3D RECONSTRUCTION USING SINGLE-PHOTON LIDAR DATA EXPLOITING THE WIDTHS OF THE RETURNS

J. Tachella, Y. Altmann, S. McLaughlin*

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

J.-Y. Tourneret

ENSEEIHT-IRIT-TéSA
University of Toulouse
Toulouse, France

ABSTRACT

Single-photon light detection and ranging (Lidar) data can be used to capture depth and intensity profiles of a 3D scene. In a general setting, the scenes can have an unknown number of surfaces per pixel (semi-transparent surfaces or outdoor measurements), high background noise (strong ambient illumination), can be acquired by systems with a broad instrumental response (non-parallel laser beam with respect to the target surface) and with possibly high attenuating media (underwater conditions). The existing methods generally tackle only a subset of these problems and can fail in a more general scenario. In this paper, we propose a new 3D reconstruction algorithm that can handle all the aforementioned difficulties. The novel algorithm estimates the broadening of the impulse response, considers the attenuation induced by scattering media, while allowing for multiple surfaces per pixel. A series of experiments performed in real long-range and underwater Lidar datasets demonstrate the performance of the proposed method.

Index Terms— Bayesian statistics, Inverse problems, Lidar, 3D reconstruction, low-photon imaging

1. INTRODUCTION

Scanning and reconstructing three-dimensional (3D) scenes has many important applications, such as autonomous driving [1], environmental monitoring [2, 3] and defence [4]. While a growing number of different 3D scanning modalities and cameras become available [5], single-photon light ranging and scanning (Lidar) offers several key advantages, such as the use of low-power lasers, long-range capabilities [6] and operation in highly attenuating environments [7], allowing for a depth precision in the order of millimetres [8]. However, recovering 3D information from single-photon measurements can be very challenging in scenarios where the number of collected photons associated with the signal of interest is very low or in the presence of strong ambient illumination. Thus, many image processing algorithms have been proposed to tackle this 3D reconstruction problem [9–20]. Each algorithm makes different assumptions on the sensed scene to reduce the complexity of the 3D reconstruction task. A summary of the capabilities of each method is presented in Table 1. The algorithms introduced in [11, 13, 18] assume very low or negligible background noise, providing poor results when the scene presents non-negligible ambient illumination. The methods proposed in [16, 19] assume only one sensed surface per pixel, thus simplifying the problem to the estimation of depth and reflectivity images. These methods provide unreliable estimates in outdoor scans, where every pixel does not necessarily contain a surface. The target detection algorithm proposed in [21] alleviates this problem by introducing a model which allows pixels without any surface. This assumption might also be insufficient in scenes with semi-transparent objects, such as camouflage [17] or windows [15], and long-range scenes (i.e., large laser footprint on target), which present multiple objects per pixel. Hence, the algorithms presented in [9, 10, 15, 17, 20] relax the assumptions on the number of surfaces per pixel, allowing an *a priori* unknown quantity. However, these methods consider a fixed instrumental response of the Lidar system, which may lead to biased estimates when a long-range surface is not orthogonal to the laser beam (e.g., [22] takes into account this effect for another single-photon gated system). Finally, only the method proposed in [18] considers the scattering effect of an underwater medium, accounting for the bias in the estimation of the target reflectivity. In this work, we present an extension of the ManiPoP algorithm recently introduced in [20] (yielding state-of-the-art reconstructions in the case of multiple surfaces per pixel) accounting for the variation of the instrumental response and also considering the effect of scattering media. The experiments performed using real Lidar datasets show that the proposed algorithm captures the aforementioned phenomena, thus improving the quality of the reconstructions in comparison to other existing methods.

*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 number EP/S000631/1 and the MOD University Defence Research Collaboration (UDRC) in Signal Processing.

	[11] (2015)	[12] (2016)	[13] (2016)	[21] (2016)	[15] (2016)	[16] (2017)	[18] (2017)	[17] (2017)	[19] (2018)	[22] (2018)	[20] (2018)	Prop. method
Single surface per pixel	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Background noise						✓			✓	✓	✓	✓
Target detection				✓	✓		✓				✓	✓
Multiple surface per pixel					✓		✓				✓	✓
Broadening of $h(t)$										✓		✓
Attenuating media								✓				✓

Table 1: Scope of recently proposed 3D reconstruction algorithms.

2. OBSERVATION MODEL

The principle of single-photon Lidar consists in constructing a histogram of time delays between emitted laser pulses and detected photon arrivals using a time-correlated single-photon detector (TCSPC). The full Lidar cube is denoted by $\boldsymbol{Z}$ and has a dimension of $N_r \times N_c \times T$, where N_r and N_c are the number of pixels in the vertical and horizontal axes and T is the number of histogram bins. The photon count recorded in pixel (i, j) and bin t is denoted as $[\boldsymbol{Z}]_{i,j,t} = z_{i,j,t} \in \{0, 1, \dots\}$. We model the objects present in the scene using a set of N_Φ marked points $\boldsymbol{\Phi} = \{(\boldsymbol{c}_n, r_n, w_n), n = 1, \dots, N_\Phi\}$, where $\boldsymbol{c}_n = (x_n, y_n, t_n)^T \in [1, N_r] \times [1, N_c] \times [1, T]$, $r_n \in \mathbb{R}_+$ and $w_n \in (1, +\infty)$ denote the spatial coordinates, intensity and width of the nth point, respectively. When the light flux received by the detector is sufficiently low, such that the probability of receiving more than one photon per laser pulse can be neglected, the number of photons in pixel (i, j) and time bin t has a Poisson distribution [11], i.e.,

$$z_{i,j,t}|(\boldsymbol{\Phi}, b_{i,j}) \sim \mathcal{P}\left(g_{i,j} s_{i,j,t} + g_{i,j} b_{i,j}\right) \tag{1}$$

where $g_{i,j} \in [0, 1]$ is the pixelwise gain of the device, $b_{i,j} \in \mathbb{R}_+$ is the background intensity due to the dark counts of the detector and the ambient illumination and $s_{i,j,t}$ is the intensity due to the surfaces. As detailed in [18, 23], the signal $s_{i,j,t}$ is expressed as

$$s_{i,j,t} = \sum_{n:(x_n,y_n)=(i,j)} r_n e^{-\alpha \Delta_b t_n} h_{w_n}(t - t_n) \tag{2}$$

where Δ_b is the bin width in metres, the exponential term accounts for the effect of scattering media with coefficient α (e.g., $\alpha \approx 0.6$ for clear water). The instrumental response of the device with width parameter w is denoted by $h_w(t)$ and modelled using a Gaussian kernel as $h_w(t) \propto \sum_k h(k) \exp(-\frac{(t-k)^2}{2(w-1)^2})$, where $h(t)$ is the instrumental response without broadening and is typically obtained during the calibration of the device. This width can be related to the angle between the laser beam and the imaged surface or to the local porosity of the object (light penetrating deeper into the object), as explained in [22]. Note that the signal model used in [20] can be recovered by assuming no attenuation (i.e., $\alpha = 0$) and no broadening of $h(t)$ (i.e., $h_w(t) = h(t)$). Assuming mutual independence between the noise realizations in different time bins and pixels, the full likelihood can be written as

$$p(\boldsymbol{Z}|\boldsymbol{\Phi}, \boldsymbol{B}) = \prod_{i=1}^{N_c} \prod_{j=1}^{N_r} \prod_{t=1}^{T} p(z_{i,j,t}|\boldsymbol{\Phi}, b_{i,j}) \tag{3}$$

where $[\boldsymbol{B}]_{i,j} = b_{i,j}$ is the background 2D image.

3. PROPOSED METHOD

Recovering the position and intensity of the objects from the raw Lidar data is an ill-posed problem, as many solutions can lead to similar data fitting errors. This problem can be tackled in a Bayesian framework, where the data generation mechanism is modelled through a set of parameters $\boldsymbol{\theta}$ that can be inferred using the available data $\boldsymbol{Z}$. The *a priori* knowledge about the unknown parameters $(\boldsymbol{\Phi}, \boldsymbol{B})$ is embedded in the prior distribution $p(\boldsymbol{\Phi}, \boldsymbol{B}|\boldsymbol{\Psi})$ depending on a set of hyperparameters $\boldsymbol{\Psi}$. Following Bayes theorem, the posterior distribution of the unknown model parameters is obtained via the relation $p(\boldsymbol{\Phi}, \boldsymbol{B}|\boldsymbol{Z}, \boldsymbol{\Psi}) \propto p(\boldsymbol{Z}|\boldsymbol{\Phi}, \boldsymbol{B}) p(\boldsymbol{\Phi}, \boldsymbol{B}|\boldsymbol{\Psi})$, where $\propto$ means "proportional to".

3.1. Prior distributions

In this work, we build upon the ManiPoP model introduced in [20], defining an additional prior distribution for the width of the points. Before introducing this new prior, we briefly summarize the priors used in [20]. A spatial point process prior was assigned to the point positions $\boldsymbol{c}_n$, promoting attraction between neighbouring points in the same surface using an area interaction process, while imposing a hard constraint on the minimum distance between two surfaces through the Strauss process. To ensure the positivity of the intensities, we used the transformation $r_n = \exp(m_n)$ and estimated the log-intensities $m_n \in \mathbb{R}$. A Gaussian Markov random field prior was also used to promote correlations among log-intensities of neighbouring points on the same surface, i.e.,

$$\boldsymbol{m}|\sigma_m^2, \beta_m, \boldsymbol{\Phi}_c \sim \mathcal{N}(\boldsymbol{0}, \sigma_m^2 \boldsymbol{P}^{-1}) \tag{4}$$

where the hyperparameter σ_m^2 controls the degree of smoothness. The precision matrix used in [20] is

$$[\boldsymbol{P}]_{n,n'} = \begin{cases} \beta_m + \sum_{\tilde{n}\in\mathcal{M}_{pp}(\boldsymbol{c}_n)} \frac{1}{d(\boldsymbol{c}_n;\boldsymbol{c}_{\tilde{n}})} & \text{if } n = n' \\ -\frac{1}{d(\boldsymbol{c}_n;\boldsymbol{c}_{n'})} & \text{if } \boldsymbol{c}_n \in \mathcal{M}_{pp}(\boldsymbol{c}_{n'}) \\ 0 & \text{otherwise} \end{cases} \tag{5}$$

where β_m is a positive hyperparameter, $d(\boldsymbol{c}_n;\boldsymbol{c}_{n'})$ denotes the Euclidean distance between points n and n' and $\mathcal{M}_{pp}(\boldsymbol{c}_{n'})$ is the set of neighbours of point n'. Finally, a gamma Markov random field [24] was used to model both positivity and spatial correlations for the background levels. An in-depth explanation of these priors can be found in [20]. In this paper, we introduce the transformation

$$w_n = e^{\tilde{w}_n} + 1, \quad \tilde{w}_n \in \mathbb{R} \tag{6}$$

to constrain the width to the interval $(1, +\infty)$[1]. Points in a small neighbourhood of a surface usually present a similar amount of broadening, as the laser beam has a similar angle of incidence on them or they present similar porosity. Thus, similarly to the log-intensity, we assign to the set of $\tilde{w_n}$ a Gaussian Markov random field prior similar to (4), which promotes spatial correlations between neighbouring widths, with hyperparameters $(\beta_{\tilde{w}}, \sigma_{\tilde{w}}^2)$ instead of (β_m, σ_m^2).

3.2. Posterior inference

We propose to estimate the point cloud using a maximum-a-posteriori estimator, i.e.,

$$\hat{\boldsymbol{\Phi}} = \arg\max_{\boldsymbol{\Phi}} p(\boldsymbol{\Phi}, \boldsymbol{B}|\boldsymbol{Z}, \Psi). \tag{7}$$

For the background level, we use the minimum mean squared error (MMSE) estimator of the background levels, i.e.,

$$\hat{\boldsymbol{B}} = \mathbb{E}\{\boldsymbol{B}|\boldsymbol{Z}, \Psi\}. \tag{8}$$

As these two estimators cannot be derived analytically, we propose to use Markov Chain Monte Carlo (MCMC) simulation methods to draw samples asymptotically distributed according to the posterior and can thus be used to compute various posterior statistics, including (7) and (8). As the number of points present in the scene is not known a priori, we use a reversible jump MCMC algorithm [26] that can handle the varying dimension nature of the spatial point process. The algorithm uses carefully designed moves to yield high acceptance rates, i.e., low execution times. These moves are inspired by the ideas presented in [20]. Due to the lack of space, we only discuss briefly the RJ-MCMC moves, referring the interested reader to [20].

[1]This is a standard choice in statistics when both spatial correlation and positivity constraints are desired [25]. Other alternatives, such as truncated Gaussian Markov random fields, result in less efficient MCMC samplers, which reduce the overall performance of the estimation.

Birth and death moves: The birth move proposes a point randomly across all possible positions, assigning a portion of the background intensity to the reflectivity of the new point. The width of the new point is proposed according to an exponential distribution. The reverse (death) move proposes the removal of a randomly chosen point.
Dilation and erosion moves: The dilation move randomly chooses an existing point and proposes a new neighbour, effectively "dilating" the current estimated surface. The reflectivity and width of the new point are sampled according to their priors. The erosion move randomly chooses a point and proposes to remove one of its neighbours.
Split and merge moves: The merge move proposes to combine two points into one, such that the new position is the average of the previous depths weighted by their respective reflectivity. Similarly, the new width is the weighted average of the previous width values. The new intensity is obtained as the sum of the previous intensities. The split move proposes to divide a point into two new ones, using auxiliary variables such that the reversible map is the merge move.
Mark move: The mark move randomly chooses a point and proposes to jointly update its log-intensity and log-width using a Gaussian random walk proposal.
Shift move: Similarly to the mark move, the shift move chooses a point at random and proposes a new position using a Gaussian proposal centred in the current estimate.
Background update: The background is sampled using a data augmentation scheme, as explained in [20].
To speed up the algorithm, an initial estimation is performed at a coarse scale (merging pixels to obtain a smaller data cube), which is used as an initial estimate at the finest scale.

4. EXPERIMENTS

4.1. Long-range imaging

The dataset presented in [6] consists of the dome of a building, imaged using terrestrial Lidar from a stand-off distance of approximately 3 kilometres. The Lidar cube has a size of $123 \times 96 \times 801$ and there are 913 recorded photons per pixel on average with a signal-to-background ratio of 1.64. In this case, the medium is air with a negligible scattering effect, i.e., $\alpha \approx 0$. Figure 1 shows the reconstruction obtained using the proposed algorithm. The estimated point widths are consistent with the orientation of the surface with respect to the incoming laser. For example, the lower part of the roof has a significant broadening w when the surface normal has a significant angle with respect to the laser. Figure 2 compares the estimated Poisson intensities obtained by ManiPoP and the proposed method for one of those pixels. The ManiPoP algorithm does not take into account the broadening of the peak, thus underestimating the reflectivity by 5%, whereas the proposed method provides an accuracy of 1%. Moreover, the estimation of the width does not significantly affect the

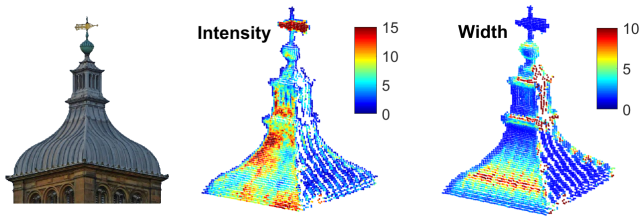

Fig. 1: Right: RGB image of the imaged dome (taken from a closer distance). The estimated point cloud intensities and widths from the college dataset are shown on the middle and left figures, respectively. The incoming laser beam is orthogonal to the left hand side of the roof.

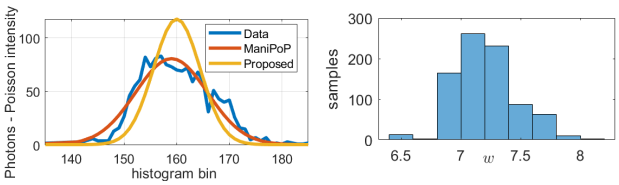

Fig. 2: Left: Estimated intensity for a pixel in the lower roof by ManiPoP and the proposed method. Right: Histogram of width samples obtained by the proposed RJ-MCMC algorithm for the same pixel.

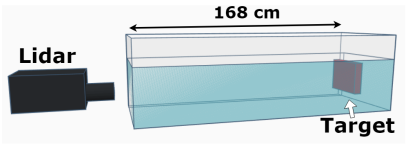

Fig. 3: The underwater measurements were taken at a standoff distance of 178 cm from the target, where 168 cm correspond to the water tank medium.

computational load, as ManiPoP requires an execution time of 174 seconds, whereas the proposed method requires 195 seconds.

4.2. Underwater imaging

The underwater scene presented in [23] is composed of a pipe inside a water tank, measured at a distance of 178 cm, as shown in Fig. 3. The measurements were repeated three times under varying concentrations of Maalox in the water, obtaining the scattering coefficients $\alpha \in [0.6, 3.9, 4.8]$. The Lidar cube has a size of $120 \times 120 \times 2500$ and the acquisition time was 100 ms in all cases. We compare the results with the ManiPoP reconstructions followed by a simple post-processing correction of the estimated intensities, i.e., by dividing them by the attenuation factor based on the estimated range. Figure 4 shows the estimated point clouds for all the values of α for both algorithms. The reconstructions obtained by the proposed algorithm have a lower variation in the estimated intensity. Moreover, in the case with highest attenuation (i.e., $\alpha = 4.8$), ManiPoP fails to recover the backplane of the scene, as its mean intensity (without the exponential term correction) is too low and the algorithm considers it as belonging to the background.

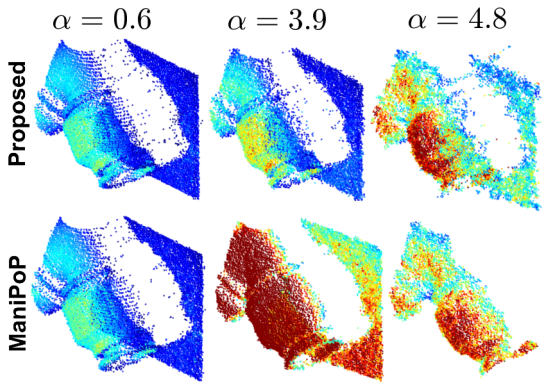

Fig. 4: The reconstructions using the proposed method are shown on the top row, whereas the reconstructions achieved by ManiPoP are shown in the lower row. The incoming laser beam is orthogonal to the backplane. All the intensities are shown in the same colormap scale.

5. CONCLUSIONS AND FUTURE WORK

We presented a new algorithm that generalizes the ManiPoP model investigated in [20], allowing for a variable width of the instrumental response and accounting for the attenuation of scattering media. These effects were incorporated within the Bayesian model and RJ-MCMC inference, thus improving the quality of the estimates in comparison to simple post-processing steps (see Section 4.2). Moreover, the proposed model refinement did not result in a significant increase of the execution time. Regarding the broadening of the instrumental response, this phenomenon is only observable when many photons per pixel are recorded and when $h(t)$ does not present a heavy tailed decay. In these cases, the proposed algorithm obtains width samples from the prior distribution, leaving the intensity and position estimation almost unmodified with respect to ManiPoP. Future work will be devoted to extending the proposed Bayesian model and the associated MCMC method to process multispectral Lidar data [27].

Acknowledgements

We would like to thank Prof. R. Lamb and Dr. A. Pawlikowska from Leonardo UK for the long-range dataset and Prof. G. S. Buller and Dr. A. Maccarone from Heriot-Watt single-photon group for the underwater data. Part of this work was conducted within the STIC-Amsud project HyperMed.

6. REFERENCES

[1] J. Hecht, "Lidar for self-driving cars," *Optics and Photonics News*, vol. 29, no. 1, pp. 26–33, 2018.

[2] C. Mallet and F. Bretar, "Full-waveform topographic lidar: State-of-the-art," *ISPRS Journal of photogrammetry and remote sensing*, vol. 64, no. 1, pp. 1–16, 2009.

[3] M. A. Canuto, F. Estrada-Belli, T. G. Garrison, S. D. Houston, M. J. Acuña, M. Kováč, D. Marken, P. Nondédéo, L. Auld-Thomas, C. Castanet, D. Chatelain, C. R. Chiriboga, T. Drápela, T. Lieskovský, A. Tokovinine, A. Velasquez, J. C. Fernández-Díaz, and R. Shrestha, "Ancient lowland maya complexity as revealed by airborne laser scanning of northern guatemala," *Science*, vol. 361, no. 6409, 2018. [Online]. Available: http://science.sciencemag.org/content/361/6409/eaau0137

[4] J. Gao, J. Sun, J. Wei, and Q. Wang, "Research of underwater target detection using a slit streak tube imaging lidar," in *Proc. Academic International Symposium on Optoelectronics and Microelectronics Technology (AISOMT)*, Harbin, China, Mar. 2012, pp. 240–243.

[5] R. Horaud, M. Hansard, G. Evangelidis, and C. Ménier, "An overview of depth cameras and range scanners based on time-of-flight technologies," *Machine Vision and Applications*, vol. 27, no. 7, pp. 1005–1020, Oct 2016. [Online]. Available: https://doi.org/10.1007/s00138-016-0784-4

[6] A. M. Pawlikowska, A. Halimi, R. A. Lamb, and G. S. Buller, "Single-photon three-dimensional imaging at up to 10 kilometers range," *Opt. Express*, vol. 25, no. 10, pp. 11 919–11 931, May 2017. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-25-10-11919

[7] A. Maccarone, A. McCarthy, X. Ren, R. E. Warburton, A. M. Wallace, J. Moffat, Y. Petillot, and G. S. Buller, "Underwater depth imaging using time-correlated single-photon counting," *Opt. Express*, vol. 23, no. 26, pp. 33 911–33 926, Dec 2015. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-23-26-33911

[8] 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. [Online]. Available: http://ao.osa.org/abstract.cfm?URI=ao-48-32-6241

[9] 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, 2007.

[10] ——, "Multilayered 3d lidar image construction using spatial models in a bayesian framework," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 6, pp. 1028–1040, June 2008.

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

[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, 2016.

[13] A. Halimi, Y. Altmann, A. McCarthy, X. Ren, R. Tobin, G. S. Buller, and S. McLaughlin, "Restoration of intensity and depth images constructed using sparse single-photon data," in *Proc. Signal Processing Conference (EUSIPCO)*, Budapest-Hungary, Sep. 2016, pp. 86–90.

[14] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Target detection for depth imaging using sparse single-photon data," in *Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, March 2016, pp. 3256–3260.

[15] D. Shin, F. Xu, F. N. Wong, J. H. Shapiro, and V. K. Goyal, "Computational multi-depth single-photon imaging," *Optics express*, vol. 24, no. 3, pp. 1873–1888, 2016.

[16] 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.

[17] A. Halimi, R. Tobin, A. McCarthy, S. McLaughlin, and G. S. Buller, "Restoration of multilayered single-photon 3d lidar images," in *Proc European Signal Processing Conference (EUSIPCO)*, Aug 2017, pp. 708–712.

[18] 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*, vol. 3, no. 3, pp. 472–484, Sept 2017.

[19] D. B. Lindell, M. OToole, and G. Wetzstein, "Single-Photon 3D Imaging with Deep Sensor Fusion," *ACM Trans. Graph. (SIGGRAPH)*, no. 4, 2018.

[20] J. Tachella, Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, J.-Y. Tourneret, and S. McLaughlin, "Bayesian 3D reconstruction of complex scenes from single-photon lidar data," *To appear in SIAM Journal on Imaging Sciences*.

[21] 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, pp. 456–467, Dec 2016.

[22] X. Ren, P. W. R. Connolly, A. Halimi, Y. Altmann, S. McLaughlin, I. Gyongy, R. K. Henderson, and G. S. Buller, "High-resolution depth profiling using a range-gated cmos spad quanta image sensor," *Opt. Express*, vol. 26, no. 5, pp. 5541–5557, Mar 2018. [Online]. Available: http://www.opticsexpress.org/abstract.cfm?URI=oe-26-5-5541

[23] A. Maccarone, A. McCarthy, X. Ren, R. E. Warburton, A. M. Wallace, J. Moffat, Y. Petillot, and G. S. Buller, "Underwater depth imaging using time-correlated single-photon counting," *Optics express*, vol. 23, no. 26, pp. 33 911–33 926, 2015.

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

[25] H. Rue and L. Held, *Gaussian Markov random fields: theory and applications*. CRC press, 2005.

[26] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, pp. 711–732, 1995.

[27] A. Wallace, C. Nichol, and I. Woodhouse, "Recovery of forest canopy parameters by inversion of multispectral lidar data," *Remote Sensing*, vol. 4, no. 2, pp. 509–531, 2012.