Robust Bayesian Target Detection Algorithm for Depth Imaging From 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 associated algorithm for depth and intensity profiling using full waveforms from time-correlated single-photon counting measurements in the limit of very low photon counts (i.e., typically less than 20 photons per pixel). The model represents each Lidar waveform as an unknown constant background level, which is combined in the presence of a target, to a known impulse response weighted by the target intensity and finally corrupted by Poisson noise. The joint target detection and depth imaging problem is expressed as a pixelwise model selection and estimation problem, which is solved using Bayesian inference. Prior knowledge about the problem is embedded in a hierarchical model that describes the dependence structure between the model parameters while accounting for their constraints. In particular, Markov random fields (MRFs) are used to model the joint distribution of the background levels and of the target presence labels, which are both expected to exhibit significant spatial correlations. An adaptive Markov chain Monte Carlo algorithm including reversible-jump updates is then proposed to compute the Bayesian estimates of interest. 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—Bayesian estimation, full waveform Lidar, poisson statistics, remote sensing, reversible jump Markov chain monte carlo, target detection.

I. INTRODUCTION

TIME-OF-FLIGHT laser detection and ranging (Lidar) based imaging systems are used to reconstruct 3-dimensional scenes in many applications, including automotive [1]–[4], environmental sciences [5], [6], architectural engineering and defence [7], [8] applications. This challenging problem consists of illuminating the scene with a train of laser pulses and analysing the distribution of the photons received by the detector to infer the presence of objects as well as their range, and radiative properties (e.g., reflectivity, observation conditions,...).

Using scanning systems, a 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. Conventionally, in the presence of objects, the recorded photon histograms are decomposed into a series of peaks whose positions can be used to infer the distance of the objects present in each region of the scene and whose amplitudes provide information about the intensity of the objects.

In this paper, we investigate the target detection problem which consists of inferring the regions or pixels of the scene where objects are present. Moreover, we propose an algorithm for applications where the flux of detected photons is small and for which classical depth imaging methods [9] usually provide unsatisfactory results in terms of range and intensity estimation. This is typically the case for free-space depth profiling on targets at very long distances based on the time-correlated single-photon counting (TCSPC) technique [9], which negotiates the trade-offs between range/intensity estimation quality, data acquisition time and output laser power. In addition, this might be extended for sparse single-photon depth imaging in turbid media, e.g., underwater depth imaging [10]. In contrast with the method proposed in [11], we consider scene observation using a scanning system whose acquisition time per pixel is fixed, thus leading to a deterministic and user-defined overall acquisition duration. As in [12], the number of detected photons can thus vary across the image pixels and some pixels can be empty (i.e., no detected photons).

In this work, we assume that the targets potentially present in the scene of interest are opaque, i.e., are composed of a single surface per pixel. As in [11], [12], we consider the potential presence of two kinds of detector events: the photons originating from the illumination laser and scattered back from the target (if present); and the 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 (if any) and the background in the photon timing histograms.

Following a classical Bayesian approach, as in [12]–[14], we express the target detection and identification problem as a pixel-wise model selection and estimation problem. More precisely, two observation models, conditioned on the presence or absence of a target (modelled by binary labels) are considered for each pixel. We then assign prior distributions to each of the unknown parameters in each model to include available information within the estimation procedure. The probabilities of target presence (or equivalently the binary labels associated with the presence/absence of target) are also assigned prior

distributions accounting for spatial organization of the objects in the scene.

Classical Bayesian estimators associated with the joint posterior cannot be easily computed due to the complexity of the model, in particular because the number of underlying parameters (number of targets) is unknown and potentially large. To tackle this problem, a Reversible-Jump Markov chain Monte Carlo (RJ-MCMC) [15], [16] method is used to generate samples according to this posterior by allowing moves between different parameter spaces. More precisely, we construct an efficient stochastic gradient MCMC (SGMCMC) algorithm [17] that simultaneously estimates the background levels and the target distances and intensity, along with the MRFs parameters.

The main contributions of this work are threefold:

1) We develop a new Bayesian algorithm for joint target detection and identification, which takes spatial correlations affecting the background levels and the target locations into account through Markovian dependencies. To the best of our knowledge, the proposed method is the first joint target detection and identification method designed for depth imaging using single-photon data.

2) An adaptive Markov chain Monte Carlo algorithm including reversible-jump updates is proposed to compute the Bayesian estimates of interest and perform Bayesian inference. This algorithm included RJ-MCMC updates and is equipped with a stochastic optimization mechanism that adjusts automatically the parameters of the Markov random fields by maximum marginal likelihood estimation, thus removing the need to set the regularization parameters, e.g., by cross-validation.

3) We show the benefits of the proposed flexible model for reconstructing a real 3D object in scenarios where the number of detected photons is very low and the background levels are significant.

The remainder of this paper in organized as follows. Section II recalls the statistical models used for depth imaging using time-of-flight scanning sensors, based on TCSPC. Section III presents the new hierarchical Bayesian model which takes into account the inherent spatial correlations between parameters of spatially close pixels. Section IV discusses the estimation of the model parameters including the detection labels using adaptive MCMC methods including reversible jump updates. 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

Assessing the presence of targets from TCSPC measurements in an unsupervised manner is a challenging problem as the detection performance highly depends on the nature of the potential targets (range and reflectivity), as well as the observation conditions (system performance, ambient illumination). In practice, the target range can be restrained to a bounded interval and the estimated reflectivity parameter can be used to assess the target presence (i.e., via thresholding). In a similar manner to [11], [18], it is possible to enhance the estimation performance of

depth and reflectivity parameters by processing simultaneously several pixels/spatial locations, thus improving the subsequent target detection. However, the nature of such sequential processes implies sub-optimal detection performance (which depends on the quality of the previous estimation steps). For this reason, we propose a Bayesian model and method allowing the joint target detection and depth/reflectivity estimation. In contrast to [11], [18], in this work we assume that the ambient noise level, which can vary among pixels, is unknown and thus needs to be estimated. This is typically the case for long range measurements where the background levels can change due to time-varying illumination conditions. In this work, the background levels are estimated from signal measured during the detection process, and not beforehand, as may be done with a detector array [19].

We consider a set of $N_{\text{row}} \times N_{\text{col}}$ observed Lidar waveforms/pixels $\mathbf{y}_{i,j} = [y_{i,j,1}, \ldots, y_{i,j,T}]^T$, $(i,j) \in \{1, \ldots, N_{\text{row}}\} \times \{1, \ldots, N_{\text{col}}\}$ where T is the number of temporal (corresponding to range) bins. To be precise, $y_{i,j,t}$ is the photon count within the tth bin of the pixel or location (i,j). Let $z_{i,j} \in \{0,1\}$ be a binary variable associated with the presence ($z_{i,j} = 1$) or absence ($z_{i,j} = 0$) of target in the pixel (i,j). In the absence of target, $y_{i,j,t}$ is assumed to be drawn from the following Poisson distribution

$$y_{i,j,t} \,|\, (z_{i,j} = 0, \boldsymbol{\theta}_{i,j}^0) \sim \mathcal{P}(b_{i,j}), \qquad (1)$$

where $b_{i,j} > 0$ stands for the background and dark photon level, which is assumed to be constant in all bins of a given pixel. The model (1) for each pixel is denoted $\mathcal{M}_0^{(i,j)}$ (0 denotes the absence of a target). Moreover $\boldsymbol{\theta}_{i,j}^0$ denotes the set of likelihood parameters in pixel (i,j) under $\mathcal{M}_0^{(i,j)}$, i.e., $\boldsymbol{\theta}_{i,j}^0 = b_{i,j} \in \mathbb{R}^+$.

Assume now the presence of a target in the pixel (i,j). Let $t_{i,j}$ be the position of that object surface at a given range from the sensor and $r_{i,j}$ its intensity. The target position is considered as a discrete variable defined on $\mathbb{T} = \{t_{\min}, \ldots, t_{\max}\}$, such that $1 \leq t_{\min} \leq t_{\max} \leq T$ (in this paper we set $(t_{\min}, t_{\max}) = (1, T)$). In that case, we obtain

$$y_{i,j,t} \,|\, (z_{i,j} = 1, \boldsymbol{\theta}_{i,j}^1) \sim \mathcal{P}(r_{i,j} g_0(t - t_{i,j}) + b_{i,j}), \qquad (2)$$

where $g_0(\cdot) > 0$ is the photon impulse response, which is assumed to be known (this response can be estimated during the imaging system calibration). In a similar fashion to (1), the model (2) for each pixel is denoted $\mathcal{M}_1^{(i,j)}$ and in a similar fashion to the background level, the target intensity in each pixel is non-negative, i.e., $r_{i,j} \geq 0$. In (2), $\boldsymbol{\theta}_{i,j}^1$ stands for the set of likelihood parameters in pixel (i,j) under $\mathcal{M}_1^{(i,j)}$, i.e., $\boldsymbol{\theta}_{i,j}^1 = [r_{i,j}, t_{i,j}, b_{i,j}] \in \mathbb{R}^+ \times \mathbb{T} \times \mathbb{R}^+$. Note that the background levels in (1) and (2) have the same physical meaning and are assumed to present the same statistical properties under $\mathcal{M}_0^{(i,j)}$ and $\mathcal{M}_1^{(i,j)}$. Thus, a single background level $b_{i,j}$, independent from the observation model, is used for each pixel.

The problem addressed in this paper consists of deciding whether a target is present ($\mathcal{M}_0^{(i,j)}$) or not ($\mathcal{M}_1^{(i,j)}$) in each pixel and of estimating the position and intensity of the targets present in the scene, from the observed data gathered in the

$N_{\text{row}} \times N_{\text{col}} \times T$ array $\mathbf{Y}$. Moreover, the background levels $b_{i,j}$ are also assumed to be unknown and need to be estimated.

The target detection problem considered in this paper can be seen as a pixel-wise model selection problem where the parameter space associated with each model is different. To be precise, under $\mathcal{M}_0^{(i,j)}$ (resp. $\mathcal{M}_1^{(i,j)}$), $\boldsymbol{\theta}_{i,j}^{(0)} \in \boldsymbol{\Theta}_0$ (resp. $\boldsymbol{\theta}_{i,j}^{(1)} \in \boldsymbol{\Theta}_1$) where $\boldsymbol{\Theta}_0 = \mathbb{R}^+$ and $\boldsymbol{\Theta}_1 = \mathbb{R}^+ \times \mathbb{T} \times \boldsymbol{\Theta}_0$. To simplify notations, the unknown parameter vector is noted $\boldsymbol{\theta}_{i,j}$ in the remainder of the paper when we do not specify whether it is included in $\boldsymbol{\Theta}_0$ or $\boldsymbol{\Theta}_1$. Estimating $\boldsymbol{\theta}_{i,j}$ is difficult using standard optimization methods since the dimensionality of the parameter vector depends on the underlying model. However, this model selection problem can be solved efficiently in a Bayesian framework by 1) performing inference for each pixel in the parameter space $\{\{0\} \times \boldsymbol{\Theta}_0\} \bigcup \{\{1\} \times \boldsymbol{\Theta}_1\}$, 2) incorporating relevant additional prior belief (through prior distributions) and 3) using RJ-MCMC methods adapted for problems whose finite dimensionality in unknown.

The next section presents a new Bayesian model for target detection accounting for spatial correlations affecting parameters of neighbouring pixels.

III. BAYESIAN MODEL FOR COLLABORATIVE TARGET DETECTION VIA MARKOVIAN DEPENDENCIES

A. Parameter Prior Distributions

1) Priors for the Background Levels: In the absence of a target (i.e., assuming (1)), gamma distributions are conjugate priors for $b_{i,j}$. Moreover, it has been shown in [12] that considering such priors in the presence of a target also simplifies the sampling procedure. In contrast to the model in [12] which assumed the background levels of the $N_{\text{row}} \times N_{\text{col}}$ pixels to be *a priori* independent, here we specify the background levels prior distribution to reflect the prior belief that background levels exhibit spatial correlations. In particular, due to the spatial organization of images, we expect the values of $b_{i,j}$ to vary smoothly from one pixel to another (as will be illustrated in Section V). In order to model this behaviour, we specify $\epsilon_{i,j}$ such that the resulting prior for the background matrix $\mathbf{B}$ such that $[\mathbf{B}]_{i,j} = b_{i,j}$ is a hidden gamma-MRF (GMRF) [20] (in a similar fashion to the intensity model in [12]). More precisely, we introduce an $(N_{\text{row}} + 1) \times (N_{\text{col}} + 1)$ auxiliary matrix $\boldsymbol{\Gamma}$ with elements $\gamma_{i,j} \in \mathbb{R}^+$ and define a bipartite conditional independence graph between $\mathbf{B}$ and $\boldsymbol{\Gamma}$ such that each $b_{i,j}$ is connected to four neighbour elements of $\boldsymbol{\Gamma}$ and vice-versa. This 1st order neighbourhood structure is similar to that depicted in [12, Fig. 2], where we notice that any given $b_{i,j}$ and $b_{i+1,j}$ are 2nd order neighbours via $\gamma_{i+1,j}$ and $\gamma_{i+1,j+1}$. We specify a GMRF prior for $\mathbf{B}, \boldsymbol{\Gamma}$ [20], and obtain the following joint prior for $\mathbf{B}, \boldsymbol{\Gamma}$

$$f(\mathbf{B}, \boldsymbol{\Gamma}|\nu) = \frac{1}{G(\nu)} \prod_{(i,j)\in\mathcal{V}_{\mathbf{B}}} b_{i,j}^{(\nu-1)} \prod_{(i',j')\in\mathcal{V}_{\boldsymbol{\Gamma}}} (\gamma_{i',j'})^{-(\nu+1)}$$
$$\times \prod_{((i,j),(i',j'))\in\mathcal{E}} \exp\left(\frac{-\nu b_{i,j}}{4\gamma_{i',j'}}\right), \quad (3)$$

where $\mathcal{V}_{\mathbf{B}} = \{1,\ldots,N_{\text{row}}\} \times \{1,\ldots,N_{\text{col}}\}$, $\mathcal{V}_{\boldsymbol{\Gamma}} = \{1,\ldots,N_{\text{row}}+1\} \times \{1,\ldots,N_{\text{col}}+1\}$, and the edge set $\mathcal{E}$ consists of pairs $((i,j),(i',j'))$ representing the connection between $b_{i,j}$ and $\gamma_{i',j'}$. In this paper, the notation $x|y$ reads "x conditioned on the value of y" and $f(x|y)$ denotes the probability distribution function of $x|y$, i.e., its probability density function (if $x|y$ is a continuous variable), its probability mass function (if $x|y$ is discrete) or its mixed density if $x|y$ contains discrete and continuous random variables. It can be seen from (3) that

$$b_{i,j}|(\boldsymbol{\Gamma},\nu) \sim \mathcal{G}\left(\nu, \frac{\epsilon_{i,j}(\boldsymbol{\Gamma})}{\nu}\right) \quad (4a)$$

$$\gamma_{i,j}|(\mathbf{B},\nu) \sim \mathcal{IG}(\nu, \nu\xi_{i,j}(\mathbf{B})) \quad (4b)$$

where $\mathcal{G}(x,y)$ (resp. $\mathcal{IG}(x,y)$) denotes the gamma (resp. inverse gamma) distribution with shape parameter x and scale parameter y, and with

$$\epsilon_{i,j}(\boldsymbol{\Gamma}) = 4\left(\gamma_{i,j}^{-1} + \gamma_{i-1,j}^{-1} + \gamma_{i,j-1}^{-1} + \gamma_{i-1,j-1}^{-1}\right)^{-1}$$
$$\xi_{i,j}(\mathbf{B}) = (b_{i,j} + b_{i+1,j} + b_{i,j+1} + b_{i+1,j+1})/4.$$

Notice that we denote explicitly the dependence on the value of the regularization parameter ν, which here controls the amount of spatial smoothness enforced by the GMRF. Following an empirical Bayesian approach, the value of ν remains unspecified and will be adjusted automatically during the inference procedure by maximum marginal likelihood estimation.

2) Priors for the Target Parameters: To reflect the absence of prior knowledge about the target ranges given $z_{i,j} = 1$, we assign each possible target depth the following uniform prior $f(t_{i,j} = t) = \frac{1}{T'}, \quad t \in \mathbb{T}$, where $T' = \text{card}(\mathbb{T})$. Note however that this prior can be adapted according to potential prior knowledge about the expected target depth distribution.

Accounting for potential spatial dependencies for the target intensities is more challenging as all pixels do not necessarily contain targets. Thus, considering fixed neighbourhood structures (as for $\mathbf{B}$) is not well adapted here. Consequently, we propose the following classical hierarchical model

$$r_{i,j}|(\alpha,\beta) \sim \mathcal{G}(\alpha,\beta), \quad \forall(i,j) \quad (5a)$$

$$\alpha|(\alpha_1,\alpha_2) \sim \mathcal{G}(\alpha_1,\alpha_2) \quad (5b)$$

$$\beta|(\beta_1,\beta_2) \sim \mathcal{IG}(\beta_1,\beta_2) \quad (5c)$$

where (α_1,α_2) and (β_1,β_2) are fixed parameters set to $(\alpha_1,\alpha_2) = (1.1,1)$ and $(\beta_1,\beta_2) = (1,1)$ to reflect the fact that the target intensities have a high probability to be in $(0,1)$. Indeed, the photon impulse response $g_0(\cdot) > 0$ estimated during the imaging system calibration can be scaled appropriately using reference targets and acquisition times. Although the model (5) does not capture the spatial dependencies between the target intensities, it translates the prior belief that the potential target intensities share similar statistical properties (through α and β).

3) Priors for the Observation Models: Finally, in a similar fashion to the background levels, it is often reasonable to expect the probability of a target to be present in a pixel to be related to the presence of targets in the neighbouring pixels (at least when considering targets larger than the spacing between pixels

as considered in Section V). To encode this prior belief, we attach the $N_{\text{row}} \times N_{\text{col}}$ detection label matrix $\mathbf{Z}$ ($[\mathbf{Z}]_{i,j} = z_{i,j}$) the following Ising model

$$f(\mathbf{Z}|c) = \frac{1}{G(c)} \exp\left[c\phi(\mathbf{Z})\right] \qquad (6)$$

where $\phi(\mathbf{Z}) = \sum_{i,j} \sum_{(i',j') \in \mathcal{V}_{i,j}} \delta(z_{i,j} - z_{i',j'})$, $\delta(\cdot)$ denotes the Kronecker delta function, and $\mathcal{V}_{i,j}$ is the set of neighbours of pixel (i,j) (in this paper we consider an 8-neighbour structure). Moreover, c is an hyperparameter that controls the spatial granularity of the Ising model and $G(c) = \sum_{\mathbf{Z} \in (0;1)^{N_{\text{row}} \times N_{\text{col}}}} \exp[c\phi(\mathbf{Z})]$. In a similar fashion to ν, c remains unspecified and will be adjusted automatically during the inference procedure by maximum marginal likelihood estimation using [17]. Due to use of the Ising model (6), we have easy access only to the probability of a target presence in the pixel (i,j) conditioned on the labels values in the neighbouring pixels, i.e., $p_{i,j} = f(z_{i,j} = 1|\mathbf{Z}_{\backslash(i,j)})$ where $\mathbf{Z}_{\backslash(i,j)}$ denotes the subset of $\mathbf{Z}$ whose element $z_{i,j}$ has been removed. Consequently, it is important to mention here that the parameters associated with a pixel (i,j) and those of the pixels in $\mathcal{V}_{i,j}$ will not be updated simultaneously. However, chessboard sampling schemes can be used to update conditionally independent sets of parameters in parallel for a more efficient implementation.

In this Section III-A, we defined priors distributions for the unknown model parameters. More precisely, priors promoting spatial correlations were used for the background levels and for the detection labels. A joint hierarchical model was proposed for the reflectivity parameters and uniform priors were used for the unknown target depth. In other words, we do not use potential spatial correlation affecting the depth or the reflectivity profiles. Although such consideration could improve the range/reflectivity estimation and possibly the detection performance (which is the main purpose of the proposed method), the model and method we propose can be applied to analyse complex target, possibly highly spatially non-smooth in terms shape and/or reflectivity profile. The next Section derives the joint posterior probability associated with the detection problem considered.

B. Joint Posterior Distribution

We can now specify the joint posterior distribution for $\mathbf{Z}, \mathbf{\Theta} = \{\boldsymbol{\theta}_{i,j}\}_{i,j}$ and $\mathbf{\Phi} = \{\mathbf{\Gamma}, \alpha, \beta\}$ given the observed waveforms $\mathbf{Y}$ and the value of the spatial regularization parameters ν and c (recall that their value will be determined by maximum marginal likelihood estimation during the inference procedure). Using Bayes theorem, and the prior independence assumptions mentioned above, the joint posterior distribution associated with the proposed Bayesian model is given by

$$f(\mathbf{Z}, \mathbf{\Theta}, \mathbf{\Phi}|\mathbf{Y}, \nu, c) \propto \left[\prod_{i,j} f(\mathbf{y}_{i,j}|z_{i,j}, \boldsymbol{\theta}_{i,j})f(\boldsymbol{\theta}_{i,j}|\mathbf{Z}, \mathbf{\Phi})\right]$$
$$\times f(\mathbf{Z}|c)f(\mathbf{\Gamma}|\nu)f(\alpha)f(\beta), \qquad (7)$$

using $f(\mathbf{\Phi}|\nu) = f(\mathbf{\Gamma}|\nu)f(\alpha)f(\beta)$. Note that for clarity the dependence of all distributions on the known fixed quantities

$(\alpha_1, \alpha_2, \beta_1, \beta_2)$ is omitted in (7) and in the remainder of the paper.

IV. BAYESIAN INFERENCE

A. Bayesian Estimators

The Bayesian model defined in Section III specifies the joint posterior density for the unknown parameters $\mathbf{Z}, \mathbf{\Theta}, \mathbf{\Gamma}, \alpha$ and β given the observed data $\mathbf{Y}$ and the parameters ν and c. 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 target detection. Here we consider the following coupled Bayesian estimators that are particularly suitable for model selection problems: the marginal maximum a posteriori (MMAP) estimator for the target presence labels

$$z_{i,j}^{MMAP} = \underset{z_{i,j} \in \{0,1\}}{\text{argmax}} f(z_{i,j}|\mathbf{Y}, \hat{\nu}, \hat{c}), \qquad (8)$$

and, conditionally on the estimated labels, 1) the minimum mean square error estimator of the background levels

$$b_{i,j}^{MMSE} = \mathrm{E}\left[b_{i,j}|z_{i,j} = \hat{z}_{i,j}^{MMAP}, \mathbf{Y}, \hat{\nu}, \hat{c}\right], \qquad (9)$$

and 2) for the pixels for which $\hat{z}_{i,j}^{MMAP} = 1$, the minimum mean square error estimator of the target intensities

$$r_{i,j}^{MMSE} = \mathrm{E}\left[r_{i,j}|z_{i,j} = 1, \mathbf{Y}, \hat{\nu}, \hat{c}\right], \qquad (10)$$

and the marginal maximum a posteriori (MMAP) estimator of the target positions

$$t_{i,j}^{MMAP} = \underset{t_{i,j} \in \mathbb{T}}{\text{argmax}} f(t_{i,j}|z_{i,j} = 1, \mathbf{Y}, \hat{\nu}, \hat{c}), \qquad (11)$$

Note that marginalising out the other unknowns (including α and β) in (8), (9), (10) and (11) automatically takes into account their uncertainty.

Computing (8) to (11) is challenging because it requires having access to the univariate marginal densities of $z_{i,j}$ and the joint marginal densities of (among others) $(b_{i,j}, z_{i,j})$, which in turn require computing the posterior (7) and integrating it over a very high-dimensional space. Fortunately, these estimators can be efficiently approximated with arbitrarily large accuracy by Monte Carlo integration. This is why we propose to compute (8) to (11) by first using an MCMC method to generate samples asymptotically distributed according to (7), and subsequently using these samples to approximate the required marginal probabilities and expectations. Here we propose an RJ-MCMC method to simulate samples from (7), as this type of MCMC method is particularly suitable for models involving hidden Markov random fields and parameter spaces of varying of dimensions [21, Chap. 10,11]. It is important to mention that a standard Gibbs sampler similar to that studied in [22] could have been considered instead of the proposed RJ-MCMC to sample from (7). However, such approach would lead to a prohibitively slow exploration of the posterior, mainly due to the weakly informative priors for the target ranges. Precisely, updating sequentially the presence labels and the models parameters

generally leads to low probabilities to move from $\mathcal{M}_0^{(i,j)}$ to $\mathcal{M}_1^{(i,j)}$ at each iteration of the sampler. The output of this algorithm are Markov chains of N_{MC} samples distributed according to the posterior distribution (7). The first N_{bi} samples of these chains correspond to the so-called *burn-in* transient period and should be discarded (the length of this period can be assessed visually from the chain plots or by computing convergence tests). The remaining $N_{\mathrm{MC}} - N_{\mathrm{bi}}$ of each chain are used to approximate the Bayesian estimators (8) to (11) as in [12], [22].

Notice that in (8) to (11), we have set $c = \hat{c}$ and $\nu = \hat{\nu}$, which denotes the maximum marginal likelihood estimator of the MRF regularisation parameters c and ν given the observed data $\mathbf{Y}$, i.e.,

$$(\hat{c}, \hat{\nu}) = \underset{c \in \mathbb{R}^+, \nu \in \mathbb{R}^+}{\operatorname{argmax}} \; f(\mathbf{Y}|c, \nu), \qquad (12)$$

This approach for specifying (c, ν) 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 [17], this strategy has several important advantages for MRF parameters with intractable conditional distributions such as c. In particular, it allows for the automatic adjustment of the value of (c, ν) for each data set (thus producing significantly better estimation results than using a single fixed value of (c, ν) for all data sets), and has a computational cost that is several times lower than that of competing approaches, such as including (c, ν) in the model (by assigning them prior distributions) and subsequently marginalising them during the inference procedure [23].

It is important to mention that the target detection problem considered here and formulated as a model selection problem can be addressed in an optimization framework, e.g., using expectation-maximization methods [24, Chap. 9]. However, in order to assess the relevance of the proposed Bayesian model while minimizing potential algorithmic convergence issues, this paper focuses on a single simulation-based algorithm. Alternative estimation strategies based on the proposed model would deserve a detailed analysis, which is out of scope of this paper.

B. Sampling Strategy

The remainder of this section provides details about the main steps of the proposed sampling strategy, summarised in Algorithm 1 below.

Each sampling iteration can be decomposed into 5 main steps (lines 4, 5, 6, (7–12) and 13 in Algorithm 1). Due to the conjugacy between (5a) and (5c), it can be easily shown that $f(\beta|\mathbf{Y}, \mathbf{Z}, \mathbf{\Theta}, \alpha) = f(\beta|\mathbf{Z}, \mathbf{\Theta}, \alpha)$ with

$$\beta \,|\, (\mathbf{Z}, \mathbf{\Theta}, \alpha) \sim \mathcal{IG}\left(\beta_1 + \alpha \|\mathbf{Z}\|_0, \beta_2 + \sum_{z_{i,j}=1} r_{i,j}\right) \quad (13)$$

where $\|\mathbf{Z}\|_0$ denotes the number of non-zero elements in $\mathbf{Z}$. In a similar fashion to [25], when $\|\mathbf{Z}\|_0 = 0$, i.e., the image does not contain any target, sampling α from its conditional distribution

---

Algorithm 1: Collaborative Target Detection.

1: *Fixed input parameters:* Lidar impulse response $g_0(\cdot)$, $(\alpha_1, \alpha_2, \beta_2, \beta_2)$, number of burn-in iterations N_{bi}, total number of iterations N_{MC}
2: *Initialization* $(t = 0)$
  Set $\mathbf{Z}^{(0)}, \mathbf{\Theta}^{(0)}, \mathbf{\Gamma}^{(0)}, \alpha^{(0)}, \beta^{(0)}, \nu^{(0)}, c^{(0)}$
3: *Iterations* $(1 \leq t \leq N_{MC})$
4: Sample $\beta^{(t)} \sim f(\beta|\mathbf{Z}^{(t)}, \mathbf{\Theta}^{(t)}, \alpha^{(t)})$ in (13)
5: Sample $\alpha^{(t)} \sim f(\alpha|\mathbf{Z}^{(t)}, \mathbf{\Theta}^{(t)}, \beta^{(t-1)})$ in (5b) or (14)
6: Sample $\mathbf{\Gamma}^{(t)} \sim f(\mathbf{\Gamma}|\mathbf{B}^{(t)}, \nu)$ in (4b)
7: **for** $i = 1 : N_{\mathrm{row}}$ **do**
8:   **for** $j = 1 : N_{\mathrm{col}}$ **do**
9:     Set $\boldsymbol{\psi}_{i,j} = [p_{i,j}, \nu^{(t-1)}, \epsilon_{i,j}(\mathbf{\Gamma}^{(t-1)}), \alpha^{(t-1)}, \beta^{(t-1)}]$, where $p_{i,j}$ is computed from (6)
10:     Update $(z_{i,j}^{(t)}, \boldsymbol{\theta}_{i,j}^{(t)})$ using Algorithm 2
11:   **end for**
12: **end for**
13: Update $(\nu^{(t)}, c^{(t)})$ using [17].
14: Set $t = t + 1$.

---

reduces to sampling from (5b). If $\|\mathbf{Z}\|_0 > 0$,

$$f(\alpha|\mathbf{Z}, \mathbf{\Theta}, \beta) \propto \alpha^{\alpha_1 - 1} \exp\left(-\frac{\alpha}{\alpha_2}\right) \prod_{(i,j), z_{i,j}=1} \frac{r_{i,j}^{\alpha-1}}{\Gamma(\alpha)\beta^\alpha} \quad (14)$$

which is strictly log-concave if $\alpha_1 \geq 1$. Consequently, α can be updated using standard Metropolis-Hastings (MH) updates or adaptive rejection sampling [26]. Here we used MH updates using a Gaussian random walk whose variance is adjusted during the early iterations of the sampler. By noting that $\mathbf{\Gamma}$ does not appear in (1) nor (2), sampling from its conditional distribution reduces to sampling from $f(\mathbf{\Gamma}|\mathbf{B}, \nu)$ in (4b).

C. RJ-MCMC Updates

To update $(z_{i,j}, \boldsymbol{\theta}_{i,j})$, we construct a transition kernel which admits

$$f(z_{i,j}, \boldsymbol{\theta}_{i,j}|\mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) =$$
$$f_0(z_{i,j}, \boldsymbol{\theta}_{i,j}|\mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) + f_1(z_{i,j}, \boldsymbol{\theta}_{i,j}|\mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) \quad (15)$$

as invariant distribution, where

$$f_0(z_{i,j}, \boldsymbol{\theta}_{i,j}|\mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) = \mathbf{1}_{\{0\} \times \mathbf{\Theta}_0}(z_{i,j}, \boldsymbol{\theta}_{i,j})$$
$$\times \frac{f(\mathbf{y}_{i,j}|z_{i,j} = 0, \boldsymbol{\theta}_{i,j}) f_0(\boldsymbol{\theta}_{i,j}|\boldsymbol{\psi}_{i,j})(1 - p_{i,j})}{f(\mathbf{y}_{i,j})}, \quad (16)$$

$$f_1(z_{i,j}, \boldsymbol{\theta}_{i,j}|\mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) = \mathbf{1}_{\{1\} \times \mathbf{\Theta}_1}(z_{i,j}, \boldsymbol{\theta}_{i,j})$$
$$\times \frac{f(\mathbf{y}_{i,j}|z_{i,j} = 1, \boldsymbol{\theta}_{i,j}) f_1(\boldsymbol{\theta}_{i,j}|\boldsymbol{\psi}_{i,j}) p_{i,j}}{f(\mathbf{y}_{i,j})}, \quad (17)$$

with $\boldsymbol{\psi}_{i,j} = [p_{i,j}, \nu, \epsilon_{i,j}, \alpha, \beta]$,

$$f_0(\boldsymbol{\theta}_{i,j}|\boldsymbol{\psi}_{i,j}) = f(b_{i,j}|\nu, \epsilon_{i,j}),$$
$$f_1(\boldsymbol{\theta}_{i,j}|\boldsymbol{\psi}_{i,j}) = f(b_{i,j}|\nu, \epsilon_{i,j}) f(r_{i,j}|\alpha, \beta) f(t_{i,j} = t),$$

and where $\mathbf{1}_{\mathcal{X}}(\cdot)$ denotes the indicator function defined on $\mathcal{X}$. The proposed update scheme, summarized in Algorithm 2, can be decomposed into four possible moves detailed in the next few paragraphs. At each iteration, depending on the current state of the chain, two possible moves are proposed. One move proposes a candidate in the current parameter space while the other move proposes a candidate in the other parameter space.

1) Fixed Dimension Moves: Let $\pi_b > 0$ be the probability of proposing a move from model $\mathcal{M}_0^{(i,j)}$ to $\mathcal{M}_1^{(i,j)}$ and $\pi_d > 0$ be the probability of proposing the reverse move from model $\mathcal{M}_1^{(i,j)}$ to $\mathcal{M}_0^{(i,j)}$. In this work we set $\pi_b = \pi_d = 0.5$.

Move in $\mathcal{M}_0^{(i,j)}$: If the current state of the chain is in $\mathcal{M}_0^{(i,j)}$, with probability $(1 - \pi_b)$, the parameter space remains the same (i.e., $z_{i,j} = 0$ remains unchanged) and $\boldsymbol{\theta}_{i,j}$ is updated according to

$$b_{i,j} | \left(\mathbf{y}_{i,j}, z_{i,j} = 0, \boldsymbol{\psi}_{i,j}\right) \sim$$

$$\mathcal{G}\left(\nu + \|\mathbf{y}_{i,j}\|_1, \left(\frac{\nu}{\epsilon_{i,j}} + T\right)^{-1}\right), \quad (18)$$

using the conjugacy of (1) and (4a).

Move in $\mathcal{M}_1^{(i,j)}$: Similarly, if the current state of the chain at the tth iteration is in $\mathcal{M}_1^{(i,j)}$ (i.e., $\boldsymbol{\theta}_{i,j}^{(t)} = [r_{i,j}^{(t)}, t_{i,j}^{(t)}, b_{i,j}^{(t)}]$), with probability $(1 - \pi_d)$, $\boldsymbol{\theta}_{i,j}$ is updated using the following three Gibbs sampling steps

$$r_{i,j}^{(t+1)} \sim f(r_{i,j} | \mathbf{y}_{i,j}, t_{i,j}^{(t)}, b_{i,j}^{(t)}, z_{i,j} = 1, \boldsymbol{\psi}_{i,j}) \quad (19a)$$

$$t_{i,j}^{(t+1)} \sim f(t_{i,j} | \mathbf{y}_{i,j}, r_{i,j}^{(t+1)}, b_{i,j}^{(t)}, z_{i,j} = 1, \boldsymbol{\psi}_{i,j}) \quad (19b)$$

$$b_{i,j}^{(t+1)} \sim f(b_{i,j} | \mathbf{y}_{i,j}, r_{i,j}^{(t+1)}, t_{i,j}^{(t+1)}, z_{i,j} = 1, \boldsymbol{\psi}_{i,j}). \quad (19c)$$

It can be shown that sampling from (19a) and (19c) can be achieved by sampling from finite mixtures of gamma distributions and that sampling from (19b) reduces to sampling from a discrete distribution defined on a finite support $\mathbb{T}$ (see [12] for details).

2) Variable Dimension Moves: Moving from $\mathcal{M}_0^{(i,j)}$ to $\mathcal{M}_1^{(i,j)}$: Assume that the current state of the chain is in $\mathcal{M}_0^{(i,j)}$, i.e., $\boldsymbol{\theta}_{i,j} = \theta_{i,j}^0$. Moving from $\mathcal{M}_0^{(i,j)}$ to $\mathcal{M}_1^{(i,j)}$ requires the construction of an appropriate proposal distribution to propose a candidate $\boldsymbol{\theta}_{i,j}^{*1} \in \boldsymbol{\Theta}_1$. Since the background level has the same physical meaning under $\mathcal{M}_0^{(i,j)}$ and $\mathcal{M}_1^{(i,j)}$, it makes sense to use its current value to propose an appropriate candidate and thus increase the acceptance probability of the move (as $\mathcal{M}_0^{(i,j)}$ and $\mathcal{M}_1^{(i,j)}$ and nested models). More precisely, the candidate is constructed as $\boldsymbol{\theta}_{i,j}^{*1} = [r^*, t^*, \theta_{i,j}^0]$ where $(r^*, t^*) \in (\mathbb{R}^+ \times \mathbb{T})$ is generated according to an arbitrary proposal distribution $q(r^*, t^* | \theta_{i,j}^0, \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j})$. Assuming that

$$q(r^*, t^* | \theta_{i,j}^0, \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) > 0, \quad \forall (r^*, t^*) \in \mathbb{R}^+ \times \mathbb{T} \quad (20)$$

leads to the following acceptance ratio

$$\alpha\left[(0, \theta_{i,j}^0), (1, \boldsymbol{\theta}_{i,j}^{*1})\right] = \min\left[1, \rho\left((0, \theta_{i,j}^0), (1, \boldsymbol{\theta}_{i,j}^{*1})\right)\right] \quad (21)$$

Algorithm 2: Reversible Jump MCMC step.

1: *Input parameters:* Current model state $z_{i,j}$ and parameter vector $\boldsymbol{\theta}_{i,j}$, Hyperparameters $\boldsymbol{\psi}_{i,j}$, Observation vector $\mathbf{y}_{i,j}$

2: Sample $u \sim \mathcal{U}_{[0,1]}(u)$ and $v \sim \mathcal{U}_{[0,1]}(v)$

3: **if** $z_{i,j} = 0$ and $u > \pi_b$ **then**

4:    Set $z_{i,j}^{out} = 0$, sample $b_{i,j}$ from (18) and set $\boldsymbol{\theta}_{i,j}^{out} = b_{i,j}$

5: **else if** $z_{i,j} = 0$ and $u \le \pi_b$ **then**

6:    Sample (r^*, t^*) from (23) and set $\boldsymbol{\theta}_{i,j}^{*(1)} = [r^*, t^*, b_{i,j}]$

7:    Compute acceptance ratio $\alpha[(0, \theta_{i,j}^0), (1, \boldsymbol{\theta}_{i,j}^{*1})]$ in (21)

8:    Sample $v \sim \mathcal{U}_{[0,1]}(v)$

9:    **if** $v < \alpha[(0, \theta_{i,j}^0), (1, \boldsymbol{\theta}_{i,j}^{*1})]$ **then**

10:     Set $z_{i,j}^{out} = 1$, $\boldsymbol{\theta}_{i,j}^{out} = \boldsymbol{\theta}_{i,j}^{*1}$

11:    **else**

12: Set $z_{i,j}^{out} = 0$, $\boldsymbol{\theta}_{i,j}^{out} = b_{i,j}$

13:    **end if**

14: **else if** $z_{i,j} = 0$ and $u > \pi_b$ **then**

15:    Set $z_{i,j}^{out} = 1$

16:    Sample $\boldsymbol{\theta}_{i,j}^{out}$ from (19)

17: **else**

18:    Set $\boldsymbol{\theta}_{i,j}^{*0} = b_{i,j}$

19:    Compute acceptance ratio $\alpha[(1, \boldsymbol{\theta}_{i,j}^1), (0, \boldsymbol{\theta}_{i,j}^{*0})]$ in (25)

20:    Sample $v \sim \mathcal{U}_{[0,1]}(v)$

21:    **if** $v < \alpha[(1, \boldsymbol{\theta}_{i,j}^1), (0, \boldsymbol{\theta}_{i,j}^{*0})]$ **then**

22:     Set $z_{i,j}^{out} = 0$, $\boldsymbol{\theta}_{i,j}^{out} = \boldsymbol{\theta}_{i,j}^{*0}$

23:    **else**

24:     Set $z_{i,j}^{out} = 1$, $\boldsymbol{\theta}_{i,j}^{out} = \boldsymbol{\theta}_{i,j}$

25:    **end if**

26: **end if**

27: *Output parameters:* $(z_{i,j}^{out}, \boldsymbol{\theta}_{i,j}^{out})$

---

where

$$\rho\left((0, \theta_{i,j}^0), (1, \boldsymbol{\theta}_{i,j}^{*1})\right) =$$

$$\frac{f_1(1, \boldsymbol{\theta}_{i,j}^{*1} | \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) \pi_d}{f_0(0, \theta_{i,j}^0 | \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) q(r^*, t^* | \theta_{i,j}^0, \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) \pi_b}. \quad (22)$$

Although the choice of the proposal distribution (20) can be arbitrary, it has a significant impact on the acceptance ratio (21) and thus on the mixing properties of the sampler. In this work we use the conditional distribution

$$f_1(r_{i,j}, t_{i,j} | \mathbf{y}_{i,j}, z_{i,j} = 1, b_{i,j}, \boldsymbol{\psi}_{i,j})$$

$$\propto f_1(1, \boldsymbol{\theta}_{i,j}^1 | \mathbf{y}_{i,j}, \boldsymbol{\psi}_{i,j}) \quad (23)$$

as proposal distribution so that the candidate (conditioned on the current value of $b_{i,j}$) lies in a region of relatively high density in $\boldsymbol{\Theta}_1$. Although the choice of the proposal could be further improved [27], this choice allows frequent moves between $\mathcal{M}_0^{(i,j)}$ and $\mathcal{M}_1^{(i,j)}$ in practice. Note that due to the high posterior correlation of the variables in $\boldsymbol{\theta}_{i,j}^1$ and because $T' = \text{card}(\mathbb{T})$ is generally large, using weakly informative proposals instead of (23) would lead to prohibitively low acceptance ratios when proposing moves from $\mathcal{M}_0^{(i,j)}$ to $\mathcal{M}_1^{(i,j)}$. Based on [12], it can

be shown that $\forall k \in \mathbb{T}$,

$$f_1(r_{i,j}, t_{i,j} = k | \mathbf{y}_{i,j}, z_{i,j} = 1, b_{i,j}, \boldsymbol{\psi}_{i,j})$$

$$= w_k \sum_{i=0}^{\|\mathbf{y}_{i,j}\|_1} w_{i,k} p_{\mathcal{G}} \left(r_{i,j}; \alpha + i, \left(\delta_k + \frac{1}{\beta} \right)^{-1} \right) \tag{24}$$

where $\forall k \in \mathbb{T}, w_k > 0, \sum_{k \in \mathbb{T}} w_k, \forall (k, i), w_{i,k} > 0, \sum_{i=0}^{\|\mathbf{y}_{i,j}\|_1}$ $w_{i,k}$, and $p_{\mathcal{G}}(\cdot; \alpha, \beta)$ denotes the probability density function of the gamma distribution with shape α and scale β. Consequently, sampling from (23) can be achieved by first sampling $t_{i,j}$ (using $f_1(t_{i,j} = k | \mathbf{y}_{i,j}, z_{i,j} = 1, b_{i,j}, \boldsymbol{\psi}_{i,j}) = w_k$ by marginalization of $r_{i,j}$) and then sampling $r_{i,j}$ from $f_1(r_{i,j} | \mathbf{y}_{i,j}, z_{i,j} = 1, t_{i,j}, \boldsymbol{\psi}_{i,j})$. For brevity, the derivation of the probabilities $\{w_k\}$ and $\{w_{i,k}\}$ is omitted here. The interested reader is invited to consult [12] for further details. Note that using (24) as a proposal ensures (20) is satisfied and that the generated samples are asymptotically drawn from (15) for $(\pi_b, \pi_d) \in (0, 1)^2$.

Moving from $\mathcal{M}_1^{(i,j)}$ to $\mathcal{M}_0^{(i,j)}$: Finally, assume now that the current state of the chain is in $\mathcal{M}_1^{(i,j)}$, i.e., $\boldsymbol{\theta}_{i,j} = \boldsymbol{\theta}_{i,j}^1$. The reverse move associated with the move from $\mathcal{M}_1^{(i,j)}$ to $\mathcal{M}_0^{(i,j)}$ first consists of considering as candidate $\boldsymbol{\theta}_{i,j}^{*0} = b_{i,j}$, where $b_{i,j}$ is the current value of the background level (under model $\mathcal{M}_1^{(i,j)}$). This candidate (and thus the move to $\mathcal{M}_0^{(i,j)}$) is accepted with probability

$$\alpha \left[(1, \boldsymbol{\theta}_{i,j}^1), (0, \boldsymbol{\theta}_{i,j}^{*0}) \right] = \min \left[1, \rho \left((1, \boldsymbol{\theta}_{i,j}^1), (0, \boldsymbol{\theta}_{i,j}^{*0}) \right) \right] \tag{25}$$

where $\rho((1, \boldsymbol{\theta}_{i,j}^1), (0, \boldsymbol{\theta}_{i,j}^{*0})) = 1/\rho((0, \boldsymbol{\theta}_{i,j}^{*0}), (1, \boldsymbol{\theta}_{i,j}^1))$

V. SIMULATION RESULTS

A. Data Acquisition

We compare the performance of the method proposed in this paper to reconstruct a depth image of a life-sized polystyrene head (see Fig. 1) located at a distance of 325 meters from a time-of-flight scanning sensor, based on TCSPC. The transceiver system and data acquisition hardware used for this work is broadly similar to that described in [9], [12], [28]–[30], and 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 841 nm, and a silicon single-photon avalanche diode (SPAD) detector. The overall system had a timing jitter of 60 ps full width at half-maximum (FWHM). The face of the head was pointing towards the scanning system housed in the roof of the lab and a flat medium density fiberboard sheet mounted behind the head target acted as a backboard. This backboard was used for calibration purposes (see Fig. 1) and prevents the detection of return photons from objects placed behind the target of interest, the head in this example. The backboard did not obstruct direct sunlight from illuminating the target, thus not significantly reducing the background levels. To assess the performance of the proposed algorithm, the measured photon histograms have been truncated in order to discard the potential peaks associated with the backplane (i.e., the backplane
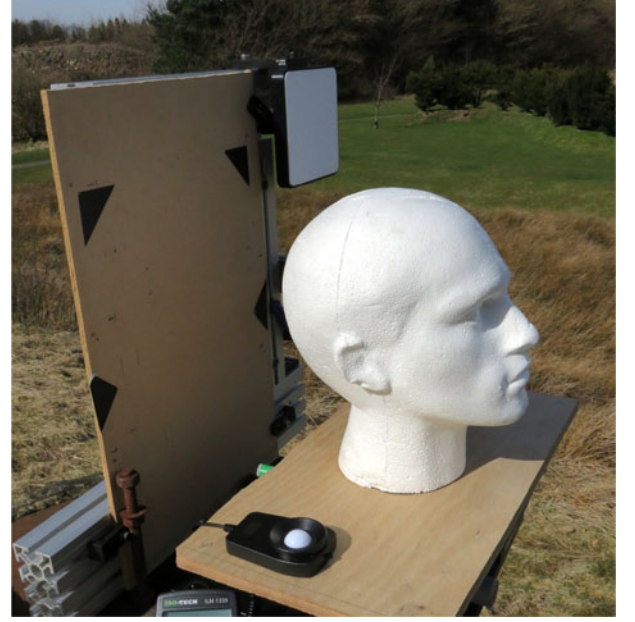

Fig. 1. Photograph showing the polystyrene head used for the experiments described here and calibration targets, including the Spectralon panel (top right corner of the fiberboard).

has been time-gated). The measurements were performed outdoors, on the Edinburgh Campus of Heriot-Watt University, in April 2015 under dry clear skies at 3 different times of day. That is, the same measurements were repeated at noon, 3 p.m. and 8 p.m. (dusk) with atmospheric conditions remaining relatively constant for the duration of each measurement. The key measurement parameters are summarized in Table I. The acquisition time per pixel in Table I is 30 ms. However, the data format of time-tagged 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 algorithm for acquisition times of 30 ms, 3 ms, 1 ms, and 300 μs per pixel, for which the number of detected photons per pixels is low.

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 325 m from laser source/detector. A long acquisition time (100 s) is considered here to reduce the impact of 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. 2. Note that the overall shape of this instrumental response, which is the important aspect for depth resolution [31], results from a combination of the laser pulse width itself and jitter from a number of sources, including detector jitter (the main contribution in this case) and jitter from the timing electronics (and other electronic components and cabling).

Table II provides details regarding the number of detected photons when varying the acquisition time. The top rows of Table II show the average number of detected photons per pixel.

TABLE I
MEASUREMENT KEY PARAMETERS

Target Stand-off Distance	≈ 325 m
Target Scene	Polystyrene head ($\approx 170 \times 285 \times 250$ mm in $W \times H \times D$ when viewed from the front) mounted on a breadboard. Backplane: MDF board.
Laser system	Supercontinuum laser source and tunable filter (NKT Photonics) fiber-coupled to the custom-designed transceiver unit
Illum. Wavelength	841 nm
Laser Repetition Rate	19.5 MHz
Illum. Power at target	$\approx 240\,\mu$W average optical power
Illum. Beam Diameter at Target	≈ 10 mm
Acquisition Mode	200×200 pixels scan centred on the head, covering an area of 285×285 mm at the scene Per-pixel acquisition time: 30 ms Total scan time: ≈ 20 minutes
Histogram bin width	2 ps
Histogram length	1500 bins (after gating)
Temporal Response of System	≈ 60 ps FWHM
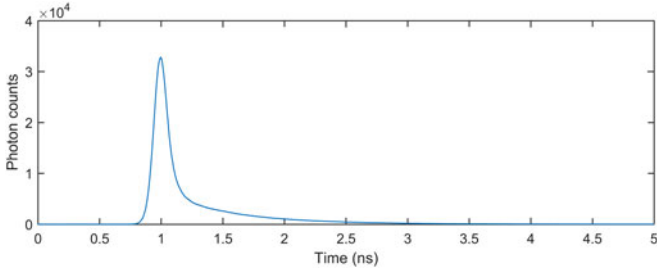
Fig. 2. Instrumental response obtained using Spectralon panel placed at 325 m from laser source/detector and for an acquisition time of 100 s (jitter ≈ 60 ps FWHM).

As expected, the number of detected photons increases linearly with exposure. The bottom rows show that the proportion of empty pixels (no detected photons) increases from noon to 8 p.m. (due to the average background levels which is higher during the day than at dusk), and that this proportion decreases when increasing the acquisition time (higher probability of detecting photons). Note that for the shortest acquisition (300 μs per pixel), less that 6 photons per pixel are detected on average, leading to a particularly challenging target detection problem.

B. Competing Method

The proposed method is compared to the standard method used for depth imaging [9] and which is divided into two steps. The first step consists of estimating $t_{i,j}$ using cross-correlation (i.e., by analysing the temporal correlation) between $\log(g_0(\cdot))$ and the photon histogram $\mathbf{y}_{i,j}$. For non-empty pixels, i.e., when

TABLE II
AVERAGE NUMBER OF DETECTED PHOTONS PER PIXEL AND PROPORTION OF EMPTY PIXELS FOR THE DIFFERENT ACQUISITIONS

		Acquisition Time			
		300 μs	1 ms	3 ms	30 ms
Av. photon counts	noon	5.6	18.5	55.5	554.6
	3 p.m.	4.1	13.7	41.0	408.9
	8 p.m.	1.2	4.9	11.6	116.0
Empty pixels (%)	noon	2.79	< 0.01	0	0
	3 p.m.	4.2	0.02	0	0
	8 p.m.	61.8	52.2	40.4	2.2

$\sum_{t=1}^{T} y_{i,j,t} > 0$, the object depth is estimated using

$$\hat{t}_{i,j,\mathrm{corr}} = \underset{\tau \in \mathbb{Z}}{\mathrm{argmax}} \sum_{t=1}^{T} y_{i,j,t} \log\left(g_0(t-\tau)\right). \qquad (26)$$

which corresponds to log-match filtering or maximum likelihood (ML) estimation under background-free ($b_{i,j} = 0$) assumption. Once the estimated time target distance $\hat{t}_{i,j,\mathrm{corr}}$ has been computed, the target intensity and the background level for each pixel are either both set to 0 (for empty pixels) or estimated using ML estimation as

$$\left(\hat{r}_{i,j,ML}, \hat{b}_{i,j,ML}\right) = \min_{\substack{r_{i,j} \geq 0 \\ b_{i,j} \geq 0}} C(\mathbf{y}_{i,j}, \hat{t}_{i,j,\mathrm{corr}}, r_{i,j}, b_{i,j}) \qquad (27)$$

with $C(\mathbf{y}_{i,j}, t_{i,j}, r_{i,j}, b_{i,j})$

$$= \sum_{t=1}^{T} y_{i,j,t} \log\left(r_{i,j} g_0\left((t-t_{i,j}) + b_{i,j}\right)\right)$$

$$- \sum_{t=1}^{T} \left[r_{i,j} g_0\left((t-t_{i,j}) + b_{i,j}\right)\right]. \qquad (28)$$

Note that (28) is convex with respect to $(r_{i,j}, b_{i,j})$ with $r_{i,j} \geq 0, b_{i,j} \geq 0$ and that (27) can be solved efficiently using constrained convex optimization methods (here we used an ADMM method similar to [32]). The proposed Bayesian algorithm has been applied with $N_{\mathrm{MC}} = 1000$ iterations, including $N_{\mathrm{bi}} = 300$ burn-in iterations. The computational complexity of the method mainly depends on the average number of detected photons per pixel and the number of admissible target depth T'. Indeed, these values have a direct impact on the number of weights to be computed in (24). For a Matlab R2014a implementation on a i7-3.0 GHz desktop computer (16 GB RAM), the average computational time is 12 hours per data set, ranging from 1.5 hour (8 p.m., shortest acquisition) to 36 hours (noon, longest acquisition). It is worth mentioning that other statistical tests, such as chi-square tests, could have been considered to perform the target detection, i.e., accept or reject the model (1) for each pixel. However, due to the statistical properties of the noise in (1) and the low background levels encountered in practice ($b_{i,j} << 1$), the distribution of the classical chi-square test statistic cannot be accurately approximated by a chi-square distribution and the detection results are generally worse than those obtained with the competing method described in this section.

TABLE III
EMPIRICAL DETECTION PERFORMANCE (PROBABILITIES IN %,
BEST RESULT IN BLUE)

			π_{00}	π_{10}	π_{01}	π_{11}
8 p.m.	3 ms	X-corr	98.9	1.1	13.3	86.7
		Prop. algo.	100	0	9.2	90.8
	1 ms	X-corr	99.5	0.5	14.1	85.9
		Prop. algo.	100	0	13.1	86.9
	0.3 ms	X-corr	99.1	0.9	15.5	84.5
		Prop. algo.	100	0	8.8	91.2
3 p.m.	3 ms	X-corr	86.0	14.0	9.7	90.3
		Prop. algo.	99.9	0.01	8.8	91.2
	1 ms	X-corr	60.7	39.3	15.9	84.1
		Prop. algo.	99.9	0.01	13.9	86.1
	0.3 ms	X-corr	62.4	37.6	37.3	62.7
		Prop. algo.	99.9	0.01	15.8	84.2
noon	3 ms	X-corr	79.9	20.1	8.9	91.1
		Prop. algo.	99.9	0.01	10.8	89.2
	1 ms	X-corr	57.4	42.6	16.9	83.1
		Prop. algo.	99.9	0.01	18.6	81.4
	0.3 ms	X-corr	59.6	40.4	39.1	60.9
		Prop. algo.	99.9	0.01	20.4	79.6

C. Target Detection

The detection performance of the algorithms is quantitatively assessed by comparing their empirical specificity π_{00} (deciding $\mathcal{M}_0^{(i,j)}$ when $\mathcal{M}_0^{(i,j)}$ is true) and sensitivity π_{11} or equivalently their empirical probability of false alarm $\pi_{10} = 1 - \pi_{00}$ and of miss $\pi_{01} = 1 - \pi_{11}$. Although the standard method does not provide target detection results directly, it is possible to infer the target presence by thresholding the estimated intensity images. In all the results presented here, we set the threshold to $\eta = 0.1$, which corresponds to an estimated target intensity 10 times smaller than that of the Spectralon panel. Moreover, we observed that this value of threshold yields a satisfactory threshold between high detection performance and high false alarm rate for the data considered. Table III compares the detection performance of the standard and proposed methods, averaged over all pixels, for the 3 sets of measurements (noon, 3 p.m. and 8 p.m.) and for acquisition times of 3 ms, 1 ms and 0.3 ms per pixel. For each data set, the results obtained by the proposed method with an acquisition time of 30 ms are used as ground truth. This table shows that the performance of the two algorithms degrade when reducing the acquisition time. However, the proposed method (as a consequence of its joint detection and estimation ability and the different spatial regularizations) generally provides lower probabilities of false alarm and miss as well as less significant performance degradation than the standard method.

For completeness, Fig. 3 depicts the posterior probabilities of target presence, i.e, $f(z_{i,j} = 1|\mathbf{Y}, \hat{\nu}, \hat{c}), \forall(i,j)$ for the observations at 8 p.m. (top), 3 p.m. (middle) and noon (bottom) and for the different per-pixel acquisition times. This figure shows that the proposed method is able to identify the central region (of high probability) where the head is located and highlights regions of high uncertainty, i.e. where $f(z_{i,j} = 1|\mathbf{Y}, \hat{\nu}, \hat{c}) \approx 0.5$, around the boundaries of the head where the detection is difficult due to the low reflectivity of the head is these regions. This
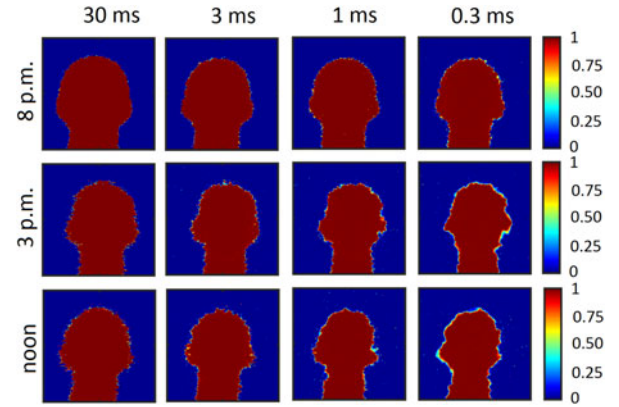

Fig. 3. Estimated posterior probability of target presence ($f(z_{i,j} = 1|\mathbf{Y}, \hat{\nu}, \hat{c}), \forall(i,j)$) for the different observation conditions considered. Red (resp. blue) regions correspond to high (resp. low) probabilities of presence.
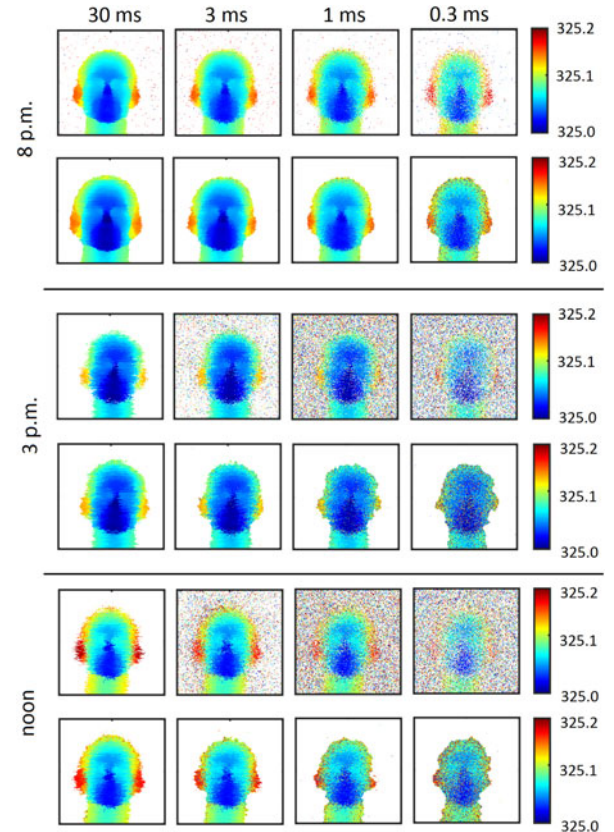

Fig. 4. Estimated depths for the 325 m target observed at 8 p.m. (top), 3 p.m. (middle) and noon (bottom) and for different per-pixel acquisition times. For each experiment, the top (resp. bottom) row is associated with the standard (resp. proposed) method.

figure also shows that the regions of high uncertainty generally broaden as the background levels increase (for fixed target reflectivity) and as the acquisition time decreases.

D. Parameter Estimation

Fig. 4 compares the estimated depth maps obtained by the standard and the proposed methods. These results show that for large acquisition times, the two methods provide similar results.
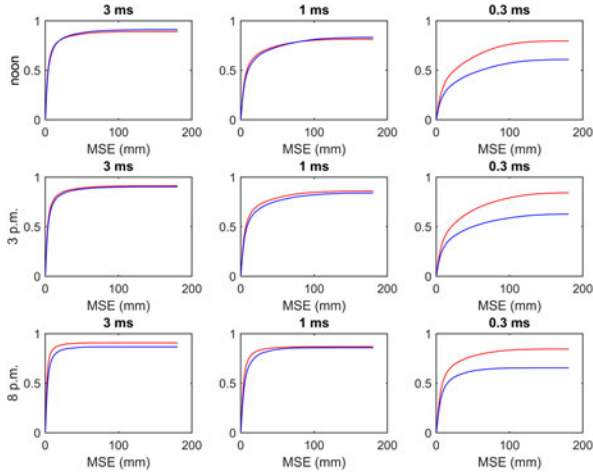
Fig. 5. Distance RMSE cdfs provided by the standard (blue) and the proposed (red) methods for the target located at 325 m at observed at noon (top), 3 p.m. (middle) and 8 p.m. (bottom).
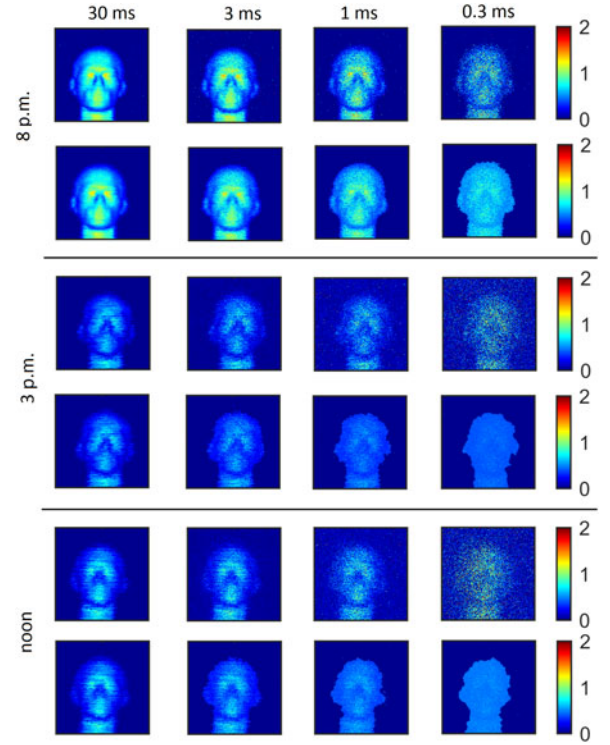

Fig. 6. Estimated intensity for the 325 m target observed at 8 p.m. (top), 3 p.m. (middle) and noon (bottom) and for different per-pixel acquisition times. For each experiment, the top (resp. bottom) row is associated with the standard (resp. proposed) method.

However, when the acquisition time decreases, the two methods start to fail in identifying the target positions, especially in pixels where no photon is detected. However, due to its better target detection ability, the proposed method provides more reliable depth images as it can more accurately detect pixels not containing a surface.

The performance of the two methods are quantitatively evaluated using the distance mean squared errors (MSEs) defined by $MSE(d_{i,j}) = \|\hat{d}_{i,j} - d_{i,j}\|_2^2$ where $\| \cdots \|_2$ denotes the ℓ_2-norm, $\hat{d}_{i,j}$ is the estimated value of $d_{i,j} = (3 \times 10^8)t_{i,j}/2$. Since the actual distances $\{d_{i,j}\}$ are unknown for the data sets considered, these values have been replaced by those estimated by the proposed method for the longest acquisition time (30 ms). Fig. 5 depict the cumulative density functions (cdfs) of the distance MSEs, defined by $F_d(\tau) = \frac{1}{N_{\text{row}}N_{\text{col}}} \sum_{i,j} \mathbf{1}_{(0,\tau)}(MSE(d_{i,j}))$ where $\mathbf{1}_{(0,\tau)}(\cdot)$ denotes the indicator function defined on $(0, \tau)$. Note that for each dataset, the cdfs are upperbounded by the sensitivity π_{11} of each method. This figure shows that for the pixels containing targets, the two methods provide similar depth estimation performance for the longer acquisition times and that the proposed method is more robust when reducing the acquisition time, thanks to its target detection ability. It is important to recall that in contrast to [12], the proposed algorithm (as the competing method) does not explicitly account for the spatial correlation of the target depths, which is why the two competing methods provide similar results for long acquisitions. Accounting for such correlations in future target detection and depth imaging methods could further improve the robustness of the method.

Fig. 6 compares the estimated intensity maps obtained by the proposed method (for the pixels containing a target) and after thresholding for the standard method ($\eta = 0.1$). These results show that the two methods provide similar results for the longest acquisition times and that the proposed method is more robust to the lack/absence of detected photons. In particular, for acquisition times shorter than 3ms per pixel, it becomes difficult to estimate accurately the intensity of the target, whatever

the background levels. By assuming that the target intensities share the same statistical properties (through (5)), the proposed method provides more homogeneous intensity images in the region containing the target than the standard method, which in turn enhance the target detection. 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).

Finally, Fig. 7 compares the background levels estimated by the two methods for the different measurements. The two top rows (8 p.m.) of Fig. 7 show that for the longer acquisition times, higher backgrounds are estimated in region of significant depth changes, which can be primarily explained by a model mismatch. In particular, due to the laser beam size and the orientation of the target surface, the peak in the photon histogram can become broader than that depicted in Fig. 2. The boundary between the head chin and neck is an even more extreme case where two peaks can be observed. Under brighter observation conditions however (middle and bottom rows of Fig. 7), these effects become negligible and the background images estimated by the proposed method are in agreement with the observation conditions. Indeed, the detected background photons correspond mainly to photons emitted by external sources (e.g., the sun) and reflected onto the targets. Thus, we can expect (assuming homogeneous ambient illumination) the background levels to be higher in pixels where more reflective surfaces are present. At 3 p.m. (middle rows), it can be observed that the background levels are generally in the head region and lower in the backplane region. It can be observed that the background levels are
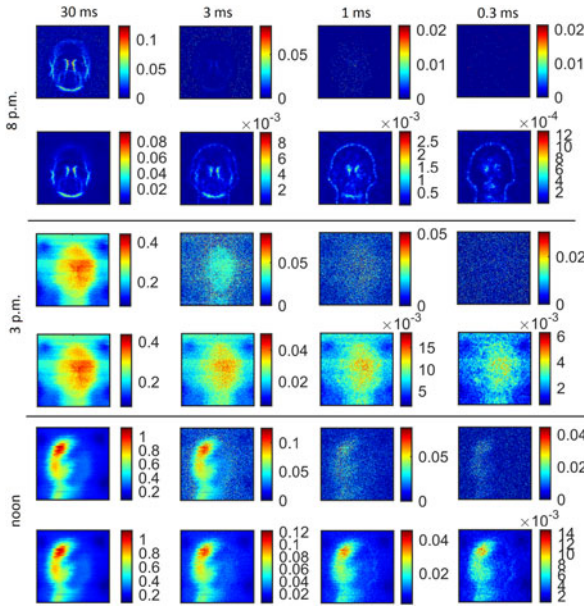
Fig. 7.   Estimated background levels for the 325 m target observed at 8 p.m. (top), 3 p.m. (middle) and noon (bottom) and for different per-pixel acquisition times. For each experiment, the top (resp. bottom) row is associated with the standard (resp. proposed) method.

particularly low in the regions where black calibration markers have been placed (see [12, Fig. 4]). Finally, the bottom rows of Fig. 7 clearly show higher background levels on the left-hand side of the head, due to the more direct sun illumination at noon and the head orientation with respect to the sun and observation directions. Note that the regions of particularly low background levels correspond to the four dark triangles mounted on the fiberboard and visible in Fig. 1. Fig. 7 also illustrates the benefits or the background model (4) used in the proposed method. Accounting for the spatial correlations of the background levels to regularize the target detection problem 1) provides more realistic background images compared to the pixel-by-pixel method and 2) enhance the target detection. Moreover, by achieving simultaneously the target detection and identification, the proposed method is more robust than the standard method whose performance highly relies on the first depth estimation step.
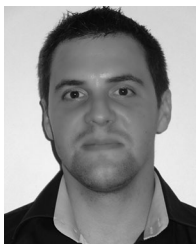
VI. Conclusion

In this paper, we have presented a Bayesian algorithm for joint target detection and depth imaging using sparse single-photon data. This problem was translated into a pixel-wise model selection problem and a Bayesian hierarchical model was proposed to describe the expected correlations between the pixels of the observed image through appropriate prior distributions. To perform Bayesian inference based on the resulting posterior distribution, we proposed a reversible-jump MCMC algorithm which allows efficient moves between the different parameter spaces. The experiments conducted on real Lidar data demonstrate the ability of the proposed method to 1) detect and 2) identify targets observed under difficult observation conditions (high and spatially variable background levels, short acquisition times), with a better accuracy than existing methods. An important property of the proposed method is its capacity to adjust

automatically the different spatial regularization parameters, thus relieving practitioners from the difficult task of setting them by cross-validations. In contrast to [12], we have not explicitly accounted for the possible correlations affecting the intensity and/or depth images. Although it is possible to apply the algorithm studied in [12] to refine the depth/intensity images after the target detection step, e.g. in an empirical Bayes fashion, it would be interesting to extend the model proposed in Section III to capture additional parameter dependency within the target detection procedure. This could be achieved by constructing a depth and/or reflectivity prior model conditioned on the values of the detection labels e.g., $f(\mathbf{T}|\mathbf{Z})$. It would also be interesting the correlate the background levels and the target reflectivities. In order to reduced the computational complexity of the target detection, it would be interesting to investigate optimization-based alternatives (e.g., expectation-maximization methods) to be compared with the proposed method in terms of accuracy (estimation performance) and robustness (convergence issues). Finally, the model considered assumed the potential presence of a single target per pixel, which might not be realistic for specific applications. Although the detection of multiple targets using sparse single-photon data is a significantly more difficult problem, extending the model to multiple targets is subject of further investigations that will be reported in subsequent papers.
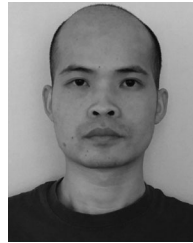
References

[1] T. Ogawa and K. Takagi, "Lane recognition using on-vehicle Lidar," in *Proc. IEEE Intell. Veh. Symp.*, 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. Veh. Veh. Syst.*, 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 Conf. Adv. Geosci.*, Mykonos, Greece, May 2012, pp. 59–60.

[6] T. Hakala, J. Suomalainen, S. Kaasalainen, and Y. Chen, "Full waveform hyper-spectral Lidar for terrestrial laser scanning," *Opt. Express*, 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, pp. 155–166, 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.*, 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. Express*, vol. 21, no. 19, pp. 22098–22113, Sep. 2013.

[10] A. Maccarone *et al.*, "Underwater depth imaging using time-correlated single-photon counting," *Opt. Express*, vol. 23, no. 26, pp. 33911–33926, Dec. 2015.

[11] A. Kirmani *et al.*, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.

[12] 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 Process.*, vol. 25, no. 5, pp. 1935–1946, May 2016.

[13] 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.

[14] 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.

[15] P. J. Green, "Reversible jump MCMC computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, Dec. 1995.

[16] C. Andrieu and A. Doucet, "Joint Bayesian model selection and estimation of noisy sinusoids via reversible jump MCMC," *IEEE Trans. Signal Process.*, vol. 47, no. 10, pp. 19–37, Oct. 1999.

[17] 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 Workshop Statist. Signal Process.*, Gold Coast, Australia, Jul. 2014, pp. 121–124.

[18] D. Shin, A. Kirmani, V. Goyal, and J. Shapiro, "Computational 3D and reflectivity imaging with high photon efficiency," in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2014, pp. 46–50.

[19] G. Gariepy *et al.*, "Single-photon sensitive light-in-fight imaging," *Nature Commun.*, vol. 6, Jan. 2015, Art. no. 6021.

[20] O. Dikmen and A. Cemgil, "Gamma Markov random fields for audio source modeling," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 3, pp. 589–601, Mar. 2010.

[21] C. P. Robert and G. Casella, *Monte Carlo Statistical Methods*, 2nd ed. New York, NY, USA: Springer-Verlag, 2004.

[22] Y. Altmann, M. Pereyra, and J. Bioucas-Dias, "Collaborative sparse regression using spatially correlated supports—Application to hyperspectral unmixing," *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 5800–5811, Dec. 2015.

[23] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, "Estimating the granularity coefficient of a Potts-Markov random field within an MCMC algorithm," *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2385–2397, Jun. 2013.

[24] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Secaucus, NJ, USA: Springer-Verlag, 2006.

[25] M. Zhou, L. Hannah, D. Dunson, and L. Carin, "Beta-negative binomial process and poisson factor analysis," in *Proc. 15th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 1462–1471.

[26] W. R. Gilks and P. Wild, "Adaptive rejection sampling for Gibbs sampling," *Appl. Stat.*, vol. 41, pp. 337–348, 1992.

[27] S. P. Brooks, P. Giudici, and G. O. Roberts, "Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions," *J. Roy. Statist. Soc. B, (Statist. Methodol.)*, vol. 65, no. 1, pp. 3–39, 2003. [Online]. Available: http://dx.doi.org/10.1111/1467–9868.03711

[28] 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.

[29] 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, Apr. 2010.

[30] 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 J. Adv. Signal Process.*, vol. 2010, 2010, Art. no. 896708.

[31] S. Pellegrini, G. S. Buller, J. M. Smith, A. M. Wallace, and S. Cova, "Laser-based distance measurement using picosecond resolution time-correlated single-photon counting," *Meas. Sci. Technol.*, vol. 11, no. 6, 2000, Art. no. 712.

[32] M. Figueiredo and J. Bioucas-Dias, "Restoration of poissonian images using alternating direction optimization," *IEEE Trans. Image Process.*, vol. 19, no. 12, pp. 3133–3145, Dec. 2010.

Ximing Ren received the B.Eng. degree in optical engineering from Beijing Institute of Technology, Beijing, China, in 2008, and the M.Eng. degree in electrical engineering from Beihang University, Beijing, in 2011. In 2011, he started his Ph.D. research with a prize scholarship from the Scottish Universities Physics Alliance in the Single-Photon Group at Heriot-Watt University, Edinburgh, U.K., studying single-photon time-of-flight imaging under the supervision of Prof. G. S. Buller. He received the Ph.D. degree in physics in November 2015. He is a Postdoctoral Research Associate in the Single-Photon Group, Heriot-Watt University. He is a member of the Optical Society of America (OSA) and a member of the Institute of Physics. He provides voluntary peer-review for IEEE PHOTONICS TECHNOLOGY LETTERS, *Chinese Optics Letters* (OSA), and *Optical Engineering* (SPIE).

Aongus McCarthy (M'03) received the B.Sc. degree from University College Galway, 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 worked in industry from 1992 until 1997 as a Design Team Leader on the development of a thermal transfer printing system. 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, Glasgow, U.K., in 1986, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 1989. He is currently a Professor of physics at Heriot-Watt University. In 2002, he co-founded Helia Photonics Ltd., Livingston, U.K. In 2015, he received an EPSRC Established Career Fellowship in Quantum Technology to research sparse photon imaging. 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, a Fellow of the Institute of Physics, U.K., and a Fellow of the Optical Society of America.

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

Steve McLaughlin (F'11) received the B.Sc. degree in electronics and electrical engineering from the University of Glasgow, , Glasgow, U.K., in 1981, and the Ph.D. degree from the University of Edinburgh, Edinburgh, U.K., in 1990. In 1988, he joined the academic staff at Edinburgh, and from 1991 until 2001, he held a Royal Society University Research Fellowship to study nonlinear signal processing techniques. In 2002, he was awarded a personal Chair in Electronic Communication Systems at the University of Edinburgh. In October 2011, he joined Heriot-Watt University as a Professor of signal processing and Head of the School of Engineering and Physical Sciences. His research interests include 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.