FAST SURFACE DETECTION IN SINGLE-PHOTON LIDAR WAVEFORMS

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) devices can be used to obtain range and reflectivity information from 3D scenes. However, reconstructing the 3D surfaces from the raw waveforms can be very challenging, in particular when the number of spurious background detections is large compared to the number of signal detections. This paper introduces a new and fast detection algorithm, which can be used to assess the presence of objects/surfaces in each waveform, allowing only the histograms where the imaged surfaces are present to be further processed. The method is compared to state-of-the-art 3D reconstruction methods using synthetic and real single-photon data and the results illustrate its benefits for fast and robust target detection using single-photon data.

Index Terms— Bayesian statistics, inverse problems, Lidar, detection, low-photon imaging and sensing

1. INTRODUCTION

Imaging systems based on time-of-flight laser detection and ranging (Lidar) are used to reconstruct 3-dimensional scenes in many applications, including automotive [1], environmental sciences [2], architectural engineering and defence [3, 4]. This modality consists of illuminating the scene with laser pulses and analyzing the distribution of the photons received by the detector to infer the presence of objects as well as their range, and radiative properties (e.g., reflectivity, observation conditions). For each pixel, associated with a different region of the scene, a histogram of time delays between the emitted pulses and the detected photon arrivals is usually recorded. Conventionally, in the presence of objects, the recorded photon histograms are decomposed into a series of peaks whose positions can be used to infer the distance of the objects present in each region of the scene and whose amplitudes provide information about the intensity of the objects.

In this work, we address the target detection problem, which aims at identifying regions or pixels of the scene where objects are present. We propose an algorithm adapted to situations where the flux of photons originally emitted by the laser source is small, where the ambient illumination level is high (i.e., low signal-to-background ratio) and for which classical depth imaging methods [5] usually provide unsatisfactory results in terms of object detection and reconstruction. In addition, this method can be easily extended to single-photon depth imaging in turbid media, e.g., underwater depth imaging [6, 7]. As in most 3D reconstruction scenarios, we assume that at most one surface can be observed in each pixel. The observation model includes two kinds of detector events [8–10]: the photons originating from the illumination laser and scattered back from the target (if present); and the background detector events originating from ambient light and dark events resulting from detector noise. Following a classical Bayesian approach, the target detection problem is first formulated as a pixel-wise model selection and estimation problem and prior distributions are assigned to each of the unknown model parameters. We also present a post-processing step to further improve the detection maps at a low additional cost. This additional step can be seen as defining the prior probabilities of target presence (or equivalently the binary labels associated with the presence/absence of targets) that account for the spatial organization of objects in the scene. In contrast to the target detection method presented in [9], where a reversible-jump Markov chain Monte Carlo (RJ-MCMC) [11, 12] method was used to generate samples according to a posterior distribution of interest, we reformulate the observation model such that the background parameters can be marginalized analytically while the other parameters (target range and reflectivity) can be marginalized from the posterior distribution using (finite sums of) one-dimensional integrals. The resulting algorithm, which relies mostly on pixel-wise, low-dimensional integrations, is thus fast and can be implemented using parallel architectures.

The remainder of this paper is organized as follows. Section 2 recalls the classical statistical model used for depth imaging using single-photon Lidar and introduces the alternative model used in this work. Section 3 details the proposed Bayesian target detection method. Simulation results conducted using real Lidar measurements are presented and discussed in Section 4. Conclusions and potential future work are finally reported in Section 5.

*This work was supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31.

978-9-0827-9703-9/19/$31.00 ©2019 IEEE

2. OBSERVATION MODEL

A histogram of photon detections with T bins is denoted by $\boldsymbol{z} = [z_1, \dots, z_T]^T \in \mathbb{Z}_+^{T\times 1}$, where $\mathbb{Z}_+ = \{0, 1, \dots\}$ is the set of positive integers. If the light flux reaching the single-photon detector is sufficiently low [13], the observed photon count in a given time bin t follows a Poisson distribution, i.e,

$$z_t|(r, t_0, b) \sim \mathcal{P}\left(rh(t - t_0) + b\right), \quad \forall t = 1, \dots, T, \quad (1)$$

where $r \in \mathbb{R}_+$ is the target intensity, $b \in \mathbb{R}_+$ is the constant background level associated with dark counts and ambient illumination and $h(t)$ is the instrumental response of the device, which is assumed to be normalized ($\sum_{t=1}^{T} h(t) = 1$). In (1), t_0 corresponds to the typical delay/time-of-flight associated with the depth of the given surface.

An equivalent model can be defined using the signal-to-background ratio (SBR), which is defined as the ratio of the useful detected photons, e.g., originally emitted by the laser source, divided by the total number of background photons in the histogram, i.e., $w = r/(bT)$. Following this alternative parametrization, the observation model (1) can be rewritten

$$z_t|(w, t_0, b) \sim \mathcal{P}\left(b\left(wTh(t - t_0) + 1\right)\right), \forall t = 1, \dots, T. \quad (2)$$

The main motivation for using (2) instead of (1) is that gamma distributions are conjugate priors for b in (2) (and not in (1)), which allows a simple marginalization of b, as will be seen in Section 3. Assuming the T observations in $\boldsymbol{z}$ are mutually independent, conditioned on their means [13], the joint likelihood can be expressed as

$$p(\boldsymbol{z}|w, t_0, b) = \prod_{t=1}^{T} p(z_t|w, t_0, b). \quad (3)$$

As can be seen from the two observation models (1) and (2), in the absence of surface in the field of view, i.e., when $r = 0$ or equivalently when $w = 0$, the observation model reduces to considering T random variables z_t drawn independently from a Poisson distribution with mean b, i.e.,

$$z_t|(w = 0, t_0, b) \sim \mathcal{P}\left(b\right). \quad (4)$$

In this work, we propose a surface detection algorithm to decide whether $w = 0$ or $w > 0$. However, the background level b, and the instant t_0 (if an object is present) are unknown in practice, which makes the detection task more difficult. The next section presents the proposed Bayesian strategy for this detection problem.

3. DETECTION STRATEGY

Adopting a Bayesian framework, we assign prior distributions to the unknown parameters using the *a priori* knowledge available about the model parameters.

3.1. Prior distributions

Similarly to previous work [9, 10, 14], independent prior distributions are assigned to the background level and target reflectivity, i.e., $p(r, b) = p(r)p(b)$. In order to model the absence ($r = 0$) or presence ($r > 0$) of a target, we use a spike and slab prior distribution [15] for the signal intensity, that is

$$p(r|u, \alpha_r, \beta_r) = u\mathcal{G}(r; \alpha_r, \beta_r) + (1 - u)\delta(r) \quad (5)$$

where $\delta(r)$ is the Dirac delta distribution centred in 0 and $u \in \{0, 1\}$ is a binary variable that indicates the presence ($u = 1$) or absence ($u = 0$) of a target. Moreover, $\mathcal{G}(r; \alpha_r, \beta_r)$ denotes a gamma density with known shape α_r and rate β_r. Note that (α_r, β_r) can usually be adjusted from calibration measurements, as the dynamic range of r is primarily guided by the laser power used, the average distance between the Lidar system and the scene, the scattering properties of the media, the efficiency of the detector and the pixel-wise acquisition time. The prior distribution for the binary label u is a Bernoulli distribution such that $p(u = 1) = \pi$ and $p(u = 0) = 1 - \pi$, where $\pi \in (0, 1)$ is the prior probability of target presence. Unless stated otherwise, we have used $\pi = 0.5$, expressing our absence of knowledge regarding this parameter.

The background level mostly depends on the amount of ambient illumination reaching the single-photon detector and is modelled as in [9, 10] with a conjugate gamma distribution $p(b|\alpha_b, \beta_b) = \mathcal{G}(b; \alpha_b, \beta_b)$ with known parameters. If limited information is available about b a weakly informative prior distribution can be defined for b (e.g., to have a heavy-tailed prior). The resulting joint prior distribution on the parametrization based on b and w can be obtained from $p(r, b)$ by applying a standard change of variables yielding

$$\begin{aligned} p(w, b|u, \phi) = (1 - u)\delta(w)\mathcal{G}(b; \alpha_b, \beta_b) \\ + uc_0(w)\mathcal{G}(b; \alpha_b + \alpha_r, \beta_b + \beta_r Tw) \end{aligned} \quad (6)$$

where $\phi = \{\alpha_r, \beta_r, \alpha_b, \beta_b\}$, $c_0(w) = (T\beta_r)^{\alpha_r} w^{\alpha_r - 1}(\beta_b + Tw\beta_r)^{\alpha_r + \alpha_b}\beta_b^{\alpha_b}/B(\alpha_r, \alpha_b)$ and $B(\cdot, \cdot)$ is the beta function. Since ϕ is known in this work, it is omitted in all the conditional distributions in the remainder of this paper. Assuming no prior knowledge on the position of the target, we assign a uniform prior for the depth, i.e., $p(t_0) = 1/T$ for any t_0 in $\{1, \dots, T\}$. However, this choice could be changed if additional information was available.

3.2. Decision rule

The proposed decision rule is based on the marginal posterior distribution of the label u, obtained by integrating out the parameters b, t_0 and w, considered here as nuisance parameters. Defining H_0 and H_1 as the absence and presence of the target respectively, the proposed decision rule is

$$p(u = 0|\boldsymbol{z}) \underset{H1}{\overset{H0}{\gtrless}} p(u = 1|\boldsymbol{z}), \quad (7)$$

where

$$p(u|\boldsymbol{z}) = \sum_{t=1}^{T} \int \int p(w, b, t_0, u|\boldsymbol{z}) db dw, \quad (8)$$

with $p(w, b, t_0, u|\boldsymbol{z}) \propto p(\boldsymbol{z}|w, b, t_0) p(w, b|u) p(t_0) p(u)$ using Bayes rule. Note that, as will be shown in Section 4, it is also possible to consider t_0 as a deterministic parameter and only marginalize (b, w), i.e., consider $p(u|\boldsymbol{z}, t_0)$ in (7), where the actual (unknown) value of t_0 is replaced by an arbitrary estimate.

3.3. Computation of marginals

In order to compute the marginal distribution $p(u|\boldsymbol{z})$ used in (7), we first integrate out the background level and target position, that is

$$p(w, u|\boldsymbol{z}) \propto \sum_{t_0=1}^{T} p(t_0) \int_0^{\infty} p(\boldsymbol{z}|w, t_0, b) p(w, b|u) p(u) db.$$

Due to the conjugacy between the observation model (2) and the prior distribution (6), the inner integral is available in closed form. The integration over the signal-to-background level is also available in closed form for $u = 0$,

$$\begin{aligned} p(u = 0|\boldsymbol{z}) &= \int p(w, u = 0|\boldsymbol{z}) dw \\ &= \frac{(1 - \pi)}{\gamma} \Gamma(\bar{\boldsymbol{z}} + \alpha_b)(T + \beta_b)^{\bar{z} + \alpha_b} \quad (9) \end{aligned}$$

where $\bar{z} = \sum_{t=1}^{T} z_t$ is the total number of photons observed and γ is a normalization constant. Finally, the marginal probability of the target being present is

$$p(u = 1|\boldsymbol{z}) = \frac{c_1}{\gamma} \int_0^{\infty} f_1(w) \sum_{t_0=1}^{T} \exp(f_2(w, t_0)) dw \quad (10)$$

with

$$\begin{aligned} f_1(w) &= w^{\alpha_r - 1} \left(\beta_b + T(1 + w(\beta_r + 1))\right)^{\bar{z} + \alpha_r + \alpha_b} \\ f_2(w, t_0) &= \sum_{t=1}^{T} z_t \log(wTh(t - t_0) + 1), \end{aligned}$$

γ is the same constant as in (10) and $c_1 = \pi \Gamma(\alpha_r) \Gamma(\bar{z} + \alpha_b + \alpha_r)(\beta_r T)^{\alpha_r}$. Since γ is shared in (9) and (10), it can be easily computed using $p(u = 0|\boldsymbol{z}) + p(u = 1|\boldsymbol{z}) = 1$. The marginal distribution (10) involves an intractable integral. However, the sum can be computed with $\mathcal{O}(T \log T)$ floating point operations using the fast Fourier transform (FFT), allowing the integral to be be numerically approximated with a quadrature method (with a computational cost of K integrand evaluations). Thus, the overall complexity is $\mathcal{O}(KT \log T)$, which is close to cross-correlation if $K << T$. Note that if t_0 is not marginalized and replaced by a point estimate instead, (10) is simplified as the sum in the integrand reduces to one term.

3.4. Spatial regularization

Histograms corresponding to neighbouring pixels generally show similar numbers of surfaces [16–18]. Thus, we proposed to refine the pixel-wise detection method to create a more homogeneous map of target presence. Such segmentation (or subsequent denoising step) can be efficiently computed by solving a total variation (TV) problem [19], that is

$$\hat{\boldsymbol{u}} = f_{\text{th}} \left(\arg\min_{\boldsymbol{v}} ||\boldsymbol{v} - \boldsymbol{y}||_2^2 + \tau ||\boldsymbol{v}||_{TV} \right) \quad (11)$$

where the input image $\boldsymbol{y}$ contains the log-ratios $y_{i,j} = \log p(u = 1|\boldsymbol{z}) - \log p(u = 0|\boldsymbol{z})$ of histogram at pixel (i, j), $||\cdot||_{TV}$ is the isotropic total variation operator, τ is a user-defined parameter which controls the impact of the TV-based denoiser ($\tau = 5$ here) and $f_{\text{th}}(\cdot)$ is a hard thresholding operation, which assigns 1 to positive inputs and 0 otherwise.

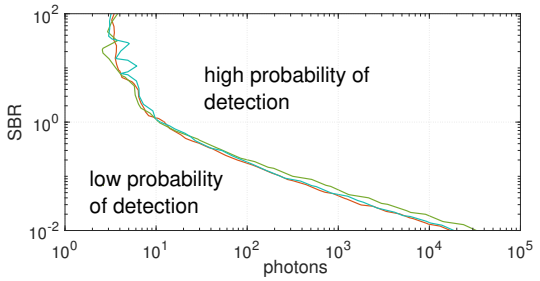

Fig. 1: Performance of the proposed detection algorithm for different numbers of photons and signal-to-background ratio. The solid lines correspond to a true positive rate of 95% with t_0 known (red), estimated by cross-correlation (blue) and marginalized (green).

4. RESULTS

In this section, we evaluate the performance of the proposed algorithm using synthetic and real Lidar datasets. In all the results presented here, we assume we know (from calibration measurements) the average number r_M of signal photons detected when observing an object of unit reflectivity under similar observation conditions as for the scene of interest. We then use this value to set $\phi = \{\alpha_r, \beta_r, \alpha_b, \beta_b\} = \{2, 2/r_M, 1, T/r_M\}$, which corresponds to a fairly informative prior for r and more weakly informative prior for b[1].

First, we evaluate the method for different values of SBR and $\bar{z}$, using a Gaussian instrumental response with standard deviation $\sigma = T/100$. Fig. 1 shows the SBR/photon counts curves for a true positive rate of 95%, without marginalizing t_0 (the true value of t_0, estimating it with the classical matched filter (see [10] for details)) and with the proposed marginalization. The probability of false alarm of the proposed detector is shown in Fig. 2. While the sensitivity of the

[1] While a thorough robustness analysis is beyond the scope of this paper, the results do not vary significantly with reasonable variations of r_M.

detector does not change significantly with the marginalization of t_0 (see Fig. 1), the probability of false alarm increases when t_0 is estimated using the standard cross-correlation. Fig. 3 depicts a map of the empirical probability of detection obtained by the proposed method (with t_0 marginalized) for various SBRs and photon counts. This figure gives an empirical bound on the minimum number of photons needed to detect a target with a given probability, for different levels of SBR, which can be used to adjust the acquisition time of the device in practice. In absence of a target, Fig. 2 shows that around 20 background detections are sufficient to correctly discard the histogram with high probability (>0.95).

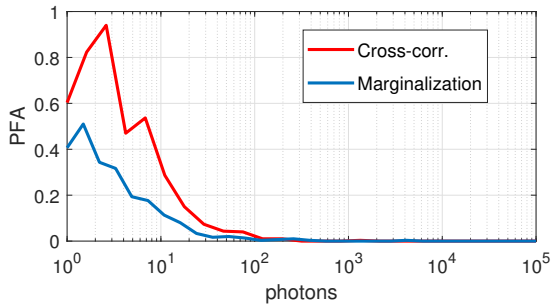

Fig. 2: Probability of false alarm achieved by the proposed method by marginalizing t_0 (blue curve) and by estimating t_0 via cross-correlation/matched filtering (red curve).

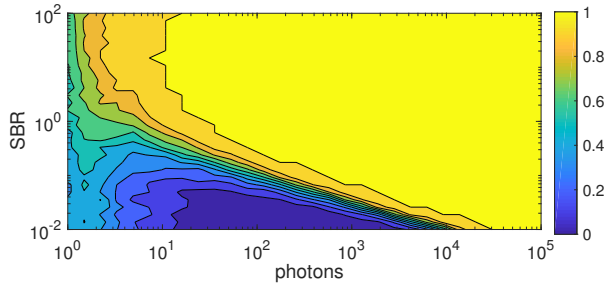

Fig. 3: True positive rate for the proposed method as a function of the SBR and total number of photons.

Secondly, we compare the proposed algorithm with two state-of-the-art single-depth [9] and multi-depth [18] detection algorithms and the standard cross-correlation with reflectivity thresholding (see [9] for details) using real Lidar dataset, which consists of a polystyrene head measured at a stand-off distance of 325 metres during midday (more details can be found in [9]). The hyperparameters of these two algorithms were chosen to obtain the best PD/PFA trade-offs. The dataset consists of 200×200 pixels with $T = 2700$ histogram bins per pixel, an approximate SBR of 0.29 with a 5th-95th percentile interval of $(0.05, 0.67)$. Table 1 shows the performance of the different detection algorithms for two different per-pixel acquisition times (3 ms and 1 ms), which correspond to average counts of 90 and 30 photons, respectively. These results show that although the proposed method is applied pixel-wise, it generally provides better results than the cross-correlation method, a significant improvement in terms of probability of false alarm (PFA) and probability of detection (PD) is obtained by using the additional denoising step presented in Section 3.4, which accounts for spatial correlation between adjacent pixels (as in [9]). This can also be confirmed visually on the detection maps depicted in Fig. 4. In contrast to the proposed method, applying a TV post-processing step directly to the cross-correlation output is challenging, as the per pixel detection probabilities are not available. Note that although our method yields a small improvement in terms of PD in the 1 ms case (see bottom row of Table 1), the corresponding PFA is significantly increased. Finally, Table 2 illustrates that the proposed method, even when combined with a post-processing step, remains very competitive from a computational point of view and can potentially be used for real-time target detection.

	PD [%]		PFA [%]	
	3 ms	1 ms	3 ms	1 ms
Altmann et al. [9]	89.2	81.4	**0.01**	0.01
ManiPoP [18]	**94.47**	84.51	0.47	**0.00**
Cross-corr.	91.1	57.4	20.1	42.6
Proposed	80.52	75.40	6.45	18.53
Proposed + TV	92.76	**94.31**	0.04	0.57

Table 1: Probabilities of false alarm (PFAs) and probabilities of detection (PDs) for the proposed algorithm and other state-of-the-art detection algorithms.

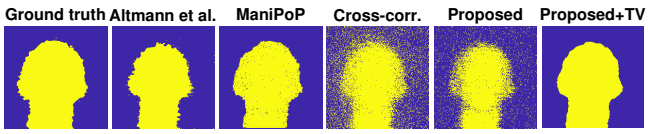

Fig. 4: Detected targets (in yellow) for a per-pixel acquisition time of 3 ms.

	3 ms	1 ms
Altmann et al. [9]	24 h	12 h
ManiPoP [18]	539 s	416 s
Cross-corr.	**1 ms (p)**	
Proposed	50 ms (p)	
Proposed + TV	50 ms (p) + 0.1 s	

Table 2: Execution time for the proposed algorithm and other state-of-the-art alternatives for per-pixel acquisition times of 3 ms and 1 ms. Cross-correlation and the proposed method are assumed to be executed in parallel, indicated by (p).

5. CONCLUSIONS AND FUTURE WORK

We presented a new fast target detection algorithm for single-photon Lidar data. Unlike other existing algorithms, the proposed method is easily parallelizable and can be used as a pre-processing step to discard histograms without useful information. This step can improve the reconstruction quality obtained by algorithms assuming one depth per pixel [20, 21], as it removes histograms without surfaces from the data cube. Moreover, it can also be used before multiple-surface-per-pixel algorithms [18, 22] to reduce the computational load. As mentioned in Section 4, the performance bounds shown in Figures 1 to 3 can be used to adjust the acquisition time depending on the minimum SBR admissible. Future work will be devoted to including the proposed method within a hierarchical model and refine the additional denoising step investigated in this work.

Acknowledgements

We thank the single-photon group led by Prof. G. S. Buller at Heriot-Watt University, for providing the real single-photon data used in this work.

6. REFERENCES

[1] P. Lindner and G. Wanielik, "3D Lidar processing for vehicle safety and environment recognition," in *Proc. Work. Comput. Intell. Vehicles and Vehicular Syst. (CIVVS'09)*, March 2009, pp. 66–71.

[2] T. Hakala, J. Suomalainen, S. Kaasalainen, and Y. Chen, "Full waveform hyper-spectral Lidar for terrestrial laser scanning," *Opt. Express*, vol. 20, no. 7, pp. 7119–7127, Mar 2012.

[3] N. Cadalli, P. J. Shargo, D. C. Munson, Jr., and A. C. Singer, "Three-dimensional tomographic imaging of ocean mines from real and simulated Lidar returns," in *Proc. SPIE, Ocean Optics: Remote Sensing and Underwater Imaging*, vol. 4488, 2002, pp. 155–166.

[4] J. Gao, J. Sun, J. Wei, and Q. Wang, "Research of underwater target detection using a slit streak tube imaging Lidar," in *Acad. Int. Symp. Optoelectronics and Microelectronics Technology (AISOMT)*, Oct 2011, pp. 240–243.

[5] A. McCarthy, X. Ren, A. D. 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. 22 098–22 113, Sep. 2013.

[6] 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.

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

[8] 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.

[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, pp. 456–467, Dec 2016.

[10] ——, "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.

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

[12] C. Andrieu and A. Doucet, "Joint Bayesian model selection and estimation of noisy sinusoids via reversible jump MCMC," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2667–2676, Oct 1999.

[13] 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.

[14] 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.

[15] J. Piironen and A. Vehtari, "Sparsity information and regularization in the horseshoe and other shrinkage priors," *Electron. J. Statist.*, vol. 11, no. 2, pp. 5018–5051, 2017.

[16] 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.

[17] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, "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.

[18] J. Tachella, Y. Altmann, X. Ren, A. McCarthy, G. Buller, S. McLaughlin, and J. Tourneret, "Bayesian 3D reconstruction of complex scenes from single-photon lidar data," *SIAM Journal on Imaging Sciences*, vol. 12, no. 1, pp. 521–550, 2019.

[19] A. Chambolle and T. Pock, "An introduction to continuous optimization for imaging," *Acta Numerica*, vol. 25, p. 161319, 2016.

[20] 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.

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

[22] J. Tachella, Y. Altmann, S. McLaughlin, and J. . Tourneret, "3D reconstruction using single-photon lidar data exploiting the widths of the returns," in *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, May 2019, pp. 7815–7819.