(n) \neq 0 \quad \Leftrightarrow \quad \mathbb{E}h(2n) \neq \mathbb{E}h(2n + 1). \quad (8)$$

Hence $h(2n)$ and $h(2n + 1)$ are drawn from a different distributions. We may obtain a large number of $w(n)$ samples by taking a difference between two noisy samples *drawn at random* from two predesignated ColorChecker patches.

For analysis, $F_{data}(w)$ was computed from the noisy DWT coefficients obtained by the above scheme. For $F_{model}(w)$, the DWT noise model derived from (1-3) have the form:

$$\begin{aligned} \alpha^{-1}w_P &\sim \text{Skellam}(f_P, f'_P) \\ w_G &\sim \mathcal{N}(\alpha(f_G - f'_G), \alpha^2(f_G + f'_G)) \\ w_H &= \alpha w_S + w_C, \quad \begin{cases} w_S \sim \text{Skellam}(f_S, f'_S) \\ w_C \sim \mathcal{N}(0, 2\sigma_C^2). \end{cases} \end{aligned} \quad (9)$$

where $\{f_P, f_G, f_S\}$ and $\{f'_P, f'_G, f'_S\}$ are parameters derived from ColorChecker patches corresponding to $h(2n)$ and $h(2n + 1)$, respectively.² The QQ plots shown in Figure 2(b-c) compare the distribution of empirical DWT coefficients against their models. Though noise models are accurate near the median, the models clearly shorten the tails.

¹We cannot infer signal-noise dependence from the QQ plots as noise samples in each plot are drawn from the same distribution with same mean/variance. Only representative examples shown due to page limit.

²*Skellam*(f_P, f'_P) is known as Irwin distribution.

3.3. Discrete Cosine Transform

DCT is defined for $k \in \{0, 1, \dots, N - 1\}$ as follows [28]:

$$d(k) = \sum_{n=0}^{N-1} \frac{h(n)}{\sqrt{N}} \cos\left(\frac{\pi}{2N}(2n - 1)k\right). \quad (10)$$

2D DCT applies (10) to horizontal and vertical directions. For quantile analysis in DCT domain, *randomly drawn* sensor measurement samples from each of N predesignated ColorChecker patches (N^2 patches for 2D DCT) are respectively assigned to $\{h(0), \dots, h(N - 1)\}$ to yield a single DCT coefficient $d(k)$ via (10). Repeating this experiment yields a large number of DCT coefficients drawn from the same distribution and $F_{data}(d)$. For $F_{model}(d)$, *randomly drawn* samples of h_P in (1) from the *models of N predesignated ColorChecker patches* are assigned to $\{h(0), \dots, h(N - 1)\}$ to compute DCT coefficient $d_P(k)$. We followed the same procedure to yield $d_G(k)$ and $d_H(k)$ from (2-3).

The resultant QQ plot in Figure 2(d) suggests that the distribution of DCT coefficients stemming from the real sensor data is well approximated by the models. The improved match is likely due to central limit theorem, which is in force as a result of DCT in (10) taking weighted average of measurements $\{h(0), \dots, h(N - 1)\}$. When N is small, the deviation of tails is still observable (not shown).

4. NOISE UNDER VARIANCE STABILIZATION

Variance stabilization (VS) is an invertible function that recovers homoscedasticity given heteroscedastic noise. VS is often combined with additive white Gaussian noise (AWGN) image denoising to address camera sensor noise. Bartlett/Anscombe VS transforms Poisson counts into a normal variable [29, 30]:

$$\eta\{h_P\} = 2\sqrt{(h_P - \beta)/\alpha + k} \sim \mathcal{N}(2\sqrt{f_P + k}, 1), \quad (11)$$

where the constant value is $k = 1/2$ for Bartlett and $k = 3/8$ for Anscombe. Poisson-Gaussian hybrid h_H in (3) is also stabilized by the generalized Anscombe transform [5, 6]:

$$\eta'\{h_H\} = 2\sqrt{(h_H - \mu_C)/\alpha + 3/8 + \sigma_C^2/\alpha^2}. \quad (12)$$

However, one can prove $\eta\{h\} \equiv \eta'\{h\}$ by the fact that $\sigma_C^2 = \alpha\mu_C - \alpha\beta$ must hold in Figure 1(b). Haar-Fisz (HF) transform is a more contemporary VS treatment that modify Haar wavelet (w) and scaling (s) coefficients [31]:

$$\gamma\{h_P\} = \frac{w_P}{\sqrt{\alpha(s_P - 2\beta)}} \sim \mathcal{N}\left(\frac{f_P - f'_P}{\sqrt{f_P + f'_P}}, 1\right). \quad (13)$$

This procedure is repeated for coarser level wavelet representations.

Figure 3 compares the normal distribution models of (11-13) against the distributions of variance stabilized sensor measurements $\eta\{h\}$ and $\gamma\{h\}$ (and their DWT/DCT). Although Anscombe transform stabilized noise variance, the tails of empirical VS coefficients $\eta\{h\}$ are clearly longer than normality. The tails of the empirical variance stabilized DWT coefficients also deviate from normal probability. By contrast, HF stabilized coefficients $\gamma\{h\}$ appear to be normally distributed, as evidenced by straight QQ plot line. However, the gentle slope of QQ plot suggests that data variance is smaller than 1. By comparison, QQ plot followed the 45° line very closely when scene was well lit (not shown). We conclude that HF VS succeeds in “Gaussianizing” the observed data, but it did not achieve homoscedasticity (i.e. variance depends on signal intensity).

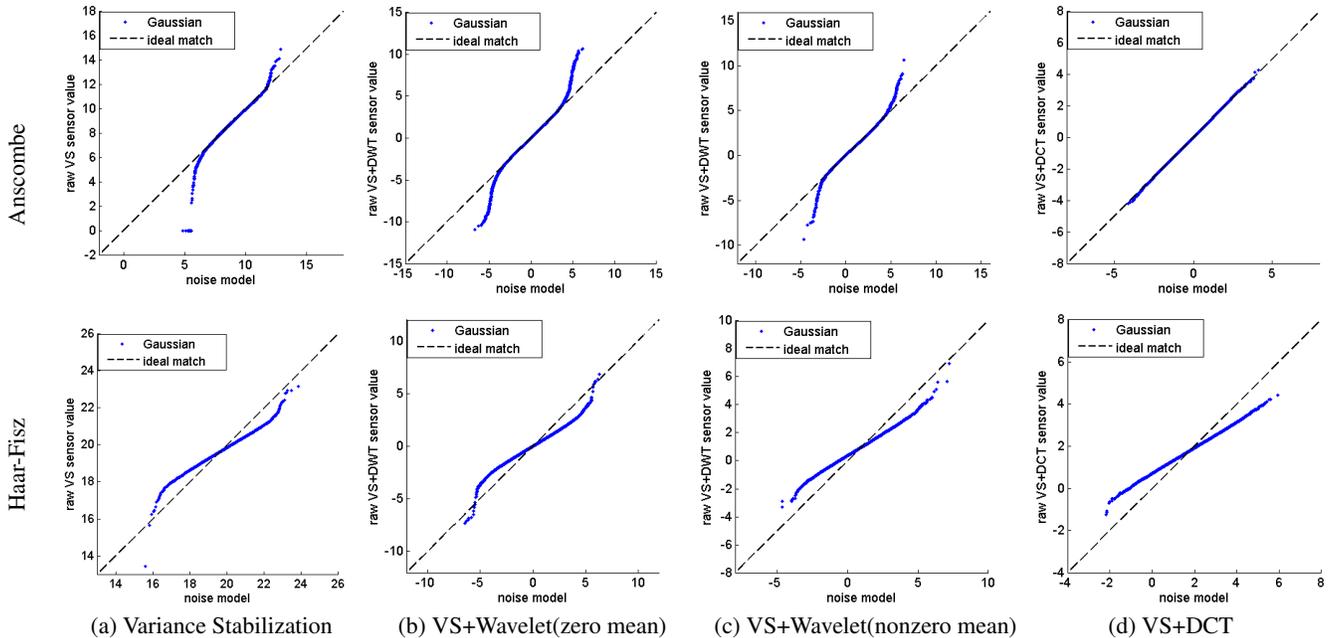


Fig. 3. QQ plot comparing the distribution of variance stabilized Canon 550D measurements to the noise models of (11-13).

5. DISCUSSIONS

Poisson, signal-dependent Gaussian, and Poisson-Gaussian distributions are commonly used to model the linear coupling between the signal strength and the noise variance in image sensors. These models provide the basis for image denoising methods that extend the imaging device’s capabilities. Our quantile analysis definitively proved that the tails of the noise models are too short to describe the actual distribution of the measurement noise. The trends we described in this paper are common among a variety of camera manufacturers, red/green/blue color components, and camera settings.

Our study is not without limitations. In Sections 3.2 and 3.3, random sampling of the pixels within Colorchecker patch were linearly combined to yield a large number of empirical DWT and DCT coefficients. This procedure is only valid if the random phenomena occurring in the spatially neighboring pixels are independent. Thermal noise, for example, is not always spatially white since electron leakage affects neighboring pixels. Hence the computed DWT/DCT coefficient noise in our study is noisier than the actual coefficients computed from an image sensor with significant leak. However, the main conclusions of this work—that the tail behavior of sensor noise is heavier than the models—remains valid. Our lab’s capabilities today do not allow for measurements with integrating spheres, which guarantees uniformity of the scene beyond our current ANOVA testing. Most commercial cameras have safety features that prevent pictures from being taken while the camera optics are removed, increasing the risks of vignetting.

What practical impact does the model mismatch play in image denoising? How should image denoising methods be improved for handling real image sensor data? Most modern image denoising methods operate in (linear) transform domain. However, we showed that the noise models also fail in Haar wavelet domain, where the model tails insufficiently account for large noise coefficients. The practical impact of the model mismatch is the undersmoothing of noise—when denoising algorithms designed with (1-3) in mind are applied to real sensor measurements, a large DWT coefficient is in-

correctly attributed to the signal since noise model does not account for it. Though improved image denoising techniques would certainly improve real sensor image denoising, the future image denoising designs would likely benefit also from incorporating a heavier tail likelihood functions—such as scale mixture of likelihood functions.

An alternatives is to combine VS transforms with AWGN image denoising. We concluded earlier that while Anscombe VS achieves homoscedasticity (i.e. noise variance decoupled from signal strength), the overall distribution profile is far from normality. In working with Anscombe VS, one would need to extend the noise model tails (e.g. scale mixture of AWGN) in order to improve denoising performance. On the other hand, Haar-Fisz VS successfully transformed sensor data into normal random variables, but the claim of the homoscedasticity could not be substantiated. Indeed, one is at risk of oversmoothing noise with conventional AWGN denoising method used in conjunction with Haar-Fisz VS, since noise variance in low light regions is lower than 1. One must first determine the true relationship between signal strength and noise variance in the Haar-Fisz transformed domain.

We also witnessed that model mismatch is less significant in DCT domain. Although it suggests that DCT-based denoising can be effective for real image sensor data, there are a few practical challenges. First, the analysis of noise heteroscedasticity in the DCT domain is far more complicated than that of wavelet noise. The combination of Anscombe VS and DCT (as was investigated in [5, 6], for example) is the only scheme in our investigation that satisfactorily yielded a homoscedastic model matching the observation distribution. Moreover, the averaging in DCT that belies the central limit theorem also results in a signal representation that is less sparse than DWT, which is disadvantageous for image denoising (though energy compaction is still in force).

6. REFERENCES

- [1] Xiaodan Jin and Keigo Hirakawa, “Approximations to camera sensor noise,” in *IS&T/SPIE Electronic Imaging*. International

- Society for Optics and Photonics, 2013, pp. 86550H–86550H.
- [2] Guy Meynants, Bart Dierickx, Dirk Uwaerts, and Jan Bogaerts, “Fixed pattern noise suppression by a differential readout chain for a radiation-tolerant image sensor,” in *proc. 2001 CCD & AIS workshop*, 2001, p. 52.
 - [3] Stephen C Cain, Majeed M Hayat, and Ernest E Armstrong, “Projection-based image registration in the presence of fixed-pattern noise,” *Image Processing, IEEE Transactions on*, vol. 10, no. 12, pp. 1860–1872, 2001.
 - [4] James R Janesick, “Photon transfer,” SPIE, 2007.
 - [5] M Makitalo and Alessandro Foi, “Poisson-gaussian denoising using the exact unbiased inverse of the generalized anscombe transformation,” in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 1081–1084.
 - [6] Markku Makitalo and Alessandro Foi, “Optimal inversion of the generalized anscombe transformation for poisson-gaussian noise,” *Image Processing, IEEE Transactions on*, vol. 22, no. 1, pp. 91–103, 2013.
 - [7] M. Raphan and E. P. Simoncelli, “Learning to be Bayesian without supervision,” in *Advances in Neural Information Processing Systems 19*, B. Schölkopf, J. Platt, and T. Hoffman, Eds., pp. 1145–1152. MIT Press, Cambridge, MA, 2007.
 - [8] Triet Le, Rick Chartrand, and Thomas J Asaki, “A variational approach to reconstructing images corrupted by Poisson noise,” *Journal of Mathematical Imaging and Vision*, vol. 27, no. 3, pp. 257–263, 2007.
 - [9] Robert Nowak and Eric D Kolaczyk, “A Bayesian multiscale framework for Poisson inverse problems,” in *Acoustics, Speech, and Signal Processing, 1999. Proceedings., 1999 IEEE International Conference on*. IEEE, 1999, vol. 3, pp. 1741–1744.
 - [10] Klaus E Timmermann and Robert D Nowak, “Multiscale modeling and estimation of Poisson processes with application to photon-limited imaging,” *Information Theory, IEEE Transactions on*, vol. 45, no. 3, pp. 846–862, 1999.
 - [11] Keigo Hiraoka and Patrick J Wolfe, “Skellam shrinkage: Wavelet-based intensity estimation for inhomogeneous poisson data,” *Information Theory, IEEE Transactions on*, vol. 58, no. 2, pp. 1080–1093, 2012.
 - [12] Florian Luisier, Thierry Blu, and Michael Unser, “Undecimated haar thresholding for poisson intensity estimation,” in *Image Processing (ICIP), 2010 17th IEEE International Conference on*. IEEE, 2010, pp. 1697–1700.
 - [13] Hugues Talbot, Harold Phelippeau, Mohamed Akil, and Stefan Bara, “Efficient poisson denoising for photography,” in *Image Processing (ICIP), 2009 16th IEEE International Conference on*. IEEE, 2009, pp. 3881–3884.
 - [14] Mário AT Figueiredo and José M Bioucas-Dias, “Restoration of Poissonian images using alternating direction optimization,” *Image Processing, IEEE Transactions on*, vol. 19, no. 12, pp. 3133–3145, 2010.
 - [15] Simon Setzer, Gabriele Steidl, and Tanja Teuber, “Deblurring Poissonian images by split Bregman techniques,” *Journal of Visual Communication and Image Representation*, vol. 21, no. 3, pp. 193–199, 2010.
 - [16] R. Willett, “Multiscale Analysis of Photon-Limited Astronomical Images,” *ASTRONOMICAL SOCIETY OF THE PACIFIC CONFERENCE SERIES*, vol. 371, pp. 247, 2007.
 - [17] H Joel Trussell and R Zhang, “The dominance of poisson noise in color digital cameras,” in *Image Processing (ICIP), 2012 19th IEEE International Conference on*. IEEE, 2012, pp. 329–332.
 - [18] Tamara Seybold, Christian Keimel, Marion Knopp, and Walter Stechele, “Towards an evaluation of denoising algorithms with respect to realistic camera noise,” in *Multimedia (ISM), 2013 IEEE International Symposium on*. IEEE, 2013, pp. 203–210.
 - [19] Tamara Seybold, Marion Knopp, Christian Keimel, and Walter Stechele, “Beyond standard noise models: Evaluating denoising algorithms with respect to realistic camera noise,” *International Journal of Semantic Computing*, vol. 8, no. 02, pp. 145–167, 2014.
 - [20] Hui Tian, Boyd Fowler, and Abbas E Gamal, “Analysis of temporal noise in cmos photodiode active pixel sensor,” *Solid-State Circuits, IEEE Journal of*, vol. 36, no. 1, pp. 92–101, 2001.
 - [21] Alessandro Foi, Mejd Trimeche, Vladimir Katkovnik, and Karen Egiazarian, “Practical poissonian-gaussian noise modeling and fitting for single-image raw-data,” *Image Processing, IEEE Transactions on*, vol. 17, no. 10, pp. 1737–1754, 2008.
 - [22] Ce Liu, Richard Szeliski, Sing Bing Kang, C Lawrence Zitnick, and William T Freeman, “Automatic estimation and removal of noise from a single image,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 2, pp. 299–314, 2008.
 - [23] Xinhao Liu, Masayuki Tanaka, and Masatoshi Okutomi, “Estimation of signal dependent noise parameters from a single image,” in *ICIP, 2013*, pp. 79–82.
 - [24] Keigo Hiraoka, “Cross-talk explained,” in *Image Processing, 2008. ICIP 2008. 15th IEEE International Conference on*. IEEE, 2008, pp. 677–680.
 - [25] Jin Xiaodan, Xu Zhenyu, and Keigo Hiraoka, “Noise parameter estimation for poisson corrupted images using variance stabilization transforms,” *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society*, vol. 23, no. 3, pp. 1329–1339, 2014.
 - [26] Florian Luisier, Thierry Blu, and Michael Unser, “Image denoising in mixed poisson-gaussian noise,” *Image Processing, IEEE Transactions on*, vol. 20, no. 3, pp. 696–708, 2011.
 - [27] Robert A Lodder and Gary M Hieftje, “Quantile analysis: a method for characterizing data distributions,” *Applied spectroscopy*, vol. 42, no. 8, pp. 1512–1520, 1988.
 - [28] Khalid Sayood and Jay C Borkenhagen, “Use of residual redundancy in the design of joint source/channel coders,” *Communications, IEEE Transactions on*, vol. 39, no. 6, pp. 838–846, 1991.
 - [29] MSo Bartlett, “The square root transformation in analysis of variance,” *Supplement to the Journal of the Royal Statistical Society*, pp. 68–78, 1936.
 - [30] Francis J Anscombe, “The transformation of poisson, binomial and negative-binomial data,” *Biometrika*, pp. 246–254, 1948.
 - [31] Piotr Fryzlewicz, Theofanis Sapatinas, and Suhasini Subba Rao, “A haar-fisz technique for locally stationary volatility estimation,” *Biometrika*, vol. 93, no. 3, pp. 687–704, 2006.