

Joint surface detection and depth estimation from single-photon Lidar data using ensemble estimators

K. Drummond

School of Engineering and Physical Sciences
Heriot-Watt University and Leonardo UK
Edinburgh, United Kingdom
kd122@hw.ac.uk

S. McLaughlin, Y. Altmann

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

A. Pawlikowska, R. Lamb

Leonardo UK
Edinburgh, United Kingdom

Abstract—This paper addresses the problem of joint surface detection and depth estimation from single-photon Lidar (SPL) data. Traditional 3D ranging methods for SPL usually perform surface detection and range estimation sequentially to alleviate the computational burden of joint detection and estimation. Adopting a Bayesian formalism, the joint detection/estimation problem is formulated as a single inference problem. To avoid the intractable integrals usually involved with variable marginalization, we consider discrete variables and the resulting problem is recast as a model selection/averaging problem. We illustrate our method for a case where the expected signal-to-background (e.g., the target reflectivity and ambient illumination level) is unknown but the proposed framework can be adapted to more complex problems where the target depth can be obtained by combining several estimators. We demonstrate the additional benefits of the proposed method in also providing a conservative approach to uncertainty quantification of the calculated depth estimates, which can be used for real time analysis. The benefits of the proposed methods are illustrated using synthetic and real SPL data for targets at up to 8.6 km.

Index Terms—Single-photon Lidar, Bayesian estimation, Detection, Ensemble estimation.

I. INTRODUCTION

Single-Photon Lidar (SPL) is a reliable Lidar technology for readily providing 3D scene reconstruction. SPL systems are based on time-correlated single-photon counting (TCSPC) [1]–[3], whereby the laser source emits a short pulse towards the scene, part of which is reflected by the target. Once an individual photon of the reflected pulse is detected, which we refer to as a “desirable” detection event, the time interval between the pulse emission and the photon detection is recorded. This time interval is the photon’s time-of-flight (ToF) and it is logged in the corresponding time-bin of a histogram. ToF histograms can be adversely affected by “undesirable” detection events, which arise due to ambient illumination and dark counts. To provide more accurate and reliable depth estimates of the target, many pulses are repeatedly emitted towards the scene to build denser histograms and improve the signal quality.

SPL has proven to have great benefits for use in a range of different fields, including autonomous vehicles [4], agriculture [5] and defence [6]. The high sensitivity of single photon detectors allows for the use of low-power, eye-safe laser sources [6]. Furthermore, the picosecond timing resolution enables greater surface-to-surface resolution at ranges up to 200 km [7]. Thanks to advances in single-photon avalanche diode (SPAD) array technology, acquisition of data can now be achieved at video rates or higher [8], [9]. Consequently,

greater interest is being focused on faster data processing to reconstruct 3D scenes as fast and as reliably as possible. Whilst great strides have been achieved along those lines, e.g., [10], important challenges such as surface detection and reliable uncertainty quantification still need to be addressed.

In this paper, we propose a novel, pixel-wise, joint detection and depth estimation method, which detects objects/surfaces in the field of view, estimates their distance and rapidly provides uncertainty measures that can then be used in more sophisticated object recognition algorithms or subsequent decision-making processes. We use a similar observation model as described previously [11]–[14]. These methods treat the unknown model parameters (e.g., the target reflectivity) as continuous variables that are classically determined using sequential or iterative processes [15]–[18]. This can be computationally intensive, especially if the target depth and reflectivity are jointly estimated. We overcome this problem by treating the reflectivity as a discrete parameter, allowing us to perform joint detection and range estimation at a fixed (and low) computational cost. We also extend the observation model from previous works by allowing non-uniform background distributions without the method being significantly more computationally intensive. This enables the analysis of data corrupted by pile-up in SPAD detectors [19].

As multiple sources of error arise when reconstructing 3D surfaces, it is ever more important to quantify the uncertainty in depth estimation. Recent works have shown it is possible to use uncertainty quantification methods for joint depth estimation and detection, e.g., [20]. However, that method is far too slow for reconstruction at real time speeds, which motivates our work.

The remainder of this paper is organized as follows. Section II recalls the statistical observation model used for SPL and describes the proposed method for joint surface detection and depth estimation. Results of simulations conducted with synthetic single-pixel histogram data and real SPL data are presented and discussed in Section III. Conclusions are finally reported in Section IV.

II. BACKGROUND THEORY

A. Observation model

In this paper, we consider a set of K photon time of arrival (ToA) values $\mathbf{y} = \{y_k\}_{k=1}^K$, such that $y_k \in (0, T)$, where it is implicitly assumed that T is the repetition period of the

laser source [21]. Indices representing pixel dependency are omitted to simplify notation. The probability density function for a photon ToA, y_k , for a given pixel is given by

$$f(y_k|d, w) = w h_0 \left(y_k - \frac{2d}{c} \right) + (1 - w) \mathcal{V}_{(0,T)}, \quad (1)$$

where d is the range of the target surface within the admissible range gate and c is the speed of light, such that $2d/c$ is the characteristic ToF associated with the illuminated target. The function $h_0(\cdot)$ in Eq. (1) is the normalised Impulse Response Function (IRF) of the Lidar system which is not required to be Gaussian nor symmetric. It is generally measured during system calibration [12], [21]. The second term $\mathcal{V}_{(0,T)}$, assumed known, represents the distribution of undesirable background photon detection events mentioned in Section I. This need not necessarily be a uniform distribution, especially in situations of high ambient illumination conditions where pile-up effects due for instance to scattering media are more prominent. The shape of $\mathcal{V}_{(0,T)}$ arises generally due to the dominating properties of the ambient illumination over the undesirable detection events caused by dark counts, which are usually constant and relatively low compared with the ambient illumination undesirable detection events. The variable w in Eq. (1) is the probability of a detection event to be a desirable detection event. This probability is related to the signal to background ratio (SBR) by $SBR = w/(1 - w)$.

When K photons are detected, and the dead-times of the SPAD detector can be neglected, the photon ToA's are mutually independent (given d and w) and the joint likelihood can be expressed as [12], [21]

$$f(\mathbf{y}|d, w) = \prod_{k=1}^K f(y_k|d, w). \quad (2)$$

Our goal is to estimate d from Eq. (1), given that w is also unknown. Moreover, we are also interested in estimating w , as it provides information about the presence ($w > 0$) or absence ($w = 0$) of a surface, as well as its reflectivity.

B. Proposed estimation strategy

We assume that d can take a finite number N_D of values, i.e., $d \in \{d_1, \dots, d_{N_D}\}$. This set is constructed from a subset of the N_T number of original non-overlapping time bins spanning $(0, T)$ and N_D is a number of non-overlapping time bins within this subset, spanning $(\tau, T - \tau)$, such that $N_D \leq N_T$ and $\tau \geq 0$, where the width of the time bins is arbitrary (usually given by the timing resolution of the SPAD used). The parameter τ is chosen to ensure that the support of $h_0(\cdot)$ is always included in $(0, T)$ (for any admissible value of d) and thus the value of τ depends on the width of the IRF peak.

Let's first assume that w is known. The probability distribution $f(d|\mathbf{y}, w)$ can be obtained from the joint likelihood Eq. (2), using Bayes theorem

$$f(d|\mathbf{y}, w) = \frac{f(\mathbf{y}|d, w)f(d)}{f(\mathbf{y}|w)}, \quad (3)$$

where $f(d)$ is a user-defined depth prior distribution, $f(\mathbf{y}|w)$ is the marginal term and $f(\mathbf{y}|w) = \sum_{j=1}^{N_D} f(\mathbf{y}|d = d_j, w)f(d =$

$d_j)$ is a tractable normalizing constant (thanks to d being discrete). This posterior can then be used easily to compute Bayesian estimators (e.g., maximum a posteriori (MAP) or minimum mean squared error (MMSE), denoted $\mu(w)$) for the depth, as well as the posterior variance, denoted $\sigma^2(w)$. Unfortunately, depth inference using Eq. (3) is challenging since w is unknown in practice and setting its value poorly can greatly impact the depth estimation.

To alleviate this issue, a classical approach consists (assuming that w is continuous) of assigning w a prior distribution and of computing the following marginal posterior distribution

$$f(d|\mathbf{y}) = \int f(d|\mathbf{y}, w)f(w|\mathbf{y}) dw, \quad (4)$$

where $f(w|\mathbf{y})$ is the marginal posterior distribution of w . Unfortunately, manipulating $f(d|\mathbf{y})$ in Eq. (4) is challenging due to the integral w.r.t. w which needs to be approximated numerically for any value of d . To overcome this difficulty, we consider the parameter w as discrete with $w \in \{w_1, w_2, \dots, w_M\}$, where we allow $w_1 = 0$ to be in the admissible set of w and M is a user determined value. Using this discretization of w , Eq. (4) becomes

$$f(d|\mathbf{y}) = \sum_{m=1}^M f(d|\mathbf{y}, w_m)f(w_m|\mathbf{y}), \quad (5)$$

which becomes tractable, provided that $f(w_m|\mathbf{y})$ can be computed easily. The marginal posterior distribution $f(d|\mathbf{y})$ can be seen as a mixture of M distributions, whose weights are given by $f(w_m|\mathbf{y})$. Thus, its mean and variance can be easily derived from the mean and variance of each of its components. More precisely, by simplifying the notation $\mu_m = \mu(w_m)$ and $\sigma_m^2 = \sigma^2(w_m)$ for the mean and variance of $f(d|\mathbf{y}, w_m)$, the mean and variance of $f(d|\mathbf{y})$ in (5) are given by

$$\bar{\mu} = \sum_{m=1}^M f(w = w_m|\mathbf{y})\mu_m, \quad (6)$$

$$\bar{\sigma}^2 = \left(\sum_{m=1}^M f(w = w_m|\mathbf{y})(\sigma_m^2 + \mu_m^2) \right) - \bar{\mu}^2. \quad (7)$$

Since w can only take a finite number of values, the marginal posterior $f(w|\mathbf{y})$ can be computed exactly using

$$f(w|\mathbf{y}) = \frac{f(\mathbf{y}|w)f(w)}{\sum_{m=1}^M f(\mathbf{y}|w = w_m)f(w = w_m)}, \quad (8)$$

and $f(\mathbf{y}|w)$ defined below Eq. (3). While in Eq. (6), we consider the marginal mean of d a posteriori (having marginalised over w), it is of course also possible to condition the estimation of d based on the marginal MAP (MMAP) estimate of w . In that case, we can first compute

$$\hat{w} = \underset{w}{\operatorname{argmax}} f(w|\mathbf{y}), \quad (9)$$

and derive the mean and variance of $f(d|\mathbf{y}, \hat{w})$, which, as will be illustrated in Section III, often has a mean close to that of $f(d|\mathbf{y})$ but a smaller variance (since it does not account for the uncertainty associated with w). It is important to notice that the computation of the summary statistics in Eqs. (6) and (7)

requires the computation of MN_D likelihoods ($f(\mathbf{y}|d_n, w_m)$). The associated cost can thus be bounded by controlling the grids of d and w . Note also that Eq. (8), as a by-product of the depth estimation procedure, allows us to perform surface detection. For instance, we can decide that a surface is present based on $f(w > \alpha|\mathbf{y})$ where α is user-defined. More complex decision criteria can also be used as the proposed method is not based on single decision criterion. Finally, this method differs from the surface detection method detailed in [11] in which the depth estimation is performed after the detection step, whereas our proposed method performs both simultaneously.

C. Depth estimation using ensemble estimators

Eq. (5) is a mixture distribution and the resulting depth estimation method can be seen as a model averaging method, whose weights are given by $f(w_m|\mathbf{y})$. These M models are simply characterized by a different value of w_m . However, more general models can also be used. Consider a set of M models denoted $\{\mathcal{M}_m\}_m$, for which we can compute $f(d|\mathbf{y}, \mathcal{M}_m)$ and its mean and variance. Robust depth estimation via model averaging can then be achieved using

$$f(d|\mathbf{y}) = \sum_{m=1}^M \pi_m f(d|\mathbf{y}, \mathcal{M}_m), \quad (10)$$

where $\pi_m \geq 0, \forall m$ and $\sum_{m=1}^M \pi_m = 1$. These weights can be either arbitrarily defined (e.g., using $1/M$ when $f(\mathcal{M}_m|\mathbf{y})$ cannot be computed, or using $\pi_m = f(\mathcal{M}_m|\mathbf{y})$ otherwise). Due to space constraints, we only illustrate our method for models parametrized by w , but the proposed approach could also be used, for instance if several IRFs or background distributions were considered simultaneously. The former would arise for instance in long-range imaging applications when the peak broadens due to reflection on surfaces which are oblique to the beam direction. The latter can arise from varying levels of back-scattered light in scattering media. Models can also differ by the priors assigned to d (and the other model parameters such as w).

III. RESULTS

We first evaluate the performance of the proposed algorithm using synthetic single-pixel data and then two real SPL datasets, provided by Leonardo UK [6]. In all results presented, the background distribution $\mathcal{V}_{(0,T)}$ is assumed to be known, regardless of whether it is a uniform distribution or not (and estimated from the data as a pre-processing step).

A. Single-Pixel analysis

First, we generate two synthetic histograms of length $T = 1500$ bins, with $K = 100$ and $K = 1000$, respectively. The real IRF data obtained from [22] is used, with histogram resolution at 2ps per bin and $FWHM = 30$ bins = 60ps, and the maximum of the peak is set to bin 746. We set the ground truth signal photon probability to $w = 0.2$. The final pixel data generated are shown in Fig. 1 (bottom). For this initial investigation, we set $M = 20$ and the values of admissible values of w are uniformly spread in $[0, 1]$. Fig. 1 (top) shows that $f(w|\mathbf{y})$ is more concentrated as K increases since more detected photons help the discrimination between signal and

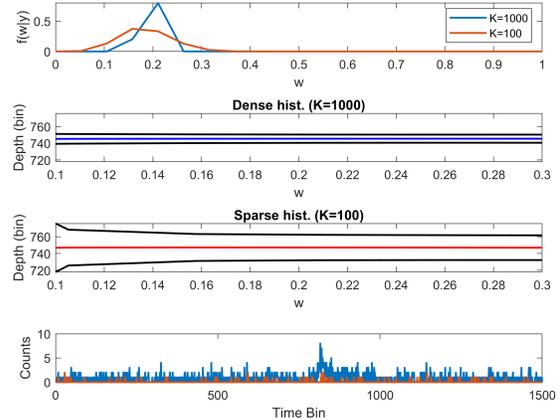


Fig. 1: Graphic results of photon dense (blue) and sparse (red) single pixel histogram data. Top: Plot of probability a posteriori $f(w|\mathbf{y})$ vs w . Middle: Plot of mean μ_m values for photon dense (Second) and sparse (Third) data, with error range plots (shown in black). Bottom: Single Pixel Histogram data plot.

| | Dense hist. ($K = 1000$) | | | Sparse hist. ($K = 100$) | | |
|---|----------------------------|-------------|-------------|----------------------------|--------------|-------------|
| | Mean of d | Var. of d | Est. of w | Mean of d | Var. of d | Est. of w |
| $f(d \mathbf{y}, \hat{w})$ using (9) | 745.44 | 2.68 | 0.21 | 747.13 | 28.55 | 0.16 |
| $f(d \mathbf{y})$ and $f(w \mathbf{y})$ | 745.43 | 2.74 | 0.20 | 747.02 | 62.50 | 0.19 |

TABLE I: Comparison of the different estimates of d (conditioned on w or not) and w (marginal MAP or marginal MMSE), for $K = 1000$ and $K = 100$. The actual value of (d, w) is $(746, 0.2)$. Note that $f(d|\mathbf{y})$ and $f(w|\mathbf{y})$ are computed using (5) and (8), respectively.

background photons. Similarly, the second and third row of Fig. 1, which depict $\mu(w = w_m)$ and $\sigma^2(w = w_m)$, illustrate how the estimated depth mean and variance using $f(d|\mathbf{y}, w)$ depend on K and on the unknown value of w (plots restricted to $w > 0.1$ below which the means degrade drastically).

Table I summarizes the different estimates of d and w for the two histograms with $K = 100$ and $K = 1000$. First, we can note that the marginal MMSE estimator of w (i.e., the mean of $f(w|\mathbf{y})$ used in the bottom row of Table I) is usually more reliable than the MMAP estimator in (9) which is more sensitive to the resolution of the w -grid (and M), especially for small values of K . Second, the MMSE depth estimates, conditioned on the MMAP estimator of w (top row) or computed from the marginal posterior $f(d|\mathbf{y})$ (bottom row) are similar. However, the estimated variance is larger in the latter case, since it incorporates the uncertainty about this unknown parameter w . This estimator is thus more conservative in terms of uncertainty quantification.

B. Real SPL data analysis

Here, we use two real SPL datasets acquired by Leonardo to illustrate the potential benefits of the proposed method. The

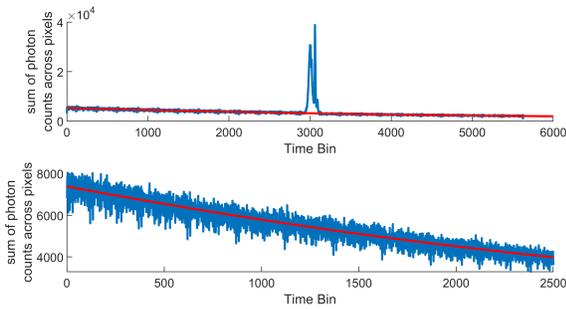


Fig. 2: Sum of photon counts (blue) across all Lidar pixels for college tower data (top) and pylon data (bottom). The red curves represent the estimated background contributions.

first SPL cube consists of 100×50 pixels and focuses on the tower of an Edinburgh college tower, taken at $\approx 3\text{km}$ range, already considered in [6]. The second dataset consists of 80×40 pixels representing a pylon, taken at $\approx 8.6\text{km}$ range, and is an example of a challenging, line-based object. These datasets primarily differ by their average SBR, much higher in the first dataset ($SBR \approx 0.22$) than in the second ($SBR \approx 0.01$). It should be noted the SBR described is not the peak SBR, where the signal/peak reaches its maximum value, but rather the ratio of desirable to undesirable detection events from all histogram temporal bins. The difference in the SBR of both datasets can be seen in Fig. 2, which depicts two histograms obtained by accumulating the histograms of all the pixels in the set. These accumulated histograms can be used to estimate an average background distribution which is not constant in these two cases. Here, we used second-order polynomials to fit $\mathcal{V}_{(0,T)}$ and compared the results obtained assuming constant background instead. The background for the SBR is equivalent to the integral of the $\mathcal{V}_{(0,T)}$ distribution curve over all the temporal bins.

For both datasets, our detection method is compared to the detection method proposed by Tachella et al. [11], assuming a constant background distribution. It is worth recalling that estimates of w can be used to estimate to the target intensity (number of signal photons $I = wK$) and the number of background photons $B = (1 - w)K$ and we used \bar{w} , the mean of $f(w|\mathbf{y})$ in these expressions, leading to \bar{I} and \bar{B} . The proposed method has been applied with larger M values than for the single pixel analysis for a more precise estimation of w . A target is assumed to be present in each pixel if and only if $f(w > w_0) > 0.5$, where w_0 is user-defined and scene dependent. Due to space constraints, the final depth variance results are not presented but their scale is of the order the timing resolution of the SPAD used.

1) *College Tower Data*: For this dataset, target detection is a relative simple task and the admissible grid of w is set using $M = 100$ equality spaced $\{w_m\}_m$ and we used $w_0 = 0.02$. The estimated probability of target presence maps are displayed in Fig. 3, which shows that for this scene, the proposed method leads to results similar to that using Tachella's method [11], irrespective of the background model

adopted. This can be explained by the fact that the peaks in the histograms can be easily identified even when assuming a constant background.

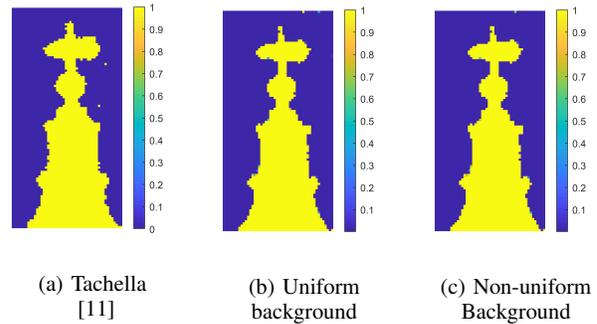


Fig. 3: College tower data comparison of probability of detection results for the method by Tachella et al. (a), and for the proposed method using uniform (b) and non uniform (c) background distributions, where $w_0 = 0.02$.

The final mean depth estimates $\bar{\mu}$, reflectivity estimates \bar{I} and background estimates \bar{B} are presented in Fig. 4.

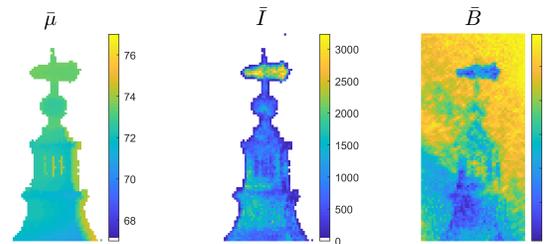


Fig. 4: Final mean depth (left), reflectivity (middle) and background (right) estimates for the tower data using the proposed method, for $w_0 = 0.02$.

2) *Pylon Data*: For this dataset for which we know the SBR is low, the admissible grid of w is set using $M = 200$ logarithmically spaced $\{w_m\}_m$, with $w_2 = 10^{-5}$, $w_M = 10^{-1}$ and $w_0 = 0.008$. The estimated target presence maps are displayed in Fig. 5, which shows that for this scene, the proposed method leads to results noisier than those obtained by Tachella et al. [11] (taking the conjugate gamma density shape parameter $\alpha_b = 100$) when we assume a uniform background distribution, and leads to improved results for detecting the pylon structure when we assume a non-uniform background distribution. In this instance, the peaks in the histograms cannot be easily identified when assuming a constant background, and so a non-uniform distribution assumption is required to obtain better results from the cross-correlation calculations.

The final mean depth, reflectivity and background estimates obtained using our proposed method are shown in Fig. 6, under the condition $w_0 = 0.008$ and $\bar{I} > 0.25$.

IV. CONCLUSION

In this paper, we proposed a novel method for joint surface detection and depth estimation from SPL data using discrete

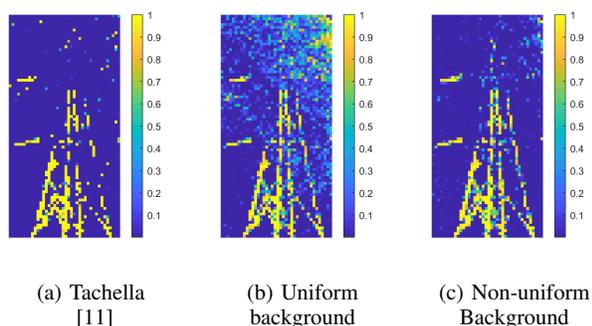


Fig. 5: Pylon data comparison of probability of detection results for Tachella et al. method (a), where $\alpha_b = 100$, and proposed method using uniform (b) and non uniform (c) background distributions, where $w_0 = 0.008$.

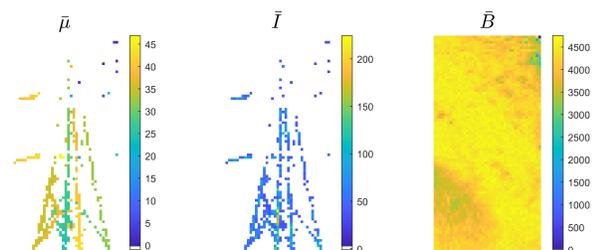


Fig. 6: Final mean depth (left), reflectivity (middle) and background (right) estimates for the pylon data using the proposed method, for $w_0 = 0.008$ and $\bar{I} > 0.25$.

variables to avoid intractable marginalizations and producing satisfactory results in the final estimates using model selection/averaging. While we illustrated the method for a case where only the expected signal-to-background is unknown, in the future we aim to adapt the framework to more complex problems where the target depth can be obtained combining several arbitrary estimators. Furthermore we plan to propose a GPU implementation to enable reliable depth estimation and uncertainty quantification at real-time speeds.

ACKNOWLEDGEMENTS

This work was supported by CENSIS and Leonardo UK. This work was supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31, the Engineering and Physical Sciences Research Council of the UK (EPSRC) Grant number EP/S000631/1 and the UK MOD University Defence Research Collaboration (UDRC) in Signal Processing.

REFERENCES

- [1] G. Buller and R. Collins, "Single-photon generation and detection," *Measurement Science and Technology*, vol. 21, no. 1, 2010.
- [2] A. McCarthy, X. Ren, A. D. Frera, N. R. Gemmill, N. J. Krichel, C. Scarcella, A. Ruggeri, A. Tosi, and G. S. Buller, "Kilometer-range depth imaging at 1550 nm wavelength using an ingaas/inp single-photon avalanche diode detector," *Opt. Express*, vol. 21, no. 19, pp. 22098–22113, Sep 2013. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-21-19-22098>
- [3] W. Becker, *Advanced Time-Correlated Single-Photon Counting Techniques*, ser. Springer Series in Chemical Physics. Springer, 2005.

- [4] J. Rapp, J. Tachella, Y. Altmann, S. McLaughlin, and V. K. Goyal, "Advances in single-photon Lidar for autonomous vehicles," *IEEE Signal Processing Magazine*, vol. 37, no. 4, 2020.
- [5] P. Yimyam and A. F. Clark, "3D reconstruction and feature extraction for agricultural produce grading," *8th International Conference on Knowledge and Smart Technology (KST)*, Chiangmai, no. 136-141, 2016.
- [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. 11919–11931, 2017.
- [7] Z.-P. Li, J.-T. Ye, X. Huang, P.-Y. Jiang, Y. Cao, Y. Hong, C. Yu, J. Zhang, Q. Zhang, C.-Z. Peng, F. Xu, and J.-W. Pan, "Single-photon imaging over 200km," *Optica*, vol. 8, no. 3, pp. 344–349, Mar 2021. [Online]. Available: <http://www.osapublishing.org/optica/abstract.cfm?URI=optica-8-3-344>
- [8] M. Entwistle, M. A. Itzler, J. Chen, M. Owens, K. Patel, X. Jiang, K. Slomkowski, and S. Rangwala, "Geiger-mode APD camera system for single-photon 3D LADAR imaging," in *Advanced Photon Counting Techniques VI*, M. A. Itzler, Ed., vol. 8375, June 2012, p. 83750D.
- [9] R. Henderson, N. Johnston, H. Chen, D. Li, G. Hungerford, R. Hirsch, P. Yip, D. McLoskey, and D. Birch, "A 192 x 128 time correlated single photon counting imager in 40nm CMOS technology," in *44th European Solid-State Circuits Conference (ESSCIRC) 2018*. IEEE, Sept. 2018.
- [10] J. Tachella, Y. Altmann, N. Mellado, A. McCarthy, R. Tobin, G. S. Buller, J.-Y. Tourneret, and S. McLaughlin, "Real-time 3D reconstruction from single-photon Lidar data using plug-and-play point cloud denoisers," *Nature Communications*, vol. 10, no. 1, Nov 2019. [Online]. Available: <http://dx.doi.org/10.1038/s41467-019-12943-7>
- [11] J. Tachella, Y. Altmann, S. McLaughlin, and J. Y. Tourneret, "Fast surface detection in single-photon Lidar waveforms," in *2019 27th European Signal Processing Conference (EUSIPCO)*, 2019, pp. 1–5.
- [12] Q. Legros, J. Tachella, R. Tobin, A. McCarthy, S. Meignen, G. S. Buller, Y. Altmann, S. McLaughlin, and M. E. Davies, "Robust 3D reconstruction of dynamic scenes from single-photon Lidar using beta-divergences," *IEEE Transactions on Image Processing*, vol. 30, p. 1716–1727, 2021. [Online]. Available: <http://dx.doi.org/10.1109/TIP.2020.3046882>
- [13] Y. Altmann, S. McLaughlin, and M. E. Davies, "Fast online 3D reconstruction of dynamic scenes from individual single-photon detection events," *IEEE Transactions on Image Processing*, vol. 29, pp. 2666–2675, 2020.
- [14] J. Rapp and V. K. Goyal, "A few photons among many: Unmixing signal and noise for photon-efficient active imaging," *IEEE Transactions on Computational Imaging*, vol. 3, no. 3, pp. 445–459, 2017.
- [15] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3D and reflectivity imaging with single-photon detectors," *IEEE Trans. Comput. Imaging*, vol. 1, no. 2, pp. 112–125, Jun 2015.
- [16] Y. Altmann and S. McLaughlin, "Range estimation from single-photon Lidar data using a stochastic EM approach," in *2018 26th European Signal Processing Conference (EUSIPCO)*, ser. European Signal Processing Conference (EUSIPCO). United States: IEEE, Dec. 2018, pp. 1112–1116.
- [17] Y. Altmann, X. Ren, A. McCarthy, G. Buller, and S. McLaughlin, "Lidar waveform-based analysis of depth images constructed using sparse single-photon data," *IEEE Transactions on Image Processing*, vol. 25, no. 5, pp. 1935–1946, May 2016.
- [18] 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>
- [19] A. Gupta, A. Ingle, A. Velten, and M. Gupta, "Photon-flooded single-photon 3D cameras," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [20] Y. Altmann, A. Maccarone, A. Halimi, A. McCarthy, G. Buller, and S. McLaughlin, "Efficient range estimation and material quantification from multispectral Lidar waveforms," in *2016 Sensor Signal Processing for Defence (SSPD)*, 2016, pp. 1–5.
- [21] Q. Legros, S. McLaughlin, Y. Altmann, S. Meignen, and M. E. Davies, "Robust depth imaging in adverse scenarios using single-photon Lidar and beta-divergences," in *2020 Sensor Signal Processing for Defence Conference (SSPD)*, 2020, pp. 1–5.
- [22] Y. Altmann, X. Ren, A. McCarthy, G. Buller, and S. McLaughlin, "Robust bayesian target detection algorithm for depth imaging from sparse single-photon data," *IEEE Transactions on Computational Imaging*, vol. 2, no. 4, pp. 456–467, Dec. 2016, arXiv admin note: text overlap with arXiv:1507.02511.