

# Unsupervised Restoration of Subsampled Images Constructed from Geometric and Binomial data

Yoann Altmann, Steve McLaughlin  
School of Engineering and Physical Sciences  
Heriot-Watt University  
Edinburgh, United Kingdom  
Email: {Y.Altmann;S.McLaughlin}@hw.ac.uk

Miles Padgett  
School of Physics and Astronomy  
University of Glasgow  
Glasgow, United Kingdom  
Email: Miles.Padgett@glasgow.ac.uk

**Abstract**—In this paper, we investigate a new imaging denoising algorithm for single-photon applications where the classical Poisson noise assumption does not hold. Precisely, we consider two different acquisition scenarios where the unknown intensity profile is to be recovered from subsampled measurements following binomial or geometric distributions, whose parameters are nonlinearly related to the intensities of interest. Adopting a Bayesian approach, a flexible prior model is assigned to the unknown intensity field and an adaptive Markov chain Monte Carlo methods is used to perform Bayesian inference. In particular, it allows us to automatically adjust the amount of regularisation required for satisfactory image inpainting/restoration. The performance of the proposed model/method is assessed quantitatively through a series of experiments conducted with controlled data and the results obtained are very promising for future analysis of multidimensional single-photon images.

## I. INTRODUCTION

Single-photon detectors (SPDs) are ubiquitous for applications where the light flux to be analysed is quantified at photonic levels. In particular, SPDs are particularly attractive for imaging applications where the light flux changes rapidly (of the order of picoseconds) or is extremely limited. For instance, the range resolution of SPD-based Lidar systems and their capability to resolve close objects depends on the ability of the detectors to accurately capture the time-of-arrival of photons emitted by fast laser sources [1]–[6]. Recent advances in fast SPDs and SPD arrays, coupled with efficient signal/image processing techniques have allowed the development of extreme imaging systems, including first photon [7] and single pixel [8], [9], and ghost [10]–[12] imaging systems, among others. Improving and investigating new systems however requires the development of statistical methods adapted to the discrete and sparse nature of the recorded data (photon counts or times of arrival).

SPDs can be classified into two groups depending on their ability to quantify a number of detected photons within an elementary time period (or depending of the light regime for which they are used). Although some detectors can be considered as photon-number resolving, in this paper we consider SPDs that can generally only distinguish no detection from at least one detection, such as single-photon avalanche diodes (SPADs), photomultiplier tubes and superconducting nanowire SPDs [13]. Although potentially not too restrictive, such limitations need to be considered when developing/applying statistical methods to analyse data recorded by non photon-number resolving SPDs or when the number of detected

photons is not reliable beyond the first detection. Indeed, although the number of photons reaching an SPD within a time period is widely assumed to be Poisson distributed (say of mean  $x$ ), the SPD saturation (e.g. due to detector dead time) can have a significant influence on the distribution of the actual photon detections.

In many imaging applications involving such non photon-revolving SPDs, images are formed by summing binary detections over several independent realizations and assuming the observed phenomenon is stationary (intensities constant over time). By ensuring that the probabilities of detection per acquisition for each pixel are small enough (as a rule-of-thumb generally lower than 5%), the actual distribution of the total number of detected photons in each pixel can be approximated by a Poisson distribution whose mean grows linearly with the number of elementary periods. This approximation becomes generally less accurate as the light flux increases and/or when the ratio between the detection interval and the dead-time duration of the detector decreases, where the dead-time period corresponds to a period following a photon detection during which the detector cannot detect additional photons.

In this work, and in contrast with most denoising methods developed for photon-limited data, we focus on applications or acquisition scenarios for which the classical Poisson observation model does not hold. This is typically the case for imaging applications using single-photon avalanche diodes (SPADs) in scenarios where the detector dead time cannot be neglected or SDPs which are not photon-number resolving (i.e., not able to accurately quantify photon detection events beyond the first detection).

Adopting a classical Bayesian approach, we consider a flexible intensity prior model able to capture correlation between intensities of neighbouring pixels. This prior model is then coupled with the observation models (or likelihood) associated with two different acquisition scenarios. While the first strategy relies on a fixed per-pixel acquisition time, the second acquisition mode, referred to as first-photon approach[7], stops after the first detection event for each pixel. A stochastic simulation method (Markov chain Monte Carlo) method is finally investigated to exploit the resulting posteriors. An important advantage of the proposed method is that it is fully unsupervised and does not require crucial parameter tuning. In particular, the parameter controlling the spatial regularisation is automatically adjusted for depending on the observed data during the early stage of the sampling process.

The remained of the paper is organized as follow. Section II presents the two observation models considered and defines our Bayesian model used for image restoration. The estimation and sampling strategies proposed to exploit the resulting posterior distributions are described in Section III. Simulation results conducted using synthetic single-photon data are discussed in Section IV and conclusions and future work are finally reported in Section V.

## II. BAYESIAN MODELS

### A. Observation models

As discussed above, we consider observation models for low-flux imaging applications where Poisson noise models are not well adapted, e.g. when detector dead time cannot be neglected. Let  $x_n$  be the unknown photon flux of interest, reaching the  $n$ th given pixel/detector over an elementary time unit (referred to as repetition period in this work). The corresponding detection rate of the  $n$ th detector can be expressed as  $\eta_n x_n + b_n$ , where  $\eta_n \in [0, 1]$  is the  $n$ th detector's quantum efficiency and  $b_n \geq 0$  stands for ambient illumination and detector dark count rate. In this work, we assume that  $(\eta_n, b_n)$  is known or previously estimated from calibration measurements. As explained in [14], for most SPAD-based single-photon imaging applications, the dead time of the detectors is of the order of the repetition period. Thus, at most one detection event is recorded per repetition period. A classical approximation consists of considering that the dead time of the detector ends at the beginning of each new repetition period, leading to

$$y_n | (x_n, t_n) \sim \text{Bin}(t_n, 1 - \exp(-\eta_n x_n - b_n)), \quad (1)$$

where  $y_n$  is the number of detection events recorded after  $t_n$  repetition periods and  $\text{Bin}(\cdot, \cdot)$  stands for the binomial distribution. Under the low-flux assumption, i.e.  $\eta_n x_n + b_n \ll 1$ , Eq. (1) can be well approximated by the classically used Poisson distribution  $y_n | (x_n, t_n) \sim \mathcal{P}(t_n(\eta_n x_n + b_n))$ .

In the first-photon imaging context, i.e., when the acquisition of each pixel is stopped after the first detection, the number of repetition periods required to record the first detection event follows a geometric distribution specified by

$$t_n | (y_n = 1, x_n) \sim \text{Geo}(1 - \exp(-\eta_n x_n - b_n)), \quad (2)$$

with  $t_n \geq 1$ . Note that for brevity, on the left-hand side of (1) and (2) and in the remainder of the paper, we omit the fixed model parameters  $\eta_n$  and  $b_n$  in the notations of the different distributions.

### B. Intensity prior model

This section describes a family of intensity prior models that can be used to regularise the intensity inpainting/denoising problem by accounting for spatial correlation of natural intensity fields, while allowing an efficient and automated procedure able to adjust the amount of spatial smoothness of the estimated intensity field. We assume that the image to be reconstructed is composed of  $N$  pixels. The unknown vectorized intensity field is gathered in the vector  $\mathbf{x} = [x_1, \dots, x_N]^T$ . To regularise the intensity estimation procedure, we first constrain each intensity  $x_n$  to belong to a discrete set of  $N_x$  user-defined intensities  $\mathcal{X} = \{x_{\min}, \dots, x_{\max}\}$ , such that  $0 \leq x_{\min} \leq x_{\max}$ .

The value of  $N_x$  and the intensity range  $\mathcal{X}$  can be selected from prior knowledge about the intensity field and computational constraints (finer grids can lead to better estimates but at a higher computational cost). In a similar manner to [6], to account for the spatial correlations between intensities of neighbouring pixels, we propose to use Markov random fields (MRFs) to define a prior distribution for  $x_n$  given the intensities in the neighbouring pixels, whose indices are gathered in  $\mathbf{x}_{\mathcal{V}(n)}$ , i.e.,  $f(x_n | \mathbf{x}_{\setminus(n)}) = f(x_n | \mathbf{x}_{\mathcal{V}(n)})$  where  $\mathbf{x}_{\mathcal{V}(n)}$  is the neighbourhood of the pixel ( $n$ ) and  $\mathbf{x}_{\setminus(n)} = \{n'\}_{n' \neq n}$ . More precisely, we propose to use the following family of discrete MRFs

$$f(\mathbf{x} | \epsilon) = G(\epsilon)^{-1} \exp[-\epsilon \phi(\mathbf{x})] \quad (3)$$

where  $\epsilon \geq 0$  is a parameter tuning the amount of correlation between pixels; the higher the value of  $\epsilon$ , the more correlated the intensities of neighbouring pixels. Moreover,  $G(\epsilon)$  is a normalisation (or partition) constant and  $\phi(\cdot)$  is an arbitrary cost function modelling correlation between neighbours. For instance, in Section III, we propose to use the following cost function

$$\phi(\mathbf{x}) = \sum_n \sum_{n' \in \mathcal{V}(n)} |x_n - x_{n'}|, \quad (4)$$

which corresponds to an anisotropic total-variation (TV) regularisation [15], [16] promoting piecewise constant intensity profiles. Note that other regularizations could also be used (e.g. isotropic TV or Gaussian MRFs [17]). Indeed, since the intensities are allowed to take a finite number of values, the choice of  $\phi(\cdot)$  does not affect the structure of the algorithm (see Section III). Several neighbourhood structures can be employed to define  $\mathcal{V}(n)$ ; here, a four pixel structure (1-order neighbourhood) will be considered in the rest of the paper for the MRF used.

Now that we have defined the observations models for any observed pixel and the intensity prior model, we can now derive the joint posterior distributions of  $\mathbf{x}$ , given the observed detection events (binomial likelihoods,  $t_n$  fixed for each observed pixel) or the observed repetition periods (first-photon imaging,  $y_n = 1$  for each observed pixel). Let  $N_{obs}$  be the set of indices of the pixels actually observed. It can be easily shown that posteriors distributions

$$f(\mathbf{x} | \mathbf{t}, \mathbf{y}, \epsilon) \propto f(\mathbf{x} | \epsilon) \prod_{n \in N_{obs}} f(y_n | x_n, t_n) \quad (5)$$

$$\propto f(\mathbf{x} | \epsilon) \prod_{n \in N_{obs}} f(t_n | x_n, y_n = 1), \quad (6)$$

with  $\mathbf{t} = \{t_n\}_{n \in N_{obs}}$  and  $\mathbf{y} = \{y_n\}_{n \in N_{obs}}$  and where  $f(y_n | x_n, t_n)$  and  $f(t_n | x_n, y_n = 1)$  are given by (1) and (2), respectively, present the same form. Thus, a unifying inference mechanism can be used to exploit the posterior distribution associated with each observation model.

## III. PROPOSED ESTIMATION STRATEGY

The proposed method consists of automatically inpainting and denoising the unknown intensity field. This task is generally difficult because the results will highly depend on the value of the regularisation parameter  $\epsilon$ , in particular when the number of detection events is small. Here we resort to an

adaptive MCMC method to compute the MMAP estimators of the target intensities, conditioned on the observed data, and given by

$$\hat{x}_n = \max_{\lambda_{x_n}} f(x_n | \mathbf{t}, \mathbf{y}, \hat{\epsilon}), \quad \forall n \in 1, \dots, N \quad (7)$$

where  $\hat{\epsilon}$  approximates the marginal maximum likelihood estimator of  $\epsilon$ , it is

$$\hat{\epsilon} = \operatorname{argmax}_{\epsilon} f(\mathbf{y} | \mathbf{t}, \epsilon) \quad \text{or} \quad \hat{\epsilon} = \operatorname{argmax}_{\epsilon} f(\mathbf{t} | \mathbf{y}, \epsilon), \quad (8)$$

depending on the observation model considered. More precisely, we use a Gibbs sampler to sample sequentially each intensity  $x_n$  from its conditional distribution  $f(x_n | \mathbf{x}_{\setminus x_n}, \mathbf{y}, \mathbf{t}, \epsilon)$ . Since  $x_n$  is assumed to take a finite number of values, we can evaluate  $f(x_n | \mathbf{x}_{\setminus x_n}, \mathbf{y}, \mathbf{t}, \epsilon)$  (up to a multiplicative constant) for each of the  $N_x$  possible values and sample  $x_n$  by drawing from a categorical distribution. The sampler, whose pseudocode is shown in Algo. 1, consists of  $N_{\text{MC}}$  iterations during each of which all the intensities  $x_n$  are updated sequentially. Note however that we can exploit the conditional independence induced by the neighbourhood structure of the MRF in (3) to update simultaneously groups of conditionally independent intensities. For instance, using a 1-order neighbourhood, only 4 successive updates are required to update all the intensities, which makes the resulting method highly parallelizable. The first  $N_{\text{bi}}$  iterations of the algorithm correspond to the burn-in period of the sampler during which the regularisation parameter  $\epsilon$  is adjusted using the method proposed in [18]. Note that this method only requires to be able to (asymptotically) sample according to (3), which is generally the case for a wide range of regularisations  $\phi(\cdot)$  since  $\mathbf{x}$  is a vector whose elements admit a finite number of discrete values. Another important property of this method is that its computational cost is similar to that of the update of  $\mathbf{x}$  and thus does not significantly increase the overall computational cost of the sampling strategy. Due to space constraints, the technical details about the update of  $\epsilon$  are not included here but the interested reader is invited to consult [18], [19] for further discussions. It is also important to mention here that such method can also be used when multiple regularizations are used, e.g. to regularize spatially and spectrally multispectral [20] or hyperspectral data [21].

It is worth noting that when the regularisation  $\phi(\cdot)$  in (3) is convex,  $f(\mathbf{x} | \mathbf{t}, \mathbf{y}, \epsilon)$  is log-concave and that  $\mathbf{x}$  could also be inferred via standard MAP estimation, using continuous intensities, and using state-of-the-art convex optimisation techniques (provided that  $\epsilon$  is properly tuned). However, our sampling method allows an efficient estimation of the intensity profile even if (5) is not log-concave. Moreover, even with the appropriate value of  $\epsilon$ , optimisation methods might suffer from slow convergence due to high uncertainty of the data, i.e. of the likelihood, especially for large values of  $\epsilon$ . By iteratively updating  $\mathbf{x}$  and  $\epsilon$  (starting with a small  $\epsilon$ ), we observed that the convergence of the sampling strategy generally is generally improved and that the final result is usually more reliable than the sequential optimization of  $\mathbf{x}$  and  $\epsilon$  via maximum a posteriori estimation (e.g. as in [22]) when the likelihood is weakly informative (e.g. low light flux).

## Proposed algorithm

- 1: Fixed input parameters: Vector of numbers of detection events  $\mathbf{y}$  and corresponding numbers of repetitions periods  $\mathbf{t}$ , number of burn-in iterations  $N_{\text{bi}}$ , total number of iterations  $N_{\text{MC}}$
- 2: Initialization ( $u = 0$ )
  - Set  $\mathbf{x}^{(0)}, \epsilon^{(0)}$
- 3: Iterations ( $1 \leq u \leq N_{\text{MC}}$ )
- 4: **for**  $n = 1 : N$  **do**
- 5: Sample  $x_n^{(u)}$  from  $f(x_n | x_1^{(u)}, \dots, x_{n-1}^{(u)}, x_{n+1}^{(u-1)}, \dots, x_N^{(u-1)}, \mathbf{y}, \mathbf{t}, \epsilon^{(u-1)})$
- 6: **end for**
- 7: **if**  $u < N_{\text{bi}}$  **then**
- 8: Update  $\epsilon^{(u)}$  using [18]
- 9: **else**
- 10: Set  $\epsilon^{(u)} = \epsilon^{(u-1)}$
- 11: **end if**
- 12: Set  $u = u + 1$ .

## IV. RESULTS

In this section, we investigate the performance of the proposed denoising technique by analysing the classical image of lena. The size of the image is  $128 \times 128$  pixels and the original image intensity is scaled so that the mean intensity  $E[x_n]$  belongs to  $\{10^{-4}, 10^{-3}, 10^{-2}, 10^{-1}, 0.5, 1\}$ . For all the results presented in this paper, we used  $(\eta_n, b_n) = (1, 0), \forall n$ . For the data generated according to (1), the same number of repetition period is chosen for all the observed pixels, i.e.  $t_n = t, \forall n$ . Moreover,  $t$  belongs to  $\{10, 50, 10^2, 10^3, 10^4\}$ . The quality of image reconstruction is quantitatively assessed using the *normalised root mean square error* (NRMSE) defined as

$$\text{NRMSE} = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2}{\|\mathbf{x}\|_2}, \quad (9)$$

where  $\mathbf{x}$  and  $\hat{\mathbf{x}}$  are the actual and estimated intensity profiles, respectively. To evaluate the consistency of the results obtained by the proposed method, we consider mean NRMSEs and associated standard deviations (stds) obtained over 20 noise realisations. We first investigate the performance of our method assuming that all the pixels are observed. Fig. 1 depicts the mean NRMSEs obtained with binomial observations for various values of  $E[x_n]$  and  $t$ . As expected, the performance generally improves as  $E[x_n]$  and  $t$  increase. Note however that for a fixed number of repetition periods  $t$ , the performance will be degraded when  $E[x_n]$  becomes significantly larger than 1 as the number of detection events per pixel will be close to  $t$  with high probability. This phenomenon is not visible in Fig. 1 since we only considered  $E[x_n] \leq 1$ . Moreover, when  $E[x_n]$  is small, a large number of repetition periods is required to ensure a sufficient number of detection events to accurately estimate the intensity profile (see pink curves in Fig. 1).

Fig. 2 compares the intensity estimation performance based on the binomial and geometric observation models considered, as a function of the mean number of repetition periods per pixel  $E[t_n]$ . The solid and dashed lines correspond to binomial data with  $t_n = t, \forall n$  ( $E[t_n] = 1$ ). The isolated error bars correspond to the geometric measurements. The colours in Fig. 2 are associated with different values of  $E[x_n]$ . Note that we did not include results obtained with binomial data for  $E[x_n] = 10^{-3}$  (pink error bar) and  $E[x_n] = 10^{-4}$  (light blue error bar) for which the NRMSEs are too large over the




Fig. 1. Mean NRMSEs (solid lines) and associated  $\pm 3$  stds confidence intervals (dashed lines) obtained using data generated according to (1), for various values of  $E[x_n]$  and  $t$ .




Fig. 2. Mean NRMSEs (solid lines) and associated  $\pm 3$  stds confidence intervals (dashed lines) obtained using data generated according to (1), for various values of  $E[x_n]$  and  $E[t_n]$ .

range of  $t$  considered. This figure shows that for  $E[x_n] \ll 1$  (e.g. when the Poisson noise assumption is well adapted), we obtain similar results using a first-photon imaging approach or binomial data with  $t \approx 1/E[x_n]$ , which corresponds to observing 1 detection event per pixel on average. If  $t > 1/E[x_n]$ , the binomial data will then provide better intensity estimates. Interestingly, the estimation performance using the geometric data is fairly constant across the range of  $E[x_n]$  considered (except for  $E[x_n] \geq 1$  where saturation will lead to  $t_n = 1$  with high probability). In that case, the NRMSE is intrinsically limited by the fact that exactly one detection event is recorded per pixel. However, it might be interesting to investigate in future work whether similar results can be obtained using “ $k$ -photons” imaging techniques where  $k$  could potentially be chosen to ensure a predefined NRMSE.




Fig. 3. Mean NRMSEs (solid lines) and associated  $\pm 3$  stds confidence intervals (dashed lines) obtained using data generated according to (1), for  $t = 10^3$  and various values of  $E[x_n]$  and  $\alpha$ .

Finally, we evaluate the robustness of the denoising method with respect to the fraction of pixels actually observed. Precisely, we randomly select  $\alpha \in \{25, 50, 80, 100\}$  percent of the original  $128 \times 128$  pixels (the selected pixels are different for each of the 20 noise realisations). Here, we only present results obtained with binomial measurements and  $t = 10^3$  but similar results have been obtained using the second model considered. The resulting NRMSEs and associated confidence intervals are depicted in Fig. 3. This figure shows that although down-sampling induces a performance degradation, this degradation reduces as the mean image intensity reduces. This observation is particularly interesting in the context of extremely low flux imaging scenarios where sparse sampling strategies could be adopted to perform color/multispectral imaging. Indeed, using single-photon mosaic filter arrays, it could be possible for instance to acquire simultaneously 4 channels/spectral bands by acquiring only 25% of the pixels for each band and reconstruct reflectivity profiles as if all the spectral bands were observed for each pixel.

## V. CONCLUSION

In this paper, we presented a new unsupervised image denoising technique for binomial and geometric observations nonlinearly related to the intensity field of interest. A discrete Markov random field was introduced to capture the spatial correlation of natural intensity fields and a unifying Bayesian inference procedure was investigated to inpainting and restoration of the unknown intensity field. The proposed model and estimation strategy is perfectly adapted for analysis of multidimensional images for which the Markov random field can be easily extended while still benefiting from interesting parallelisation properties during the inference process. Future work includes the consideration of more complex (including non-local) regularisations and the generalization of the method to continuous intensity fields which will require the use of efficient simulation methods adapted to high-dimensional and non-standard posterior distributions.

## REFERENCES

- [1] A. McCarthy, R. J. Collins, N. J. Krichel, V. Fernández, A. M. Wallace, and G. S. Buller, "Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting," *Appl. Opt.*, vol. 48, no. 32, pp. 6241–6251, Nov. 2009.
- [2] N. J. Krichel, A. McCarthy, and G. S. Buller, "Resolving range ambiguity in a photon counting depth imager operating at kilometer distances," *Opt. Express*, vol. 18, no. 9, pp. 9192–9206, April 2010.
- [3] A. M. Wallace, J. Ye, N. J. Krichel, A. McCarthy, R. J. Collins, and G. S. Buller, "Full waveform analysis for long-range 3d imaging laser radar," *EURASIP Journal on Advances in Signal Processing*, vol. 2010, no. 1, p. 896708, 2010.
- [4] A. McCarthy, X. Ren, A. D. Frera, N. R. Gemmill, N. J. Krichel, C. Scarcella, A. Ruggeri, A. Tosi, and G. S. Buller, "Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector," *Opt. Express*, vol. 21, no. 19, pp. 22098–22113, Sept. 2013.
- [5] A. Maccarone, A. McCarthy, X. Ren, R. E. Warburton, A. M. Wallace, J. Moffat, Y. Petillot, and G. S. Buller, "Underwater depth imaging using time-correlated single-photon counting," *Opt. Express*, vol. 23, no. 26, pp. 33911–33926, Dec 2015.
- [6] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform based analysis of depth images constructed using sparse single-photon data," *IEEE Trans. Image Processing*, vol. 25, no. 5, pp. 1935–1946, May 2016.
- [7] A. Kirmani, D. Venkatraman, D. Shin, A. Colao, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [8] Y. Zhang, M. P. Edgar, B. Sun, N. Radwell, G. M. Gibson, and M. J. Padgett, "3d single-pixel video," *Journal of Optics*, vol. 18, no. 3, p. 035203, 2016.
- [9] M.-J. Sun, M. P. Edgar, D. B. Phillips, G. M. Gibson, and M. J. Padgett, "Improving the signal-to-noise ratio of single-pixel imaging using digital microscanning," *Opt. Express*, vol. 24, no. 10, pp. 10476–10485, May 2016.
- [10] R. S. Aspden, D. S. Tasca, R. W. Boyd, and M. J. Padgett, "Epr-based ghost imaging using a single-photon-sensitive camera," *New Journal of Physics*, vol. 15, no. 7, p. 073032, 2013.
- [11] R. S. Aspden, N. R. Gemmill, P. A. Morris, D. S. Tasca, L. Mertens, M. G. Tanner, R. A. Kirkwood, A. Ruggeri, A. Tosi, R. W. Boyd, G. S. Buller, R. H. Hadfield, and M. J. Padgett, "Photon-sparse microscopy: visible light imaging using infrared illumination," *Optica*, vol. 2, no. 12, pp. 1049–1052, Dec. 2015. [Online]. Available: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-2-12-1049>
- [12] P. A. Morris, R. S. Aspden, J. E. C. Bell, R. W. Boyd, and M. J. Padgett, "Imaging with a small number of photons," *Nat. Commun.*, 2015.
- [13] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov, "Invited review article: Single-photon sources and detectors," *Review of Scientific Instruments*, vol. 82, no. 7, 2011.
- [14] J. Rapp and V. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," *IEEE Transactions on Computational Imaging*, 2017, to appear.
- [15] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.
- [16] A. Chambolle, "An algorithm for total variation minimization and applications," *J. of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 89–97, 2004.
- [17] Y. Altmann, R. Aspden, M. Padgett, and S. McLaughlin, "A bayesian approach to denoising of single-photon binary images," *IEEE Trans. Comput. Imaging*, 2017, to appear.
- [18] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tourneret, "Maximum marginal likelihood estimation of the granularity coefficient of a Potts-Markov random field within an MCMC algorithm," in *Proc. IEEE-SP Workshop Stat. and Signal Processing*, Gold Coast, Australia, July 2014.
- [19] Y. Altmann, M. Pereyra, and J. Bioucas-Dias, "Collaborative sparse regression using spatially correlated supports - application to hyper-spectral unmixing," *IEEE Trans. Image Processing*, vol. 24, no. 12, pp. 5800–5811, Dec. 2015.
- [20] Y. Altmann, A. Maccarone, A. McCarthy, G. Newstadt, G. Buller, S. McLaughlin, and A. Hero, "Robust spectral unmixing of sparse multispectral lidar waveforms using gamma markov random fields," *IEEE Trans. Comput. Imaging*, 2017, to appear.
- [21] Y. Altmann, S. McLaughlin, and A. Hero, "Robust linear spectral unmixing using anomaly detection," *IEEE Trans. Comput. Imaging*, vol. 1, no. 2, pp. 74–85, June 2015.
- [22] M. Pereyra, J. M. Bioucas-Dias, and M. A. T. Figueiredo, "Maximum-a-posteriori estimation with unknown regularisation parameters," in *Proc. European Signal Processing Conf. (EUSIPCO)*, Aug. 2015, pp. 230–234.