



Heriot-Watt University
Research Gateway

Lidar Waveform-Based Analysis of Depth Images Constructed Using Sparse Single-Photon Data

Citation for published version:

Altmann, Y, Ren, X, McCarthy, A, Buller, GS & McLaughlin, S 2016, 'Lidar Waveform-Based Analysis of Depth Images Constructed Using Sparse Single-Photon Data', *IEEE Transactions on Image Processing*, vol. 25, no. 5, pp. 1935-1946. <https://doi.org/10.1109/TIP.2016.2526784>

Digital Object Identifier (DOI):

[10.1109/TIP.2016.2526784](https://doi.org/10.1109/TIP.2016.2526784)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

IEEE Transactions on Image Processing

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Lidar Waveform-Based Analysis of Depth Images Constructed Using Sparse Single-Photon Data

Yoann Altmann, *Member, IEEE*, Ximing Ren, Aongus McCarthy, *Member, IEEE*,
Gerald S. Buller, and Steve McLaughlin, *Fellow, IEEE*

Abstract—This paper presents a new Bayesian model and algorithm used for depth and reflectivity profiling using full waveforms from the time-correlated single-photon counting measurement in the limit of very low photon counts. The proposed model represents each Lidar waveform as a combination of a known impulse response, weighted by the target reflectivity, and an unknown constant background, corrupted by Poisson noise. Prior knowledge about the problem is embedded through prior distributions that account for the different parameter constraints and their spatial correlation among the image pixels. In particular, a gamma Markov random field (MRF) is used to model the joint distribution of the target reflectivity, and a second MRF is used to model the distribution of the target depth, which are both expected to exhibit significant spatial correlations. An adaptive Markov chain Monte Carlo algorithm is then proposed to perform Bayesian inference. This algorithm is equipped with a stochastic optimization adaptation mechanism that automatically adjusts the parameters of the MRFs by maximum marginal likelihood estimation. Finally, the benefits of the proposed methodology are demonstrated through a series of experiments using real data.

Index Terms—Remote sensing, full waveform Lidar, Poisson statistics, Bayesian estimation, Markov chain Monte Carlo.

I. INTRODUCTION

RECONSTRUCTION of 3-dimensional scenes using time-of-flight light detection and ranging (Lidar) systems is a challenging problem encountered in many applications, including automotive [1]–[4], environment sciences [5], [6], architectural engineering and defence [7], [8]. This problem consists of illuminating the scene with a train of laser pulses and analyzing the distribution of the photons reflected from the targets to infer the range, as well as reflectivity information about the scene. Using scanning systems, an histogram of time delays between the emitted pulses and the detected photon arrivals is usually recorded for each pixel, associated with a different region of the scene. The recorded photon histograms are classically decomposed into a series of peaks whose

positions can be used to infer the distance of the object(s) present in each region of the scene and whose amplitudes provide information about the reflectivity of the objects.

In this paper, we consider applications where the flux of detected photons is small and for which classical methods usually provide unsatisfactory results in terms of range and reflectivity estimation. This is typically the case when the acquisition time or the laser source power are small relative to the range of the target(s) [9].

Recently, Kirmani *et al.* [10], investigated a new method for reconstructing 3-dimensional scenes under low photon flux conditions by registering the time of arrival of the first photon in each pixel. Based on an appropriate statistical model relating the time of arrival to the target distance and reflectivity, they proposed different processing steps, to handle ambient noise (background photons) and the spatial regularity of the scene to obtain estimated target distances and reflectivities. In contrast with the method proposed in [10], we consider a scanning system whose acquisition time per pixel is fixed, thus leading to a deterministic and user-defined overall acquisition duration. Consequently, the number of detected photons can be larger than one for some pixels, whereas some pixels may be empty, i.e., contain no detected photons.

In this work, we assume that the targets of interest are opaque, i.e., are composed of a single surface per pixel and that at least a surface is present in each scanned pixel. Generalization to more complex objects will be discussed in the conclusion of this paper. As in [10], we consider the presence of two kinds of detector events: the “useful” photons originating from the illumination laser and scattered back from the target; and those background detector events originating from ambient light and the “dark” events resulting from detector noise. The proposed method aims to estimate the respective contributions of the actual target and the background in the photon timing histograms. More precisely, it also allows the estimation of the distance and reflectivity of the surface associated with each pixel, together with the average background levels, within a single estimation procedure.

Adopting a Bayesian framework as in [11] and [12], we assign prior distributions to the unknown model parameters to include available information (such as parameter constraints) within the estimation procedure. In particular, Markov random fields (MRFs) are introduced to model spatial correlations for the target distances and reflectivities. The joint posterior distribution of the unknown parameter vector is then derived, based on the Poisson statistical properties of the observed data.

Manuscript received May 7, 2015; revised December 1, 2015 and January 22, 2016; accepted January 28, 2016. Date of publication February 8, 2016; date of current version March 18, 2016. This work was supported by the Engineering and Physical Sciences Research Council under Grant EP/J015180/1, Grant EP/K015338/1, and Grant EP/M01326X/1. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Aleksandra Pizurica.

The authors are with the School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, U.K. (e-mail: y.altmann@hw.ac.uk; x.ren@hw.ac.uk; a.mccarthy@hw.ac.uk; g.s.buller@hw.ac.uk; s.mclaughlin@hw.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2016.2526784

Since classical Bayesian estimators cannot be easily computed from this joint posterior, a Markov chain Monte Carlo (MCMC) method is used to generate samples according to this posterior. More precisely, we construct an efficient stochastic gradient MCMC (SGMCMC) algorithm [13] that simultaneously estimates the background levels and the target distances and reflectivity, along with the MRFs parameters.

The main contributions of this work are

- 1) We develop new Bayesian reflectivity and depth models taking into account spatial correlations through Markovian dependencies. These flexible models are embedded within the observation model for full waveform Lidar-based low photon count imaging
- 2) An adaptive Markov chain Monte Carlo algorithm is then proposed to compute the Bayesian estimates of interest and perform Bayesian inference. This algorithm is equipped with a stochastic optimization adaptation mechanism that automatically adjusts the parameters of the Markov random fields by maximum marginal likelihood estimation, thus removing the need to set the regularization parameters by cross-validation.

The benefits of the proposed flexible model for reconstructing a real 3D object are illustrated through experiments conducted using real single-photon data for scenarios where the number of detected photons is very low. Specifically, we demonstrate the ability of the proposed algorithm to handle empty pixels as well as background levels.

The remainder of this paper is organized as follows. Section II recalls the classical statistical model used for depth imaging using time-of-flight scanning sensors, based on TCSPC. Section III presents the proposed Bayesian model for low photon count depth imaging, which takes into account the inherent spatial correlations between pixels. The estimation of the Bayesian model parameters using adaptive MCMC methods is discussed in Section IV. Simulation results conducted using an actual time-of-flight scanning sensor are presented and discussed in Section V. Finally, conclusions and potential future work are reported in Section VI.

II. PROBLEM FORMULATION

Consider a set of $N_{\text{row}} \times N_{\text{col}}$ observed Lidar waveforms $\mathbf{y}_{i,j} = [y_{i,j,1}, \dots, y_{i,j,T}]^T$, $(i, j) \in \{1, \dots, N_{\text{row}}\} \times \{1, \dots, N_{\text{col}}\}$ where T is the number of temporal (corresponding to range) bins. These waveforms are associated with $N_{\text{row}} \times N_{\text{col}}$ regions (or pixels) of the scene, which is assumed to be regularly sampled. To be precise, $y_{i,j,t}$ is the photon count within the t th bin of the pixel or location (i, j) . Let $t_{i,j}$ be the characteristic time of flight associated with an object surface at a given range from the sensor (estimating $t_{i,j}$ or the associated range are equivalent). According to [11], each photon count $y_{i,j,t}$ is assumed to be drawn from the following Poisson distribution

$$y_{i,j,t} \sim \mathcal{P}(r_{i,j} g_0(t - t_{i,j}) + b_{i,j}) \quad (1)$$

where $g_0(\cdot) > 0$ is the photon impulse response, $r_{i,j}$ denotes the reflectivity of the target and $b_{i,j}$ stands for the background and dark photon level, which is constant in all bins of

a given pixel. The photon impulse response $g_0(\cdot)$ is assumed to be known and estimated during the imaging system calibration.

Due to physical considerations, the target reflectivity is assumed to satisfy the following positivity constraints

$$r_{i,j} \geq 0, \quad \forall i, j. \quad (2)$$

Similarly, the average background level in pixel satisfies $b_{i,j} \geq 0, \forall i, j$. The problem addressed in this paper consists of estimating the depth/range (*i.e.*, $t_{i,j}$) of the targets, their reflectivity $r_{i,j}$ and the background levels $b_{i,j}$ from the observed data gathered in the $N_{\text{row}} \times N_{\text{col}} \times T$ array \mathbf{Y} . The next section studies the proposed Bayesian model to estimate the unknown parameters in (1) while ensuring the positivity constraints mentioned above.

III. BAYESIAN MODEL

The unknown parameter vector associated with (1) contains the surfaces reflectivity \mathbf{R} , the range of the target \mathbf{T} and the background levels \mathbf{B} (satisfying positivity constraints), where $[\mathbf{R}]_{i,j} = r_{i,j}$, $[\mathbf{T}]_{i,j} = t_{i,j}$ and $[\mathbf{B}]_{i,j} = b_{i,j}$. This section summarizes the likelihood and the parameter priors associated with these parameters.

A. Likelihood

Assuming the Poisson noise realizations in all bins and for all wavelengths are independent, Eq. (1) leads to

$$P(\mathbf{Y}|\mathbf{R}, \mathbf{B}, \mathbf{T}) = \prod_{i,j} \prod_{t=1}^T \frac{\lambda_{i,j,t}^{y_{i,j,t}}}{y_{i,j,t}!} \exp^{-\lambda_{i,j,t}} \quad (3)$$

where $\lambda_{i,j,t} = r_{i,j} g_0(t - t_{i,j}) + b_{i,j}$.

B. Prior for the Target Ranges

To account for the spatial correlations between the target distances in neighboring pixels, we propose to use a Markov random field as a prior distribution for $t_{i,j}$ given its neighbors $\mathbf{T}_{\mathcal{V}(i,j)}$, *i.e.*, $f(t_{i,j}|\mathbf{T}_{\setminus(i,j)}) = f(t_{i,j}|\mathbf{T}_{\mathcal{V}(i,j)})$, where $\mathcal{V}(i, j)$ is the neighborhood of the pixel (i, j) and $\mathbf{T}_{\setminus(i,j)} = \{t_{i',j'}\}_{(i',j') \neq (i,j)}$. More precisely, this paper focuses on the following discrete MRF

$$f(\mathbf{T}|c) = \frac{1}{G(c)} \exp[-c\phi(\mathbf{T})], \quad \forall t_{i,j} \in \{1, \dots, T\}, \quad (4)$$

where $c \leq 0$ is a parameter tuning the amount of correlation between pixels, $G(c)$ is a normalization (or partition) constant and where $\phi(\cdot)$ is an arbitrary cost function modeling the correlation between neighbors. In this work we propose to use the following cost function

$$\phi(\mathbf{T}) = \sum_{i,j} \sum_{(i',j') \in \mathcal{V}(i,j)} |t_{i,j} - t_{i',j'}|, \quad (5)$$

which corresponds to a *total-variation* (TV) regularization [14], [15] promoting piecewise constant depth image. Moreover, the higher the value of c , the more correlated the neighboring pixels. Several neighborhood structures can be employed to define $\mathcal{V}(i, j)$. Fig. 1 shows two examples of neighborhood structures. Here, the eight pixel structure

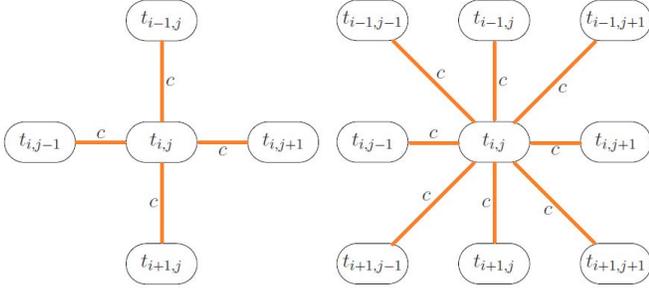


Fig. 1. 4-pixel (left) and 8-pixel (right) neighborhood structures. The pixel considered appears as a black circle whereas its neighbors are depicted in white.

(or 2-order neighborhood) will be considered in the rest of the paper for the depth parameters, as it provides in practice smoother depth images.

Accounting for the spatial correlation between depth values in neighbouring pixels makes particularly sense when the spacing between adjacent pixels/locations is small compared to the size of observed object(s). Increasing the spacing between pixels (i.e., reducing the spatial resolution of the scene sampling) generally reduces the correlation between neighbouring pixels but this case can be handled by (4) by allowing the regularization parameter c to tend to 0 (i.e. depths a priori independently and uniformly distributed). In this work, we proposed to use a TV regularization which will allow some sharp depth transitions, as it occurs when observing distinct objects having significantly different ranges (as in Section V). However, if the actual depth profile is expected to exhibit only smooth variations (e.g., observation of a single object with regular surface), different cost functions (e.g., quadratic functions) might lead to more accurate estimated depth profiles. However, since the depths are assumed to take a finite number of values, changing the cost function would not significantly change the estimation procedure presented in Section IV.

C. Prior for the Target Reflectivity

In the absence of background, i.e., if $b_{i,j} = 0$, gamma distributions are conjugate priors for each reflectivity parameter $r_{i,j}$. Consequently, it seems natural to consider gamma priors for the reflectivity. As will be shown later, gamma priors are still conjugate distributions when $b_{i,j} > 0$. We propose to assign $r_{i,j}$ the following gamma prior

$$r_{i,j} \sim \mathcal{G}\left(\alpha_0, \frac{\alpha_{i,j}}{\alpha_0}\right) \quad (6)$$

where $\alpha_{i,j} > 0$ is a local parameter related to the prior mean and mode of $r_{i,j}$ and $\alpha_0 > 0$ is a global parameter which controls the shape of the distribution tails and thus the prior deviation of $r_{i,j}$ from $\alpha_{i,j}$. Hierarchical Bayesian models generally allow the construction of elaborate prior models in which parameters can be related through the introduction of additional parameters which generally belong to higher levels in the Bayesian hierarchical model. Here we specify (6) to reflect the prior belief that reflectivity exhibit

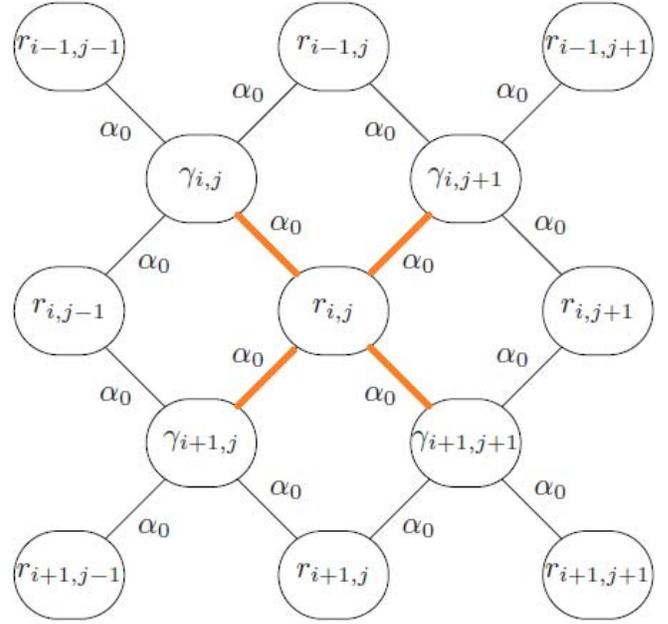


Fig. 2. Proposed 1st order GMRF neighborhood structure $\forall(i, j) \in \mathcal{V}_{\mathbf{R}}$. We set $r_{i,j} = 0.1, \forall(i, j) \notin \mathcal{V}_{\mathbf{R}}$.

spatial correlations. In particular, due to the spatial organization of images, we expect the values of $r_{i,j}$ to vary smoothly from one pixel to another. In order to model this behaviour, we specify $\alpha_{i,j}$ such that the resulting prior for \mathbf{R} is a hidden gamma-MRF (GMRF) [16].

More precisely, we introduce an $(N_{\text{row}} + 1) \times (N_{\text{col}} + 1)$ auxiliary matrix $\mathbf{\Gamma}$ with elements $\gamma_{i,j} \in \mathbb{R}^+$ and define a bipartite conditional independence graph between \mathbf{R} and $\mathbf{\Gamma}$ such that each $r_{i,j}$ is connected to four neighbor elements of $\mathbf{\Gamma}$ and vice-versa. This 1st order neighbourhood structure is depicted in Fig. 2, where we notice that any given $r_{i,j}$ and $r_{i+1,j}$ are 2nd order neighbors via $\gamma_{i+1,j}$ and $\gamma_{i+1,j+1}$. Following the general GMRF model proposed in [16] and specified here by the neighbouring structure depicted in Fig. 2, we assign a GMRF prior for $(\mathbf{R}, \mathbf{\Gamma})$, and obtain the following joint prior for $(\mathbf{R}, \mathbf{\Gamma})$

$$\begin{aligned} f(\mathbf{R}, \mathbf{\Gamma} | \alpha_0) &= \frac{1}{G(\alpha_0)} \prod_{(i,j) \in \mathcal{V}_{\mathbf{R}}} r_{i,j}^{(\alpha_0-1)} \\ &\times \prod_{(i',j') \in \mathcal{V}_{\mathbf{\Gamma}}} (\gamma_{i',j'})^{-(\alpha_0+1)} \\ &\times \prod_{((i,j),(i',j')) \in \mathcal{E}} \exp\left(\frac{-\alpha_0 r_{i,j}}{4\gamma_{i',j'}}\right), \end{aligned} \quad (7)$$

where $\mathcal{V}_{\mathbf{R}} = \{1, \dots, N_{\text{row}}\} \times \{1, \dots, N_{\text{col}}\}$, $\mathcal{V}_{\mathbf{\Gamma}} = \{1, \dots, N_{\text{row}} + 1\} \times \{1, \dots, N_{\text{col}} + 1\}$, and the edge set \mathcal{E} consists of pairs $((i, j), (i', j'))$ representing the connection between $r_{i,j}$ and $\gamma_{i',j'}$. It can be seen from (7) that

$$r_{i,j} | \mathbf{\Gamma}, \alpha_0 \sim \mathcal{G}\left(\alpha_0, \frac{\alpha_{i,j}(\mathbf{\Gamma})}{\alpha_0}\right) \quad (8a)$$

$$\gamma_{i,j} | \mathbf{R}, \alpha_0 \sim \mathcal{IG}\left(\alpha_0, \alpha_0 \beta_{i,j}(\mathbf{R})\right) \quad (8b)$$

where

$$\alpha_{i,j}(\mathbf{\Gamma}) = 4 \left(\gamma_{i,j}^{-1} + \gamma_{i-1,j}^{-1} + \gamma_{i,j-1}^{-1} + \gamma_{i-1,j-1}^{-1} \right)^{-1}$$

$$\beta_{i,j}(\mathbf{R}) = (r_{i,j} + r_{i+1,j} + r_{i,j+1} + r_{i+1,j+1}) / 4.$$

In addition to their flexibility, the one of the main motivations for considering GMRFs here is the fact that they ease the sampling strategy and thus the inference process using the conjugacy of (8a) and (1) (see Eq. (18)), while introducing spatial dependencies between the reflectivity coefficients.

Notice that we denote explicitly the dependence on the value of α_0 , which here acts as a regularization parameter that controls the amount of spatial smoothness enforced by the GMRF. Following an empirical Bayesian approach, the value of α_0 remains unspecified and will be adjusted automatically (together with the depth parameter c) during the inference procedure by maximum marginal likelihood estimation.

Finally, it is worth mentioning that this reflectivity model has similarities with the depth model (4) in the sense that spatial correlation is used to regularize the parameter estimation problem. However, the depth model (4) follows a segmentation approach in which the target depths are assumed (and constrained) to take values in a finite set. This leads to a piecewise constant depth representation which is usually sufficient for most applications, due to the depth resolution of the recent Lidar-based imaging systems.

The reflectivity model proposed in this paper provides a spatially smooth representation of the reflectivities that is possibly more realistic than a piece-wise constant representation (that would result from a TV-based reflectivity regularization), as it does not enforce the reflectivities to take a finite number of possible values.

D. Prior for the Background Levels

In a similar fashion to the reflectivities, when $r_{i,j} = 0$, gamma distributions are conjugate priors for $b_{i,j}$ and the following Gamma priors

$$b_{i,j} \sim \mathcal{G}(\eta, \nu) \quad (9)$$

are assigned to $b_{i,j}$, where $\eta > 0$ and $\nu > 0$ are fixed hyperparameters. In order to reflect the lack of prior knowledge about the background levels, (η, ν) is arbitrarily set to obtain a weakly informative prior $((\eta, \nu) = (1, 10)$ in all results presented in this paper). However, (η, ν) can be easily adapted if additional information, e.g., observation conditions, is available. It could also be estimated as in [17], however the choice $(\eta, \nu) = (1, 10)$ made here has a limited impact on the estimation performance.

Assuming prior independence between the average background levels of the different pixels, we obtain

$$f(\mathbf{B}|\eta, \nu) = \prod_{i,j} f(b_{i,j}|\eta, \nu). \quad (10)$$

E. Joint Posterior Distribution

We can now specify the joint posterior distribution for $\mathbf{T}, \mathbf{B}, \mathbf{R}$ and $\mathbf{\Gamma}$ given the observed waveforms \mathbf{Y} and the

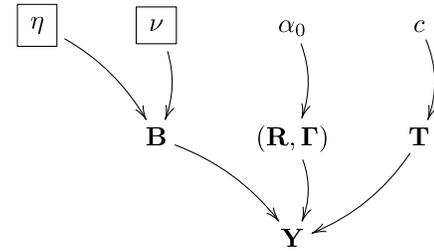


Fig. 3. Graphical model for the proposed Bayesian model (fixed quantities appear in boxes).

value of the spatial regularization parameters c and α_0 (recall that their value will be determined by maximum marginal likelihood estimation during the inference procedure). Using Bayes' theorem, and assuming prior independence between \mathbf{T} , $(\mathbf{\Gamma}, \mathbf{R})$ and \mathbf{B} , the joint posterior distribution associated with the proposed Bayesian model is given by

$$f(\mathbf{T}, \mathbf{B}, \mathbf{R}, \mathbf{\Gamma}|\mathbf{Y}, \alpha_0, \eta, \nu) \propto f(\mathbf{Y}|\mathbf{R}, \mathbf{B}, \mathbf{T})f(\mathbf{B}|\eta, \nu)f(\mathbf{T}|\alpha_0)f(\mathbf{R}, \mathbf{\Gamma}|\alpha_0). \quad (11)$$

For illustration, Fig. 3 depicts the directed acyclic graph (DAG) summarising the structure proposed Bayesian model (recall that $\mathbf{R}, \mathbf{\Gamma}$ have a bi-partite neighbourhood structure, which is illustrated in the graphical model of Fig. 2).

IV. BAYESIAN INFERENCE

A. Bayesian Estimators

The Bayesian model defined in Section III specifies the joint posterior density for the unknown parameters $\mathbf{T}, \mathbf{B}, \mathbf{R}$ and $\mathbf{\Gamma}$ given the observed data \mathbf{Y} and the parameters c and α_0 . This posterior distribution models our complete knowledge about the unknowns given the observed data and the prior information available. In this section we define suitable Bayesian estimators to summarize this knowledge and perform depth imaging. More precisely, we propose to use the following two Bayesian estimators: 1) the minimum mean square error estimator (MMSE) of the reflectivity matrix

$$\hat{\mathbf{R}}_{\text{MMSE}} = \mathbb{E}[\mathbf{R}|\mathbf{Y}, \hat{c}, \hat{\alpha}_0], \quad (12)$$

where the expectation is taken with respect to the marginal posterior density $f(\mathbf{R}|\mathbf{Y}, c, \alpha_0)$ (by marginalizing \mathbf{T}, \mathbf{B} and $\mathbf{\Gamma}$ this density takes into account their uncertainty); 2) the maximum a posteriori (MAP) estimator of target ranges

$$\hat{\mathbf{T}}_{\text{MAP}} = \underset{\mathbf{T}}{\operatorname{argmax}} f(\mathbf{T}|\mathbf{Y}, \hat{c}, \hat{\alpha}_0), \quad (13)$$

which is particularly adapted to estimate discrete parameters. Notice that in (12) and (13), we have set $c = \hat{c}$ and $\alpha_0 = \hat{\alpha}_0$, which denotes the maximum marginal likelihood estimator of the MRF regularisation parameters c and α_0 given the observed data \mathbf{Y} , i.e.,

$$(\hat{c}, \hat{\alpha}_0) = \underset{c \in \mathbb{R}^+, \alpha_0 \in \mathbb{R}^+}{\operatorname{argmax}} f(\mathbf{Y}|c, \alpha_0), \quad (14)$$

This approach for specifying (c, α_0) is taken from the empirical Bayes framework in which parameters with unknown values are replaced by point estimates computed from observed

data (as opposed to being fixed a priori or integrated out of the model by marginalization). As explained in [13], this strategy has several important advantages for MRF parameters with intractable conditional distributions such as (c, α_0) . In particular, it allows for the automatic adjustment of the value of (c, α_0) for each image (thus producing significantly better estimation results than using a single fixed value of (c, α_0) for all data sets), and has a computational cost that is several times lower than that of competing approaches, such as including (c, α_0) in the model and subsequently marginalising them during the inference procedure [18].

B. Bayesian Algorithm

Computing the estimators (12) and (13) is challenging because it involves calculating expectations with respect to posterior marginal densities, which in turn require evaluating the full posterior (11) and integrating it over a very high-dimensional space. Computing $(\hat{c}, \hat{\alpha}_0)$ is also difficult because it involves solving an intractable optimisation problem, (because it is not possible to evaluate the marginal likelihood $f(\mathbf{Y}|c, \alpha_0)$ or its gradient $\nabla f(\mathbf{Y}|c, \alpha_0)$). Here we adopt the approach proposed in [13] for the estimation of a Potts model hyperparameter and design a stochastic optimisation and simulation algorithm to compute (12) and (13) simultaneously. That is, we construct a stochastic gradient Markov chain Monte Carlo (SGMCMC) algorithm that simultaneously estimates $(\hat{c}, \hat{\alpha}_0)$ and generates a chain of N_{MC} samples $\{\mathbf{R}^{(n)}, \mathbf{T}^{(n)}, \mathbf{B}^{(n)}, \mathbf{\Gamma}^{(n)}\}_{n=1}^{N_{MC}}$ asymptotically distributed according to the marginal density $f(\mathbf{T}, \mathbf{B}, \mathbf{R}, \mathbf{\Gamma}|\mathbf{Y}, \hat{c}, \hat{\alpha}_0)$ (this algorithm is summarised in Algo. 1 below). Once the samples have been generated, the estimators (12) and (13) are approximated by Monte Carlo integration [19, Ch. 10], i.e.,

$$\hat{\mathbf{R}}_{MMSEj} = \frac{1}{N_{MC} - N_{bi}} \sum_{n=N_{bi}+1}^{N_{MC}} \mathbf{R}^{(n)}, \quad (15)$$

and

$$(\hat{t}_{i,j})_{MAPj} = \underset{k \in \{1, \dots, T\}}{\operatorname{argmax}} \sum_{n=N_{bi}+1}^{N_{MC}} \delta(t_{i,j}^{(n)} - k) \quad (16)$$

where the samples from the first N_{bi} iterations (corresponding to the transient regime or burn-in period) are discarded and where $\delta(\cdot)$ denotes the Kronecker delta function. The main steps of this algorithm are detailed in below.

1) *Sampling the Target Ranges*: It can be seen from (11) that

$$f(t_{i,j} = t | \mathbf{Y}, \mathbf{T}_{\setminus i,j}, \mathbf{B}, \mathbf{R}, \mathbf{\Gamma}, c, \alpha_0, \eta, \nu) \propto f(\mathbf{Y}_{i,j} | t_{i,j} = t, r_{i,j}, b_{i,j}) f(t_{i,j} = t | \mathbf{T}_{\setminus (i,j)}). \quad (17)$$

Consequently, sampling the target ranges can be achieved by sampling sequentially each depth from its conditional distribution, i.e., by drawing randomly from $\{1, \dots, T\}$ with known probabilities. In our experiments we used a Gibbs sampler implemented using a colouring scheme such that many depths can be updated in parallel (9 steps required when considering a 2-order neighborhood structure).

2) *Sampling the Reflectivity Coefficients*: Similarly, from (11) we obtain

$$f(\mathbf{R}|\mathbf{Y}, \mathbf{T}, \mathbf{B}, \mathbf{\Gamma}, c, \alpha_0, \eta, \nu) = \prod_{i,j} f(r_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0) \quad (18)$$

i.e., the elements of \mathbf{R} are a posteriori independent (conditioned on the other parameters) and can thus be updated simultaneously. Moreover,

$$f(r_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0) \propto r_{i,j}^{\alpha_0-1} \exp^{-\frac{\alpha_0 r_{i,j}}{a_{i,j}(\mathbf{\Gamma})}} \prod_{t=1}^T \lambda_{i,j,t}^{y_{i,j,t}} \exp^{-\lambda_{i,j,t}}, \quad (19)$$

with $\lambda_{i,j,t} = r_{i,j} g_0(t - t_{i,j}) + b_{i,j}$. By noticing that $\prod_{t=1}^T \lambda_{i,j,t}^{y_{i,j,t}}$ is a polynomial function of $r_{i,j}$, whose degree is $\sum_{t=1}^T y_{i,j,t}$ (since $g_0(t - t_{i,j}) > 0, \forall t$), it turns out that (19) can be expressed as a finite mixture of gamma distributions whose weights and parameters are known (see Appendix for the derivation of (19)). Sampling $r_{i,j}$ from its conditional distribution finally reduces to randomly selecting one of components of the mixture and then sampling from the corresponding gamma distribution. Note that the auxiliary variables in $\mathbf{\Gamma}$ do not appear in the likelihood (3) and that sampling $\mathbf{\Gamma}$ from its conditional distribution reduces to sampling from (8b).

3) *Sampling the Background*: Similarly analysis applies when sampling the background levels. Precisely,

$$f(\mathbf{B}|\mathbf{Y}, \mathbf{T}, \mathbf{R}, \mathbf{\Gamma}, c, \alpha_0, \eta, \nu) = \prod_{i,j} f(b_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, r_{i,j}, \eta, \nu) \quad (20)$$

i.e., the elements of \mathbf{B} are a posteriori independent (conditioned on the other parameters) and can thus be updated simultaneously. Moreover,

$$f(b_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, r_{i,j}, \eta, \nu) \propto b_{i,j}^{\eta-1} \exp^{-\frac{b_{i,j}}{\nu}} \prod_{t=1}^T \lambda_{i,j,t}^{y_{i,j,t}} \exp^{-\lambda_{i,j,t}}, \quad (21)$$

which can also be expressed as a finite mixture of gamma distributions.

4) *Updating the MRF Regularization Parameters c and α_0* : If marginal likelihood $f(\mathbf{Y}|c, \alpha_0)$ was tractable we could update (c, α_0) from one MCMC iteration to the next by using a classic gradient descent step

$$\begin{aligned} \alpha_0^{(n+1)} &= \alpha_0^{(n)} + \zeta_n \frac{\partial}{\partial \alpha_0} \log f(\mathbf{Y}|c^{(n)}, \alpha_0^{(n)}), \\ c^{(n+1)} &= c^{(n)} + \zeta_n \frac{\partial}{\partial c} \log f(\mathbf{Y}|c^{(n)}, \alpha_0^{(n)}) \end{aligned} \quad (22)$$

with $\zeta_n = n^{-3/4}$, such that $\alpha_0^{(n)}$ (resp. $c^{(n)}$) converges to $\hat{\alpha}_0$ (resp. \hat{c}) as $n \rightarrow \infty$. However, this gradient has two levels of intractability, one due to the marginalization of $(\mathbf{T}, \mathbf{B}, \mathbf{R}, \mathbf{\Gamma})$ and another one due to the intractable normalizing constant of the gamma MRF. We address this difficulty by following the approach proposed in [13]; that is, by replacing $\nabla \log f(\mathbf{Y}|c^{(n)}, \alpha_0^{(n)})$ with estimators computed with the samples generated by the MCMC algorithm at iteration n , and

Algorithm 1 Proposed MCMC Algorithm

-
- 1: Fixed input parameters: Lidar impulse response $g_0(\cdot)$, number of burn-in iterations N_{bi} , total number of iterations N_{MC}
 - 2: Initialization ($n = 0$)
 - Set $\mathbf{R}^{(0)}, \mathbf{T}^{(0)}, \mathbf{B}^{(0)}, \mathbf{\Gamma}^{(0)}, c^{(0)}, \alpha_0^{(0)}$
 - 3: Iterations ($1 \leq n \leq N_{\text{MC}}$)
 - 4: Sample $\mathbf{T}^{(n)}$ from (17)
 - 5: Sample $\mathbf{R}^{(n)}$ from (18)
 - 6: Sample $\mathbf{B}^{(n)}$ from (20)
 - 7: Sample $\mathbf{\Gamma}^{(n)}$ from (8b)
 - 8: **if** $n < N_{\text{bi}}$ **then**
 - 9: Sample $(\mathbf{R}', \mathbf{\Gamma}') \sim \mathcal{K}_1(\mathbf{R}, \mathbf{\Gamma} | \mathbf{R}^{(n)}, \mathbf{\Gamma}^{(n)}, \alpha_0^{(n-1)})$
 - 10: Sample $\mathbf{T}' \sim \mathcal{K}_2(\mathbf{T} | \mathbf{T}^{(n)}, c^{(n-1)})$
 - 11: Set $\alpha_0^{(n)} = \mathcal{P}_{[0, A_n]} \left(\alpha_0^{(n-1)} + \xi_n \left[\Lambda(\mathbf{R}^{(n)}, \mathbf{\Gamma}^{(n)}) - \Lambda(\mathbf{R}', \mathbf{\Gamma}') \right] \right)$
 where $\Lambda(\mathbf{R}, \mathbf{\Gamma}) = \sum_{(i,j) \in \mathcal{V}_{\mathbf{R}}} \log(r_{i,j}) - \sum_{(i',j') \in \mathcal{V}_{\mathbf{T}}} \log(\gamma_{i',j'}) - \sum_{((i,j),(i',j')) \in \mathcal{E}} \frac{r_{i,j}}{4\gamma_{i',j'}}$
 - 12: Set $c^{(n)} = \mathcal{P}_{[0, C_n]} \left(c^{(n-1)} + \xi_n \left[\phi(\mathbf{T}') - \phi(\mathbf{T}^{(n)}) \right] \right)$
 - 13: **else**
 - 14: Set $c^{(n)} = c^{(n-1)}$
 - 15: Set $\alpha_0^{(n)} = \alpha_0^{(n-1)}$
 - 16: **end if**
 - 17: Set $n = n + 1$.
 - 18: Output $\{\mathbf{R}^{(n)}, \mathbf{T}^{(n)}\}_{n=1}^{N_{\text{MC}}}$.
-

two sets of auxiliary variables. More precisely, we generate $(\mathbf{R}', \mathbf{\Gamma}') \sim \mathcal{K}_1(\mathbf{R}, \mathbf{\Gamma} | \mathbf{R}^{(n)}, \mathbf{\Gamma}^{(n)}, \alpha_0^{(n-1)})$ with an MCMC kernel \mathcal{K}_1 with target density (7) (in our experiments we used a Gibbs sampler implemented using a colouring scheme such that all the elements of \mathbf{R}' and $\mathbf{\Gamma}'$ are generated in parallel). We also generate $\mathbf{T}' \sim \mathcal{K}_2(\mathbf{T} | \mathbf{T}^{(n)}, c^{(n-1)})$ with an MCMC kernel \mathcal{K}_2 with target density (4). The updated value $\alpha_0^{(n)}$ (resp. $c^{(n)}$) is then projected onto an interval $[0, A_n]$ (resp. $[0, C_n]$, see (10:) and (11:) in Algo 1) to guarantee the positivity constraints $c, \alpha_0 \in \mathbb{R}^+$ and the stability of the stochastic optimization algorithm (we have used $A_t = C_t = 20$).

It is worth mentioning that if it was possible to simulate the auxiliary variables $(\mathbf{R}', \mathbf{\Gamma}')$ (resp. \mathbf{T}') exactly from (7) (resp. (4)), then the estimator of $\nabla \log f(\mathbf{Y} | c^{(n)}, \alpha_0^{(n)})$ used in Algo. 1 would be unbiased and as a result $(c^{(n)}, \alpha_0^{(n)})$ would converge exactly to $(\hat{c}, \hat{\alpha}_0)$. However, exact simulation from (7) and (4) is not computationally feasible and therefore we resort to the MCMC kernels \mathcal{K}_1 and \mathcal{K}_2 and obtain a biased estimator of $\nabla \log f(\mathbf{Y} | c^{(n)}, \alpha_0^{(n)})$ that drives $(c^{(n)}, \alpha_0^{(n)})$ to a neighbourhood of $(\hat{c}, \hat{\alpha}_0)$ [13]. We have found that computing this biased estimator is significantly less expensive than alternative approaches, (e.g., using an approximate Bayesian computation algorithm as in [18]), and that it leads to very accurate depth and reflectivity results.

V. SIMULATION RESULTS

A. Data Acquisition

We propose comparing the performance of the proposed method to reconstruct a depth image of a life-sized polystyrene head located at a distance of 40m from a time-of-flight scanning sensor, based on TCSPC. The transceiver system and



Fig. 4. Photograph showing the polystyrene head used for the experiments described here and calibration targets, including the Spectralon panel (top right corner).

data acquisition hardware used for this work is broadly similar to that described in [9] and [20]–[22], which was previously developed at Heriot-Watt University. For the measurements reported in this section, the optical path of the transceiver was configured to operate with a fiber-coupled illumination wavelength of 841nm, and a silicon single-photon avalanche diode (SPAD) detector. The overall system had a jitter of 95ps full width at half-maximum (FWHM). The measurements have been performed outdoors, on the Edinburgh Campus of Heriot-Watt University, in November 2014 under dry clear skies and with atmospheric conditions remaining relatively constant for the duration of the measurement. The key measurement parameters are summarized below:

- Target stand-off distance: $\approx 40\text{m}$
- Target scene: polystyrene head ($\approx 170 \times 285 \times 250\text{mm}$ in $W \times H \times D$ when viewed from the front) mounted on a breadboard. Backplane: MDF board (see Fig. 4)
- Laser system: Supercontinuum laser source and tunable filter (NKT Photonics) fiber-coupled to the custom-designed transceiver unit
- Laser repetition Rate: 19.5MHz
- Illumination power at target: $\approx 240\mu\text{W}$ average optical power
- Illumination beam diameter at target: $\approx 1\text{mm}$
- Acquisition mode: 142×142 pixels scan centered on the head, covering an area of $285 \times 285\text{mm}$ at the scene
- Per-pixel acquisition time: 30 ms
- Total scan time: ≈ 10 minutes
- Histogram bin width: 16ps
- Histogram length: 586 bins (after gating)

Although the acquisition time per pixel is 30ms, the data format of timed events allows the construction of photon timing histograms associated with shorter acquisition times, after measurement, as the system records the time of arrival of each detected photon. Here, we evaluate our algorithms for acquisition times of 30ms, 6ms, 3ms, $600\mu\text{s}$, $300\mu\text{s}$ and $60\mu\text{s}$ per pixel.

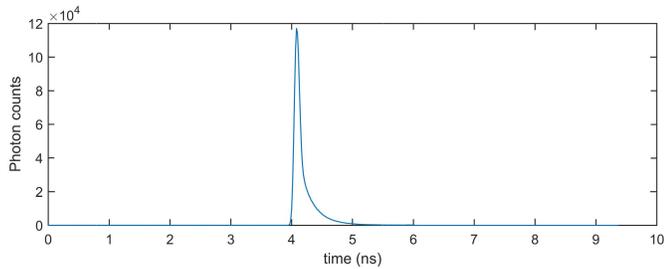


Fig. 5. Instrumental response obtained using spectralon panel placed at 40m from laser source/detector and for an acquisition time of 60s (jitter \approx 95ps FWHM).

TABLE I

AVERAGE NUMBER OF DETECTED PHOTONS PER PIXEL AND PROPORTION OF EMPTY PIXELS AS A FUNCTION OF THE ACQUISITION TIME

	60 μ s	300 μ s	600 μ s	3ms	6ms	30ms
Av. Photon counts	0.8	4.2	8.4	42.0	83.7	418.6
Empty pixels (%)	48.7	7.5	1.7	≥ 0.1	0	0

The instrumental impulse response $g_0(\cdot)$ is estimated from preliminary experiments by analysing the distribution of photons reflected onto a Spectralon panel (a commercially available Lambertian scatterer), placed at 40m from laser source/detector. A long acquisition time (60s) is considered here to reduce the impact of the photon count variability and a pre-processing step is used to remove the constant background in the measured response. The resulting instrumental impulse response is depicted in Fig. 5.

Table I provides details regarding the amount of detected photons when varying the acquisition time. The second row of I shows the average number of detected photons per pixel, ranging from about 0.85 photons for a 60 μ s acquisition time to about 420 for a 30ms acquisition time. As expected, the number of detected photons increases linearly with exposure. The bottom row shows that almost 49% of the pixels do not contain any detected photons for a 60 μ s acquisition time and that this proportion decreases when increasing the acquisition time.

B. Estimation Performance

The proposed method is first compared to the method classically used for depth imaging [9] and which is divided into two steps. The first step consists of estimating $t_{i,j}$ using cross-correlation between $g_0(\cdot)$ and the photon histogram $y_{i,j}$. The object depth is estimated using

$$\hat{t}_{i,j,\text{corr}} = \underset{\tau \in \mathbb{Z}}{\operatorname{argmax}} \sum_{t=1}^T y_{i,j,t} g_0(t - \tau). \quad (23)$$

Once the estimated time target distance $\hat{t}_{i,j,\text{corr}}$ has been computed, the target reflectivity is estimated using maximum likelihood (ML) estimation (assuming that $b_n = 0$) as

$$\hat{r}_{i,j,ML} = \frac{\sum_{t=1}^T y_{i,j,t}}{\sum_{t=1}^T g_0(t - \hat{t}_{i,j,\text{corr}})}. \quad (24)$$

When the background level is relatively low compared to the maximum value of $r_{i,j} g_0(t - t_{i,j})$, the ML reflectivity estimates (conditioned on the previously estimated depths) provide satisfactory results and are thus considered as the comparative method in the remainder of this paper. Moreover, we also compared the proposed method to denoising methods applied to the standard depth and reflectivity images. Precisely, due to the potential presence of empty pixels, a nearest neighbor interpolation has been applied to the depth and reflectivity images estimated by the standard method. We then applied two different denoising techniques. The first method consists of applying a median filter whose size has been chosen by cross-validation (5 \times 5 here). The second method is the Block-Matching and 3D filtering (BM3D) algorithm [23] whose noise variance parameter has also been fixed from preliminary runs. The proposed algorithm has been applied with $N_{MC} = 1000$ iterations, including $N_{bi} = 200$ burn-in iterations.

Fig. 6 (top) compares the estimated depth maps obtained by the proposed and competing methods. These results show that for large acquisition times, the four methods provide similar results. However, when the acquisition time decreases, the cross-correlation method starts to fail in identifying the target ranges, especially in pixels where no photon is detected in a pixel, indicating that the other methods seem more robust to the absence of signal in some pixels. The median and BM3D filters seem to better identify the edges of the head than the proposed method, however they generally provide less accurate estimated depths in the center of the head, where more subtle depth variations occur.

Fig. 6 (bottom) compares the estimated reflectivity maps obtained by the four methods. These results show that the proposed and the standard methods provide similar results for the longest acquisition times and the two additional filters tend to oversmooth the reflectivity images. The proposed method seems more robust than the other methods to the lack/absence of detected photons, especially those including filters. In particular, for the 60 μ s acquisition time, few photons are detected in the pixels around the head and the proposed algorithm provides a more reliable reflectivity image due to consideration of spatial correlations, in contrast to the standard method which processes the pixels independently. Note that for each experiment, the Spectralon response $g_0(\cdot)$ is scaled to account for the acquisition time (e.g., amplitude divided by ten between the 30ms and 3ms experiments). Note also that for some pixels, the reflectivities estimated by the four methods can exceed 1. This point will be discussed further in the conclusions. For completeness, Fig. 7 compares the depth/reflectivity estimation results obtained by the four methods for an acquisition time of 300 μ s.

In addition to the depth and reflectivity maps, the proposed method also estimates the average background level in each pixel, depicted in Fig. 8. This figure shows that for the longer acquisition times, higher backgrounds are estimated at the boundary between the head and the backplane, which can be explained by the presence of two peaks in the histograms of detected photons. Due to the laser beam size, some photons are scattered from the head whereas others are scattered from the backplane and thus arrive later onto the detector.

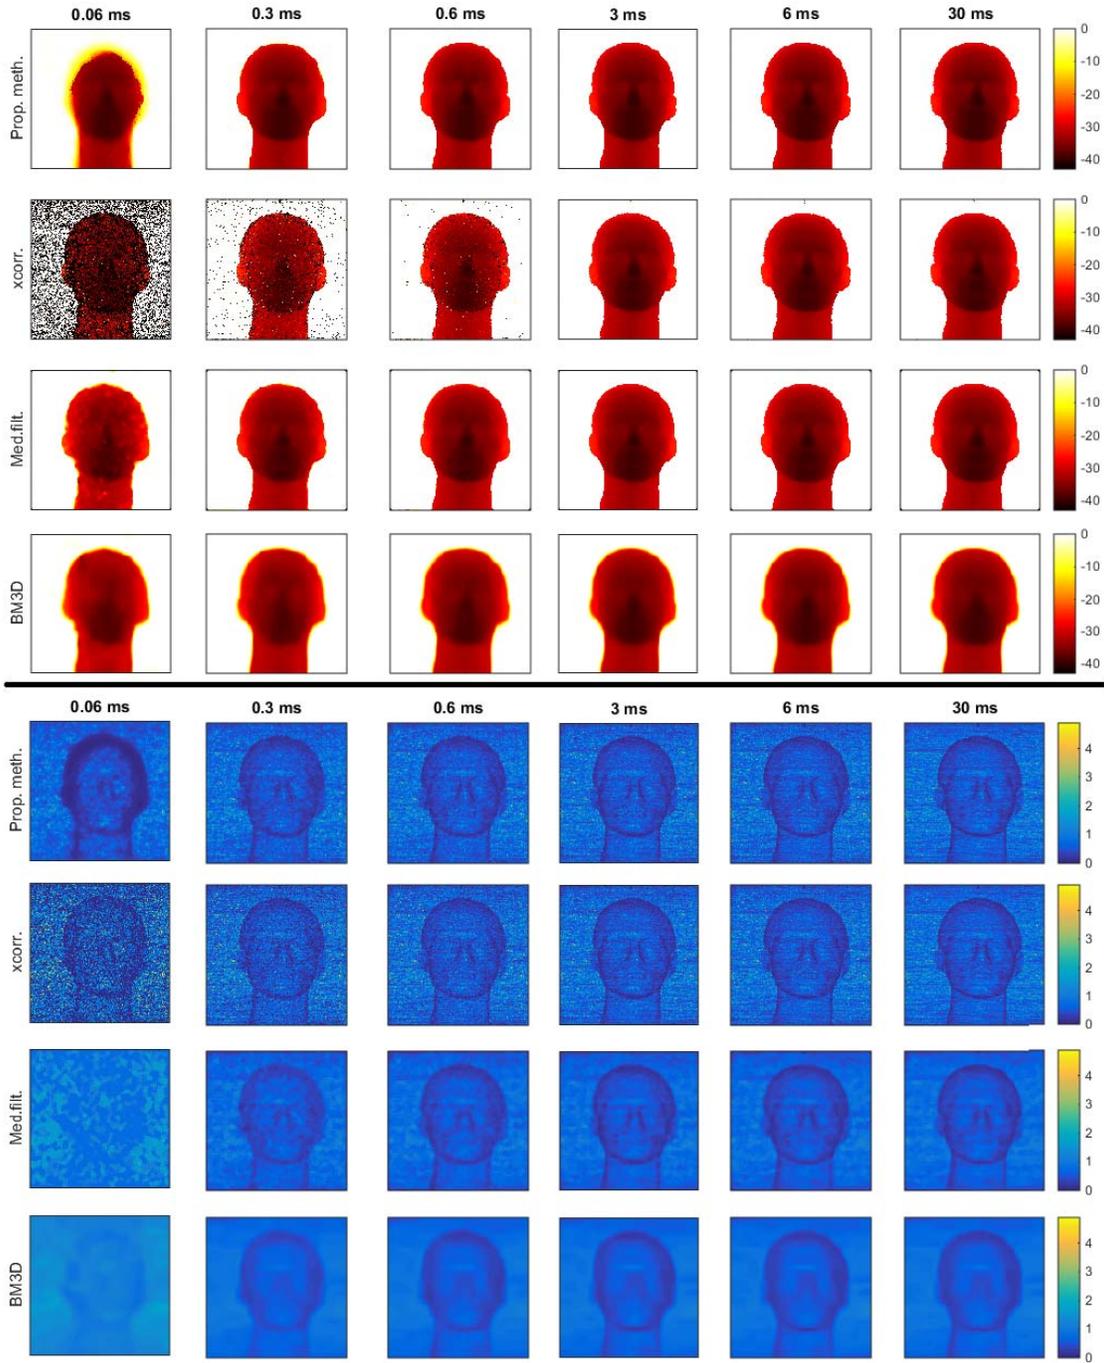


Fig. 6. Top rows: depth maps for different per-pixel acquisition times, estimated by the proposed Bayesian algorithm and competing methods. Distances shown are in centimeters and the reference distance is the distance of the backplane. Black pixels correspond to pixels where no photon are detected and for which the cross-correlation method cannot identify the target distance. Bottom rows: reflectivity maps for different per-pixel acquisition times, estimated by the proposed Bayesian algorithm and the competing methods.

When the number of detected photons decreases, the amplitudes of the two peaks decrease, which makes the detection of multiple peaks more difficult.

The performance of the methods are quantitatively evaluated using the distance and reflectivity mean squared errors (MSEs) defined by

$$MSE(d_{i,j}) = \left\| \hat{d}_{i,j} - d_{i,j} \right\|_2^2 \quad (25)$$

where $\|\cdot\|_2$ denotes the ℓ_2 -norm, $\hat{d}_{i,j}$ is the estimated value of $d_{i,j} = 3 \times 10^8 t_{i,j} / 2$ and

$$MSE(r_{i,j}) = \left\| \hat{r}_{i,j} - r_{i,j} \right\|_2^2 \quad (26)$$

where $\hat{r}_{i,j}$ is the estimated value of $r_{i,j}$. Note that $\{d_{i,j}\}$ and $\{r_{i,j}\}$ are unknown for the data set considered. Consequently we replace these values by those estimated by the standard method for the longest acquisition time (30ms).

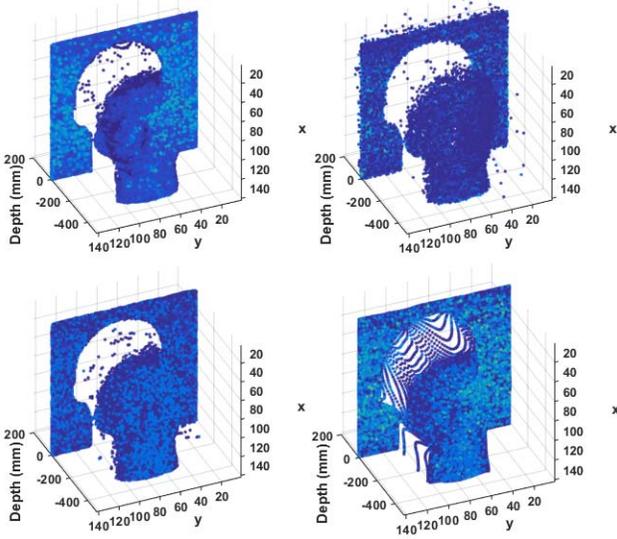


Fig. 7. Depth/reflectivity reconstruction of the target, estimated by the proposed Bayesian algorithm (top left) and standard method (top right) method, and the standard method followed by median filtering (bottom left) or BM3D (bottom right), for a $300\mu s$ acquisition time. The colours represent the target reflectivity (dark blue for low reflectivity coefficients).

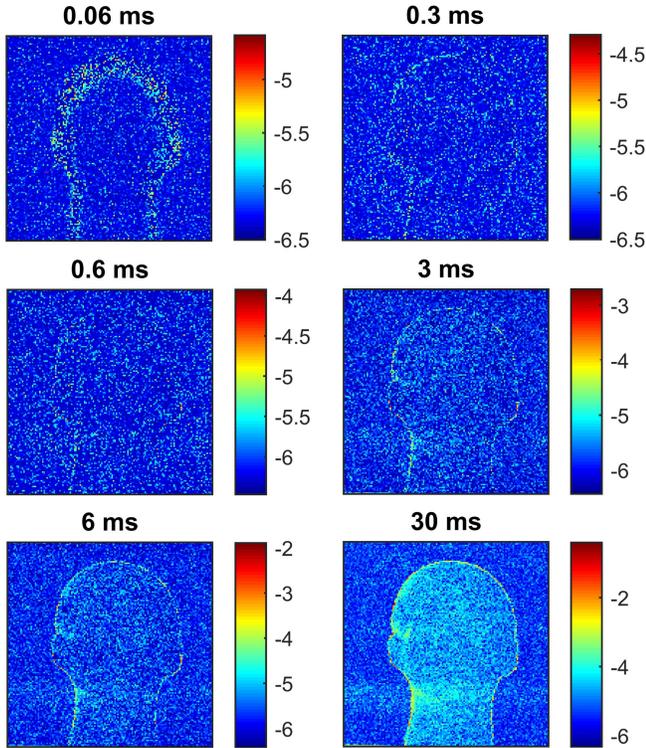


Fig. 8. Background level maps $(\log(b_{i,j}))_{i,j}$ for different acquisition time, estimated by the proposed Bayesian algorithm.

Figs. 9 and 10 depict the cumulative density functions (cdfs) of the distance and reflectivity MSEs, defined by

$$F_d(\tau) = \frac{1}{N_{\text{row}}N_{\text{col}}} \sum_{i,j} \mathbf{1}_{(0,\tau)}(MSE(d_{i,j})) \quad (27)$$

$$F_r(\tau) = \frac{1}{N_{\text{row}}N_{\text{col}}} \sum_{i,j} \mathbf{1}_{(0,\tau)}(MSE(r_{i,j})) \quad (28)$$

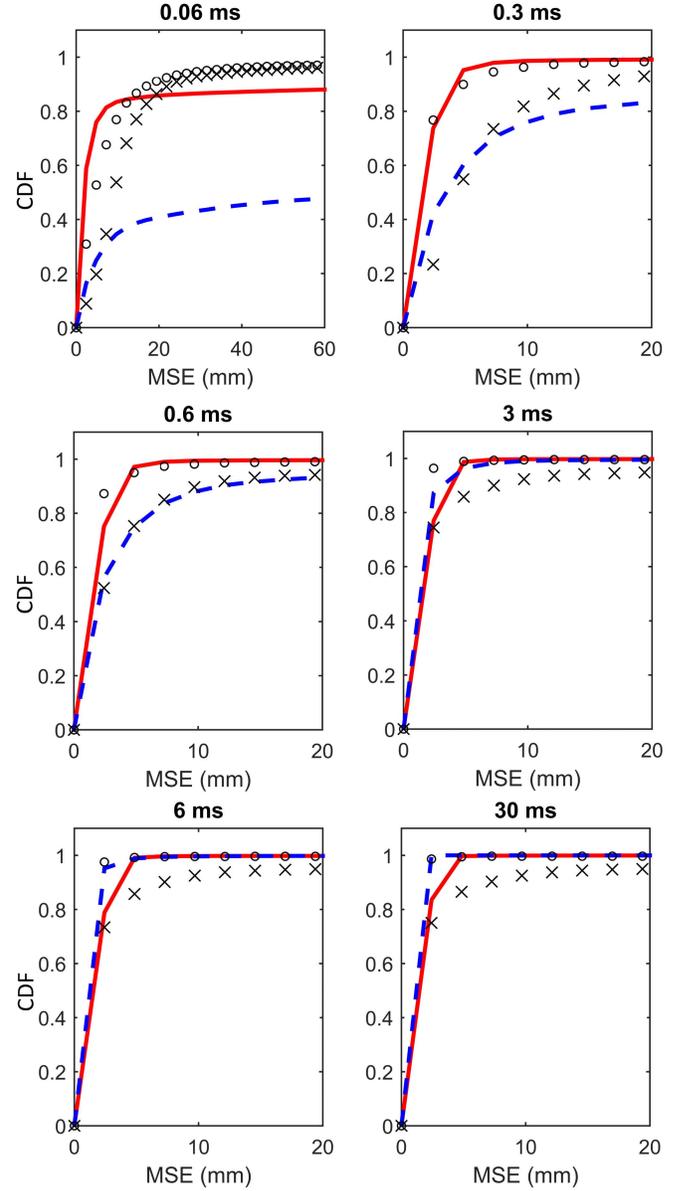


Fig. 9. Distance RMSE cdfs provided by the standard (dashed blue lines), standard+median filter (black circles), standard+BM3D (black crosses) and the proposed (red lines) methods for the target located at 40m.

where $\mathbf{1}_{(0,\tau)}(\cdot)$ denotes the indicator function defined on $(0, \tau)$. Figs. 9 and 10 show that the proposed method is more robust than the other methods when reducing the acquisition time and provide more consistent results in terms of depth and reflectivity estimation, except for the depth estimation with an acquisition time of $60\mu s$ due to the strong smoothness penalization around the boundaries of the head. These figures also highlight the ability of the proposed method to process pixels for which no photons are detected, as the cdfs are upper-bounded by the proportion of pixels that can be processed by each method.

Finally, Table II compares the computational costs of the four methods to process the whole image (142×142 pixels, $T = 586$), for the different acquisition durations. Due to the use of the MCMC method, the proposed method is

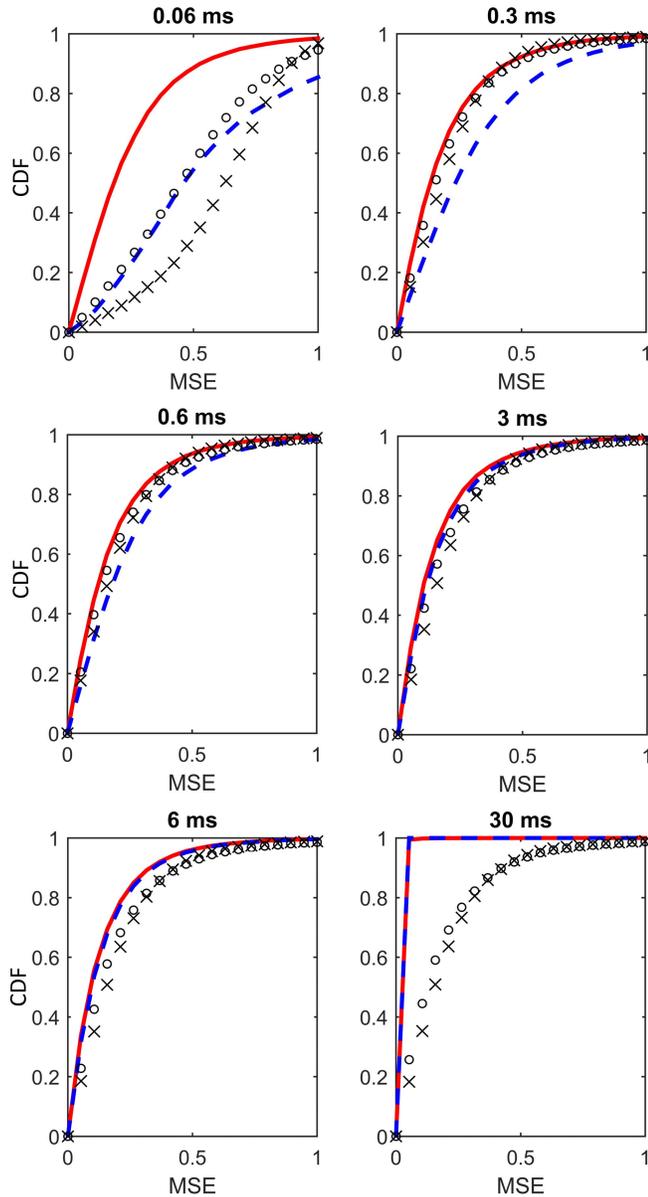


Fig. 10. Reflectivity RMSE cdfs provided by the standard (dashed blue lines), standard+median filter (black circles), standard+BM3D (black crosses) and the proposed (red lines) methods for the target located at 40m.

significantly more computationally demanding than the standard methods. However, this cost must be balanced by the performance improvement in terms of depth and reflectivity estimation, when the number of detected photons is low. When the flux of detected photons is large enough, the consideration of spatial correlations has a limited impact on the estimation performance, as long as the background levels are low compared to the amplitudes of the peaks associated with actual targets. It is interesting to note the computational cost of the Bayesian increases with the acquisition time, in contrast to the standard method. This is mainly due to the MCMC steps used to update the reflectivity coefficients and the background levels which require the computation of polynomial coefficients whose number depends on the number of detected photons

TABLE II
PROCESSING TIME (IN MINUTES)

	60 μ s	300 μ s	600 μ s	3ms	6ms	30ms
X-corr+MLE	1	1	1	1	1	1
X-corr+MLE + median filter	1	1	1	1	1	1
X-corr+MLE + BM3D	1	1	1	1	1	1
Prop. method	113	123	131	152	197	347

in each pixels (see Appendix). Note however that due the conditional independence between the reflectivity parameter, they can be updated in a parallel manner. The same observation holds for \mathbf{B} and $\mathbf{\Gamma}$. Similarly, the depth parameters can be updated using checkerboard pixel clustering.

VI. CONCLUSION

In this paper, we proposed a new Bayesian model for Lidar-based low photon count imaging of single-layered targets. In the Bayesian framework, prior distributions were assigned to the unknown target depths and reflectivity to account for the intrinsic correlations between neighboring pixels. An adaptive Markov chain Monte Carlo method was then developed to estimate the model unknown parameters, including the spatial regularization parameters, thus relieving practitioners from setting these parameters by cross-validation. The model and method were validated using real Lidar data and the results showed the benefits of the proposed approach compared to the classical method used when the number of detected photons is low.

In the paper, we assumed that the beam associated with a given pixel is incident on a single surface. This assumption is reasonable for small beam sizes, compared to the target distance and when the scene is composed of locally continuous surfaces. When the beam encounters multiple surfaces, one peak will be considered as principal surface, depending on its amplitude and on the Lidar returns in the neighboring pixels. The remaining peaks will be considered as part of the background noise. Considering returns from multiple surfaces is an interesting problem already addressed in [5], [11], and [12] for applications where the number of detected photons are significantly higher. It would be interesting to extend this work for the low-photon imaging problem.

Since the model considered in this paper assumed the presence of a target in each pixel, the proposed method will tend to process empty pixels (i.e., containing no photon) using the neighboring pixels, which might be inaccurate for non-locally continuous surfaces (such as wire fence). Accounting for the absence of target in some pixels is currently under investigation.

In the results presented in Section V, some of the estimated reflectivity coefficients were significantly greater than one, even for long acquisition time and even when scaling the instrumental response. Thus, it can be difficult to directly link the estimated reflectivities to the actual target reflectivity values as these estimation errors do not seem to be only due

to estimation errors when extracting the instrumental impulse response. Thus constraining the reflectivity coefficients to be less than one might not be sufficient to provide accurate reflectivity estimates. As studied in [24]–[26], atmospheric perturbations can have a significant impact on the distribution of the detected photons, specially for long-range targets. Although such scintillation effects will have a small impact on the depth estimation, accounting for them will be necessary in future work to improve the reflectivity estimation. Finally, extension of the imaging processing approach outlined to different sparse photon imaging scenarios, for example imaging using the effects of entangled photon pairs [27] is also worthy of future investigation.

APPENDIX ON THE CONDITIONAL DISTRIBUTION OF THE REFLECTIVITY COEFFICIENTS

The conditional distribution $f(r_{i,j}|y_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0)$ can be expressed (up to a multiplicative constant) as

$$f(r_{i,j}|y_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0) \propto r_{i,j}^{\alpha_0-1} \exp^{-\frac{\alpha_0 r_{i,j}}{a_{i,j}(\mathbf{\Gamma})}} \exp^{-\sum_{t=1}^T \lambda_{i,j,t}} \prod_{t=1}^T \lambda_{i,j,t}^{y_{i,j,t}}, \quad (29)$$

where $\sum_{t=1}^T \lambda_{i,j,t} = u_{i,j} + r_{i,j} v_{i,j}$,

$$u_{i,j} = \sum_{t=1}^T b_{i,j}$$

$$v_{i,j} = \sum_{t=1}^T g_0(t - t_{i,j}),$$

and

$$P(r_{i,j}) = \prod_{t=1}^T \lambda_{i,j,t}^{y_{i,j,t}} = \prod_{t=1}^T (r_{i,j} g_0(t - t_{i,j}) + b_{i,j})^{y_{i,j,t}} \quad (30)$$

is a polynomial function of $r_{i,j}$. Since $g_0(t - t_{i,j}) > 0, \forall t, -b_{i,j}/g_0(t - t_{i,j})$ is a root of $P(r_{i,j})$ if $y_{i,j,t} > 0$. Moreover, this root is of multiplicity $y_{i,j,t}$ and the polynomial order is thus $O_{i,j} = \sum_{t=1}^T y_{i,j,t}$. Let

$$P(r_{i,j}) = \sum_{k=0}^{O_{i,j}} \epsilon_k r_{i,j}^k, \quad (31)$$

be the polynomial expansion of $P(r_{i,j})$, whose coefficients $\{\epsilon_k\}$ can be obtained from the polynomial roots. From Eq. (29), we obtain

$$f(r_{i,j}|y_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0) \propto \sum_{k=0}^{O_{i,j}} \epsilon_k r_{i,j}^{\alpha_0+k-1} \exp^{-r_{i,j} \left(\frac{\alpha_0}{a_{i,j}(\mathbf{\Gamma})} + v_{i,j} \right)}. \quad (32)$$

which can be expressed as the following mixture of $O_{i,j} + 1$ gamma distributions $f(r_{i,j}|y_{i,j}, t_{i,j}, b_{i,j}, \mathbf{\Gamma}, \alpha_0)$

$$= \sum_{k=0}^{O_{i,j}} w_k \mathcal{G} \left(r_{i,j}; \alpha_0 + k, \left(\frac{\alpha_0}{a_{i,j}(\mathbf{\Gamma})} + v_{i,j} \right)^{-1} \right), \quad (33)$$

with

$$w_k \propto \epsilon_k \frac{\Gamma(\alpha_0 + k)}{\left(\frac{\alpha_0}{a_{i,j}(\mathbf{\Gamma})} + v_{i,j} \right)^{\alpha_0+k}}, \quad \forall k, \quad (34)$$

where $\Gamma(\cdot)$ denotes the Gamma function and $\sum_{k=0}^{O_{i,j}} w_k = 1$.

REFERENCES

- [1] T. Ogawa and K. Takagi, "Lane recognition using on-vehicle LIDAR," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2006, pp. 540–545.
- [2] P. Lindner and G. Wanielik, "3D LIDAR processing for vehicle safety and environment recognition," in *Proc. IEEE Workshop Comput. Intell. Vehicles Veh. Syst. (CIVVS)*, Mar. 2009, pp. 66–71.
- [3] S. Matzka, A. M. Wallace, and Y. R. Petillot, "Efficient resource allocation for attentive automotive vision systems," *IEEE Trans. Intell. Transp. Syst.*, vol. 13, no. 2, pp. 859–872, Jun. 2012.
- [4] Y. Kang, C. Roh, S.-B. Suh, and B. Song, "A lidar-based decision-making method for road boundary detection using multiple Kalman filters," *IEEE Trans. Ind. Electron.*, vol. 59, no. 11, pp. 4360–4368, Nov. 2012.
- [5] D. Ramirez *et al.*, "Developing hyperspectral LiDAR for structural and biochemical analysis of forest data," in *Proc. EarSel Symp.*, Mykonos, Greece, May 2012, p. 333.
- [6] T. Hakala, J. Suomalainen, S. Kaasalainen, and Y. Chen, "Full waveform hyperspectral LiDAR for terrestrial laser scanning," *Opt. Exp.*, vol. 20, no. 7, pp. 7119–7127, Mar. 2012.
- [7] N. Cadalli, P. J. Shargo, D. C. Munson, Jr., and A. C. Singer, "Three-dimensional tomographic imaging of ocean mines from real and simulated lidar returns," *Proc. SPIE*, vol. 4488, no. 155, pp. 155–166, Jan. 2002.
- [8] J. Gao, J. Sun, J. Wei, and Q. Wang, "Research of underwater target detection using a slit streak tube imaging lidar," in *Proc. Acad. Int. Symp. Optoelectron. Microelectron. Technol. (AISOMT)*, Oct. 2011, pp. 240–243.
- [9] A. McCarthy *et al.*, "Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector," *Opt. Exp.*, vol. 21, no. 19, pp. 22098–22113, Sep. 2013.
- [10] A. Kirmani *et al.*, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [11] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, "Bayesian analysis of lidar signals with multiple returns," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 12, pp. 2170–2180, Dec. 2007.
- [12] A. M. Wallace *et al.*, "Design and evaluation of multispectral LiDAR for the recovery of arboreal parameters," *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 8, pp. 4942–4954, Aug. 2014.
- [13] 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 Statist. Signal Process.*, Gold Coast, VIC, Australia, Jul. 2014, pp. 121–124.
- [14] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D, Nonlinear Phenomena*, vol. 60, nos. 1–4, pp. 259–268, Nov. 1992.
- [15] A. Chambolle, "An algorithm for total variation minimization and applications," *J. Math. Imag. Vis.*, vol. 20, nos. 1–2, pp. 89–97, 2004.
- [16] O. Dikmen and A. T. Cemgil, "Gamma Markov random fields for audio source modeling," *IEEE Trans. Audio, Speech, Language Process.*, vol. 18, no. 3, pp. 589–601, Mar. 2010.
- [17] M. Zhou, L. A. Hannah, D. B. Dunson, and L. Carin, "Beta-negative binomial process and Poisson factor analysis," in *Proc. 15th AISTATS*, 2012, pp. 1–10.
- [18] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "Estimating the granularity coefficient of a Potts-Markov random field within a Markov chain Monte Carlo algorithm," *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2385–2397, Jun. 2013.
- [19] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York, NY, USA: Springer-Verlag, 2004.
- [20] 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.
- [21] N. J. Krichel, A. McCarthy, and G. S. Buller, "Resolving range ambiguity in a photon counting depth imager operating at kilometer distances," *Opt. Exp.*, vol. 18, no. 9, pp. 9192–9206, Apr. 2010.

- [22] A. M. Wallace, J. Ye, N. Krichel, A. McCarthy, R. J. Collins, and G. S. Buller, "Full wave form analysis for long-range 3D imaging laser radar," *EURASIP J. Adv. Signal Process.*, no. 1, p. 896708, 2010.
- [23] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [24] P. W. Milonni, J. H. Carter, C. G. Peterson, and R. J. Hughes, "Effects of propagation through atmospheric turbulence on photon statistics," *J. Opt. B, Quantum Semiclass. Opt.*, vol. 6, no. 8, p. S742, 2004.
- [25] I. Capraro *et al.*, "Impact of turbulence in long range quantum and classical communications," *Phys. Rev. Lett.*, vol. 109, p. 200502, Nov. 2012.
- [26] M. Henriksson and L. J. Sjöqvist, "Scintillation index measurement using time-correlated single-photon counting laser radar," *Opt. Eng.*, vol. 53, no. 8, p. 081902, 2014.
- [27] M. Malik, O. S. Magaña-Loaiza, and R. W. Boyd, "Quantum-secured imaging," *Appl. Phys. Lett.*, vol. 101, p. 241103, Dec. 2012.



Yoann Altmann (S'10–M'14) was born in Toulouse, France, in 1987. He received the Engineering degree in electrical engineering from the Ecole Nationale Supérieure d'Electrotechnique, d'Electronique, d'Informatique, d'Hydraulique et des Télécommunications, Toulouse, and the M.Sc. degree in signal processing from the National Polytechnic Institute of Toulouse (INP Toulouse), Toulouse, in 2010, and the Ph.D. degree from INP Toulouse, in 2013. Since 2014, he has been with Heriot-Watt University, Edinburgh, as a

Post-Doctoral Researcher. He conducts his research with the Institute of Sensors, Signals and Systems, School of Engineering and Physical Sciences. His current research activities focus on statistical signal and image processing, with a particular interest in Bayesian inverse problems with applications to remote sensing and biomedical imaging.



Ximing Ren received the B.Eng. degree in optical engineering from the Beijing Institute of Technology, China, in 2008, the M.Eng. degree in electrical engineering from Beihang University, China, in 2011, and the Ph.D. degree in physics with a prize scholarship from the Scottish Universities Physics Alliance, Single-Photon Group, Heriot-Watt University, studying single-photon time-of-flight imaging under the supervision of Prof. G. S. Buller in 2015. He is currently a Post-Doctoral Research Associate with the Single-Photon Group, Heriot-Watt

University. He is a member of the Optical Society of America (OSA) and the Institute of Physics. He is providing voluntary peer-review for the IEEE PHOTONICS TECHNOLOGY LETTERS, *Chinese Optics Letters of OSA*, and *Optical Engineering of SPIE*.



Aongus McCarthy (M'03) received the B.Sc. degree from University College Galway, Ireland, in 1989, the Diploma degree in electronics engineering from the Institute of Technology, Carlow, Ireland, in 1990, the B.Sc. degree in physical optoelectronics from Essex University, Essex, U.K., in 1991, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 2002. He was with the industry before his Ph.D. degree. He is currently a Research Fellow with the School of Engineering and Physical



Sciences, Heriot-Watt University. His research interests include optical and optomechanical system design, time-of-flight depth imaging, single-photon counting technologies, and microscope systems. He is a member of the Optical Society of America and the IEEE Photonics Society.

Gerald S. Buller received the B.Sc. (Hons.) degree in natural philosophy from the University of Glasgow, in 1986, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 1989. In 2002, he co-founded Helia Photonics Ltd., Livingston, U.K. He is currently a Professor of Physics with Heriot-Watt University. His current research interests include single-photon detection methods, particularly at infrared wavelengths, and applications involving the use of single-photons. These applications include single-photon

depth imaging and quantum communications protocols. He is a fellow of the Royal Society of Edinburgh and the U.K. Institute of Physics, and a member of the Optical Society of America. In 2015, he received an EPSRC Established Career Fellowship in Quantum Technology to research sparse photon imaging.



Steve McLaughlin (F'11) was born in Clydebank, U.K., in 1960. He received the B.Sc. degree in electronics and electrical engineering from the University of Glasgow, in 1981, and the Ph.D. degree from the University of Edinburgh, in 1990. From 1981 to 1984, he was a Development Engineer with the industry involved in the design and simulation of integrated thermal imaging and fire control systems. From 1984 to 1986, he worked on the design and development of high frequency data communication systems. In 1986, he joined the Department of Electronics and Electrical Engineering, University of Edinburgh, as a Research Fellow, where he studied the performance of linear adaptive algorithms in high noise and nonstationary environments. In 1988, he joined the Academic Staff Member at Edinburgh, and from 1991 to 2001, he held a Royal Society University Research Fellowship to study nonlinear signal processing techniques. In 2002, he was a Personal Chair in Electronic Communication Systems with the University of Edinburgh. In 2011, he joined Heriot-Watt University as a Professor of Signal Processing and the Head of the School of Engineering and Physical Sciences. His research interests lie in the fields of adaptive signal processing and nonlinear dynamical systems theory and their applications to biomedical, energy, and communication systems. He is a fellow of the Royal Academy of Engineering, the Royal Society of Edinburgh, and the Institute of Engineering and Technology.

depth imaging and quantum communications protocols. He is a fellow of the Royal Society of Edinburgh and the U.K. Institute of Physics, and a member of the Optical Society of America. In 2015, he received an EPSRC Established Career Fellowship in Quantum Technology to research sparse photon imaging.