Bayesian Restoration of Reflectivity and Range Profiles from Subsampled Single-Photon Multispectral Lidar Data

Y. Altmann, R. Tobin, A. Maccarone, X. Ren, A. McCarthy, G. S. Buller and S. McLaughlin
School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, U.K.
Email: {Y.Altmann, rt2, A.Maccarone, Ximing.Ren, A.McCarthy, G.S.Buller, S.McLaughlin}@hw.ac.uk

***Abstract*—In this paper, we investigate the recovery of range and spectral profiles associated with remote three-dimensional scenes sensed via single-photon multispectral Lidar (MSL). We consider different spatial/spectral sampling strategies and compare their performance for similar overall numbers of detected photons. For a regular spatial grid, the first strategy consists of sampling all the spatial locations of the grid for each of the wavelengths. Conversely, the three other strategies consist, for each spatial location, of acquiring a reduced number of wavelengths, chosen randomly or in a deterministic manner. We propose a fully automated computational method, adapted for the different sampling strategies in order to recover the target range profile, as well as the reflectivity profiles associated with the different wavelengths. The performance of the four sampling strategies is illustrated using a single photon MSL system with four wavelengths. The results presented demonstrate that although the first strategy usually provides more accurate results, the subsampling strategies do not exhibit a significant performance degradation, particularly for extremely photon-starved data (down to one photon per pixel on average).**

I. Introduction

In recent years, single-photon timing has emerged as a candidate technology for high-resolution three-dimensional profiling [1], and the performance of the approach has been demonstrated in a number of field trials [2]–[4]. Time-correlated single-photon counting (TCSPC) is a statistical sampling technique which records the arrival time of detected photons with respect to the emitted laser pulse or absolute time. In addition to the high timing resolution of such systems (some picoseconds), the high sensitivity of single-photon detectors can mean detection over longer ranges and/or the use of lower power laser sources. However, a potential drawback of single photon counting is that the integration times required for accurate depth measurement can be too long for rapid depth profiling. Even with single-detector scanning systems, significant reductions in acquisition time have been demonstrated by application of advanced computational imaging approaches, such as first-photon imaging [5] or single-photon data analysis in the extremely photon-starved regime [6], [7], which have allowed depth images to be reconstructed with very few photon returns.

Single-photon approaches have been used for demonstrations of multispectral depth imaging for target identification [8], [9], quantification [10] and for the measurement of the physiological parameters of foliage [11]. In most cases, range and spectral information was assumed to be available for each wavelength. Recently, the implementation of mosaic filters, or multispectral filter arrays, to produce simultaneous multispectral images on a single frame read-out has been pursued by several groups [12], [13]. In such cases, a single wavelength (out of the L wavelengths) is acquired at each spatial location and an interpolation scheme is thus required to recover the L full intensity profiles. Moreover, subsampling strategies are clearly interesting for scanning systems where the acquisition time is proportional to the number of pixels scanned.

In this work, we propose a Bayesian computational method to compare the performance of sampling strategies based on subsampled MSL data to the classical full scan approach. We also consider the case where the sampling patterns are independent, i.e. where multiple wavelengths can be acquired at some spatial locations while other locations are not sensed. Adopting a Bayesian approach, prior distributions are assigned to the unknown parameters involved in the observation model. In particular, Markov random fields (MRFs) are used to capture the intrinsic spatial correlation affecting the range and intensity profiles of natural scenes. Markov chain Monte Carlo (MCMC) methods are then used to sample posterior distributions of interest and estimate the unknown depth and reflectivity profiles. Moreover, the resulting estimation strategy is able to automatically adjust the spatial regularisation parameters which are usually difficult to adjust without ground truth depth and intensity profiles, in particular when considering extremely sparse and uncertain measurements.

The remainder in this paper is organised as follows. Section II introduces the observation model associated with MSL returns for single-layered objects. Section III presents the Bayesian model associated with the range and reflectivity estimation problem considered and Section IV describes the MCMC methods used to sample from the posterior distributions of interest and subsequently approximate appropriate Bayesian estimators. Results of experiments conducted on real MSL data are discussed in Section V and conclusions are reported in Section VI.

II. Observation model

We consider a 4-D array $\mathbf{Y}$ of Lidar waveforms of dimension $N_{\text{row}} \times N_{\text{col}} \times L \times T$, where N_{row} and N_{col} stands

ISBN 978-0-9928626-7-1 © EURASIP 2017

for the number of rows and columns of the regular spatial sampling grid (in the transverse plane), L is the total number of spectral bands or wavelengths and T is the number of temporal (corresponding to range) bins. Let $\mathbf{y}_{i,j,\ell} = [\mathbf{Y}]_{i,j,\ell,:} = [y_{i,j,\ell,1}, \ldots, y_{i,j,\ell,T}]^T$ be the Lidar waveform obtained in the pixel (i,j) using the ℓth wavelength. The element $y_{i,j,\ell,t}$ is the photon count within the tth bin of the ℓth spectral band (assuming that a waveform is actually observed for this wavelength and spatial location). For each pixel, the set of observed wavelengths in denoted $\mathcal{L}_{i,j}$. Pixels/wavelengths that are not observed satisfy $y_{i,j,\ell,t} = 0, \forall t$. In this work, we assume that for each pixel, the detected photons result from two main contributions: 1) from direct path reflections of the photons originally emitted by the laser sources onto the surface of the object of interest or 2) dark photon counts and ambient illumination (assumed to be stationary in time but potentially non-stationary spatially). Moreover, we assume that the laser beam, for each pixel, encounters a single surface which is assumed to be locally orthogonal to the beam direction. This is typically the case for short to mid-range (up to dozens of metres) depth imaging where the divergence of the laser source(s) can be neglected. Let $d_{i,j}$ be the position of an object surface at a given range from the sensor, whose mean spectral signature is denoted as $\boldsymbol{\lambda}_{i,j} = [\lambda_{i,j,1}, \ldots, \lambda_{i,j,L}]^T$. According to [6], [14], each photon count $y_{i,j,\ell,t}$ is assumed to be drawn from the following Poisson distribution

$$y_{i,j,\ell,t} | \lambda_{i,j,\ell}, t_{i,j} \sim \mathcal{P}\left(\lambda_{i,j,\ell} g_\ell(t - t_{i,j}) + b_{i,j,\ell}\right) \tag{1}$$

where $g_\ell(\cdot)$ is the instrumental response of the system, which consists of a peak whose shape differs between wavelength channels (see [9] for examples of impulse responses). 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 instrumental responses $\{g_\ell(\cdot)\}$ are assumed to be known, as occurs when they can be accurately estimated during imaging system calibration. $\{b_{i,j,\ell}\}$ represent background illumination levels. It is important to stress that in this work, we consider applications where the observed objects consist of a single visible surface per pixel. We do not consider cases where the photons can penetrate through objects or be reflected from multiple surfaces. This assumption allows the estimation of the target spectral responses to be reduced to a two spatial dimensions problem. Thus, the problem addressed here is to estimate the range of the targets (for all the image pixels) and estimate the target spectral responses. For brevity, we assume that the background levels are known (and negligible in Section V). In practice, these parameters can often be estimated from calibration measurements using a pre-processing step similar to the reflectivity estimation strategy as discussed in Section IV-A.

III. Bayesian model

A. Likelihood

Using the classical assumption that the detected photon counts, conditioned on their mean in all channels, are conditionally independent, the joint likelihood of $\mathbf{Y}$ can be expressed as

$$f(\mathbf{Y}|\boldsymbol{\Lambda}, \mathbf{B}, \mathbf{T}) = \prod_{i,j} \prod_{t, \ell \in \mathcal{L}_{i,j}} f_{\mathcal{P}}(y_{i,j,\ell,t}; \lambda_{i,j,\ell} g_{i,j,\ell}(t - t_{i,j}) + b_{i,j,\ell}), \tag{2}$$

where $\boldsymbol{\Lambda} = \{\boldsymbol{\lambda}_{i,j}\}_{i,j}$, $\mathbf{B} = \{\mathbf{b}_{i,j}\}_{i,j}$ and $\mathbf{T}$ is a matrix gathering the target ranges. Note that in Eq. (2), $f_{\mathcal{P}}(\cdot;\lambda)$ denotes the probability mass function of the Poisson distribution with mean λ.

B. Prior distributions

Each target position is 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 $(t_{min}, t_{max}) = (301, T - 300)$ (for reasons discussed later) and the temporal resolution of the grid is set to the resolution of the single-photon detection (i.e., 2 picoseconds (ps) in Section V). As in [6], to account for the spatial correlations between the neighbouring pixels, we propose to use a 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 $f(\mathbf{T}|\epsilon) = G(\epsilon)^{-1} \exp\left[-\epsilon \phi(\mathbf{T})\right]$, where $\epsilon \geq 0$ is a parameter tuning the amount of correlation between ranges in adjacent pixels, $G(\epsilon)$ is a normalisation constant and where $\phi(\mathbf{T}) = \sum_{i,j} \sum_{(i',j') \in \mathcal{V}(i,j)} |t_{i,j} - t_{i',j'}|$. In a convex optimisation framemork, this MRF corresponds to a total-variation (TV) regularisation [15], [16] promoting piecewise constant depth image. Moreover, the higher the value of ϵ, the more correlated the ranges of neighbouring pixels. Here, a four-pixel structure is used to define the neighbourhood structure of the all the MRFs considered.

In a similar fashion to the depth parameters, we use TV-based priors for the target spectral signatures in $\boldsymbol{\Lambda}$. Precisely, we assume that the elements of $\boldsymbol{\Lambda}$ take value in arbitrarily defined finite sets of discrete values (typically $(0,1)$ for the target reflectivity) and define the following prior models

$$f(\boldsymbol{\Lambda}_\ell | \gamma_\ell) = G(\gamma_\ell)^{-1} \exp\left[-\gamma_\ell \phi(\boldsymbol{\Lambda}_\ell)\right], \tag{3}$$

where $\boldsymbol{\Lambda}_\ell$ stands for the target reflectivities associated with the ℓth spectral channel. The prior models in (3) rely on L regularisation parameters $\{\gamma_\ell\}$ which control the amount of spatial correlation between reflectivities of adjacent pixels. This strategy allows us to account for spatial correlations between reflectivity parameters in order to improve estimation performance. Moreover, as will be discussed in the next paragraph, the proposed algorithm is able to automatically adjust $\{\gamma_\ell\}$ and thus estimate the amount a spatial correlation between pixels.

ISBN 978-0-9928626-7-1 © EURASIP 2017

IV. Estimation strategy

Simultaneously estimating $\mathbf{\Lambda}$ and $\mathbf{T}$ from the observed waveforms $\mathbf{Y}$ is challenging mainly due to the multimodal nature of the likelihood $f(\mathbf{Y}|\mathbf{\Lambda},\mathbf{B},\mathbf{T})$. In a similar fashion to [5], [10], [17], we simplify the problem by estimating sequentially $\mathbf{\Lambda}$ and $\mathbf{T}$, using weak assumptions which can often be satisfied in practice. The two estimation steps are detailed in what follows.

A. Target reflectivity estimation

The first step of the proposed method consists of estimating the target reflectivity within each pixel and for each spectral band, given known or previously estimated background levels. This step is generally difficult because the target ranges are unknown. Nevertheless, if we consider $\tilde{y}_{i,j,\ell} = \sum_{t=1}^{T} y_{i,j,\ell,t}$, i.e., the sum of the photons detected (within a single spectral band) in a given observed pixel, Eq. (2) leads to

$$\tilde{y}_{i,j,\ell}|\lambda_{i,j,\ell}, t_{i,j}, b_{i,j,\ell} \sim \mathcal{P}\left(\lambda_{i,j,\ell}\tilde{g}_{i,j,\ell} + Tb_{i,j,\ell}\right), \tag{4}$$

where $\tilde{g}_{i,j,\ell} = \sum_{t=1}^{T} g_\ell(t - t_{i,j})$ is the (discretised) integral of the instrumental response associated with an object with unitary reflectivity and located at $t_{i,j}$ from the detector. In this work, we assume that $\tilde{g}_{i,j,\ell} = \tilde{g}_\ell, \forall t_{i,j} \in \mathbb{T}$, i.e., that the integral of the instrumental response does not depend on the distance of the target. This is typically the case in practice when the admissible set of target ranges are far enough from the extreme bins of the recorded histograms, i.e. when the peaks associated with target returns are not truncated. For instance in Section V, the support of the impulse responses are smaller than 300 bins, what motivated the range constraints in Section III-B. In such cases, Eq. (4) yields

$$\tilde{y}_{i,j,\ell}|\lambda_{i,j,\ell}, \tilde{g}_\ell, b_{i,j,\ell} \sim \mathcal{P}\left(\lambda_{i,j,\ell}\tilde{g}_\ell + Tb_{i,j,\ell}\right), \tag{5}$$

which does not depend on the distance of the target. We can thus define the following posterior distribution for $\mathbf{\Lambda}_\ell$

$$f(\mathbf{\Lambda}_\ell|\widetilde{\mathbf{Y}},\mathbf{B},\gamma_\ell) \propto f(\mathbf{\Lambda}_\ell|\gamma_\ell)\prod_{i,j}\prod_{\ell\in\mathcal{L}_{i,j}} f(\tilde{y}_{i,j,\ell}|\lambda_{i,j,\ell},\tilde{g}_\ell,b_{i,j,\ell}), \tag{6}$$

which can be exploited to estimate $\mathbf{\Lambda}_\ell$, with $\widetilde{\mathbf{Y}} = \{\tilde{y}_{i,j,\ell}\}$. Here we resort to an adaptive Markov chain Monte Carlo (MCMC) similar to that proposed in [6], [18] to compute the marginal maximum *a posteriori* (MMAP) estimators of the target intensities given by

$$\hat{\lambda}_{i,j,\ell} = \max_{\lambda_{i,j,\ell}} f(\lambda_{i,j,\ell}|\widetilde{\mathbf{Y}},\mathbf{B},\hat{\gamma}_\ell), \quad \forall(i,j,\ell) \tag{7}$$

where $\hat{\gamma}_\ell$ approximates the marginal maximum likelihood estimator of γ_ℓ. Precisely, for each wavelength, a Gibbs sampler is used to sample according to $f(\mathbf{\Lambda}_\ell|\widetilde{\mathbf{Y}},\mathbf{B},\gamma_\ell)$ and the regularization parameter γ_ℓ is updated during the early iterations (burn-in) period of the sampler (the interested reader is invited to consult [6], [18] for further details about the sampling strategy adopted here). Here we resort to a simulation-based algorithm here in order to automatically adjust the regularisation parameters γ_ℓ. However, it is important to mention that $f(\mathbf{\Lambda}_\ell|\widetilde{\mathbf{Y}},\mathbf{B},\gamma_\ell)$ is log-concave and that $\mathbf{\Lambda}$ could also be inferred via standard maximum *a posteriori* (MAP) estimation using state-of-the-art convex optimisation techniques (provided that $\boldsymbol{\gamma} = [\gamma_1,\ldots,\gamma_L]$ is properly tuned). A similar estimation procedure can be used to estimate the L full background profiles from (potentially sparsely sampled) MSL waveforms that do not contain photons originally emitted by the laser. Such waveforms can be obtained by switching off the laser sources or by considering parts of the longer waveforms which do not contain target returns. Using TV-based priors for the background levels, it is possible to estimate $\mathbf{B}$ via MMAP estimation while adjusting automatically the spatial regularisation parameters. Due to space constraints, this step is however not further detailed in this paper.

B. Target range estimation

The second step of the proposed method consists of recovering the target ranges, given that the background levels and target reflectivities are known (or have been previously estimated). In contrast to the distributions in 6 considered to estimate the target intensities, the conditional distribution $f(\mathbf{T}|\mathbf{Y},\mathbf{\Lambda},\mathbf{B},\epsilon)$ is generally highly multimodal and thus cannot be maximised (w.r.t. $\mathbf{T}$) efficiently using convex optimisation tools. Again, we use an efficient MCMC method similar to that used in [6], [18] to estimate simultaneously $\mathbf{T}$ (via MMAP estimation) and ϵ.

The resulting algorithm, which estimate sequentially $\mathbf{\Lambda}$ and $\mathbf{T}$, together with the associated regularisation parameters is thus fully automated and does not require practitioners to adjusted critical parameters. In addition to its computational efficiency (highly parallelizable) and robustness (w.r.t. convergence issues), the proposed method can also be used to provide *a posteriori* measures of uncertainty associated with each estimation step. The interested reader is invited to consult [10] for examples of use of such measures for ranging assessment.

V. Results

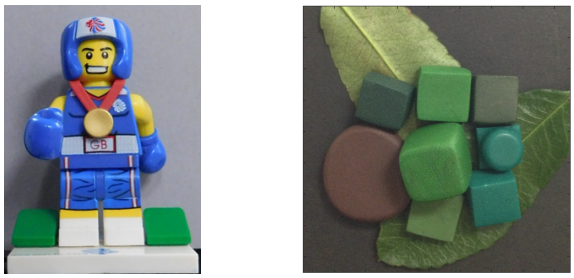

(a) (b)

Fig. 1. RGB photographs of the front view of the first (a) and second (b) target (approximately 45×45 mm) which were located at a stand-off distance of 1.8 m from the transceiver unit. The depth variations are of approximately 30 mm (Target 1) and 20 mm (Target 2).

In this section, we evaluate the performance of the algorithm proposed in Section IV for ranging and reflectivity estimation from multispectral Lidar data using data acquired by an imaging system with $L = 4$ wavelengths, similar to that detailed in [9]. We propose comparing the performance

ISBN 978-0-9928626-7-1 © EURASIP 2017

of 4 different sampling strategies to analyze the depth and reflectivity profiles of two scenes (Fig. 1). The first target is a lego figurine place in front of a dark backboard (Fig. 1 (a)) while the second target (Fig. 1 (b)) is composed of different objects made of Fimo clay fixed on tree leaves and mounted onto a dark backboard. The two targets were placed at a distance of 1.8 m from the imaging system. The measurements were performed indoor to control ambient illumination. To investigate the effect of the sampling strategy on the recovered depths and reflectivities, we consider measurements performed under dark conditions, with a negligible contribution from background ambient illuminations. The scenes were scanned using a regular spatial grid of 200×200 pixels using two different sets of wavelengths. For the first target, we used $473, 532, 589$ and 640nm, which corresponds to red, green, yellow and blue colours. For the second target, we selected $L = 4$ regularly spaced wavelengths ($500, 580, 660$ and 740nm) out of the 33 wavelengths considered in [9]. The histograms consist of $T = 3000$ bins of 2 ps, which represents a depth resolution of 300μm per bin. The instrumental impulse responses $g_\ell(\cdot)$ were estimated from preliminary experiments by analyzing the distribution of photons reflected onto a commercially available Lambertian scatterer. The ground truth depths and reflectivity parameters have been obtained from complete 200×200 pixels scans (also acquired with negligible ambient illumination) for each wavelength and long per-pixel acquisition times to reduce uncertainty cause by the Poisson noise statistics. We consider the following sampling strategies

1) Scenario #1: full scans
2) Scenario #2: regular subsampling, without overlap
3) Scenario #3: random subsampling, without overlap
4) Scenario #4: random subsampling, with overlap.

For Scenarios 2 to 4, 25% of the pixels are sensed for each wavelength. This ratio corresponds to a gain of 75% in terms of acquisition time if a scanning system processing the four wavelengths sequentially was used. Scenario 2 is obtained by defining a four-colour checker-board of size 200×200. For the Scenario 3, for each pixel, the observed wavelength is randomly chosen by drawing uniformly in $\{1, 2, 3, 4\}$, while $L = 4$ independent random masks are created for the last scenario. In that case, some pixels are observed through several wavelengths while others are not sensed. For each target, the power of the laser source and the acquisition times were adjusted from preliminary runs so that similar photon counts are returned by the target for the four wavelengths and so that the four scenarios lead similar per-pixel photons counts on average. For instance, for 10 detected photons per pixel on average, Scenario #1 corresponds to 4×2.5 photons on average, while Scenarios #2 and #3 correspond to 1×10 photons on average.

The proposed algorithm has been applied with $N_{\text{MC}} = 5000$ sampler iterations for each step, including $N_{\text{bi}} = 2000$ burn-in iterations, which corresponds to a processing time of about 30 minutes for each scene, for a Matlab R2014a implementation on a i7-3.0 GHz desktop computer (16GB RAM). The computational cost of the four scenarios are similar as the likelihood contributions to be computed for each observed data can be computed only once at the first iteration for each admissible intensity and range (constrained to be a fixed grids).

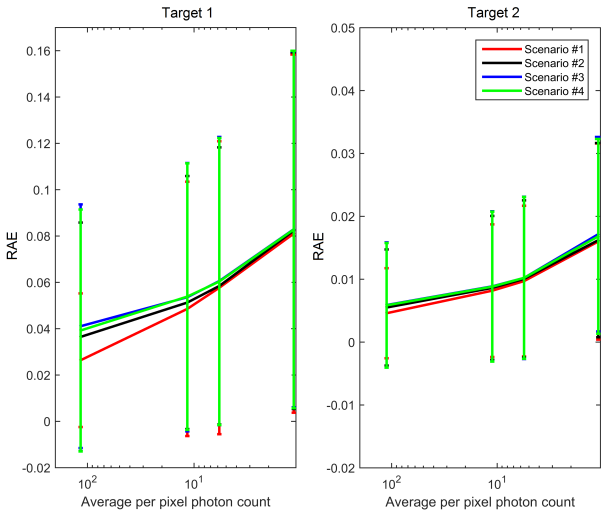

Fig. 2. Mean RAEs obtained for Target 1 (left) and Target 2 (right) using the different scenarios as a function of the average per-pixel photon counts. The error bars represent the $\pm$ one standard deviation intervals.

The performance of the four imaging strategies in terms of reflectivity estimation are compared using the reflectivity absolute error (RAE) defined by $RAE_{i,j} = \sum_{\ell=1}^{L} |\lambda_{i,j,\ell} - \hat{\lambda}_{i,j,\ell}|$, where $\lambda_{i,j,\ell}$ (resp. $\hat{\lambda}_{i,j,\ell}$) is the actual (resp. estimated) target reflectivity in the ℓth band of the pixel (i, j). Fig. 2 compares the mean RAEs ($\pm$ standard deviation) obtained with the four methods as the average number of detected photons decreases. The performance of all methods degrades (higher mean RAEs and larger variances) as the data quality degrades. Yet, with almost 1 photon per pixel, the mean RAE of all method remains below 0.1 for the first target and 0.02 for the second target. Note that smaller RAEs are generally obtained for the second target which consists of large homogeneous reflectivity regions for which the TV-based regularisation is particularly well adapted. Moreover, for higher detected photon counts, the full scan approach (red curves in Fig. 2) outperforms the other strategies which suffers from the reduced number of observed pixels (25%) for each wavelength, in particular for the first target which includes finer spectral features.

Similarly, the ranging performance of the sampling strategies is measured using the depth AE (DAE) defined by $DAE_{i,j} = |t_{i,j} - \hat{t}_{i,j}|$, 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 DAEs obtained using the four scenarios are depicted in Fig. 3. With more than 10^2 photons per pixel on average, the mean DAEs obtained are smaller than 500μm for all methods and for the two targets. In contrast to the estimated intensities, the range profiles obtained by the the four methods with larger photon counts are relatively similar. Indeed, in such cases the overall number of detected photons in each pixel is large enough to provide sufficiently accurate range estimates,

ISBN 978-0-9928626-7-1 © EURASIP 2017

even if a single band is observed. The full scan approach is generally more robust since all the wavelengths are observed in each pixel. Using a reduced number of wavelengths per pixel generally increases the probability of weak target returns while the full scan approach only requires at least one of the L reflectivity parameters to be high in order to obtain an accurate depth estimate. Consequently, spreading the energy in each pixel across multiple wavelengths generally improves the ranging performance. Note that the DAEs associated with the second target are generally smaller than those obtained with the first target due to 1) smaller local depth variations (range profile of Target 2 locally constant) and 2) the peak of the instrumental responses $g_\ell(\cdot)$ which broadens as the wavelength decreases, leading to less accurate depth estimates (see [9] for examples of impulse responses associated with the imaging system considered).

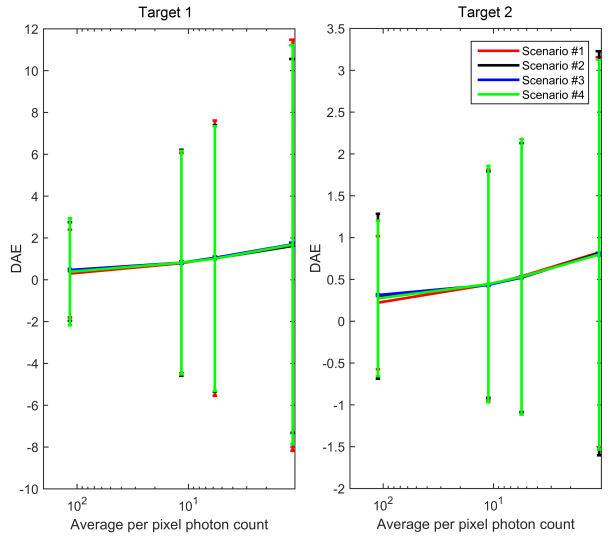

Fig. 3. Mean DAEs obtained for Target 1 (left) and Target 2 (right) using the different scenarios as a function of the average per-pixel photon counts. The error bars represent the $\pm$ one standard deviation intervals.

VI. Conclusion

In this paper, we compared the ranging and reflectivity estimation performance of several sampling strategies for 3D scenes analysis using MSL waveforms. The results obtained show that the lack of spectral information induced by the acquisition of a reduced number of wavelengths per pixel can be mitigated to some extent by incorporating spatial correlation in particular when the overall number of detected photons reduces. For higher photon counts, the full scan approach generally performs better than the other approaches using the TV-based spatial regularisations used in this work. However, this observation should be mitigated by the fact the performance of sparse sampling approaches could be further improved using other spatial regularisations. Moreover, the performance degradation induced by such approaches reduces in the presence of ambient illumination sources (results not presented here due to space constraints), which introduces more uncertainty than that the lack of observed pixels.

References

[1] G. S. Buller and A. M. Wallace, "Ranging and three-dimensional imaging using time-correlated single-photon counting and point-by-point acquisition," *IEEE J. Sel. Topics in Quantum Electronics*, vol. 13, no. 4, pp. 1006–1015, July 2007.

[2] M. A. Albota, R. M. Heinrichs, D. G. Kocher, D. G. Fouche, B. E. Player, M. E. O'Brien, B. F. Aull, J. J. Zayhowski, J. Mooney, B. C. Willard, and R. R. Carlson, "Three-dimensional imaging laser radar with a photon-counting avalanche photodiode array and microchip laser," *Appl. Opt.*, vol. 41, no. 36, pp. 7671–7678, Dec 2002.

[3] J. J. Degnan, "Photon-counting multikilohertz microlaser altimeters for airborne and spaceborne topographic measurements," *Journal of Geodynamics*, vol. 34, no. 34, pp. 503 – 549, 2002.

[4] A. McCarthy, X. Ren, A. Della Frera, N. R. Gemmell, N. J. Krichel, C. Scarcella, A. Ruggeri, A. Tosi, and G. S. Buller, "Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector," *Opt. Express*, vol. 21, no. 19, pp. 22098–22113, Sept. 2013.

[5] A. Kirmani, D. Venkatraman, D. Shin, A. Colao, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.

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

[7] A. Halimi, A. Maccarone, A. McCarthy, S. McLaughlin, and G. S. Buller, "Object depth profile and reflectivity restoration from sparse single-photon data acquired in underwater environments," *IEEE Trans. Comput. Imaging*, 2017, to appear.

[8] G. S. Buller, R. D. Harkins, A. McCarthy, P. A. Hiskett, G. R. MacKinnon, G. R. Smith, R. Sung, A. M. Wallace, R. A. Lamb, K. A. Ridley, and J. G. Rarity, "Multiple wavelength time-of-flight sensor based on time-correlated single-photon counting," *Review of Scientific Instruments*, vol. 76, no. 083112, 2005.

[9] Y. Altmann, A. Maccarone, A. McCarthy, G. S. Buller, and S. Mclaughlin, "Joint spectral clustering and range estimation for 3d scene reconstruction using multispectral lidar waveforms," in *Proc. European Signal Processing Conf. (EUSIPCO)*, Budapest, Hungary, Sept. 2016.

[10] 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," Tech. Rep., Heriot-Watt University, Oct. 2016, Available online, http://yoannaltmann.weebly.com/publications.html.

[11] A. M. Wallace, A. McCarthy, C. J. Nichol, Ximing Ren, S. Morak, D. Martinez-Ramirez, I. H. Woodhouse, and G. S. Buller, "Design and evaluation of multispectral lidar for the recovery of arboreal parameters," *IEEE Trans. Geoscience and Remote Sensing*, vol. 52, no. 8, pp. 4942–4954, Aug 2014.

[12] P.-J. Lapray, X. Wang, J.-B. Thomas, and P. Gouton, "Multispectral filter arrays: Recent advances and practical implementation," *Sensors*, vol. 14, no. 11, pp. 21626–21659, 2014.

[13] A. V. Kanaev, M. R. Kutteruf, M. K. Yetzbacher, M. J. Deprenger, and K. M. Novak, "Imaging with multi-spectral mosaic-array cameras," *Appl. Opt.*, vol. 54, no. 31, pp. F149–F157, Nov. 2015.

[14] 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.

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

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

[17] 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.

[18] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tourneret, "Maximum marginal likelihood estimation of the granularity coefficient of a Potts-Markov random field within an MCMC algorithm," in *Proc. IEEE-SP Workshop Stat. and Signal Processing*, Gold Coast, Australia, July 2014.

ISBN 978-0-9928626-7-1 © EURASIP 2017