

IMAGE RECONSTRUCTION AND THRESHOLD DESIGN FOR QUANTA IMAGE SENSORS

Omar A. Elgendy and Stanley H. Chan

School of ECE and Dept of Statistics, Purdue University, West Lafayette, IN 47907.

ABSTRACT

Quanta Image Sensor (QIS) has been envisioned as a candidate solution for next generation image sensors. We provide two new contributions to the signal processing aspects of QIS. First, we develop an image reconstruction algorithm to recover the underlying images from the QIS data, which is a massive array of binarized Poisson random variables. The new algorithm supersedes existing methods by enabling arbitrary threshold level. Second, we present a threshold design scheme to adaptively update the threshold level for optimal image reconstruction. We discuss the existence of a phase transition in determining the optimal threshold. Experimental results on tone-mapped high dynamic range images validates the effectiveness of the threshold scheme and the image reconstruction algorithm.

Index Terms— Image reconstruction, quanta image sensors, ADMM, high dynamic range, quantization map.

1. INTRODUCTION

Quanta image sensors (QIS) is a collective name referring to a class of solid-state image sensors proposed for next generation camera after CMOS [1]. Having single photon sensitivity [2, 3, 4] and small pitch size (500nm pitch [5]), QIS has the potential to achieve high spatial resolution (e.g., giga-pixel cameras [6]), high frame rate (e.g., 10^3 frames/sec [7]), and high dynamic range [4, 8]. However, because of the small full well capacity QIS, a pixel observed by a QIS is a result of a quantized Poisson process — it is 1 when the photon count is above certain threshold, and 0 when below the threshold.

This paper addresses two signal processing questions of QIS: (1) Given a massive array of 1s and 0s, how can one reconstruct the underlying image from the observations? (2) How to adaptively design the threshold level for each pixel, or a group of nearby pixels? Addressing these two questions have important implications for the following reasons:

On the reconstruction level, without a viable image reconstruction algorithm one will not be able to use QIS as an imaging sensor, which obviously defeats the purpose of making QIS. However, image reconstruction algorithms for QIS are relatively young compared to other conventional reconstruction techniques (e.g., image denoising and deblurring). To date, the two major existing methods for QIS image reconstruction are the gradient descent algorithm which solves a maximum likelihood estimation (MLE) [9] and the alternating direction method of multiplier (ADMM) algorithm which solves a maximum-a-posteriori problem (MAP) [10]. Yet, both have limitations: The former has slow convergence and the latter does not handle threshold greater than 1. There is a recent approach using deep learning [11], but we do not consider here because it is an approximation and requires training. The first goal of this paper is to present an algorithm which has fast convergence and is able to handle any threshold level greater than 1.

E-mail: {oelgendy,stanleychan}@purdue.edu.



(a) With un-optimized threshold (b) With proposed threshold

Fig. 1. Reconstructed images from simulated QIS data using the proposed ADMM algorithm, (a) when the data is acquired using a constant threshold map; (b) when the data is acquired using an optimized threshold map.

On the threshold design level, we lack an adaptive threshold scheme to control the dynamic range of an image — we want a high threshold to prevent all-one measurements under a bright scene and a low threshold to prevent all-zeros under a dark scene. Adaptive threshold is implementable using a multi-bit QIS [2], which currently supports 3 bits (max 8 levels). The method we present here is more general as it works for any number of bits. It is also different from the time-sequential Markov chain approach which updates the threshold [12] or the duty cycle [13], and [14] which conditionally resets the photon counter.

In order to highlight the significance of the results presented in this paper, we show in Figure 1 a reconstructed result using our proposed ADMM algorithm with an un-optimized threshold map and a designed threshold map. It is evident from the figure that a fully designed threshold has significant performance gain compared to the unoptimized threshold.

The rest of the paper is organized as follows. In Section 2, we provide a quick review of the QIS imaging model and present the MLE formulation for a general threshold greater than 1. In section 3, we discuss the ADMM algorithm for solving the generalized MLE problem. The adaptive threshold design is presented in Section 4 and experimental results are presented in Section 5.

2. PROBLEM FORMULATION

In this section we provide a brief review of the QIS imaging model. Interested readers can refer to [9] for more details.

2.1. QIS Imaging Model

We represent the incoming light field by its expansion coefficients $\mathbf{c} = [c_1, \dots, c_N]^T$ with $0 \leq c_n \leq 1$. These N coefficients are the quantities we try to recover. When the incoming light arrives at the QIS, the QIS uses M uniformly distributed pixels to read the incoming light. The ratio $K \stackrel{\text{def}}{=} M/N$ defines the *over-sampling factor* of

the QIS. Without loss of generality, we assume K is an integer. Accumulated over a short period of time and integrated over the space, we define $\boldsymbol{\theta} = [\theta_1, \dots, \theta_M]^T$ as the amount of light arriving at the QIS pixel, and it is given by

$$\boldsymbol{\theta} = \alpha \mathbf{G} \mathbf{c},$$

where $\alpha > 0$ is a constant gain factor that multiplies with c_n to generate the actual photon intensity, and $\mathbf{G} \in \mathbb{R}^{M \times N}$ is a matrix capturing the projection and up-sampling of the sensor.

Given $\boldsymbol{\theta}$, the number of photons Y_m observed at the m -th QIS pixel follows a Poisson process, meaning that

$$\mathbb{P}(Y_m = y_m; \theta_m) = \frac{\theta_m^{y_m} e^{-\theta_m}}{y_m!}. \quad (1)$$

Letting $q > 0$ be a quantization threshold, the final observed pixel value at the m -th QIS pixel is a truncated version of Y_m :

$$B_m = \begin{cases} 1, & \text{if } Y_m \geq q \\ 0, & \text{if } Y_m < q. \end{cases}$$

Consequently, B_m is a Bernoulli random variable with probabilities

$$\mathbb{P}(B_m = b_m; \theta_m, q) = \begin{cases} \sum_{k=0}^{q-1} \frac{\theta_m^k e^{-\theta_m}}{k!}, & \text{if } b_m = 0, \\ \sum_{k=q}^{\infty} \frac{\theta_m^k e^{-\theta_m}}{k!}, & \text{if } b_m = 1. \end{cases} \quad (2)$$

While the probability defined in (2) is valid for all ranges of q and θ , it would be easier to represent the right hand side of (2) using the incomplete Gamma function, defined as

$$\Gamma(q, \theta) \stackrel{\text{def}}{=} \int_{\theta}^{\infty} t^{q-1} e^{-t} dt, \quad \text{for } \theta > 0, q \in \mathbb{N}.$$

Then, we can show the following results.

Proposition 1. [15]. *The probabilities in (2) can be expressed as*

$$\mathbb{P}(B_m = b_m; \theta_m, q) = \begin{cases} \Gamma(q, \theta_m) / \Gamma(q), & \text{if } b_m = 0, \\ 1 - \Gamma(q, \theta_m) / \Gamma(q), & \text{if } b_m = 1, \end{cases} \quad (3)$$

where $\Gamma(q) = (q-1)!$ for $q \in \mathbb{N}$ is the standard gamma function.

2.2. Maximum Likelihood Estimation

The goal of image reconstruction for QIS is to recover \mathbf{c} from the observed binary bit pattern $\mathbf{B} \stackrel{\text{def}}{=} [B_1, \dots, B_M]^T$. To this end, we consider the maximum likelihood estimation (MLE) framework, which aims to solve

$$\hat{\mathbf{c}} = \underset{\mathbf{c}}{\text{argmin}} F(\boldsymbol{\theta}), \quad \text{subject to } \boldsymbol{\theta} = \alpha \mathbf{G} \mathbf{c}, \quad (4)$$

where

$$F(\boldsymbol{\theta}) = - \sum_{m=1}^M \log \mathbb{P}(B_m = b_m; \theta_m, q) \quad (5)$$

is the negative log-likelihood of the data. Extending the MLE to maximum-a-posteriori (MAP) estimation by incorporating priors is possible [10], but we shall not dive into this option here.

To solve (4), we consider the alternating direction method of multipliers (ADMM). However, existing ADMM algorithm for QIS image reconstruction has only been successfully developed for the case of $q = 1$ and $\alpha = 1$ [10]. Our focus here is how to extend the existing ADMM for $q > 1$ and $\alpha > 1$.

3. ADMM ALGORITHM FOR SOLVING MLE

In this section we discuss how to solve the MLE problem in (4) using the ADMM algorithm [16]. Our focus here is the modification required to accommodate the case of $q > 1$ and $\alpha > 1$ for the original ADMM algorithm presented in [10].

Inspecting (4), we note that it is an equality constrained optimization. Therefore, we can formulate its augmented Lagrangian function as

$$\mathcal{L}(\mathbf{c}, \boldsymbol{\theta}, \mathbf{z}) = F(\boldsymbol{\theta}) - \mathbf{z}^T (\boldsymbol{\theta} - \alpha \mathbf{G} \mathbf{c}) + \frac{\rho}{2} \|\boldsymbol{\theta} - \alpha \mathbf{G} \mathbf{c}\|^2, \quad (6)$$

and solve the optimization problem via an iterative approach

$$\mathbf{c}^{(k+1)} = \underset{\mathbf{c}}{\text{argmin}} \mathcal{L}(\mathbf{c}, \boldsymbol{\theta}^{(k)}, \mathbf{z}^{(k)}), \quad (7a)$$

$$\boldsymbol{\theta}^{(k+1)} = \underset{\boldsymbol{\theta}}{\text{argmin}} \mathcal{L}(\mathbf{c}^{(k+1)}, \boldsymbol{\theta}, \mathbf{z}^{(k)}), \quad (7b)$$

$$\mathbf{z}^{(k+1)} = \mathbf{z}^{(k)} - \rho (\boldsymbol{\theta}^{(k+1)} - \alpha \mathbf{G} \mathbf{c}^{(k+1)}). \quad (7c)$$

Since $F(\boldsymbol{\theta})$ is convex, convergence of (7a)-(7c) is guaranteed under appropriate conditions [17].

For this ADMM algorithm, (7a) is a quadratic minimization which can be solved in closed-form by exploiting the structure of \mathbf{G} [10]. The challenge of the algorithm is (7b). Substituting (3) into (5), solving (7b) is equivalent to solving

$$\min_{\boldsymbol{\theta}} \sum_{m=1}^M \left[-z_m \theta_m + \frac{\rho}{2} (\theta_m - d_m)^2 - \log \left((1 - b_m) \frac{\Gamma(q, \theta_m)}{\Gamma(q)} + b_m \left(1 - \frac{\Gamma(q, \theta_m)}{\Gamma(q)} \right) \right) \right], \quad (8)$$

where $\mathbf{d} = [d_1, \dots, d_M]^T$ with $\mathbf{d} = \alpha \mathbf{G} \mathbf{c}$.

To solve (8), we recognize that it is a sum of M separable functions. Therefore, (8) is minimized when each individual term in the sum is minimized. Letting $w_m = z_m + \rho d_m$, the first order optimality returns us the following result.

Proposition 2. *The optimal solution θ_m of (8) satisfies the equations*

$$\begin{cases} e^{-\theta_m} \theta_m^{q-1} / \Gamma(q, \theta_m) = w_m - \rho \theta_m, & \text{if } b_m = 0, \\ e^{-\theta_m} \theta_m^{q-1} / (\Gamma(q) - \Gamma(q, \theta_m)) = \rho \theta_m - w_m & \text{if } b_m = 1. \end{cases} \quad (9)$$

Proof. The first order derivative of the incomplete gamma function is

$$\frac{\partial}{\partial \theta} \Gamma(q, \theta) = -e^{-\theta} \theta^{q-1}.$$

Differentiating the m th term in (8) and setting to zero yields (9). \square

From Proposition 2, it remains to solve (9). However, since (9) is a transcendental equation, we must adopt a numerical approach to solve the equation. Our proposed solution relies on building a look up table (offline) for $L+1$ values of w_m distributed uniformly in the interval $[w_{\min}, w_{\max}]$ with a step $\Delta w = (w_{\max} - w_{\min})/L$. Then, the solution at any value of w is obtained by a simple linear interpolation.

Remark: Because of the nonlinearity of the incomplete Gamma function $\Gamma(q, \theta)$, when building the look up table a solution may lie in a region close to discontinuity. To mitigate this issue, we use a bisection to determine an approximate interval in which the solution must be contained.

4. ADAPTIVE THRESHOLD DESIGN

In this section we discuss the threshold design methodology. For simplicity, we consider one unit space with a constant light field c so that $N = 1$, $M = K$ and $\theta_k = \theta$ for all $k = 1, \dots, K$.

4.1. Optimal Threshold

To start with, we define the signal-to-noise ratio of the MLE solution as

$$\text{SNR}(\theta, q) \stackrel{\text{def}}{=} 10 \log_{10}(\theta^2 / \mathbb{E}[(\hat{\theta}_{\text{ML}} - \theta)^2]),$$

where the expectation is taken over the random observations B_1, \dots, B_K . For large over-sampling ratio K , we can approximate the SNR as follows.

Proposition 3. [13]. As $K \rightarrow \infty$,

$$\text{SNR}(\theta, q) \approx 10 \log_{10}(\theta^2 I(\theta, q)) + 10 \log_{10} K,$$

where $I(\theta, q)$ is the Fisher Information of the distribution in (3).

Using SNR as a measure of the estimate's quality, we propose to determine a q which maximizes the SNR for a given θ . To this end, we show that the Fisher Information of the distribution in (3) can be expressed as follows, a result similar to the one in [12].

Proposition 4. The Fisher Information $I(\theta, q)$ of the distribution in (3) under a threshold q is:

$$I(\theta, q) = \frac{e^{-2\theta} \theta^{2q-2}}{\Gamma(q, \theta) (\Gamma(q) - \Gamma(q, \theta))}. \quad (10)$$

Proof. It is easy to check that the probability distribution $\mathbb{P}(B; \theta, q)$ satisfies the regularity conditions [18]. Therefore, $I(\theta, q)$ can be derived by taking second-order derivative

$$I(\theta, q) = \mathbb{E} \left[\frac{\partial^2}{\partial \theta^2} \log \mathbb{P}(B; \theta, q) \right],$$

yielding the desired result. \square

Maximizing the Fisher Information involves maximizing a non-linear function as described in (10), which clearly requires a numerical solver. However, it is possible to derive a reasonably tight lower bound on $\log(\theta^2 I(\theta, q))$ as follows.

Proposition 5. The Fisher Information is lower bounded by

$$\log(\theta^2 I(\theta, q)) \geq -2\theta + 2q \log \theta - 2 \log \Gamma(q) \stackrel{\text{def}}{=} L(\theta, q). \quad (11)$$

Proof. The lower bound is established by showing that $\Gamma(q, \theta) \leq \Gamma(q)$, and $\Gamma(q) - \Gamma(q, \theta) \leq \Gamma(q)$. \square

Consequently, we can derive the optimal threshold q under $L(\theta, q)$.

Proposition 6. The optimal q^* is

$$q^* = \underset{q}{\text{argmax}} L(\theta, q) = \lfloor \theta \rfloor + 1. \quad (12)$$

Proof. Write $L(\theta, q) = -2\theta + 2 \sum_{k=1}^{q-1} \log(\theta/k) + 2 \log \theta$. Then we can show that the sum is decreasing for $q > \lfloor \theta \rfloor + 1$, and increasing for $q < \lfloor \theta \rfloor + 1$. Thus, maximum is obtained at $q = \lfloor \theta \rfloor + 1$. \square

The result of Proposition 6 matches well with intuition: For dark images with a small θ , a small q is needed to ensure that not all B 's are 0. Similarly, for bright images with a large θ , q should also be large so that not all B 's are 1. The optimal threshold q should therefore match with the light intensity θ .

4.2. Phase Transition of Optimal Threshold

While Proposition 6 provides important theoretical justification of the optimal threshold q , it is however not useful because it requires knowing the underlying pixel intensity θ . Therefore, what remains to be solved is how to update q in the absence of θ . We first show a proposition. The proof is deferred to a follow up paper.

Proposition 7. Fix q , and let θ_0 be the ground truth unknown pixel value. Let K_0 and K_1 be the numbers of zeros and ones in the K observed pixels, respectively, and let $\gamma(K, q, \theta_0) = K_0 / (K_0 + K_1)$ be the proportion of zeros. Assume that $c = \theta$ (i.e., $\alpha = 1$, and $G = 1$). Then, the optimal $\hat{\theta}$ of the MLE in (4) satisfies

$$\Psi(q, \hat{\theta}) \stackrel{\text{def}}{=} \frac{\Gamma(q, \hat{\theta})}{\Gamma(q)} = \gamma(K, q, \theta_0). \quad (13)$$

The term $1 - \gamma$ (which is $K_1 / (K_0 + K_1)$) is called the *bit density* [4]. γ has a few important properties. Most critically, γ is a random variable depending on K , θ_0 and q . The mean and variance are $\mathbb{E}[\gamma] = \Psi(q, \theta_0)$, and $\text{Var}[\gamma] = \Psi(q, \theta_0)(1 - \Psi(q, \theta_0))/K$, respectively. Therefore, for a true pixel value θ_0 (which is unknown to us) and a given q , the right hand side of (13) is a realization of the random variable γ . Consequently, solving $\hat{\theta}$ in (13) is equivalent to finding a $\hat{\theta}$ which equates the deterministic quantity $\Psi(q, \hat{\theta})$ and the realization $\gamma(K, q, \theta_0)$. As an example at the extreme case when $K \rightarrow \infty$, we can show by law of large number that $\gamma(K, q, \theta_0) \xrightarrow{P} \Psi(q, \theta_0)$. In this case, (13) becomes $\Psi(q, \hat{\theta}) = \Psi(q, \theta_0)$ and hence the solution is $\hat{\theta} = \theta_0$, independent of the choice of q .

The situation is, however, slightly more complicated in practice because $\gamma(K, q, \theta_0)$ is the result of a *measurement*. Thus, finding $\hat{\theta}$ requires taking the inverse: $\hat{\theta} = \Psi^{-1}(\gamma, q)$, which has to be done numerically. This could be problematic as Ψ does not allow numerical inversion when $q \rightarrow \infty$ or $q \rightarrow 0$. To see this, we show in Figure 2 the function $\Psi(q, \theta)$ as a function of q . On the same figure, we also show the estimated value of $\hat{\theta}/\theta_0$ as a function of q . ($\hat{\theta}/\theta_0 = 1$ means that $\hat{\theta} = \theta_0$, i.e., correct estimate.) The result of this figure demonstrates an important *phase transition* behavior: When q is below certain level (59 in this example), $\hat{\theta}/\theta_0 \rightarrow \infty$; When q is above certain level (85 in this example), $\hat{\theta}/\theta_0 \rightarrow 0$. However, for q residing inside the interval, we can obtain an *exact* $\hat{\theta}$ (subject to some minor numerical approximation error).

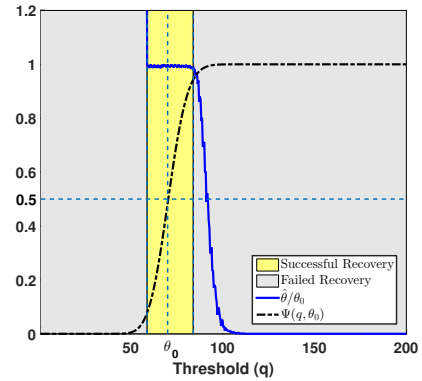


Fig. 2. Phase transition of the estimated solution and its relationship with $\Psi(q, \theta)$. The gray region is where it is impossible to recover θ , whereas the yellow region is where we can have perfect recovery.



(a) Recovery using $q = 10$ (b) Recovery using $q = 30$ (c) Optimized threshold (d) Recovery using (c)
Fig. 3. Two examples of using unoptimized threshold $q = \{10, 30\}$, and the optimized threshold. Gain factor is $\alpha = 50$.

Algorithm 1 Threshold Update using Bisection

Require: Given T binary measurements $\{\mathbf{Y}(t)\}_{t=1}^T$, where $\mathbf{Y}(t) = [Y_1(t), \dots, Y_K(t)]^T$.

Require: Initial thresholds q_a and q_b such that $q_a < q_b$.

for $t = 1, \dots, T$ **do**

 Let $q_c = (q_a + q_b)/2$.

 Quantize $\mathbf{B}_i(t) = \mathbb{I}(\mathbf{Y}(t) > q_i)$, for $i = \{a, b, c\}$.

 Compute $\gamma_i = 1 - \mathbf{1}^T \mathbf{B}_i(t)/K$, for $i = \{a, b, c\}$.

 If $\gamma_c \geq 0.5$, then set $q_a = q_c$ and $\gamma_a = \gamma_c$; Else, set $q_b = q_c$ and $\gamma_b = \gamma_c$.

 Break when $|\gamma_c - 0.5| < \text{tol}$.

end for

return q_c

4.3. Adaptive Threshold Update Scheme

The phase transition result we presented above offers insight about how to update the threshold. Since a perfect recovery can be obtained as long as q lies within certain interval, it remains to find one q which can always satisfy this criteria. Observing Figure 2, one candidate is at $\Psi(q, \theta) = 0.5$. In other words, for an unknown but fixed θ_0 , we should adjust the threshold q until $\gamma(K, q, \theta_0) = 0.5$. Then, the estimated θ is $\hat{\theta} = \Psi^{-1}(0.5, q)$.

In practice, in order to obtain $\gamma(K, q, \theta_0) = 0.5$ we can exploit the spatial-temporal over-sampling property of the QIS. Suppose there are K pixels in a unit space, at every time instant t we collect K binary measurements $\mathbf{B}(t) = [B_1(t), \dots, B_K(t)]^T$. If the proportion of zeros is below 0.5, we increase q ; otherwise we decrease q . We then repeat the process for the $(t + 1)$ th measurement $\mathbf{B}(t + 1)$ until the number of zeros stabilizes. In updating q for the $t + 1$ th measurement, we can update by a fixed step size, or other root-finding schemes. Algorithm 1 summarizes the threshold update scheme using a bisection method.

The threshold update scheme discussed in Algorithm 1 is defined for every K binary measurements. In the actual implementation, we can relax the requirement by using more than K measurements, e.g., define a common threshold for a group of pixels. Moreover, in the beginning of this section we made simplifications to the problem. For general situations where $N > 1$, $K = M/N$, $\alpha > 1$ and for a general \mathbf{G} , Algorithm 1 can still be applied, but the ADMM algorithm is needed to reconstruct the image.

5. EXPERIMENTAL RESULTS

We evaluate the proposed ADMM algorithm and the threshold design methodology using tone-mapped high-dynamic range (HDR) images downloaded from [19] and [20]. We convert and normalize the RGB channels of the tone-mapped HDR images to define the ground truth $\mathbf{c} \in [0, 1]^{3N}$. For this experiment, we set \mathbf{G} as an up-sampling operation followed by a box-car filtering [9]. The gain factor α is set as $\alpha = 50$. The over-sampling ratio is $K = 5$ along both the horizontal and vertical directions. Given \mathbf{c} , α and \mathbf{G} , we compute $\boldsymbol{\theta} = \alpha \mathbf{G} \mathbf{c}$ and generate Poisson observations $\mathbf{Y} \stackrel{\text{iid}}{\sim} \text{Poisson}(\boldsymbol{\theta})$.

On the left two columns of Figure 3 we show the reconstructed images using a constant threshold throughout the entire image. The results show that by using an un-optimized threshold map, the reconstructed images are poor. In particular, if q is low, dark regions are recovered but bright regions turn into 1, and vice versa for high values of q . This behavior is similar to a conventional HDR image.

On the right two columns of Figure 3 we show the designed threshold maps and the reconstructed images. For simplicity we only show the green channel of the threshold map, although in practice there are three separate threshold maps. In computing the threshold map, we share a common threshold for non-overlapping blocks of $K \times K$ pixels. This is equivalent to sharing one q for $K \times K$ c_n 's. We run the bisection method in Algorithm 1 for $T = 10$ frames to obtain the threshold map, and the ADMM algorithm to recover the final image. It can be seen that the recovered image captures both the dark and bright regions.

6. CONCLUSION

We presented solutions to two signal processing questions for QIS. The first solution addresses the image reconstruction problem for a general threshold $q > 1$. We derived an ADMM algorithm using the incomplete Gamma function to solve the problem. The second solution addresses the problem of designing optimal thresholds for QIS. We discussed the optimal threshold criteria and showed a phase transition of the threshold. We demonstrated the application of QIS for high-dynamic range imaging. Considering the high spatial resolution and high speed of QIS, we anticipate that novel applications of QIS beyond HDR imaging will be developed in a near future.

Acknowledgement: We thank Professor Eric Fossum and his students for useful discussions about multi-bit QIS.

7. REFERENCES

- [1] E. R. Fossum, "What to do with sub-diffraction-limit (SDL) pixels?—A proposal for a gigapixel digital film sensor (DFS)," in *Proc IEEE Workshop CCDs Adv. Image Sensors*, Sep. 2005, pp. 214–217.
- [2] J. Ma and E. R. Fossum, "Quanta image sensor jot with sub 0.3e- r.m.s. read noise and photon counting capability," *IEEE Electron Device Letters*, vol. 36, no. 9, pp. 926–928, Sep. 2015.
- [3] W. H. P. Pernice, C. Schuck, O. Minaeva, M. Li, G. N. Goltsman, A. V. Sergienko, and H. X. Tang, "High-speed and high-efficiency travelling wave single-photon detectors embedded in nanophotonic circuits," *Nature Communications*, vol. 3, pp. 1325–, Dec. 2012.
- [4] E. R. Fossum, "Modeling the performance of single-bit and multi-bit quanta image sensors," *IEEE J. Electron Devices Soc.*, vol. 1, no. 9, pp. 166–174, Sep. 2013.
- [5] J. Ma and E. R. Fossum, "A pump-gate jot device with high conversion gain for a quanta image sensor," *IEEE J. Electron Devices Soc.*, vol. 3, no. 2, pp. 73–77, Mar. 2015.
- [6] E. R. Fossum, "The Quanta Image Sensor (QIS): Concepts and Challenges," in *Proc OSA Topical Mtg Computational Optical Sensing and Imaging*, Jul. 2011, paper JTUE1.
- [7] S. Masoodian, A. Rao, M. Jiaju, K. Odame, and E. R. Fossum, "A 2.5 pj/b binary image sensor as a pathfinder for quanta image sensors," *IEEE Trans. on Electron Devices*, vol. 63, no. 1, pp. 100–105, Jan. 2016.
- [8] E. R. Fossum, "Multi-bit quanta image sensors," in *Proc. Int. Image Sens. Workshop*, Jun. 2015, pp. 292–295.
- [9] F. Yang, Y. M. Lu, L. Sbaiz, and M. Vetterli, "Bits from photons: Oversampled image acquisition using binary poisson statistics," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1421–1436, Apr. 2012.
- [10] S. H. Chan and Y. M. Lu, "Efficient image reconstruction for gigapixel quantum image sensors," in *Proc IEEE Global Conf. on Signal and Information Processing (GlobalSIP'14)*, Dec 2014, pp. 312–316.
- [11] T. Remez, O. Litany, and A. Bronstein, "A picture is worth a billion bits: Real-time image reconstruction from dense binary pixels," Available online at: <http://arxiv.org/abs/1510.04601>, Oct. 2015.
- [12] C. Hu and Y. M. Lu, "Adaptive time-sequential binary sensing for high dynamic range imaging," in *Proc. SPIE*, 2012, vol. 8375, pp. 83750A1 – A11.
- [13] Y. M. Lu, "Adaptive sensing and inference for single-photon imaging," in *Proc 47th Annual Conf. on Information Sciences and Systems (CISS'13)*, March 2013, pp. 1–6.
- [14] T. Vogelsang, D. G. Stork, and M. Guidash, "Hardware validated unified model of multibit temporally and spatially over-sampled image sensor with conditional reset," *Journal of Electronic Imaging*, vol. 23, no. 1, pp. 013021–1–013021–13, Jan-Feb 2014.
- [15] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*, Number 55. Courier Corporation, 1964.
- [16] S. H. Chan, R. Khoshabeh, K. B. Gibson, P. E. Gill, and T. Q. Nguyen, "An augmented Lagrangian method for total variation video restoration," *IEEE Trans. Image Process.*, vol. 20, no. 11, pp. 3097–3111, Nov. 2011.
- [17] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [18] T. S. Ferguson, *A Course in Large Sample Theory*, vol. 49, Chapman & Hall London, 1996.
- [19] F. Xiao, J. M. DiCarlo, P. B. Catrysse, and B. A. Wandell, "High dynamic range imaging of natural scenes," in *Proc S&T/SID 10th Color and Imaging Conference*, Scottsdale, Arizona, Nov. 2002, pp. 337–342.
- [20] I. Sprow, D. Kuepper, Z. Barańczuk, and P. Zolliker, "Image quality assessment using a high dynamic range display," in *Proc. 12th Congress of the International Colour Association, Newcastle upon Tyne, England*, Jul. 2013.