



Heriot-Watt University
Research Gateway

On fast object detection using single-photon lidar data

Citation for published version:

Tachella, J, Altmann, Y, McLaughlin, S & Tourneret, J-Y 2019, On fast object detection using single-photon lidar data. in D Van De Ville, D Van De Ville, M Papadakis & YM Lu (eds), *Wavelets and Sparsity XVIII.*, 111380T, Proceedings of SPIE, vol. 11138, SPIE, SPIE Optical Engineering + Applications 2019, San Diego, United States, 11/08/19. <https://doi.org/10.1117/12.2527685>

Digital Object Identifier (DOI):

[10.1117/12.2527685](https://doi.org/10.1117/12.2527685)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Wavelets and Sparsity XVIII

Publisher Rights Statement:

Copyright 2019 Society of PhotoOptical Instrumentation Engineers (SPIE). One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this publication for a fee or for commercial purposes, and modification of the contents of the publication are prohibited.

Proceedings Volume 11138, Wavelets and Sparsity XVIII; 111380T (2019)
<https://doi.org/10.1117/12.2527685>

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

On fast object detection using single-photon lidar data

Tachella, Julian, Altmann, Yoann, McLaughlin, Stephen, Tourneret, Jean-Yves

Julian Tachella, Yoann Altmann, Stephen McLaughlin, Jean-Yves Tourneret, "On fast object detection using single-photon lidar data," Proc. SPIE 11138, Wavelets and Sparsity XVIII, 111380T (9 September 2019); doi: 10.1117/12.2527685

SPIE.

Event: SPIE Optical Engineering + Applications, 2019, San Diego, California, United States

On fast object detection using single-photon lidar data

Julian Tachella^a, Yoann Altmann^a, Stephen McLaughlin^a, and Jean-Yves Tournet^b

^aSchool of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, United Kingdom

^bENSEEIH-IRIT-TéSA, University of Toulouse, Toulouse, France

ABSTRACT

Light detection and ranging (Lidar) systems based on single-photon detection can be used to obtain range and reflectivity information from 3D scenes with high range resolution. However, reconstructing the 3D surfaces from the raw single-photon waveforms is challenging, in particular when a limited number of photons is detected and when the ratio of spurious background detection events is large. This paper reviews a set of fast detection algorithms, 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 original method we recently proposed is extended here using a multiscale approach to further reduce the computational complexity of the detection process. The proposed methods are compared to state-of-the-art 3D reconstruction methods using synthetic and real single-photon data and the results illustrate their benefits for fast and robust target detection.

Keywords: Bayesian statistics, inverse problems, Lidar, detection, low-photon imaging and sensing

1. INTRODUCTION

Imaging systems based on time-of-flight light 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. As in our recent work,⁵ our aim is to address situations where 1) the photon flux originally emitted by the laser source is small, 2) where the ambient illumination level is high (i.e., low signal-to-background ratio) and 3) for which classical depth imaging methods⁶ usually provide unsatisfactory results in terms of object detection and reconstruction. As will be shown hereafter, this method can be easily extended to single-photon depth imaging in scattering media, e.g., for underwater depth imaging.^{7,8} As in most 3D reconstruction scenarios, we assume that at most one surface can be observed in each pixel although in practice, the algorithm is able to identify regions where at least one surface is visible.

The observation model includes two kinds of detection events;⁹⁻¹¹ 1) the background detection events originating from ambient (e.g., solar) illumination and dark counts from detector; and 2) the detected photons originating from the pulsed illumination and scattered back from the target (if present). Adopting a classical Bayesian approach, the pixel-wise target detection problem is formulated as a model selection problem, whereby prior distributions are assigned to each of the unknown model parameters, which are subsequently marginalized out. While in previous work⁵ we presented a post-processing step to further improve the detection maps at a low additional cost, in this work we adopt an alternative multi-resolution approach which generally reduces the number of tests to be performed and thus the computational complexity of the detection process. By aggregating

Corresponding author: Yoann Altmann: E-mail: Y.Altmann@hw.ac.uk.

Lidar waveforms in neighbouring pixels or patches, the resulting signal-to-noise ratio is expected to increase if objects are present in that region, which in turn increases locally the probability of detection. Similarly, this approach allows also large object-free regions to be identified at a lower computational cost.

In contrast to the target detection method presented in,¹⁰ which uses a reversible-jump Markov chain Monte Carlo (RJ-MCMC)^{12,13} method to generate samples according to a posterior distribution of interest and approximate the pixel-wise marginal model evidences, we reformulate the observation model such that the marginal evidences can be obtained using a more tractable marginalization procedure. More Precisely, by processing (sums of) waveforms independently, 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 on low-dimensional integration, is thus fast and can be partly 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 the alternative model parametrization used in⁵ and 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.

2. OBSERVATION MODEL

Let $\mathbf{z} = [z_1, \dots, z_T]^T \in \mathbb{Z}_+^{T \times 1}$ be a histogram of photon detection events with T temporal bins, where $\mathbb{Z}_+ = \{0, 1, \dots\}$ is the set of positive integers. Assuming that the light flux reaching the single-photon detector is sufficiently low,¹⁴ the observed photon count in a given time bin t follows the Poisson distribution

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

where $r \in \mathbb{R}_+$ is the target intensity, which combines the optical power of the pulsed source, the reflectivity of the target and the quantum efficiency of the detector). Moreover, $b \in \mathbb{R}_+$ models the background level associated with dark counts and ambient illumination, assumed constant over the T temporal bins 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), $t_0 \in 1, \dots, T$ corresponds to the typical delay/time-of-flight associated with the depth/range of the given object of the scene.

In,⁵ we defined an equivalent model using the signal-to-background ratio (SBR), which is defined as the ratio of the expected number of useful detected photons, e.g., originally emitted by the laser source, divided by the expected total number of background photons in the histogram, i.e., $w = r/(bT)$. Following this alternative parametrization, the observation model (1) can be rewritten

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

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 enables us to marginalize b analytically, as will be seen in Section 3. Assuming the T observations in \mathbf{z} are mutually independent, conditioned on their means,¹⁴ the joint likelihood can be expressed as

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

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

$$z_t | (w = 0, t_0, b) \sim \mathcal{P}(b), \quad \forall t = 1, \dots, T. \quad (4)$$

This motivates our surface detection algorithm which reduces to deciding whether $w = 0$ or $w > 0$, given \mathbf{z} . However, the background level b , and the object range and reflectivity (if an object is present) are unknown in practice. Thus, the detection task is difficult since these parameters, considered as nuisance parameters, need to be either estimated or marginalized. The next section details the proposed Bayesian strategy for this detection problem.

3. DETECTION STRATEGY

As mentioned above, we adopt a Bayesian approach and assign prior distributions to the unknown parameters using the *a priori* information available and propose to marginalize them to simplify the model selection problem.

3.1 Prior distributions

Similarly to previous work,^{5,10,11,15} 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 object, we use a Bernoulli-gamma distribution¹⁶ for the signal intensity, i.e.,

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

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 in practice, the parameters (α_r, β_r) can be adjusted from calibration measurements, since 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, as mentioned in Section 2.

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. In the results presented in Section 4, we have used $\pi = 0.5$, expressing our absence of knowledge regarding this parameter but this parameter can be user-defined.

The background level mostly depends on the amount of ambient illumination reaching the single-photon detector and is modelled as in^{10,11} with a gamma distribution $p(b|\alpha_b, \beta_b) = \mathcal{G}(b; \alpha_b, \beta_b)$ with known parameters. If a limited amount of 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 using the parametrization based on b and w can be obtained from $p(r, b|u, \alpha_r, \beta_r, \alpha_b, \beta_b) = p(r|u, \alpha_r, \beta_r)p(b|\alpha_b, \beta_b)$ by applying a standard change of variables yielding

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

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. An interesting property of the joint prior distribution (6) is that is that the conditional prior distribution $p(b|w, u, \phi)$ is a mixture of two gamma distributions, which is a conjugate prior for the model (2). Since ϕ is known in this work, it is omitted in all the conditional distributions in the remainder of this paper.

Finally, 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 can be changed if additional information is available, without affecting the computational complexity of the method.

3.2 Decision rule

Since the value of w is directly related to the value of u (i.e. $w = 0$ if $u = 0$ and $w > 0$ if $u = 1$ in (6)), the proposed decision rule is based on the marginal posterior distribution of the binary label u , obtained by integrating out the parameters b , t_0 and w , considered as nuisance parameters. Defining H_0 and H_1 as the absence and presence of the target respectively, the proposed pixel-wise decision rule is

$$p(u = 0|\mathbf{z}) \underset{H_1}{\overset{H_0}{\geq}} p(u = 1|\mathbf{z}), \quad (7)$$

where

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

with $p(w, b, t_0, u|\mathbf{z}) \propto p(\mathbf{z}|w, b, t_0)p(w, b|u)p(t_0)p(u)$ using Bayes rule. The next section discussed the computation of the integrals involved in (8).

3.3 Computation of marginals

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

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

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

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

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

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

with

$$\begin{aligned} f_1(w) &= w^{\alpha_r - 1} (\beta_b + T(1 + w(\beta_r + 1)))^{\bar{z} + \alpha_r + \alpha_b} \\ f_2(w, t_0) &= \sum_{t=1}^T z_t \log(wTh(t - t_0) + 1) \\ c_1 &= \pi \Gamma(\alpha_r) \Gamma(\bar{z} + \alpha_b + \alpha_r) (\beta_r T)^{\alpha_r}, \end{aligned}$$

and where γ is the same constant as in (10). Since γ is shared in (9) and (10), it can be easily computed using $p(u = 0|\mathbf{z}) + p(u = 1|\mathbf{z}) = 1$. 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 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 that of the classical cross-correlation method if $K \ll T$. 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 (see⁵ for details about this alternative test, not further discussed here).

3.4 Total variation (existing) approach

As mentioned above, in our recent work,⁵ we proposed to refine the pixel-wise detection method to create a more homogeneous map of target presence by solving the following total variation (TV) problem¹⁷

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

where the input image \mathbf{y} contains the log-ratios $y_{i,j} = \log p(u = 1|\mathbf{z}) - \log p(u = 0|\mathbf{z})$ 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.

While this approach can efficiently denoise the detection map, it still requires the test in (7) to be computed for each individual pixel, which might become computationally intensive for large images (and large values of T) due to the evaluation of the integrals (10). Thus, we propose an alternative solution described below to reduce the number of integral evaluations.

3.5 Multiscale approach

Histograms corresponding to neighbouring pixels generally show similar numbers of surfaces.^{18–20} Moreover, for a fixed SBR, the detection performance depends on the number of photons collected in the histogram. Thus, similarly to the multiscale approach of ManiPoP²⁰ or the unmixing algorithm in²¹, we integrate histograms in super-pixels (windows of 2×2 pixels), yielding approximately 4 times more photons and a similar SBR. In this way, we can improve the detection performance, while reducing the number of tests to be evaluated by a factor of 4 (see Fig. 1). The worst case scenario corresponds to having only one pixel out of four which contains a target, where the SBR of the super-pixel is roughly four times smaller than at the finer scale. In such cases, the probability of target presence in the super-pixel is close to 0.5 (i.e., neither presence or absence of target is certain) and a more informed decision can potentially be made in the finer scale. Hence, super-pixels with $p(u = 1) \leq 1 - \alpha$ and $p(u = 1) \geq \alpha$ are left uncertain, and reprocessed as 4 individual pixels in the finer scale. This strategy starts at a coarse level of S scales, and is repeated until a decision has been made in all super-pixels or the finest scale is reached. The confidence level α should be adjusted by the practitioner depending on the application. The number of scales should be set according to the size of the image and the expected detection detail. In all our experiments, we set $\alpha = 0.05$ and $S = 4$.

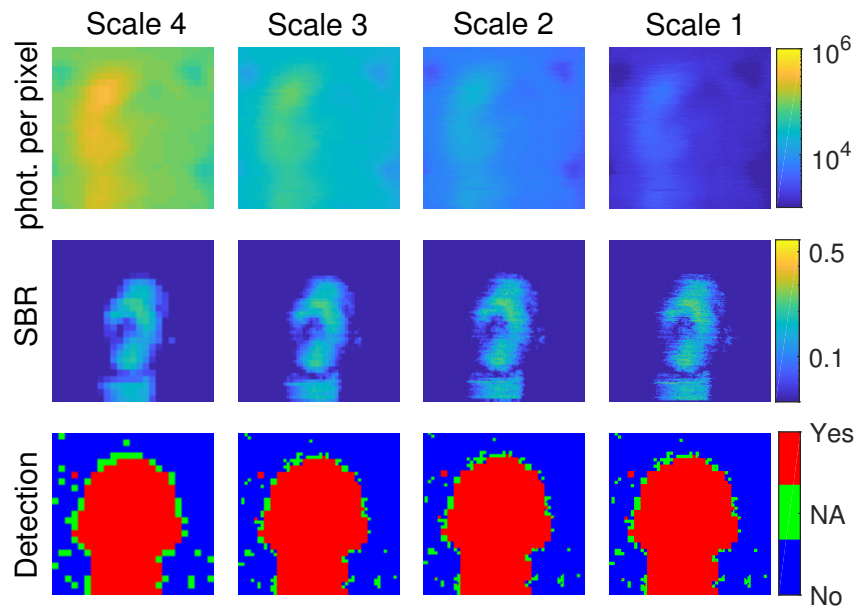


Figure 1. Integrating histograms within small windows of 2×2 pixels increases approximately 4 times the photons per histogram (top row), while reducing the amount of pixels to be processed 4 times. In most regions of the image, the SBR of the integrated waveforms does not change significantly (middle row), hence increasing the detection performance. The bottom row illustrates the proposed detection strategy, starting from the coarse scale (Scale 4) and refining the detection results sequentially in the uncertain regions (in green) using finer scales.

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 that r_M , the average number of signal photons detected when observing an object of unit reflectivity under similar observation conditions as for the scene of interest, is known. In practice, this can be obtained from calibration measurements. 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^* .

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

4.1 Synthetic data

First, we compare the three detection procedures based in the probabilities defined in (9) and (10). Comparative results with other existing methods will be presented in Section 4.2. More precisely, here we consider

- *1 scale*: Decision rule in (7) applied independently to each pixel of the scene using (9) and (10).
- *1 scale + TV*: Computation of the probabilities (9) and (10) pixel-wise, TV-denoising and thresholding procedure as described in Section 3.4.
- Multiscale: Coarse-to-fine detection procedure as described in Section 3.5, using 4 scales.

These three approaches are first compared using a set of 128×128 waveforms composed of $T = 1000$ temporal bins. The impulse response used is depicted in Fig. 2 and is based on the actual impulse response of the system used in.¹⁰ Note that to generate the synthetic data, the original impulse response has been decimated by a factor 2 and note also that only 350 temporal bins are used in Fig. 2, for visualization purpose only.

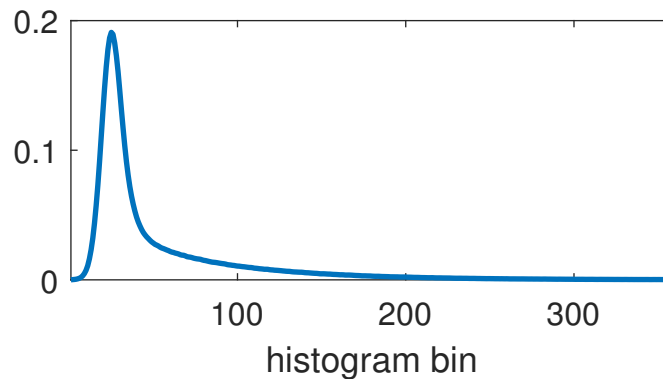


Figure 2. Impulse response used to generate the synthetic data in Section 4.1.

We generated a synthetic scene composed of a rectangular plane in the central region of the field of view and the corresponding reflectivity, background, and depth profiles generated are depicted in Fig. 3. This figure shows that the rectangular shape does not present a constant depth/range profile but rather smooth variations mimicking a direction of observation that differs from the local normal to the surface. While the background levels vary in the vertical direction, the reflectivity profile changes in the horizontal direction, which allows us to vary the SBR across pixels, as well as the overall number of detected photons. With the parameters reported in Fig. 3, we obtain an average of 7.2 photons per pixel and an average SBR of 0.13 (e.g., 13% of signal photons).




Figure 3. Synthetic lidar dataset. From left to right: reflectivity, per-pixel background photon-count, per-pixel signal-to-background ratio and ground truth depth.

Fig. 4 depicts the detection results obtained using the three approaches mentioned above. First, it can be noted that the *1-scale* approach provides a noisy detection map, with a high false alarm rate and a low detection rate in the low SBR region of the object. This can also be observed from in the results in Table 1 which reports the empirical probabilities of false alarm (PFA) and detection (PD). Conversely, the two approaches using local

information to regularize the detection problem provide less noisy detection maps and higher PDs. Moreover, while the TV regularization can lead to underestimating the object size (in particular in the low SBR region), the *multiscale* approach seems slightly more robust. Interestingly, Table 1 also shows that the proposed *multiscale* approach provides a reduction by 88% of the number of tests to be computed, when compared to the *1 scale* and *1 scale + TV* approaches, thus leading a significantly reduced computational time.

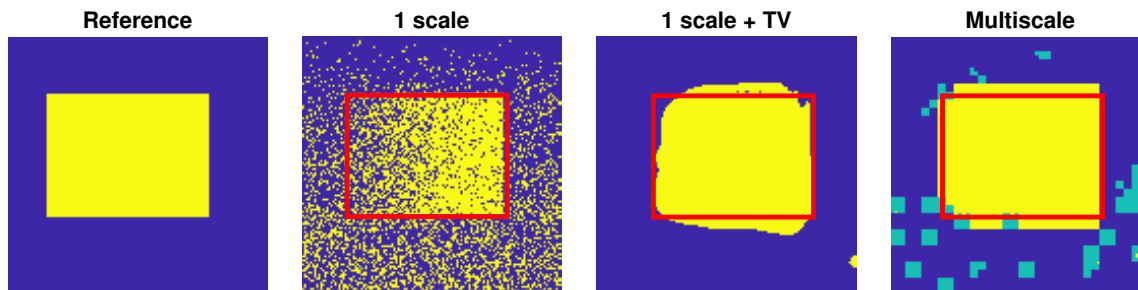


Figure 4. Target detection performance of the compared methods, target presence is indicated in yellow, while target absence is indicated in blue. The proposed multiscale detection algorithm does not reach a decision for some pixels (depicted in greenish blue). To highlight the performance the evaluated algorithms, the red bounding boxes denote the limits of the ground truth object.

Table 1. Probabilities of detection (PD), false alarm (PFA) and number of test evaluations per pixel for the evaluated algorithms on the synthetic dataset. For the multiscale approach, the uncertain regions are counted as regions where objects are present.

	PD [%]	PFA [%]	Average number of tests per pixel
1 scale	65.6	15.8	1
1 scale + TV	84.3	5.9	1
Multiscale	95.7	12.8	0.12

4.2 Real data

In this section, we compare the proposed detection strategies to two state-of-the-art single-depth¹⁰ and multi-depth²⁰ detection algorithms and the standard cross-correlation with reflectivity thresholding (see¹⁰ for details) using a real Lidar dataset, which consists of a polystyrene head measured at a stand-off distance of 325 metres during midday (more details about the acquisition can be found in¹⁰). The dataset consists of 200×200 pixels with $T = 2691$ histogram bins per pixel, an approximate SBR of 0.29 with a 5th-95th percentile interval of (0.05, 0.67). Four different per-pixel acquisition times were considered (30ms, 3ms, 1ms and 0.3ms), which corresponds to decreasing per-pixel photon counts (900 to 9 on average) while keeping the average SBR constant. Even using the longest acquisition time (30ms per pixel), the boundary of the head is difficult to clearly identify as the SBR is low due to the local orientation of the head surface. Thus, the ground truth presence map is not available and hyperparameters of the algorithms presented in^{10,20} were chosen to obtain the best visual PD/PFA trade-offs. The intensity threshold used after applying the simple cross-correlation method has been set similarly by visual inspection.

Fig. 5 depicts the detection results obtained using the three existing algorithms and the three variations of our approach, for the four acquisition times. These results first show that although *1 scale* is applied pixel-wise, it generally provides better results than the cross-correlation method which is more affected by background counts. A significant improvement in terms of probability of false alarm (PFA) and probability of detection (PD) is obtained by accounting for spatial correlation between adjacent pixels. While the method in Altmann et al.¹⁰ and MANIPOP²⁰ are prone to underestimating the size of the head, the proposed *1 scale + TV* and *multiscale* approaches seem more robust. Moreover, one of their main advantages is their computational cost, as

the method in Altmann et al.¹⁰ required 1 to 2 days per image, MANIPOPOP required 5 to 12 minutes per image, whereas the three variations of the proposed detection strategy were performed in less than 0.5s each.[†]

Table 2. Average number of test evaluations (normalized by the number of pixels in the image) for different acquisition times.

		Algorithm	
		1 scale	multiscale
Acq. time	30 ms	1	0.040
	3 ms	1	0.034
	1 ms	1	0.077
	0.3 ms	1	0.245

Finally, Table 2 illustrates how the proposed *multiscale* approach allows a general reduction of the number of tests performed to construct the detection maps in Fig. 5. While the *1 scale* and *1 scale + TV* approaches require a test per pixel, only 4% of that maximum number is used for the longest acquisition time. This gain generally reduces as the acquisition time decreases since the data become more uncertain. Still, for the shortest acquisition time (most difficult scenario), the original overall number of tests is divided by 4.

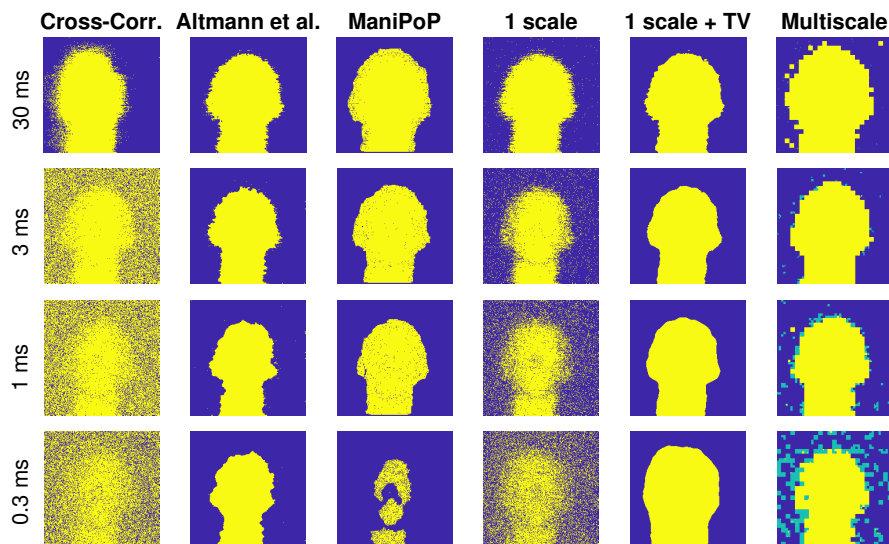


Figure 5. Target detection performance of the compared methods, target presence is indicated in yellow, while target absence is indicated in blue. The proposed multiscale detection algorithm does not reach a decision for some pixels (depicted in greenish blue).

5. CONCLUSION

We presented a set of new fast target detection algorithms for single-photon Lidar data. Unlike other existing algorithms, the proposed methods are easily parallelizable and can be used as a pre-processing step to discard histograms without useful information. We extended our previous work by proposing a novel multiscale approach in order to reduce the number of tests to be computed, which can be extremely relevant to process large images rapidly. The detection strategy can then be used to improve the reconstruction quality obtained by algorithms assuming one depth per pixel,^{21,22} as it removes histograms without surfaces from the data cube. Moreover, it can also be used before applying intensive multiple-surface-per-pixel algorithms^{20,23} to reduce their computational load.

[†]All the experiments were performed using a Matlab R2018a implementation on a i7-3.0 GHz desktop computer (16GB RAM).

ACKNOWLEDGMENTS

This work was partly supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31. The authors would also like to thank the single-photon group led by Prof. G. S. Buller at Heriot-Watt University (<http://www.single-photon.com/>), for providing the real single-photon data used in this work.

REFERENCES

- [1] Lindner, P. and Wanielik, G., “3D Lidar processing for vehicle safety and environment recognition,” in [*Proc. Work. Comput. Intell. Vehicles and Vehicular Syst. (CIVVS'09)*], 66–71 (March 2009).
- [2] Hakala, T., Suomalainen, J., Kaasalainen, S., and Chen, Y., “Full waveform hyper-spectral Lidar for terrestrial laser scanning,” *Opt. Express* **20**, 7119–7127 (Mar 2012).
- [3] Cadalli, N., Shargo, P. J., Munson, Jr., D. C., and Singer, A. C., “Three-dimensional tomographic imaging of ocean mines from real and simulated Lidar returns,” in [*Proc. SPIE, Ocean Optics: Remote Sensing and Underwater Imaging*], **4488**, 155–166 (2002).
- [4] Gao, J., Sun, J., Wei, J., and Wang, Q., “Research of underwater target detection using a slit streak tube imaging Lidar,” in [*Acad. Int. Symp. Optoelectronics and Microelectronics Technology (AISOMT)*], 240–243 (Oct 2011).
- [5] Tachella, J., Altmann, Y., McLaughlin, S., and Tourneret, J.-Y., “Fast surface detection in single-photon Lidar waveforms,” in [*Proc European Signal Processing Conference (EUSIPCO)*], (Sept. 2019). to appear.
- [6] McCarthy, A., Ren, X., Frera, A. D., Gemmell, N. R., Krichel, N. J., Scarcella, C., Ruggeri, A., Tosi, A., and Buller, G. S., “Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector,” *Opt. Express* **21**, 22098–22113 (Sept. 2013).
- [7] Maccarone, A., McCarthy, A., Ren, X., Warburton, R. E., Wallace, A. M., Moffat, J., Petillot, Y., and Buller, G. S., “Underwater depth imaging using time-correlated single-photon counting,” *Opt. Express* **23**, 33911–33926 (Dec 2015).
- [8] Halimi, A., Maccarone, A., McCarthy, A., McLaughlin, S., and Buller, G. S., “Object depth profile and reflectivity restoration from sparse single-photon data acquired in underwater environments,” *IEEE Trans. Comput. Imaging* **3**, 472–484 (Sept 2017).
- [9] Kirmani, A., Venkatraman, D., Shin, D., Colao, A., Wong, F. N. C., Shapiro, J. H., and Goyal, V. K., “First-photon imaging,” *Science* **343**(6166), 58–61 (2014).
- [10] Altmann, Y., Ren, X., McCarthy, A., Buller, G. S., and McLaughlin, S., “Robust Bayesian target detection algorithm for depth imaging from sparse single-photon data,” *IEEE Trans. Comput. Imaging* **2**, 456–467 (Dec 2016).
- [11] Altmann, Y., Ren, X., McCarthy, A., Buller, G. S., and McLaughlin, S., “Lidar waveform-based analysis of depth images constructed using sparse single-photon data,” *IEEE Trans. Image Process.* **25**(5), 1935–1946 (2016).
- [12] Green, P. J., “Reversible jump MCMC computation and Bayesian model determination,” *Biometrika* **82**, 711–732 (Dec. 1995).
- [13] Andrieu, C. and Doucet, A., “Joint Bayesian model selection and estimation of noisy sinusoids via reversible jump MCMC,” *IEEE Transactions on Signal Processing* **47**, 2667–2676 (Oct 1999).
- [14] McCarthy, A., Collins, R. J., Krichel, N. J., Fernández, V., Wallace, A. M., and Buller, G. S., “Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting,” *Appl. Opt.* **48**, 6241–6251 (Nov 2009).
- [15] Ren, X., Connolly, P. W. R., Halimi, A., Altmann, Y., McLaughlin, S., Gyongy, I., Henderson, R. K., and Buller, G. S., “High-resolution depth profiling using a range-gated cmos spad quanta image sensor,” *Opt. Express* **26**, 5541–5557 (Mar 2018).
- [16] Piironen, J. and Vehtari, A., “Sparsity information and regularization in the horseshoe and other shrinkage priors,” *Electron. J. Statist.* **11**(2), 5018–5051 (2017).
- [17] Chambolle, A. and Pock, T., “An introduction to continuous optimization for imaging,” *Acta Numerica* **25**, 161319 (2016).

- [18] Altmann, Y., Ren, X., McCarthy, A., Buller, G. S., and McLaughlin, S., “Target detection for depth imaging using sparse single-photon data,” in [*Proc. IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*], 3256–3260 (March 2016).
- [19] Hernandez-Marin, S., Wallace, A. M., and Gibson, G. J., “Multilayered 3d Lidar image construction using spatial models in a bayesian framework,” *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**, 1028–1040 (June 2008).
- [20] Tachella, J., Altmann, Y., Ren, X., McCarthy, A., Buller, G., McLaughlin, S., and Tourneret, J., “Bayesian 3D reconstruction of complex scenes from single-photon lidar data,” *SIAM Journal on Imaging Sciences* **12**(1), 521–550 (2019).
- [21] Rapp, J. and Goyal, V. K., “A few photons among many: Unmixing signal and noise for photon-efficient active imaging,” *IEEE Trans. Comput. Imaging* **3**, 445–459 (Sept 2017).
- [22] Lindell, D. B., OToole, M., and Wetzstein, G., “Single-Photon 3D Imaging with Deep Sensor Fusion,” *ACM Trans. Graph. (SIGGRAPH)* **37**(4) (2018).
- [23] Tachella, J., Altmann, Y., McLaughlin, S., and Tourneret, J. ., “3D reconstruction using single-photon lidar data exploiting the widths of the returns,” in [*Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*], 7815–7819 (May 2019).