



HAL
open science

Real-time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers

Julian Tachella, Yoann Altmann, Nicolas Mellado, Aongus Mccarthy, Rachel Tobin, Gerald Stuart Buller, Jean-Yves Tournéret, Stephen Mclaughlin

► **To cite this version:**

Julian Tachella, Yoann Altmann, Nicolas Mellado, Aongus Mccarthy, Rachel Tobin, et al.. Real-time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers. *Nature Communications*, 2019, 10, pp.4984. 10.1038/s41467-019-12943-7 . hal-02306826

HAL Id: hal-02306826

<https://hal.science/hal-02306826>

Submitted on 19 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Real-Time 3D reconstruction from single-photon lidar data using plug-and-play point cloud denoisers

Julián Tachella¹, Yoann Altmann¹, Nicolas Mellado², Aongus McCarthy¹, Rachael Tobin¹, Gerald S. Buller¹, Jean-Yves Tourneret³ and Stephen McLaughlin¹

¹School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK.

²IRIT, CNRS, University of Toulouse, Toulouse, France

³ENSEEIH-IRIT-TeSA, University of Toulouse, Toulouse, France.

Single-photon lidar has emerged as a prime candidate technology for depth imaging through challenging environments. Until now, a major limitation has been the significant amount of time required for the analysis of the recorded data. Here we show a new computational framework for real-time three-dimensional (3D) scene reconstruction from single-photon data. By combining statistical models with highly scalable computational tools from the computer graphics community, we demonstrate 3D reconstruction of complex outdoor scenes with processing times of the order of 20 ms, where the lidar data was acquired in broad daylight from distances up to 320 metres. The proposed method can handle an unknown number of surfaces in each pixel, allowing for target detection and imaging through cluttered scenes. This enables robust, real-time target reconstruction of complex moving scenes, paving the way for single-photon lidar at video rates for practical 3D imaging applications.

Introduction

Reconstruction of three-dimensional (3D) scenes has many important applications, such as autonomous navigation [1], environmental monitoring [2] and other computer vision tasks [3]. While geometric and reflectivity information can be acquired using many scanning modalities (e.g., RGB-D sensors [4], stereo imaging [5] or full waveform lidar [2]), single-photon systems have emerged in recent years as an excellent candidate technology. The time-correlated single-photon counting (TCSPC) lidar approach offers several advantages: the high sensitivity of single-photon detectors allows for the use of low-power, eye-safe laser sources; and the picosecond timing resolution enables excellent surface-to-surface resolution at long range (hundreds of metres to kilometres) [6]. Recently, the TCSPC technique has proved successful at reconstructing high resolution three-dimensional images in extreme environments such as through fog [7], with cluttered targets [8], in highly scattering underwater media [9], and in free-space at ranges greater than 10km [6]. These applications have demonstrated the potential of the approach with relatively slowly scanned optical systems in the most challenging optical scenarios, and image reconstruction provided by post-processing of the data. However, recent advances in arrayed SPAD technology now allow rapid acquisition of data [10, 11], meaning that full-field three-dimensional (3D) image acquisition can be achieved at video rates, or higher, placing a severe bottleneck on the processing of data.

Even in the presence of a single surface per transverse pixel, robust 3D reconstruction of outdoor scenes is challenging due to the high ambient (solar) illumination and the low signal return from the scene. In these scenarios, existing approaches are either too slow or not robust enough and thus do not allow rapid analysis of dynamic scenes, such as road activity monitoring (e.g., pedestrian detection), and subsequent automated decision-making processes. In this work, we propose a new algorithm structure, differing significantly from existing approaches, to meet speed, robustness and scalability requirements. Existing computational imaging approaches can generally be divided into two families of methods. The first family assumes the presence of a single surface per observed pixel, which greatly simplifies the reconstruction problem as classical image reconstruction tools can be used to recover the range and reflectivity profiles. These algorithms address the 3D reconstruction by using some prior knowledge about these images. For instance, the algorithms introduced in [12, 13] propose a hierarchical Bayesian model and compute estimates using samples generated by appropriate Markov Chain Monte Carlo (MCMC) methods. Despite providing robust 3D reconstructions with limited user supervision (where limited critical parameters are user-defined), these intrinsically iterative methods suffer from a high computational cost (several hours per reconstructed image). Faster alternatives based on convex optimisation tools and spatial regularisation, such as [14, 15, 16], have been proposed for 3D reconstruction but they often require supervised parameter tuning and still need to run several seconds to minutes to converge for a single image. The authors in [17] proposed a parallel optimization algorithm, but still reported reconstruction times of the order of seconds. Even the recent algorithm [18] based on convolutional neural network (CNN) to estimate the scene depth, using a passive 2D image as additional information, does not meet real-time requirements after training.

Although the single-surface per pixel assumption greatly simplifies the reconstruction problem, it does not hold for complex scenes, for example with cluttered targets, and long-range scenes with larger target footprints. Hence, a second family of methods has been proposed to handle multiple surfaces per pixel

[19, 15, 20, 21]. In this context, 3D reconstruction is significantly more difficult as the number of surfaces per pixel is not a priori known. The earliest methods [21] were based on Bayesian models and so-called reversible-jump MCMC methods (RJ-MCMC) and were mostly designed for single-pixel analysis. Faster optimisation-based methods have also been proposed [19, 15], but the recent ManiPoP algorithm [20] combining RJ-MCMC updates with spatial point processes has been shown to provide more accurate results with a similar computational cost.

To the best of our knowledge, all state-of-the-art multiple-surface algorithms involve execution times of the order of seconds or minutes, which are prohibitively long for any real-time processing of lidar array data. Here we propose a new algorithm that can process dozens of frames per second, achieving state-of-the-art reconstructions in the general multiple-surface per pixel setting. The novel method efficiently models the target surfaces as two dimensional manifolds embedded in a 3D space. This is achieved using manifold modelling and point cloud denoising tools from the computer graphics community (see [22] for a complete survey). A typical computer graphics pipeline for 3D reconstruction consists of a 2-step process using a simple maximum likelihood algorithm to find a rough (initial) point cloud estimate from the data and then a second step using a point cloud denoising algorithm capable of rapidly processing millions of points. This strategy is efficient when the initial point cloud is dense and moderately noisy but it does not provide satisfactory results here due to the general poor quality of the initial point cloud (extracted from single-photon lidar waveforms) and the relatively small number of pixels of current SPAD arrays - e.g., the 32×32 array of the Kestrel Princeton Lightwave camera used in this study. Moreover, this strategy does not take into account *a priori* information on the observation model, such as presence of dead pixels [23, 24] or compressive sensing strategies [25, 26]. On the other hand, most image processing approaches work with the full lidar waveforms and use accurate observation models but their denoising tools are not tailored for point clouds, as mentioned above.

In this work we propose a new inference scheme which benefits from the best of each strategy. This 3D reconstruction algorithm works directly on the lidar waveforms, making use of off-the-shelf computer graphics point cloud denoisers for distributed surfaces. We extend and adapt the ideas of plug-and-play priors [27, 28, 29, 30] and regularisation by denoising [31, 32], which have recently appeared in the image processing community, to point cloud restoration. The proposed algorithm iterates gradient descent steps, which take into account the observation model, and denoising steps, which benefit from powerful point cloud denoising techniques from the computer graphics literature (e.g., [33]). The resulting algorithm achieves real-time processing of 3D lidar video due to the intrinsic parallel architecture of the gradient evaluation steps and of the point cloud denoising strategy. The proposed method can be easily implemented for general purpose graphical processing units (GPGPU) [34], and thus is compatible with use in modern embedded systems (e.g., self-driving cars).




Figure 1: Illustration of a single-photon lidar dataset containing two surfaces (the man behind camouflage from [15]). The graph on the left shows the histogram of a given pixel. The limited number of collected photons and the high background level makes the reconstruction task very challenging. In this case, processing the pixels independently yields poor results, but they can be improved by considering *a priori* knowledge about the scene’s structure.

Results

Observation model

A lidar data cube of $N_r \times N_c$ pixels and T histogram bins is denoted by \mathbf{Z} , where the photon-count recorded in pixel (i, j) and histogram bin t is $[\mathbf{Z}]_{i,j,t} = z_{i,j,t} \in \mathbb{Z}_+ = \{0, 1, 2, \dots\}$. We represent a 3D point cloud by a set of N_{Φ} 3D points $\Phi = \{(\mathbf{c}_n, r_n) \mid n = 1, \dots, N_{\Phi}\}$, where $\mathbf{c}_n \in \mathbb{R}^3$ is the point location in real-world coordinates and $r_n \in \mathbb{R}_+$ is the intensity (unnormalised reflectivity) of the point. A point \mathbf{c}_n is mapped into the lidar data cube according to the function $f(\mathbf{c}_n) = [i, j, t_n]^T$, which takes into account the camera parameters of the lidar system, such as depth resolution and focal length, and other characteristics, such as super-resolution or spatial blurring. For ease of presentation, we also denote the set of lidar depths values by $\mathbf{T} = [t_1, \dots, t_{N_{\Phi}}]^T$ and the set of intensity values by $\mathbf{R} = [r_1, \dots, r_{N_{\Phi}}]^T$. Under the classical assumption [35, 14] that the incoming light flux incident on the TCSPC detector is very low, the observed photon-counts can be accurately modelled by a linear mixture of signal and background photons corrupted by Poisson noise. More precisely, the data likelihood which models how the observations \mathbf{Z} relate to the model parameters can be expressed as

$$z_{i,j,t} | (\mathbf{T}, \mathbf{R}, b_{i,j}) \sim \mathcal{P} \left(\sum_{N_{i,j}} g_{i,j} r_n h_{i,j}(t - t_n) + g_{i,j} b_{i,j} \right) \quad (1)$$

where $t \in \{1, \dots, T\}$, $h_{i,j}(\cdot)$ is the known (system-dependent) per-pixel temporal instrumental response, $b_{i,j}$ is the background level in present in pixel (i, j) and $g_{i,j}$ is a scaling factor that represents the gain/sensitivity

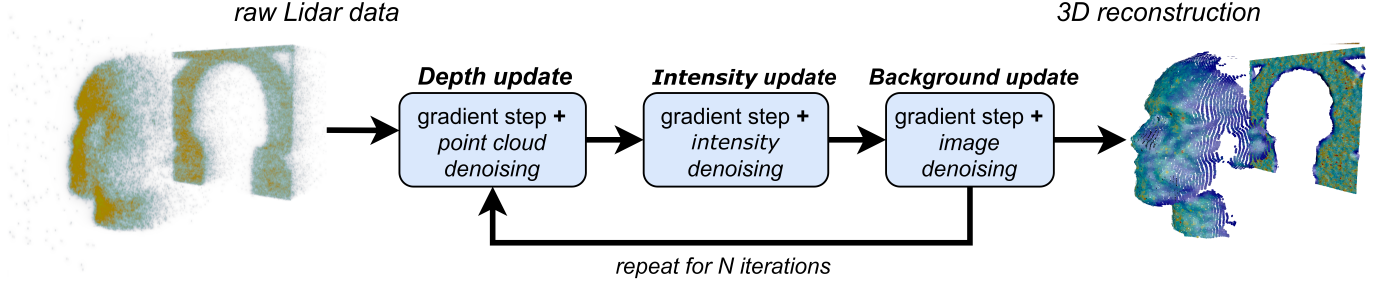


Figure 2: Block diagram of the proposed real-time framework. The algorithm iterates between depth, intensity and background updates, applying a gradient step followed by a denoiser. Each step can be processed very quickly in parallel, resulting in a low total execution time.

of the detector. The set of indices $\mathcal{N}_{i,j}$ correspond to the points (c_n, r_n) that are mapped into pixel (i, j) . Figure 1 shows an example of a collected depth histogram.

Assuming mutual independence between the noise realizations in different time bins and pixels, the negative log-likelihood function associated with the observations $z_{i,j,t}$ can be written as

$$g(\mathbf{T}, \mathbf{R}, \mathbf{B}) = - \sum_{i=1}^{N_c} \sum_{j=1}^{N_r} \sum_{t=1}^T \log p(z_{i,j,t} | \mathbf{T}, \mathbf{R}, b_{i,j}) \quad (2)$$

where $p(z_{i,j,t} | \mathbf{T}, \mathbf{R}, b_{i,j})$ is the probability mass associated with the Poisson distribution in (1). This function contains all the information associated with the observation model and its minimisation equates to maximum likelihood estimation (MLE). However, MLE approaches are sensitive to data quality and additional regularisation is required, as discussed below.

Reconstruction algorithm

The reconstruction algorithm follows the general structure of PALM [36], computing proximal gradient steps on the blocks of variables \mathbf{T} , \mathbf{R} and \mathbf{B} . Each update first adjusts the current estimates with a gradient step taken with respect to the log-likelihood (data-fidelity) term $g(\mathbf{T}, \mathbf{R}, \mathbf{B})$, followed by an off-the-shelf denoising step, which plays the role of a proximal operator [37]. While the gradient step takes into account the single-photon lidar observation model (i.e., Poisson statistics, presence of dead pixels, compressive sensing, etc.), the denoising step profits from off-the-shelf point cloud denoisers. A summary of each block update is presented below, whereas an in-detail explanation of the full algorithm can be found in the supplementary notes.

Depth update: A gradient step is taken with respect to the depth variables \mathbf{T} and the point cloud Φ is denoised with the algebraic point set surfaces (APSS) algorithm [33, 38] working in the real-world coordinate system. APSS fits a smooth continuous surface to the set of points defined by \mathbf{T} , using spheres as local primitives. The fitting is controlled by a kernel, whose size adjusts the degree of low-pass filtering of the surface. In contrast to conventional depth image regularisation/denoisers, the point cloud denoiser

can handle an arbitrary number of surfaces per pixel, regardless of the pixel format of the lidar system. Moreover, all of the 3D points are processed in parallel, equating to very low execution times.

Intensity update: In this update, the gradient step is taken with respect to \mathbf{R} , followed by a denoising step using the manifold metrics defined by Φ in real-world coordinates. In this way, we only consider correlations between points within the same surface. A low-pass filter is applied using the nearest neighbours of each point, as in [39]. This step also processes all the points in parallel, only accounting for local correlations. After the denoising step, we remove the points with intensity lower than a given threshold, which is set as the minimum admissible reflectivity (normalized intensity).

Background update: In a similar fashion to the intensity and depth updates, a gradient step is taken with respect to \mathbf{B} . Here, the proximal operator depends on the characteristics of the lidar system. In bi-static raster-scanning systems, the laser source and single-photon detectors are not co-axial and background counts are not necessarily spatially correlated. Consequently, no spatial regularisation is applied to the background. In this case, the denoising operator reduces to the identity, i.e., no denoising. In mono-static raster-scanning systems and lidar arrays, the background detections resemble a passive image. In this case, spatial regularisation is useful to improve the estimates. Thus, we replace the proximal operator with an off-the-shelf image denoising algorithm. Specifically, we choose a simple denoiser based on the fast Fourier transform (FFT), which has low computational complexity.

Large raster-scan scene results

A life-sized polystyrene head was scanned at a stand-off distance of 40 metres using a raster-scanning lidar system (a detailed explanation of the scene can be found in [12]). The data cuboid has size $N_r = N_c = 141$ pixels and $T = 4613$ bins, with a binning resolution of 0.3 mm. A total acquisition time of 1 ms was used for each pixel, yielding a mean of 3 photons per pixel with a signal-to-background ratio of 13. The scene consists mainly of one surface per pixel, with 2 surfaces per pixel around the borders of the head. Figure 3 shows the results for the proposed method, the standard maximum-likelihood estimator and two state-of-the-art algorithms assuming a single [16] or multiple [20] surfaces per pixel. Within a maximum error of 4 cm, the proposed method finds 96.6% of the 3D points, which improves the results of cross-correlation [35], which finds 83.46%, and also performs slightly better than [16] and ManiPoP [20], which find 95.2% and 95.23%, respectively. The most significant difference is the processing time of each method: the novel algorithm only takes 13 ms to process the entire frame, whereas [16] and ManiPoP require 201 s and 37 s, respectively. Whereas a parallel implementation of cross-correlation will almost always be faster than a regularised algorithm (requiring only 1 ms per lidar frame), the execution time of the proposed method only incurs a small overhead cost while significantly improving the reconstruction quality of single-photon data. The performance of the algorithm was also validated in other raster-scanned scenes, which are detailed in the supplementary material.




Figure 3: Comparison of 3D reconstruction methods. Reconstruction results of (a) cross-correlation, (b) Rapp and Goyal [16], (c) ManiPoP [20] and (d) the new method. The colour bar scale depicts the number of returned photons from the target assigned to each 3D point. Cross-correlation does not include any regularisation, yielding noisy estimates, whereas the results of [16], ManiPoP and the proposed method show structured point clouds. The method of [16] correlates the borders of the polystyrene head and the backplane (as it assumes 1 surface per pixel), whereas ManiPoP and the proposed method do not promote correlations between them.

3D Dynamic scenes results

To demonstrate the real-time processing capabilities of the proposed algorithm, we acquired, using the Kestrel Princeton Lightwave camera, a series of 3D videos with a single-photon array of $N_r = N_c = 32$ pixels and $T = 153$ histogram bins (binning resolution of 3.75 cm), which captures 150,400 binary frames per second. As the pixel resolution of this system is relatively low, we followed a super-resolution scheme, estimating a point cloud of $N_r = N_c = 96$ pixels. This can be easily achieved by defining an undersampling operation in $f(\cdot)$, which maps a window of 3×3 points in the finest resolution (real-world coordinates) to a single pixel in the coarsest resolution (lidar coordinates). As processing a single lidar frame with the novel method takes 20 ms, we integrated the binary acquisitions into 50 lidar frames per second (i.e., real-time acquisition and reconstruction). At this frame rate, each lidar frame is composed of 3008 binary frames.

Figure 4 shows the imaging scenario, which consists of two people walking between a camouflage net and a backplane at a distance of approximately 320 metres from the lidar system. Each frame has approximately 900 photons per pixel, where 450 photons are due to target returns and the rest are related to dark counts or ambient illumination from solar background. Most pixels present two surfaces, except for those in the left and right borders of the camouflage, where there is only one return per pixel. A maximum number of three surfaces per pixel can be found in some parts of the contour of the human targets. The complete 3D video, together with a reference RGB video, is included in the supplementary material. Additional 3D videos, including indoor scenes, can also be found in the supplementary material.

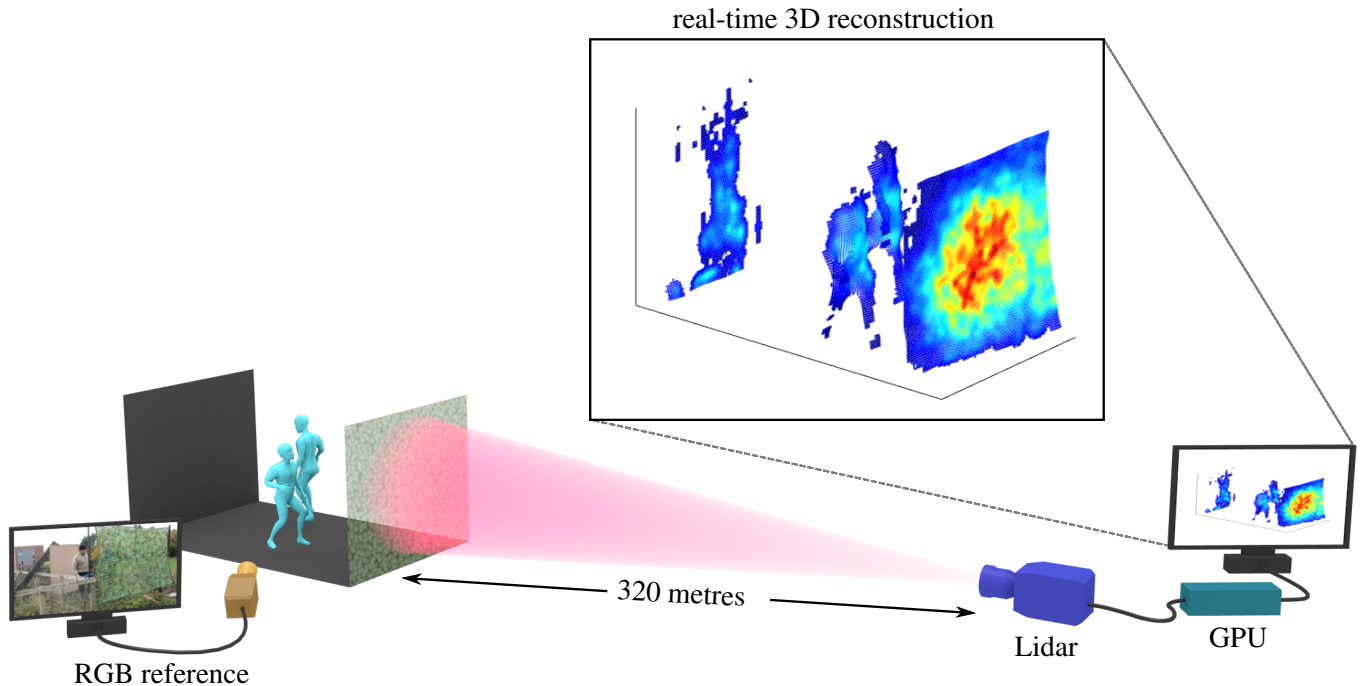


Figure 4: Schematic of the 3D imaging experiment. The scene consists of two people walking behind a camouflage net at a stand-off distance of 320 metres from the lidar system. An RGB camera was positioned a few metres from the 3D scene and used to acquire a reference video. The proposed algorithm is able to provide real-time 3D reconstructions using a GPU. As the lidar presents only $N_r = N_c = 32$ pixels, the point cloud was estimated in a higher resolution of $N_r = N_c = 96$ pixels. The full 3D video can be found in the supplementary material.

Discussion

We have proposed a new real-time 3D reconstruction algorithm, which is able to obtain reliable estimates of distributed scenes using very few photons and/or in the presence of spurious detections. The proposed method does not make any strong assumptions about the 3D surfaces to be reconstructed, allowing an unknown number of surfaces to be present in each pixel. We have demonstrated similar or better reconstruction quality than other existing methods, while improving the execution speed by a factor up to 10^5 . We have also demonstrated reliable real-time 3D reconstruction of scenes with multiple surfaces per pixel at long distance (320 m) and high frame rates (50 frames per second) in daylight conditions. Minimal operating conditions (i.e., minimum signal-to-background ratio and photons per pixel required to ensure good reconstruction with high probability) are discussed in the supplementary material. The novel algorithm combines *a priori* information on the observation model (sensor statistics, dead pixels, sensitivity of the detectors, etc.) with powerful point cloud denoisers from computer graphics literature, outperforming methods based solely on computer graphics or image processing techniques. Moreover, we have shown that the observation model can be easily modified to perform super-resolution. It is worth noting that the proposed model could

also be applied to other scenarios, e.g., involving spatial deblurring due to highly scattering media. While we have chosen the APSS denoiser, the generality of our formulation allows us to use many point cloud (depth and intensity) and image (background) denoisers as building blocks to construct other variants. In this way, we can control the trade-off between reconstruction quality and computing speed. A discussion of other choices can be found in the supplementary material. Finally, we observe that the proposed framework can also be easily extended to other 3D reconstruction settings, such as sonar [40] and multispectral lidar [26].

Methods

3D Reconstruction algorithm The reconstruction algorithm has been implemented on a graphics processing unit (GPU) to exploit the parallel structure of the update rules. Both the initialization and gradient steps process each pixel independently in parallel, whereas the point cloud and intensity denoising steps process each world-coordinates pixel in parallel, making use of the GPU shared memory to gather information of neighbouring points. The background denoising step is performed using the CuFFT library [34]. A detailed explanation of the implementation can be found in the supplementary notes. The algorithm was implemented using the parallel programming language CUDA C++ and all the experiments were performed using an NVIDIA Xp GPU. The surface fitting was performed using the Patate library [41].

Figure 5 shows the execution time per frame as a function of the total number of pixels and the mean active bins per pixel (i.e., the number of bins that have one or more photons) for the mannequin head dataset of fig. 3. For image sizes smaller than 150×150 , the algorithm has approximately constant execution time, due to the completely parallel processing of pixels. Larger images yield an increased execution time, as a single GPU does not have enough processors to handle all pixels at the same time (and other memory read/write constraints). As the per-pixel computations are not parallelized, the algorithm shows an approximately linear dependence with the mean number of active bins per pixel. A detailed analysis of complexity and memory requirements is discussed in the supplementary material.




Figure 5: Execution time of the proposed method as a function of the number of (a) lidar pixels having a mean of 4 active bins per pixel, and (b) histogram bins, for an array of 141×141 pixels. All the steps involved in the reconstruction algorithm can process the pixel information in parallel, the total execution time does not increase significantly when more pixels are considered. However, as the pixel-wise operations are not fully parallel, there is a dependency on the number of active bins present in the lidar frame.

Imaging set-up (dynamic scenes) Our system used a pulsed fibre laser (by BKTel, HFL-240am series) as the source for the flood illumination of the scene of interest. This had a central wavelength of 1550 nm and a spectral full width half maximum (FWHM) of approximately 9 nm. The output fibre from the laser module was connected to a reflective collimation package and the exiting beam then passed through a beam expander arrangement consisting of a pair of lenses. The lenses were housed in a zoom mechanism that enabled the diameter of the illuminating beam at the scene of interest to be adjusted to match the field of view of the camera.

We used a camera with a 32×32 array of pixels for the depth and intensity measurements reported here. This camera (by Princeton Lightwave Incorporated, Kestrel model) had an InGaAs/InP SPAD detector array with the elements on a $100 \mu\text{m}$ square pitch, resulting in an array with active area dimensions of approximately 3.2×3.2 mm. At the operating wavelength of 1550 nm, the elements in the array had a quoted photon detection efficiency of approximately 25% and a maximum mean dark count rate of approximately 320 kcps. The camera was configured to operate with 250 ps timing bins, a gate duration of 40 ns, and a frame rate of 150 kHz (this was close to the expected maximum frame rate of the camera). The camera provided this 150 kHz electrical clock signal for the laser, and the average optical output power from the laser at this repetition rate was approximately 220 mW and the pulse duration was approximately 400 ps. The camera recorded data continuously to provide a stream of binary frames at a rate of 150,400 binary frames per second.

An $f/7$, 500 mm effective focal length lens (designed for use in the 900 to 1700 nm wavelength region) was attached to the camera to collect the scattered return photons from the scene. This resulted in a field of view of approximately 0.5 arc degrees. As these measurements were carried out in broad daylight, a set of high performance passive spectral filters was mounted between the rear element of the lens and the sensor of the camera in order to minimise the amount of background light detected.

Our optical setup was a bistatic arrangement - the illuminating transmit channel and the collecting receive channel had separate apertures, i.e., the two channels were not co-axial. This configuration was used in order to avoid potential issues that could arise in a co-axial (monostatic) system due to back reflections from the optical components causing damage to the sensitive focal plane array. The parallax inherent in the bistatic optical configuration meant that a slight re-alignment of the illumination channel, relative to the receive (camera) channel, was required for scenes at different distances from the system.

Code availability A cross-platform executable file containing the real-time method will be made freely available upon publication of this work in the repository <https://gitlab.com/Tachella/real-time-single-photon-lidar>. The software requires an NVIDIA GPU with compute capability 5.0 or higher.

Data availability The lidar data will be made freely available upon publication of of this work in the repository <https://gitlab.com/Tachella/real-time-single-photon-lidar>.

Acknowledgements This work was supported by the Royal Academy of Engineering under the Research Fellowship scheme RF201617/16/31 and by the Engineering and Physical Sciences Research Council (EPSRC) Grants EP/N003446/1, EP/M01326X/1, EP/K015338/1, and EP/S000631/1, and the MOD University Defence Research Collaboration (UDRC) in Signal Processing. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. We also thank Bradley Schilling of the US Army RDECOM CERDEC NVESD and his team for their assistance with the field trial (camouflage netting data). Finally, we thank Robert J. Collins for his help with the measurements made of the scene at 320 metres, and David Vanderhaeghe and Mathias Paulin for their help on volumetric rendering.

Author contributions J.T. performed the data analysis, developed and implemented the computational reconstruction algorithm; Y.A. and N.M. contributed to the development of the reconstruction algorithm; A.M., R.T. and G.S.B. developed the experimental set-up; A.M. and R.T. performed the data acquisition and developed the long-range experimental scenarios; Y.A., G.S.B., J.-Y.T. and S.M. supervised and planned the project. All authors contributed to writing the manuscript.

References

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

- [2] C. Mallet and F. Bretar, “Full-waveform topographic lidar: State-of-the-art,” *ISPRS Journal of photogrammetry and remote sensing*, vol. 64, no. 1, pp. 1–16, 2009.
- [3] R. Horaud, M. Hansard, G. Evangelidis, and C. M  nier, “An overview of depth cameras and range scanners based on time-of-flight technologies,” *Machine Vision and Applications*, vol. 27, pp. 1005–1020, Oct 2016.
- [4] S. Izadi, D. Kim, O. Hilliges, D. Molyneaux, R. Newcombe, P. Kohli, J. Shotton, S. Hodges, D. Freeman, A. Davison, and A. Fitzgibbon, “Kinectfusion: Real-time 3D reconstruction and interaction using a moving depth camera,” in *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, UIST ’11, (New York, NY, USA), pp. 559–568, ACM, 2011.
- [5] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [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, pp. 11919–11931, May 2017.
- [7] R. Tobin, A. Halimi, A. McCarthy, M. Laurenzis, F. Christnacher, and G. S. Buller, “Three-dimensional single-photon imaging through obscurants,” *Opt. Express*, vol. 27, pp. 4590–4611, Feb 2019.
- [8] R. Tobin, A. Halimi, A. McCarthy, X. Ren, K. J. McEwan, S. McLaughlin, and G. S. Buller, “Long-range depth profiling of camouflaged targets using single-photon detection,” *Optical Engineering*, vol. 57, no. 3, pp. 1 – 10 – 10, 2017.
- [9] 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, pp. 33911–33926, Dec. 2015.
- [10] 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*, vol. 8375, p. 83750D, International Society for Optics and Photonics, 2012.
- [11] R. K. Henderson, N. Johnston, H. Chen, D. D. Li, G. Hungerford, R. Hirsch, D. McLoskey, P. Yip, and D. J. S. Birch, “A 192128 time correlated single photon counting imager in 40nm CMOS technology,” in *Proceedings IEEE 44th European Solid State Circuits Conference (ESSCIRC)*, pp. 54–57, Sep. 2018.
- [12] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, “Lidar waveform-based analysis of depth images constructed using sparse single-photon data,” *IEEE Transactions on Image Processing*, vol. 25, no. 5, pp. 1935–1946, 2016.
- [13] 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 Transactions on Computational Imaging*, vol. 2, pp. 456–467, Dec. 2016.
- [14] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, “Photon-efficient computational 3-D and reflectivity imaging with single-photon detectors,” *IEEE Transactions on Computational Imaging*, vol. 1, no. 2, pp. 112–125, 2015.
- [15] A. Halimi, Y. Altmann, A. McCarthy, X. Ren, R. Tobin, G. S. Buller, and S. McLaughlin, “Restoration of intensity and depth images constructed using sparse single-photon data,” in *Proc. 24th European Signal Processing Conference (EUSIPCO), Budapest, Hungary*, pp. 86–90, Aug. 2016.
- [16] 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.
- [17] F. Heide, S. Diamond, D. B. Lindell, and G. Wetzstein, “Sub-picosecond photon-efficient 3D imaging using single-photon sensors,” *Scientific reports*, vol. 8, no. 1, p. 17726, 2018.
- [18] D. B. Lindell, M. OToole, and G. Wetzstein, “Single-Photon 3D Imaging with Deep Sensor Fusion,” *ACM Trans. Graph. (SIGGRAPH)*, no. 4, 2018.
- [19] D. Shin, F. Xu, F. N. Wong, J. H. Shapiro, and V. K. Goyal, “Computational multi-depth single-photon imaging,” *Optics express*, vol. 24, no. 3, pp. 1873–1888, 2016.

- [20] 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.
- [21] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, “Bayesian analysis of lidar signals with multiple returns,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 12, pp. 2170–2180, 2007.
- [22] M. Berger, A. Tagliasacchi, L. M. Seversky, P. Alliez, G. Guennebaud, J. A. Levine, A. Sharf, and C. T. Silva, “A survey of surface reconstruction from point clouds,” *Computer Graphics Forum*, vol. 36, no. 1, pp. 301–329, 2017.
- [23] D. Shin, F. Xu, D. Venkatraman, R. Lussana, F. Villa, F. Zappa, V. K. Goyal, F. N. Wong, and J. H. Shapiro, “Photon-efficient imaging with a single-photon camera,” *Nature communications*, vol. 7, p. 12046, 2016.
- [24] Y. Altmann, R. Aspden, M. Padgett, and S. McLaughlin, “A bayesian approach to denoising of single-photon binary images,” *IEEE Transactions on Computational Imaging*, vol. 3, pp. 460–471, Sept 2017.
- [25] M.-J. Sun, M. P. Edgar, G. M. Gibson, B. Sun, N. Radwell, R. Lamb, and M. J. Padgett, “Single-pixel three-dimensional imaging with time-based depth resolution,” *Nature communications*, vol. 7, p. 12010, 2016.
- [26] Y. Altmann, R. Tobin, A. Maccarone, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, “Bayesian restoration of reflectivity and range profiles from subsampled single-photon multispectral lidar data,” in *25th European Signal Processing Conference (EUSIPCO), Kos Island, Greece*, pp. 1410–1414, Aug. 2017.
- [27] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, “Plug-and-play priors for model based reconstruction,” in *Proceedings Global Conference on Signal and Information Processing (GlobalSIP), Austin, USA*, pp. 945–948, IEEE, 2013.
- [28] S. H. Chan, “Algorithm-induced prior for image restoration,” *arXiv preprint arXiv:1602.00715*, 2016.
- [29] S. Sreehari, S. V. Venkatakrisnan, B. Wohlberg, G. T. Buzzard, L. F. Drummy, J. P. Simmons, and C. A. Bouman, “Plug-and-play priors for bright field electron tomography and sparse interpolation,” *IEEE Transactions on Computational Imaging*, vol. 2, pp. 408–423, Dec. 2016.
- [30] S. H. Chan, X. Wang, and O. A. Elgendy, “Plug-and-play admm for image restoration: Fixed-point convergence and applications,” *IEEE Transactions on Computational Imaging*, vol. 3, pp. 84–98, March 2017.
- [31] Y. Romano, M. Elad, and P. Milanfar, “The little engine that could: Regularization by denoising (RED),” *SIAM Journal on Imaging Sciences*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [32] E. T. Reehorst and P. Schniter, “Regularization by denoising: Clarifications and new interpretations,” *arXiv preprint arXiv:1806.02296*, 2018.
- [33] G. Guennebaud and M. Gross, “Algebraic point set surfaces,” *ACM Trans. Graph.*, vol. 26, July 2007.
- [34] J. Sanders and E. Kandrot, *CUDA by example: an introduction to general-purpose GPU programming*. Addison-Wesley Professional, 2010.
- [35] 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, pp. 22098–22113, Sep 2013.
- [36] J. Bolte, S. Sabach, and M. Teboulle, “Proximal alternating linearized minimization for nonconvex and nonsmooth problems,” *Mathematical Programming*, vol. 146, pp. 459–494, Aug 2014.
- [37] N. Parikh and S. Boyd, “Proximal algorithms,” *Foundations and Trends in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [38] G. Guennebaud, M. Germann, and M. Gross, “Dynamic sampling and rendering of algebraic point set surfaces,” *Computer Graphics Forum*, vol. 27, no. 2, pp. 653–662, 2008.
- [39] J. B. Tenenbaum, V. d. Silva, and J. C. Langford, “A global geometric framework for nonlinear dimensionality reduction,” *Science*, vol. 290, no. 5500, pp. 2319–2323, 2000.

- [40] Y. Petillot, I. T. Ruiz, and D. M. Lane, “Underwater vehicle obstacle avoidance and path planning using a multi-beam forward looking sonar,” *IEEE Journal of Oceanic Engineering*, vol. 26, pp. 240–251, April 2001.
- [41] N. Mellado, G. Ciaudo, S. Boyé, G. Guennebaud, and P. Barla, “Patate library.” <http://patate.gforge.inria.fr/>, 2013. Accessed: 2019-04-12.