

Stochastic EM algorithm for fast analysis of single waveform multi-spectral Lidar data

Q. Legros*, S. McLaughlin, Y. Altmann
School of Engineering and Physical Sciences
Heriot-Watt University
Edinburgh, United Kingdom
qll1@hw.ac.uk

S. Meignen†
Laboratoire Jean Kuntzmann
University Grenoble Alpes
Grenoble, France

Abstract—This paper addresses the problem of estimating spectral and range profiles from single-photon Lidar waveforms associated with single surfaces in presence of an unknown background. A single Lidar waveform per pixel is considered, whereby a single detector is used to acquire information simultaneously at multiple wavelengths. A novel Bayesian approach is developed to perform the estimation of model parameters in a reduced computational time. This is achieved by transforming an EM-based algorithm recently proposed into a stochastic EM algorithm, which is computationally more attractive. The reconstruction performance and computational complexity of our approach are assessed through a series of experiments using synthetic data under different observation scenarios. The obtained results demonstrate a significant speed-up compared to the state-of-the-art method, without significant degradation of the estimation quality.

Index Terms—Multispectral imaging, 3D imaging, Single-photon Lidar, Bayesian estimation.

I. INTRODUCTION

Light detection and ranging (Lidar) systems have received considerable interest recently for their ability to reconstruct a 3D scene at high resolution [1], [2]. By sending light pulses onto a target and measuring the time-of-flight (ToF) of the reflected photons, it is possible to estimate the depth profile of an entire scene. More precisely, the ToFs of the recorded photons are gathered to form histograms of photon detection events. Additional detection events caused by ambient illumination and detector dark counts are generally uniformly distributed over the histogram. Single-band Lidar (SBL) allows the reconstruction of high-resolution 3D scenes, but it only captures one intensity value per pixel, associated with the laser source wavelength. Multispectral Lidar (MSL) circumvents such an issue, as a single imaging system [3] is used to collect simultaneously spectral and spatial features [4]. However, the acquisition of such signals requires an overall longer acquisition time (sequential acquisition) or more expensive imaging systems (parallel acquisition) [5], [6].

A recent single-waveform MSL (SW-MSL) approach has been developed in [7], to reconstruct the depth and spectral profiles while keeping the acquisition time and complexity

of the imaging system of the same order as SBL. It consists of registering a single histogram containing spectral information from multiple wavelengths for each spatial location, via wavelength-time coding (see Fig 1). Although the temporal delays between wavelengths are known, the overall waveform shape varies as the amplitude of its different peaks depends on the spectral signature of the surface observed. In [7], a full Bayesian method was proposed to analyze SW-MSL signals, which suffers from a prohibitive computational time (about 15 hours to reconstruct a single 3D scene of 200×200 pixels and 4 wavelengths). In [8], we recently proposed a faster alternative where the model parameters are estimated sequentially. More precisely, a mixture model was used as it naturally fits Expectation-Maximization (EM) based algorithms. The resulting algorithm first estimates the mixture weights and subsequently the depth profile. The spectral signatures are computed from the mixture weights using a post-processing step.

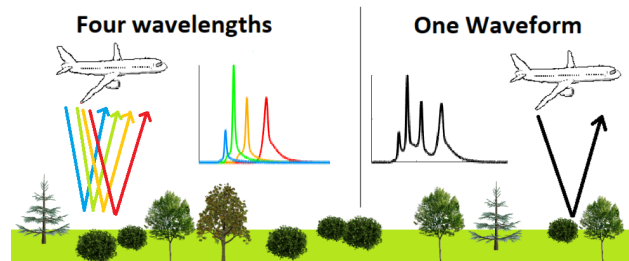


Fig. 1. Example of acquisition of four MSL waveforms, using one wavelength per waveform, and SW-MSL capturing simultaneously information from four wavelengths.

In this paper, we extend the work presented in [8] in two aspects. First, we use the traditional Poisson observation noise model and aim at directly estimating the spectral response of each surface, without resorting to the mixtures of distributions used in [8]. The reflectivity profiles are directly assigned prior models, without using additional transforms. Second, we propose a more computationally attractive algorithm, namely a stochastic EM algorithm [9], [10], to estimate the model parameters from single-photon Lidar waveforms. We also provide an additional solution to further accelerate the estimation

*This work was supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31 and by the Engineering and Physical Sciences Research Council (EPSRC) (grant EP/T00097X/1)

†This work was also supported by the IDEX of University Grenoble Alpes.

process, by sub-sampling the grid of admissible depth position. We show that this new method yields a significant speed-up without degradation of the estimation performance when compared to our previous methods.

The main contributions of the paper are:

- A new Bayesian model based on SW-MSL observation model assuming Poisson noise.
- A generalized stochastic EM algorithm allowing fast analysis of SW-MSL, extending the work done in [8].
- An efficient method for accelerating the sequential inference process, relying on approximating (via subsampling) the likelihood function.

The remainder of this paper is organized as follows. Section II introduces the problem addressed in this work. In Section III the new SW-MSL observation model and the different prior distributions associated with the unknown model parameters are defined. The estimation strategy and computational improvement proposal are discussed in Section IV, and the performance of the proposed method is then assessed in Section V on synthetic SW-MSL data. Conclusions and future work are finally reported in Section VI.

II. PROBLEM FORMULATION

Consider a 2D observation array \mathbf{Y} of SW-MSL signals $\mathbf{y}_n = [\mathbf{Y}]_{n,:} = [y_{n,1}, \dots, y_{n,T}]^\top$ of size $N \times T$, where N is the number of pixels or spatial locations and T is the length of the ToF histograms. The element $y_{n,t}$ is the photon count in the n th pixel in the t th temporal bin. The spectral response of the object in the n th pixel is denoted by $\mathbf{r}_n = [r_{n,1}, \dots, r_{n,L}]^\top \in \mathbb{R}_+^L$, where L stands for the number of spectral bands, and $b_n \geq 0$ gathers all the background contributions in that pixel. In the low-flux regime, and in the presence of a single surface per pixel, $y_{n,t}$ follows the Poisson distribution [7]

$$y_{n,t} | \mathbf{r}_n, b_n, t_n \sim \mathcal{P} \left(b_n + \sum_{\ell=1}^L r_{n,\ell} g_\ell(t - t_n) \right), \quad (1)$$

where t_n is the characteristic ToF of photons emitted by the laser source and reflected by an object at distance d_n of the sensor and $\{g_\ell(\cdot)\}_\ell$ are the instrumental/impulse response functions (IRFs) associated with the L wavelengths, and whose shape can differ between spectral channels. These IRFs are assumed to be known as they are generally estimated during the system calibration. Moreover, the integral $G_\ell = \sum_{t=1}^T g_\ell(t - t_n)$ is constant for any admissible t_n .

The problem addressed in this paper is a computationally attractive estimation of $\mathbf{t} = \{t_n\}_n$ and $(\mathbf{R}, \mathbf{b}) = (\{\mathbf{r}_n\}_n, \{b_n\}_n)$ in presence of possibly high background illumination.

III. BAYESIAN MODEL

A. Likelihood

Assuming that the photon counts in all bins and histograms are mutually independent (conditioned on their mean), the joint likelihood can be expressed as

$$p(\mathbf{Y} | \mathbf{R}, \mathbf{b}, \mathbf{t}) = \prod_n p(\mathbf{y}_n | \mathbf{r}_n, t_n, b_n). \quad (2)$$

Since maximum likelihood estimation (MLE) based on Eq. (2) is generally sensitive to noise, the problem has to be regularized in order to perform robust estimation as discussed below.

B. Prior model

Depth and spectral profiles are often spatially correlated in natural scenes. As in [7], [11]–[13] we model separately the depth and reflectivity profiles and neglect potential correlation between them to keep the inference process tractable.

Depth prior model: The range profile of the target is assumed to belong to a discrete grid of admissible positions $\mathbb{T} = \{t_{min}, \dots, t_{max}\}$ such that $1 < t_{min} < t_{max} < T$ as in [14], [15]. As in [11], [12], we define a prior model for \mathbf{t} which preserves sharp edges, i.e., we use the following total-variation (TV) Markov random field (MRF),

$$p(\mathbf{t} | \epsilon) = \exp[-\epsilon \|\mathbf{t}\|_{TV}], \quad (3)$$

where the fixed (user-defined) hyperparameter ϵ , which mostly depends on the range resolution of the single-photon detector, determines the level of correlation between the depth parameters of neighboring pixels. Here we fixed $\epsilon = 0.05$ for all the experiments performed in this paper.

Reflectivity prior model: The reflectivity profile is also assumed to present spatial correlation as neighboring pixels are likely to belong to the same surface. We propose a hierarchical model based on the gamma distribution as in [16]. In a similar fashion to [8] we assume that the pixels can be clustered into C distinct and known classes and that the pixels within each class share the same prior model. The joint prior model for \mathbf{R} is then given by

$$f(\mathbf{R} | \Theta) = \prod_{c=1}^C \prod_{n \in I_c} \prod_{\ell=1}^L \mathcal{G}(r_{n,\ell}; k_{c,\ell}, \theta_{c,\ell}), \quad (4)$$

where $\mathcal{G}(\cdot; k_{c,\ell}, \theta_{c,\ell})$ is the probability density function of the gamma distribution with respective shape and scale parameters $(k_{c,\ell}, \theta_{c,\ell})$, $\Theta = \{k_{c,\ell}, \theta_{c,\ell}\}_{c,\ell}$ and I_c is the set of pixel indices in the c th group.

Weakly informative independent truncated gamma (resp. inverse-gamma) hyper-prior models are finally assigned to the shape (resp. scale) parameters as follows

$$f(\Theta) = \prod_{c,\ell} \mathcal{G}(k_{c,\ell}; 2, 0.5) \mathbf{1}_{(1,\infty)}(k_{c,\ell}) \mathcal{IG}(\theta_{c,\ell}; 1.01, 0.5). \quad (5)$$

In practice the C classes are often unknown. Here they are set as in [8] by applying patch-based k-means clustering on \mathbf{R} after a few iterations of the algorithm using a single class. Note that the gamma prior distribution was used to ensure the positivity of \mathbf{R} and \mathbf{b} , but other prior distributions could have been used. In all the experiments performed in this work, the segmentation was done on the third iteration of the algorithm, as the current estimates were sufficiently informative to cluster the reflectivity image in our experiments, and patches of size (3×3) are clustered in $C = 7$ classes.

Background prior model: The background image may present spatial structure but can also be significantly different

from \mathbf{R} . In the high-background regime, the likelihood is relatively informative and the background image \mathbf{b} does not need to be strongly regularized. For this reason, we simply assign a separable prior (with respect of \mathbf{b})

$$f(\mathbf{b}|\alpha, \beta) = \prod_n \mathcal{G}(b_n|\alpha, \beta), \quad (6)$$

and (α, β) is assigned the same prior model $f(\alpha, \beta)$ as in (5).

IV. ESTIMATION STRATEGY

Using the likelihood and prior models defined in Section III, we obtain

$$f(\mathbf{R}, \mathbf{b}, \mathbf{t}, \Phi|\mathbf{Y}) \propto p(\mathbf{Y}|\mathbf{R}, \mathbf{b}, \mathbf{t})p(\mathbf{t})f(\mathbf{R}, \mathbf{b}|\Phi)f(\Phi), \quad (7)$$

where $\Phi = \{\Theta, \alpha, \beta\}$, $f(\mathbf{R}, \mathbf{b}|\Phi) = f(\mathbf{R}|\Theta)f(\mathbf{b}|\alpha, \beta)$ and where all the fixed model parameters have been omitted in the conditional distributions to simplify notations.

The likelihood in Eq.(2) is multimodal (because of \mathbf{t}) and the joint estimation of \mathbf{t} and $(\mathbf{R}, \mathbf{b}, \Phi)$ is challenging within a reasonable computational time. Thus, we propose as in [8], [13] to decompose the estimation process into successive problems, by first considering \mathbf{t} as a nuisance parameter which is marginalized to simplify the estimation of $(\mathbf{R}, \mathbf{b}, \Phi)$ such as

$$\begin{aligned} (\hat{\mathbf{R}}, \hat{\mathbf{b}}, \hat{\Phi}) &= \operatorname{argmax}_{\mathbf{R}, \mathbf{b}, \Phi} \sum_{\mathbf{t}} f(\mathbf{R}, \mathbf{b}, \mathbf{t}, \Phi|\mathbf{Y}) \\ &= \operatorname{argmax}_{\mathbf{R}, \mathbf{b}, \Phi} f(\mathbf{R}, \mathbf{b}, \Phi|\mathbf{Y}). \end{aligned} \quad (8)$$

The depth \mathbf{t} is then estimated via marginal maximum a posteriori (MMAP) estimation using the estimates from Eq. (8)

$$\hat{t}_n = \operatorname{argmax}_{t_n} \sum_{\mathbf{t}_{\setminus n}} p(\mathbf{t}|\mathbf{Y}, \hat{\mathbf{R}}, \hat{\mathbf{b}}, \hat{\Phi}), \quad \forall n, \quad (9)$$

where $\tilde{\mathbf{t}}_{\setminus n}$ stands for the vector $\tilde{\mathbf{t}}$ whose n th element has been removed. The next sections detail how Eqs. (8) and (9) are obtained.

A. Estimation of $(\mathbf{R}, \mathbf{b}, \Phi)$:

The estimation of $(\mathbf{R}, \mathbf{b}, \Phi)$ can be performed as in [13] using a stochastic EM (SEM) algorithm capable of handling the intractable expectation step induced by the MRF prior in (3). However, here we adopt a generalized SEM (GSEM) approach to obtain a more computationally attractive inference process. The main modifications with respect to the algorithm presented in [8] are twofold. Let $(\mathbf{R}^i, \mathbf{b}^i, \Phi^i)$ be the current estimate of $(\mathbf{R}, \mathbf{b}, \Phi)$. Instead of using a mean field-like approximation [17] in the E-step, we simulate a set of N_t samples $\{\tilde{\mathbf{t}}^1, \dots, \tilde{\mathbf{t}}^{N_t}\}$ via Markov chain Monte Carlo (MCMC) sampling from $p(\mathbf{t}|\mathbf{Y}, \mathbf{R}^i, \mathbf{b}^i, \Phi^i)$ such that

$$\tilde{\mathbf{t}}^m = \operatorname{argmax}_{\mathbf{t}_n} \sum_{\mathbf{t}_{\setminus n}} p(\mathbf{t}|\mathbf{Y}, \mathbf{R}^i, \mathbf{b}^i, \Phi^i). \quad (10)$$

Note that at each iteration, the Gibbs sampler used is hot-started with the samples generated at the previous iteration to improve the convergence speed. Sampling from

$p(\mathbf{t}|\mathbf{Y}, \mathbf{R}^i, \mathbf{b}^i, \Phi^i)$ is, in addition to the likelihood computation, one of the most computationally expensive step of the algorithm, even if the sampler is initialised properly, as it requires inverting a cumulative density function (CDF) defined on \mathbb{T} for each pixel. To reduce this cost, it is possible to downsample \mathbb{T} by a factor T_s and compute the CDFs on a coarser grid without affecting significantly the quality of the N_t simulated depth samples as well as the overall estimation quality.

After generation of the N_t samples, the cost function to be maximized in the M-step becomes

$$Q(\mathbf{R}, \mathbf{b}, \Phi|\mathbf{R}^i, \mathbf{b}^i, \Phi^i) =$$

$$\frac{1}{N_t} \sum_{m=1}^{N_t} [\log p(\mathbf{Y}|\mathbf{R}, \mathbf{b}, \tilde{\mathbf{t}}^m)] + \log f(\mathbf{R}, \mathbf{b}|\Phi)f(\Phi). \quad (11)$$

Instead of exactly maximizing this function at each iteration, in a generalized EM fashion [9], we used a few iterations of gradient ascent (resp. Newton-Raphson) to update (\mathbf{R}, \mathbf{b}) (resp. Φ) sequentially. A gradient ascent was preferred to second order methods for (\mathbf{R}, \mathbf{b}) to avoid the inversion of Hessian matrices, which are heavy to compute and whose complexity depends on L . We fixed the gradient step size to 10^{-3} as it gave the best trade-off between speed and performance in our experiments. When the number of signal photons is not so small compared to the background counts, our GSEM algorithm converges quickly and an average of 5 iterations are usually sufficient for the successive estimates to oscillate around fixed values. We assumed the convergence is reached when the relative error between two successive estimates of \mathbf{R} falls below 10^{-2} . The final estimates $(\hat{\mathbf{R}}, \hat{\mathbf{b}}, \hat{\Phi})$ are obtained by averaging the next 5 successive estimates of $(\mathbf{R}, \mathbf{b}), \Phi$ after convergence.

B. Depth estimation:

After convergence of GSEM algorithm proposed to estimate $(\mathbf{R}, \mathbf{b}, \Phi)$, the final step of the algorithm consists of estimating \mathbf{t} using (9). A MMAP estimates is computed from a variational approximation of $p(\mathbf{t}|\mathbf{Y}, \hat{\mathbf{R}}, \hat{\mathbf{b}}, \hat{\Phi})$ using the samples generated during the last iteration of the GSEM algorithm. As such solution is computationally expensive for the reasons mentioned above, we fixed $N_t = 1$ for all the experiments unless otherwise stated. As in [8], this approximation can be expressed as a product of independent distributions over the N pixels, whose maximum can be computed easily.

V. RESULTS

We assess the performance of the proposed approach using synthetic SW-MSL data, generated with $N = 200 \times 200$ pixels depth and reflectivity profiles obtained from real data in [7], with $L = 4$ wavelengths (473, 532, 589 and 640nm). The number of temporal bins is $T = 1500$ (bin width of 2ps) and a spatially constant background was added to all of the waveforms. We set $[t_{min}, t_{max}] = [301, 900]$ to ensure the integrals of the IRFs remain constant over the admissible object range. Different illumination scenarios have been generated to assess the performance of the proposed

method by changing the mean signal photon counts $MSC = \frac{1}{N} \sum_{n,\ell} r_{n,\ell} G_\ell$ and the average signal to background ratio $SBR = \frac{1}{N} \sum_{n=1}^N \frac{1}{T b_n} \sum_{\ell=1}^L r_{n,\ell} G_\ell$. The proposed approach, referred as C-Gamma (cluster-Gamma), is compared to the C-Dirichlet method proposed in [8]. To assess the speed-up/performance trade-off when sub-sampling the admissible depth grid or when generating more depth samples, we considered several versions of our method with a downsampling factor $T_s \in [4, 10]$. We also performed experiments with C-Gamma where $N_t = 100$ samples are generated in the S-step. These approaches are referred as $C\text{-Gamma}\{T_s, N_t\}$, when either T_s or N_t differ from 1.

We assess the reflectivity estimation quality using the mean squared error (MSE) defined as $\frac{1}{N} \sum_{n=1}^N \|r_n - \hat{r}_n\|_2^2$, where r_n (resp. \hat{r}_n) is the actual (resp. estimated) spectral response of the n th pixel. The ranging performance is then quantified using the CDF of the depth absolute error $h_n = |t_n - \hat{t}_n|, \forall n$, where \hat{t}_n is the estimate of the actual depth parameter t_n in the n th pixel. The reconstruction performance of the reflectivity and depth profiles obtained with $MSC \in [10, 44]$ photons are assessed in Fig. 2 and Fig. 3. Note that modifying T_s does not significantly change the reflectivity reconstruction performance for $N_t = 1$, so $C\text{-Gamma}\{4, 1\}$ is omitted in Fig. 2. Increasing T_s leads to estimating the reflectivity profile with a coarser depth profile, but as T_s remains lower than the delays between each IRF, and more importantly as long as the resolution of the coarse grid resolution stays smaller than the width of the IRFs, this does not impact \hat{R} as shown in Fig. 2. Although this affect the final depth reconstruction performance, the loss is not important for $T_s \leq 4$.

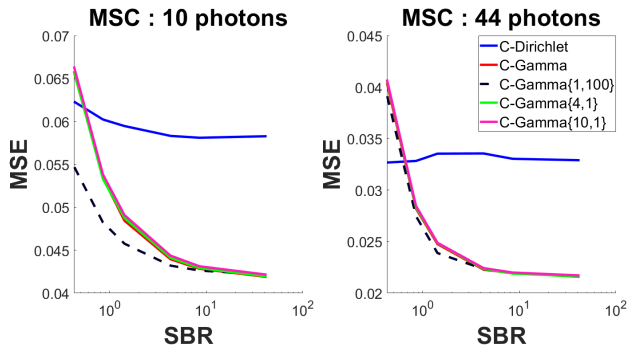


Fig. 2. MSE of \hat{R} obtained with the different competing methods as function of the SBR. The left (resp. right) subplot corresponds to MSC equals to 10 (resp. 44) photons.

	10/ ∞	10/0.046	44/ ∞	44/0.426
C-Gamma	104s	132s	350s	493s
C-Gamma{4, 1}	43s	49s	55s	69s
C-Gamma{10, 1}	34s	36s	38s	42s
C-Dirichlet	530s	592s	629s	662s

TABLE I

COMPUTATIONAL COST OF THE COMPETING APPROACHES FOR SYNTHETIC SW-MSL DATA ANALYSIS. IN THE TOP ROW, THE FIRST (RESP. SECOND) VALUE STANDS FOR THE MSC (RESP. SBR).




Fig. 3. CDFs of the depth absolute error, with high background ($SBR = 0.426$). The left (resp. right) subplot corresponds to MSC equals to 10 (resp. 44) photons.

The reflectivity reconstruction performance of C-Gamma and $C\text{-Gamma}\{1, 100\}$ differ only at low SBR in short imaging cases, explained by the accurate depth reconstruction of $C\text{-Gamma}\{1, 100\}$ in the same scenario that outperforms C-Dirichlet. Although the performance of $C\text{-Gamma}\{1, 100\}$ are better than those of C-Gamma, they differ only slightly and the high complexity of $C\text{-Gamma}\{1, 100\}$ makes it unattractive. The similarity between those results show that the proposed approach does not require perfect depth sampling at each iteration, as the depth profile converges quickly to a satisfactory estimate. The depth reconstruction performance of C-Dirichlet is only outperformed by $C\text{-Gamma}\{1, 100\}$ when $MSC = 10$. The impact of T_s is more significant on the depth estimates than on the reflectivity reconstruction. Indeed, increasing T_s only slightly degrades the estimation quality, and this is not significant for $T_s = [1, 4]$. Although the depth estimate with C-Dirichlet is more accurate than that of C-Gamma, their performance remain close and C-Gamma enables better reflectivity estimation for $SBR \geq 0.426$. Moreover, the loss is balanced by the significantly lower computational cost of our approach. Although we did not explore a broad range of values for T_s , the performance loss is negligible, especially given the associated speed-up.

Finally, Table I reports the computational time associated with the proposed model, compared to the method in [8]. The computational time of $C\text{-Gamma}\{1, 100\}$ is omitted here as it is computationally unattractive. All the experiments we performed in this work have been run on Matlab R2019a with a intel i7-8700k CPU @ 3.70GHz. Although the computational gain is implementation dependent and could still be improved via parallel programming, we noticed that a significant speed-up is obtained by the proposed method when compared to [8], for all illumination scenarios. The gain in term of speed is even more significant when increasing T_s from 1 to 4, but not much so when T_s goes from 4 to 10. Note that the estimation performance and speed-up associated with T_s depend on the data, and that the admissible set of values T_s can take depends on the histogram length, as well as on the temporal width of the impulse response. In general, C-Gamma

provides similar depth as the method in [8] but more accurate spectral signatures reconstruction and is about 10 times faster.

VI. CONCLUSION

In this paper, a novel algorithm has been presented to estimate the spectral and depth profile of a target from SW-MSL return. The method works in presence of non negligible background photon level and for short acquisition scenarios (2.5 signal photons per pixel per wavelength, on average). The alternative estimation of the model parameters simplifies the inference process and makes it computationally tractable. Moreover, the non-local C-Gamma prior model, similarly built as in [8], allows to split the problem pixel wise, and the concatenation of the admissible depth position reduces markedly the tractability of the likelihood computation. The significant computational gain compared to the slight performance degradation when increasing T_s motivates the search of an optimal sampling value. Also, this could be done by splitting the problem using for instance alternating direction method of multipliers, or by generalizing the proposed EM-based algorithm. While our new approach benefits from a significantly reduced computational cost, it is limited by the presence of high background level. Moreover, the starved photon regime has not been studied in our experiments. Thus, further studies should focus on adapting the method to extreme acquisitions scenarios.

REFERENCES

- [1] T. Bosch, "Laser ranging: a critical review of usual techniques for distance measurement," *Opt. Eng.*, vol. 40, no. 1, 2001.
- [2] P. Vines, K. Kuzmenko, J. Kirdoda, D. Dumas, M. Mirza, R. Millar, D. Paul, and G. Buller, "High performance planar Ge-on-Si single-photon avalanche diode detectors," *Nature Communications*, Sept. 2018.
- [3] P. S. Chhabra, A. Maccarone, A. McCarthy, A. M. Wallace, and G. S. Buller, "Analysis of foliage penetrating photon counting Lidar data for underwater mine counter measures," in *MTS/IEEE Oceans Conference*, June 2017.
- [4] X. Zhao, S. Shi, J. Yang, W. Gong, J. Sun, B. Chen, K. Guo, and B. Chen, "Active 3D imaging of vegetation based on multi-wavelength fluorescence Lidar," *Sensors*, vol. 20, no. 3, 2020.
- [5] W. Gong, J. Sun, S. Shi, J. Yang, L. Du, B. Zhu, and S. Song, "Investigating the potential of using the spatial and spectral information of multispectral Lidar for object classification," *Sensors*, vol. 15, no. 9, pp. 21 989–22 002, Sept. 2015.
- [6] J. Manzanera, A. Garc, C. Pascual, R. Gimeno, S. Mart, T. Tokola, and R. Valbuena, "Fusion of airborne Lidar and multispectral sensors reveals synergic capabilities in forest structure characterization," *GIScience & Remote Sensing*, vol. 53, Sept. 2016.
- [7] X. Ren, Y. Altmann, R. Tobin, A. McCarthy, S. McLaughlin, and G. S. Buller, "Wavelength-time coding for multispectral 3D imaging using single-photon Lidar," *Opt. Exp.*, vol. 26, no. 23, pp. 30 146–30 161, Nov. 2018.
- [8] Q. Legros, S. Meignen, S. McLaughlin, and Y. Altmann, "EM-based approach to 3D reconstruction from single-waveform multispectral Lidar data," *arXiv preprint arXiv:1912.06092*, 2019.
- [9] G. McLachlan and T. Krishnan, *The EM algorithm and extensions*. John Wiley & Sons, 2007, vol. 382.
- [10] G. Celeux, D. Chauveau, and J. Diebolt, "Stochastic versions of the EM algorithm: an experimental study in the mixture case," *Journal of Stat. Comp. and Simulation*, vol. 55, no. 4, pp. 287–314, 1996.
- [11] 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.
- [12] 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, p. 094006, Aug. 2017.
- [13] Y. Altmann and S. McLaughlin, "Range estimation from single-photon Lidar data using a stochastic EM approach," in *26th European Signal Processing Conference (EUSIPCO)*, Rome, Italy, Sept. 2018, pp. 1112–1116.
- [14] 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.
- [15] 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.
- [16] Y. Altmann, A. Maccarone, A. McCarthy, S. McLaughlin, and G. Buller, "Spectral classification of sparse photon depth images," *Optics Express*, vol. 26, no. 5, pp. 5514–5530, Feb. 2018.
- [17] 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.