RANGE ESTIMATION FROM SINGLE-PHOTON LIDAR DATA USING A STOCHASTIC EM APPROACH

Yoann Altmann and Stephen McLaughlin

School of Engineering and Physical Sciences
Heriot Watt University
Edinburgh, UK

ABSTRACT

This paper addresses the problem of estimating range profiles from single-photon waveforms in the photon-starved regime, with a background illumination both high and unknown a priori such that the influence of nuisance photons cannot be neglected. We reformulate the classical observation model into a new mixture model, adopt a Bayesian approach and assign prior distributions to the unknown model parameters. First, the range profile of interest is marginalised from the Bayesian model to estimate the remaining model parameters (considered as nuisance parameters) using a stochastic EM algorithm. The range profile is then estimated via Monte Carlo simulation, conditioned on the previously estimated nuisance parameters. Results of simulations conducted with controlled data demonstrate the possibility to maintain satisfactory range estimation performance in high-background scenarios with less than 10 signal photons per pixel on average.

Index Terms— Single-photon Lidar, Bayesian estimation, mixture model, Stochastic Expectation-Maximisation

1. INTRODUCTION

Time-of-flight light detection and ranging (Lidar) systems are increasingly used in many applications (e.g., robotics, environmental sciences and defence) to reconstruct 3-dimensional structures. By illuminating the scene with a train of laser pulses and analysing the distribution of the photons reflected from the targets, it is possible to infer range and reflectivity information about the scene. In this work, we assume that for each pixel, the detected photons result either from direct path reflections of the photons originally emitted by the laser source onto a single surface of the scene or from dark photon counts and ambient illumination (background photons). These assumptions generally hold for short to mid-range (up to dozens of metres) depth imaging where the divergence of the laser source can be neglected. The presence of high and unknown background levels in the photon-starved regime, which occurs for instance in short measurements performed in daylight conditions, make methods relying on the classical log-matched filter approach [1, 2] collapse. In [1], we proposed a Markov chain Monte Carlo (MCMC) method to estimate the range profile in such scenarios via maximum *a posteriori* (MAP) estimation. However, this approach suffers from a high computational cost induced by the random simulation process itself and the need for accept-reject procedures or complex operations (depending on the implementation) to update the target reflectivity and background profiles. Less computationally intensive alternatives such as those adopted in [3, 4] consist of splitting this original problem into smaller problems which are easier to solve. For instance, in [3], the background profile was estimated first using convex optimisation, followed by a similar algorithm for estimating the reflectivity parameters and finally the depth profile. While such approaches are valid when data acquired without using active illumination is available, they become impractical when the background changes between the calibration and the actual measurements (dynamic scene or moving sensor).

Here, we propose a new ranging algorithm which does not require knowledge of the background levels in advance, for applications where these nuisance parameters cannot be extracted directly from the data (as e..g., done in [3]). To address the multimodality of the likelihood function w.r.t. the range parameters, we propose to marginalise these parameters from the model when estimating the model parameters associated with target reflectivity and background levels. This is achieved by reformulating the observation model as a mixture model for which Bayesian inference can be performed efficiently using Expectation-Maximisation (EM)-based [5] algorithms with a reduced computational cost compared to MCMC alternatives [1]. The remainder of the paper is organised as follows. Section 2 introduces the new formulation of the observation model and associated Bayesian model used for range estimation. The proposed estimation strategy is detailed in Section 3 and Section 4 discusses results of experiments conducted with controlled Lidar waveforms to illustrate the performance of the proposed method. Conclusions and future work are finally reported in Section 5.

Part of this work was supported by the Royal Academy of Engineering under the Research Fellowship Scheme.

ISBN 978-90-827970-1-5 © EURASIP 2018

2. BAYESIAN MODEL

We consider a 3-D array $\mathbf{Y}$ of single wavelength Lidar waveforms of dimension $N_{\text{row}} \times N_{\text{col}} \times T$, where N_{row} (resp. N_{col}) stands for the number of rows (resp. columns) of the regular spatial sampling grid (in the transverse plane), and T is the number of temporal (corresponding to range) bins. Let $\mathbf{y}_{i,j} = [\mathbf{Y}]_{i,j,:} = [y_{i,j,1}, \ldots, y_{i,j,T}]^T$ be the Lidar waveform obtained in the pixel (i,j). The element $y_{i,j,t}$ is the photon count within the tth time bin of the histogram and $\bar{y}_{i,j} = \sum_{t=1}^{T} y_{i,j,t}$ is the total number of photons detected in that pixel.

2.1. Poisson noise model

Let $d_{i,j}$ be the position of an object surface at a given range from the sensor, whose reflectivity is denoted by $0 \leq r_{i,j} \leq 1$. According to [1,6], in the low-flux regime, each photon count $y_{i,j,t}$ is assumed to be drawn from the Poisson distribution

$$y_{i,j,t}|(r_{i,j}, b_{i,j}, t_{i,j}) \sim \mathcal{P}\left(r_{i,j} g\left(t - t_{i,j}\right) + b_{i,j}\right), \qquad (1)$$

where $g(\cdot)$ is the instrumental response of the system, whose scale depends on the integration time and whose shape is assumed to be known, as occurs when it can be accurately estimated during imaging system calibration. In Eq. (1), $t_{i,j}$ is the characteristic time-of-flight of photons emitted by a pulsed laser source and reaching the detector after being reflected by a target at range $d_{i,j}$ ($d_{i,j}$ and $t_{i,j}$ are linearly related in free-space propagation). Moreover the parameters $\{b_{i,j}\}_{i,j}$ gather dark counts and background illumination levels. As explained in [2], in the low-flux regime, the model in Eq. (1) is an accurate observation model when the probability of multiple detections in each illumination period is negligible and when detector dead-time can be neglected . Here, we further assume without loss of generality that the integral of the instrumental response $G = \sum_{t=1}^{T} g\left(t - t_{i,j}\right)$ is constant for any admissible value of $t_{i,j}$. This assumption holds in practice as the length T of the observation window is generally chosen such that the photon returns associated with surfaces of interest are not cropped.

As discussed in the introduction, the estimation of $\mathbf{T} = \{t_{i,j}\}_{i,j}$, assuming that $\mathbf{R} = \{r_{i,j}\}_{i,j}$ and $\mathbf{B} = \{b_{i,j}\}_{i,j}$ in Eq. (1) are also unknown is a difficult problem. In this work we reformulate Eq. (1) to estimate $\mathbf{T}$ more efficiently. Note that once $\mathbf{T}$ has been estimated, it is straightforward to estimate $(\mathbf{R}, \mathbf{B})$ e.g., via MAP estimation using (1) and convex optimisation tools such as [7, 8], provided that a log-concave prior model is assigned to$(\mathbf{R}, \mathbf{B})$. Consequently, in this work we mainly focus on estimating $\mathbf{T}$.

2.2. Alternative observation model

Let $s_{i,j}^p \in (1, T)$ denote the time of arrival (ToA) of the pth photon detected in the pixel (i,j), with $1 \leq p \leq \bar{y}_{i,j}$. In the low-flux regime, this photon was originally emitted by the laser source with probability $w_{i,j} = r_{i,j}G/(r_{i,j}G + Tb_{i,j})$ and is a background photon with probability $1 - w_{i,j}$. Thus, an alternative observation model associated with the time of arrival of each detected photon can be expressed, $\forall p \in (1, \ldots, \bar{y}_{i,j})$, as $f(s_{i,j}^p|w_{i,j}, t_{i,j}) = w_{i,j} f_s\left(s_{i,j}^p|t_{i,j}\right) + (1 - w_{i,j}) f_b(s_{i,j}^p)$ with $f_s\left(s_{i,j}^p|t_{i,j}\right) = g\left(s_{i,j}^p - t_{i,j}\right)/G$ and $f_b(s_{i,j}^p) = 1/T$. Note that we used $f_s(\cdot)$ and $f_b(\cdot)$ to denote the distributions of the ToAs of "signal" and "background" photons, respectively. Assuming mutual independence between the observed ToAs (conditioned on $\mathbf{W} = \{w_{i,j}\}_{i,j}$ and $\mathbf{T}$) yields the following joint likelihood

$$f(\mathbf{S}|\mathbf{W}, \mathbf{T}) \quad = \quad \prod_{i,j} \prod_{p=1}^{\bar{y}_{i,j}} f(s_{i,j}^p|w_{i,j}, t_{i,j}). \qquad (2)$$

In a similar fashion to Eq. (1), the likelihood in Eq. (2) only depends on $\mathbf{T}$ and $(\mathbf{R}, \mathbf{B})$ (through $\mathbf{W}$). However, as will be shown in Section 3, this new formulation allows us to use efficient EM-based algorithms to simplify the estimation of $\mathbf{T}$ (and $\mathbf{W}$). The remainder of this section discusses the prior models assigned to $(\mathbf{T}, \mathbf{W})$ to complete our Bayesian model.

2.3. Parameter prior models

As in [1, 9], the target positions are assumed to be discrete variables defined on $\mathbb{T} = \{t_{min}, \ldots, t_{max}\}$, such that $1 \leq t_{min} \leq t_{max} \leq T$. As in [1, 10], to account for the spatial correlations between the neighbouring pixels, we propose to use a discrete Markov random field (MRF) as a prior distribution for $t_{i,j}$ given its neighbours $\mathbf{T}_{\mathcal{V}(i,j)}$, where $\mathcal{V}(i,j)$ is the neighbourhood of the pixel (i,j) and $\mathbf{T}_{\setminus(i,j)} = \{(i',j')\}_{(i',j') \neq (i,j)}$. More precisely, we propose the following discrete MRF $f(\mathbf{T}|\epsilon) \propto \exp\left[-\epsilon\phi(\mathbf{T})\right]$, where $\epsilon \geq 0$ is a parameter tuning the amount of correlation between ranges in adjacent pixels, and where $\phi(\mathbf{T})$ corresponds to a total-variation (TV) regularisation [11, 12] promoting piecewise constant depth image. Here, a four-pixel structure is used to define the neighbourhood structure of $\phi(\mathbf{T})$. For brevity, we assume that ϵ fixed, but it could also be estimated using the methods presented in [13, 14].

The elements of $\mathbf{W}$ represent the expected ratios of detected photons originally emitted by the laser source. In most natural images, the reflectivity and background profiles are expected to present (correlated) spatial structures. To account for this prior belief, we propose to assign $\mathbf{W}$ the following smoothness promoting, proper and log-concave prior distribution

$$f(\mathbf{W}|\gamma^2) \propto \exp\left[-\frac{\|\mathbf{D}\mathbf{w}\|_2^2}{2\gamma^2}\right] \mathbf{1}_{(0,1)^N}(\mathbf{w}) \qquad (3)$$

where $\mathbf{w}$ represents the vectorized matrix $\mathbf{W}$, γ^2 is a fixed hyperparameter and $\mathbf{D}$ is the block circulant matrix of the 2D Laplacian filter.

ISBN 978-90-827970-1-5

2.4. Joint Posterior distribution

Now that we have defined the new observation model and assigned prior distributions to the unknown model parameters, we can use the Bayes' rule to derive the joint posterior distribution of $(\mathbf{W}, \mathbf{T})$ given by

$$f(\mathbf{W}, \mathbf{T}|\mathbf{S}, \epsilon, \gamma^2) \propto f(\mathbf{S}|\mathbf{W}, \mathbf{T}) f(\mathbf{W}|\gamma^2) f(\mathbf{T}|\epsilon) \quad (4)$$

assuming prior independence between $\mathbf{W}$ and $\mathbf{T}$. The next section details the algorithm used to infer $(\mathbf{W}, \mathbf{T})$.

3. ESTIMATION STRATEGY

Estimating jointly $(\mathbf{W}, \mathbf{T})$ (or equivalently the original parameters in (1)) is difficult due to the shape of the likelihoods (1) and (2) which are generally multimodal with respect to $t_{i,j}$. Thus, we propose to first estimate $\mathbf{W}$ by marginalising over $\mathbf{T}$ which is considered at first as a set of nuisance parameters. After computing the marginal maximum a posteriori (MMAP) estimator of $\mathbf{W}$

$$\widehat{\mathbf{W}} = \underset{\mathbf{W}}{\operatorname{argmax}} \; f(\mathbf{W}|\mathbf{S}, \epsilon, \gamma^2), \quad (5)$$

where $f(\mathbf{W}|\mathbf{S}, \epsilon, \gamma^2) = \sum_{\mathbf{T}} f(\mathbf{W}, \mathbf{T}|\mathbf{S}, \epsilon, \gamma^2)$, we infer the range profile by analysing the posterior distribution $f(\mathbf{T}|\mathbf{S}, \widehat{\mathbf{W}}, \epsilon, \gamma^2)$ which can be exploited more easily than the full posterior distribution (4). More precisely, in a similar fashion to [1], the range profile is estimated by MMAP estimation, i.e.,

$$\hat{t}_{i,j} = \underset{t_{i,j}}{\operatorname{argmax}} f(t_{i,j}|\mathbf{S}, \widehat{\mathbf{W}}, \epsilon, \gamma^2), \forall (i,j). \quad (6)$$

3.1. Estimation of W

The EM-like [5] algorithms are particularly adapted to solve problem involving (discrete) latent variables such as that involved in (5) . However, the standard EM algorithm cannot be used here since the Expectation (E) step is intractable due to the Markovian nature of $f(\mathbf{T}|\epsilon)$. Although variational approximations could be used (as in [13]), such EM-like approaches become computationally expensive for large values of $\Delta_t = t_{max} - t_{min}$. Indeed, the derivative of the expected log-likelihood w.r.t $w_{i,j}$ involved in the Maximisation (M) step of the EM algorithm consists of a sum of $\Delta_t \bar{y}_{i,j}$ terms. Thus, we resort to the stochastic EM (SEM) algorithm [15] summarised in Algo. 1. At the kth iteration, the classical E-step is replaced by the simulation of an auxiliary depth profile $\mathbf{T}^{(k)}$ according to $f(\mathbf{T}|\mathbf{S}, \mathbf{W}^{(k-1)}, \epsilon, \gamma^2)$. This is achieved using a 2-step Gibbs sampler accounting for the first-order neighbourhood structure of $f(\mathbf{T}|\epsilon)$ (see Section 3.2). The M-step then reduces to maximising the cost function

$$\begin{aligned} C_{\mathbf{T}^{(k)}}(\mathbf{W}) &= \log(f(\mathbf{W}, \mathbf{T}^{(k)}|\mathbf{S}, \epsilon, \gamma^2)) + c \quad (7) \\ &= \sum_{i,j} \sum_{p=1}^{\bar{y}_{i,j}} \left[\log(f(s_{i,j}^p | w_{i,j}, t_{i,j}^{(k)})) \right] - \frac{\|\mathbf{D}\mathbf{w}\|_2^2}{2\gamma^2}, \end{aligned}$$

subject to $\mathbf{W} \in (0,1)^N$, where c is a constant. It can be easily shown that $C_{\mathbf{T}^{(k)}}(\mathbf{W})$ is a concave function and its maximisation w.r.t. $\mathbf{W}$ can thus be achieved using state-of-the-art convex optimisation methods. Here we used an alternating direction method of multipliers (ADMM) method, which is not detailed here due to space constraints. The SEM algorithm is stopped after an arbitrarily fixed number of iterations N_{iter} and $\widehat{\mathbf{W}}$ is obtained by averaging the generated samples $\mathbf{W}^{(k)}{}_k$ in a similar fashion to MCMC methods, after removing the first samples corresponding to the transient regime of the algorithm.

ALGORITHM 1

SEM algorithm

1: Fixed input parameters: ϵ, γ^2, number of burn-in iterations N_{bi}, total number of iterations N_{iter}.
2: Initialization $(k = 0)$
3: Set $\mathbf{W}^{(0)}, \mathbf{T}^{(0)}$
4: **for** $k = 1, \ldots N_{\text{iter}}$ **do**
5: Stochastic Expectation step: Sample $\mathbf{T}^{(k)} \sim f(\mathbf{T}|\mathbf{S}, \mathbf{W}^{(k-1)}, \epsilon, \gamma^2)$
6: Maximization step: solve $\mathbf{W}^{(k)} = \underset{\mathbf{W} \in (0,1)^N}{\operatorname{argmax}} \; C_{\mathbf{T}^{(k)}}(\mathbf{W})$
7: **end for**
8: Set $\widehat{\mathbf{W}} = 1/(N_{\text{iter}} - N_{\text{bi}}) \sum_{k=N_{\text{bi}}+1}^{N_{\text{iter}}} \mathbf{W}^{(k)}$

3.2. Estimation of T

Once $\widehat{\mathbf{W}}$ has been estimated, the solution of Eq. (6) can be approximated *via* Monte Carlo integration by simulating from $f(\mathbf{T}|\mathbf{S}, \widehat{\mathbf{W}}, \epsilon, \gamma^2)$. This can be achieved sampling sequentially each depth from its conditional distribution $f(t_{i,j}|\mathbf{S}, \widehat{\mathbf{W}}, \mathbf{T}_{\mathcal{V}(i,j)}, \epsilon, \gamma^2)$, i.e., by drawing randomly from discrete (with finite support) distributions. In a similar fashion to [1], we use a Gibbs sampler implemented using a colouring scheme such that many depths can be updated in parallel (2 steps required when considering a 1-order neighbourhood structure). The solution in Eq. (6) for each pixel is then approximated by the most frequently generated value. Note that the chain generated can also be used to compute *a posteriori* confidence regions for each depth parameter.

4. RESULTS

We investigate the performance of the proposed method by analysing synthetic data based on actual measurements recorded with a single-photon Lidar system developed at Heriot-Watt University. Precisely, we use as instrumental response $g(\cdot)$ the response of the system presented in [3] and depicted in red in Fig. 1 (top). The reference reflectivity and range profiles ($\mathbf{R}$ and $\mathbf{T}$) used in this section are those of the scene considered in [3] and which consists of a

ISBN 978-90-827970-1-5 © EURASIP 2018

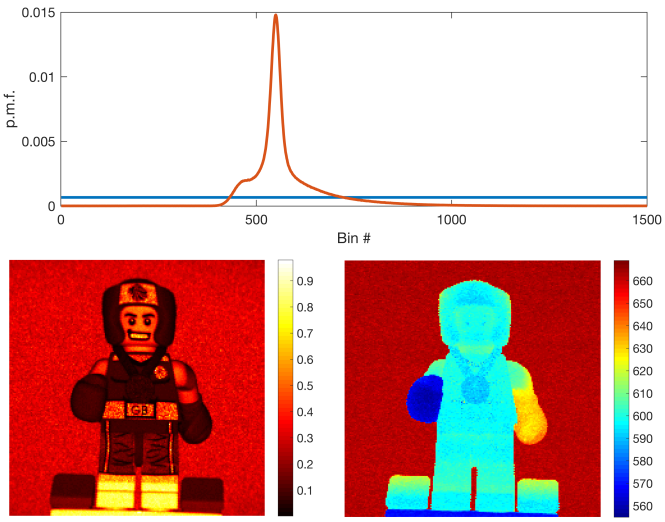

Fig. 1. Top: Probability mass functions (p.m.f.) of the time of arrival of a signal photon (red) for $t_{i,j} = 550$ and a background photon (blue). Bottom: Actual reflectivity map (left) and range map (right, in bins) associated with the simulated data.

mini-figurine placed in front of a backboard. The reflectivity and range profiles which consist of 200×200 pixels are depicted in the bottom left and right subplots of Fig. 1, respectively. At the wavelength considered (532nm), the scene includes surfaces of low reflectivity (< 0.1) which are thus more difficult to detect, in particular in the photon-starved regime. We generated synthetic Lidar waveforms according to (1) with $T = 1500$ temporal bins, each bin representing a 2ps time interval and thus a range resolution of 300μm (free-space light propagation). Without loss of generality, spatially uniform background levels $\mathbf{B}$ are considered for all the generated data. To vary the quality of the measurements, we introduce two parameters, α and β, such that $g\left(s^p_{i,j} - t_{i,j}\right) = \alpha f_s\left(s^p_{i,j}|t_{i,j}\right)$ and $b_{i,j} = \alpha\beta/T$. The values of α and β are chosen in $\alpha \in \{10; 20; 30; 50; 100; 200; 500\}$ and $\beta \in \{0.05; 0.1; 0.3; 0.5\}$. While α controls the overall amount of photons detected in each waveform, β (together with $r_{i,j}$) controls the fraction of background photons detected. Precisely, for a given pixel associated with a surface reflectivity defined by $r_{i,j}$, using (1), the expected number of detected photons is given by $\mathrm{E}[\bar{y}_{i,j}] = \alpha(r_{i,j} + \beta)$ and the expected fraction of signal photons is $w_{i,j} = r_{i,j}/(r_{i,j} + \beta)$. With the parameters considered, the mean fraction of signal photons $\mathrm{E}[w_{i,j}]$ (averaged over all the pixels) is equal to 0.80, 0.69, 0.45 and 0.34 with $\beta = 0.05$, $\beta = 0.1$, $\beta = 0.3$ and $\beta = 0.5$, respectively. The average photon counts per pixel for the different values of (α, β) are presented in Table 1. In the most challenging scenario ($(\alpha, \beta) = (10, 0.5)$), on average, 2.6 signal photons are detected per pixel out of 7.9 (i.e., 34% of the detected photons). Based on the Poisson statistics of the observations, we define the peak signal-to-noise ratio

		α						
		500	200	100	50	30	20	10
β	0.5	392.8	157.1	78.6	39.3	23.6	15.7	7.9
	0.3	292.7	117.1	58.6	29.3	17.6	11.7	5.9
	0.1	192.8	77.1	38.6	19.3	11.6	7.7	3.9
	0.05	167.8	67.1	33.6	16.8	10.1	6.7	3.3

Table 1. Mean photon counts ($\mathrm{E}[\alpha(r_{i,j}+\beta)]$) for the different values of (α, β) considered.

(PSNR) as $PSNR_{i,j} = \dfrac{(\alpha r_{i,j} f_M)^2}{\alpha r_{i,j} f_M + \alpha\beta/T}$, where f_M is the maximum value of $f_s\left(\cdot|t_{i,j}\right)$. The mean PSNRs averaged over all the pixels are presented in Table 2. For values of α equal or smaller than 200, i.e., in the photon-starved regime, the average PSNRs are smaller than 1 and the peak variance of the data is larger that the power of the peak energy on average over all the pixels. Moreover, although the value of β does not significantly affect the PSNR in the range considered, it affects the fraction of nuisance photons whose increase yields a more ill-posed problem.

		α						
		500	200	100	50	30	20	10
β	0.5	1.96	0.79	0.39	0.20	0.12	0.08	0.04
	0.3	2.02	0.81	0.40	0.20	0.12	0.08	0.04
	0.1	2.09	0.83	0.42	0.21	0.12	0.08	0.04
	0.05	2.10	0.84	0.42	0.21	0.13	0.08	0.04

Table 2. Average PSNRs for the different values of (α, β) considered.

The regularisation parameters of the model have been set to $(\epsilon, \gamma^2) = (0.07, 1)$ from preliminary runs and were kept unchanged for all the values of (α, β) considered. Optimisation of these parameters is out of scope of this paper but it is worth noting that ϵ could also be adjusted using methods such as those proposed in [13] or [14]. We set $(t_{min}, t_{max}) = (301, T - 600)$ based on shape of the impulse response and the expected dynamic range of the range profile ($T = 1500$) and the temporal resolution of the grid is set to the resolution of the single-photon detection (i.e., 2ps). The SEM algorithm has been applied with $N_{\text{iter}} = 100$ iterations while the Gibbs sampler used to estimate the range profile was stopped after 600 iterations (including 100 burn-in iterations). The structure of the SEM algorithm is similar to the MCMC algorithm proposed in [1] and further extended in [10, 16] (similar prior models). However the number of iterations and the computational cost of the SEM is significantly reduced because the MCMC algorithm requires intensive computations to update $\mathbf{R}$ and $\mathbf{B}$ while the SEM uses high dimensional convex optimisation tools to update $\mathbf{W}$. Due to its high computation complexity (2 hours to more than a day depending on the number of detected photons, whereas the proposed SEM+Gibbs sampler algorithm takes about 15 minutes on a Macbook pro, 16 GB of RAM, running Matlab 2017b)

ISBN 978-90-827970-1-5

and space constraints, we did not include results obtained the MCMC algorithm developed in [1]. Note however, that similar estimation performances have been observed in practice (but at a higher computational cost).

We quantitatively assess the ranging performance of the proposed method using the depth absolute error (AE) defined by $AE_{i,j} = |t_{i,j} - \hat{t}_{i,j}|$, where $t_{i,j}$ (resp. $\hat{t}_{i,j}$) is the actual (resp. estimated) range associated with the target in the pixel (i, j). The mean AEs (MAEs) in bins obtained with the different values of (α, β) are depicted in Fig. 2. As expected, the MAEs increase with β, i.e., when the background level increases and the MAEs decrease when α increases. Still, even with the highest value of β considered, the MAEs remain below 3 for $\alpha \geq 30$ (approximately 8 signal photons per pixel), which corresponds to an average error of 0.9mm for the range parameters.

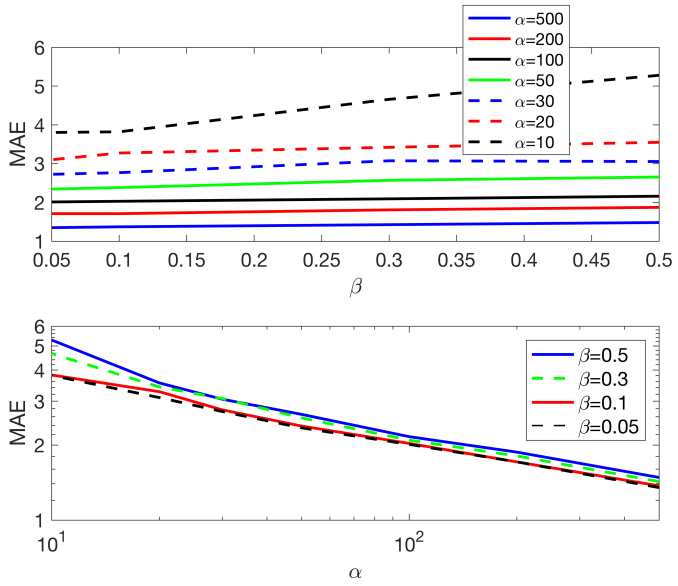

Fig. 2. Depth MAE (in bins) as a function of β (top) and α (bottom).

5. CONCLUSION

We presented a novel a novel algorithm for range profile estimation from single-photon Lidar waveforms in the photon-starved regime. To address the problem of high background levels generally degrading the performance of ranging algorithms, we reformulated the classical observation model based on Poisson statistics as a mixture model whose parameters are assigned prior distributions, enabling the use of EM-like algorithms for robust and faster estimation using Bayesian inference. The results demonstrate good estimation performance with high background levels (66% of the detected photons) and extremely uncertain data (PSNR of 0.04), the main advantage of the method being to present a significantly reduced computational cost compared to full Monte Carlo simulation methods, while preserving the possibility to compute measures of uncertainty about the estimated range profile. Future work include a more comprehensive performance analysis and comparisons with recent alternative methods (e.g., [2, 3]). The estimation of the hyper-parameters controlling the impact of the prior distributions and generalisation to the analysis of multispectral Lidar are currently under investigation.

6. REFERENCES

[1] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform based analysis of depth images constructed using sparse single-photon data," *IEEE Trans. Image Processing*, vol. 25, no. 5, pp. 1935–1946, May 2016.

[2] J. Rapp and V. K. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," *IEEE Trans. Comput. Imaging*, vol. 3, no. 3, pp. 445–459, Sept. 2017.

[3] R. Tobin, Y. Altmann, X. Ren, A. McCarthy, R. A. Lamb, S. McLaughlin, and G. S. Buller, "Comparative study of sampling strategies for sparse photon multispectral lidar imaging: towards mosaic filter arrays," *Journal of Optics*, vol. 19, no. 9, pp. 094006, 2017.

[4] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3-d and reflectivity imaging with single-photon detectors," *IEEE Transactions on Computational Imaging*, vol. 1, no. 2, pp. 112–125, June 2015.

[5] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society, series B*, vol. 39, no. 1, pp. 1–38, 1977.

[6] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, "Bayesian analysis of lidar signals with multiple returns," *IEEE Trans. Patt. Anal. Mach. Intell.*, vol. 29, no. 12, pp. 2170–2180, Dec 2007.

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

[8] Z. T. Harmany, R. F. Marcia, and R. M. Willett, "This is SPIRAL-TAP: Sparse Poisson intensity reconstruction algorithms - theory and practice," *IEEE Transactions on Image Processing*, vol. 21, no. 3, pp. 1084–1096, March 2012.

[9] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Robust Bayesian target detection algorithm for depth imaging from sparse single-photon data," *IEEE Trans. Comput. Imaging*, vol. 2, no. 4, Dec. 2016.

[10] Y. Altmann, A. Maccarone, A. McCarthy, S. McLaughlin, and G. S. Buller, "Spectral classification of sparse photon depth images," *Opt. Express*, vol. 26, no. 5, pp. 5514–5530, March 2018.

[11] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Phys. D*, vol. 60, no. 1-4, pp. 259–268, Nov. 1992.

[12] A. Chambolle, "An algorithm for total variation minimization and applications," *J. of Mathematical Imaging and Vision*, vol. 20, no. 1-2, pp. 89–97, 2004.

[13] G. Celeux, F. Forbes, and N. Peyrard, "EM procedures using mean field-like approximations for Markov model-based image segmentation," *Pattern Recognition*, vol. 36, no. 1, pp. 131 –144, 2003.

[14] F. Forbes and G. Fort, "Combining Monte Carlo and mean-field-like methods for inference in hidden Markov random fields," *IEEE Trans. Image Processing*, vol. 16, no. 3, pp. 2385–2397, March 2007.

[15] G. Celeux and J. Diebolt, "The SEM algorithm: A probabilistic teacher algorithm derived from the EM algorithm for the mixture problem," *Computational Statistics Quarterly*, vol. 2, pp. 73–82, 1985.

[16] Y. Altmann, A. Maccarone, A. McCarthy, G. Newstadt, G. S. Buller, S. McLaughlin, and A. Hero, "Robust spectral unmixing of sparse multispectral lidar waveforms using gamma Markov random fields," *IEEE Trans. Comput. Imaging*, vol. 3, no. 4, pp. 658–670, Dec. 2017.

ISBN 978-90-827970-1-5 © EURASIP 2018