

An image feature-based approach to improving SPAD Flash LiDAR imaging through fog

Joyce Mau^a, Jochen Trumf^b, Geoff Day^a, and Dennis Delic^a

^aDefence Science Technology Group, Edinburgh, Australia

^bAustralian National University, Canberra, Australia

ABSTRACT

Fog is a difficult medium to image through using Single-Photon Avalanche Diode (SPAD) based Light Detection and Ranging (LiDAR) systems because of its light scattering properties. Scattering significantly decreases the signal-to-noise ratio of photon returns, making it difficult to reconstruct meaningful images for target detection. In this paper, an image feature-based approach for reconstructing SPAD LiDAR images of a single target is proposed. Geometric characteristics of the target are used in the algorithm to differentiate between target and background photon returns. Combinations of different features such as Fourier shape descriptors and apparent target size are used to improve performance. To validate the algorithm, a 32×32 silicon SPAD array Flash LiDAR system operating at 532nm is used for collecting images through fog. Simple geometric shapes are placed indoors in a dark tunnel 44.6m from the sensor with fog decreasing the visibility in steps down to 12m. The proof-of-concept algorithm achieves good localization performance at a fog level of 1.4 attenuation lengths.

Keywords: Single-photon avalanche diode, SPAD, LiDAR, imaging through obscurant, fog, image feature-based approach

1. INTRODUCTION

As the military continues to adopt the use of autonomous platforms and unmanned systems, there is an increasing need for sensors that have low size, weight and power (SWaP) requirements. A sensor that is showing promise for this application is the Single Photon Avalanche Diode (SPAD). It is a single-photon sensor that can be used with a laser as a Light Detection and Ranging (LiDAR) system. The SPAD LiDAR system has been shown to provide good performance in long range imaging¹ where there are sparse photon returns.² However, one current challenge SPAD LiDAR has is sensing an environment affected by obscurant, where fog is one of the most challenging obscurants of interest. The light scattering properties of fog significantly decrease the signal-to-noise ratio of photon returns, making it difficult to perform target localization and reconstruct meaningful images for target detection.

The challenge of reconstructing SPAD LiDAR imagery of targets obscured by fog is an active area of research problem.²⁻⁸ Currently, the most common solution is to model the probability of photon return timing values to determine the most likely depth value. The light scattering properties of fog cause similar imaging challenges to environments with low signal-to-background ratio or sparse photon returns. Therefore, photon return modeling is a popular choice because it has been successful in these similar environments.^{2,9-11} For example, the Unmixing algorithm⁹ is a state-of-the-art algorithm for reconstructing SPAD images with low signal-to-background ratios and has been shown to successfully reconstruct SPAD images of objects in fog.⁵

Much of the literature assumes that the empirical photon return distribution can be modeled by a Poisson distribution.⁵⁻⁸ The differences lie in the optimisation and implementation methods used to numerically converge to a set of acceptable Poisson parameter values. Common optimization methods used are Bayesian models⁵ and

Further author information:

Joyce Mau: E-mail: joyce.mau@defence.gov.au

Jochen Trumf: E-mail: jochen.trumpf@anu.edu.au

Geoff Day: Email: geoff.day@defence.gov.au

Dennis Delic: Email: dennis.delic@defence.gov.au

maximum likelihood estimation.⁷ Some algorithms also use spatial and temporal correlation between pixels to better estimate these parameters.^{5,6,9}

The current state-of-the-art algorithm for fog obscured environments is the Median-based Multi-scale Restoration of 3D images (M2R3D) algorithm.⁵ The M2R3D algorithm also assumes the photon return timing values can be modeled as a Poisson distribution and uses a hierarchical Bayesian model to estimate the Poisson parameters.⁵ In addition, it also uses spatial processing techniques to improve the spatial quality of the SPAD images.⁵ The M2R3D algorithm is able to partially reconstruct a 32×32 SPAD array image at a fog level of 5.5 attenuation lengths (fully reconstructed at 4.5) at an operating wavelength of 1550nm for multiple stationary objects at 50m away using data acquired over 1s, which is assumed to correspond to 150,000 frames given the system operates at 150 kHz.⁵ It is also able to reconstruct a 10 frames per second video of an indoor foggy scene with multiple moving targets at 50m away. The fog level in this scene is approximately 3 attenuation lengths.⁵

Only a few papers have deviated from using a Poisson distribution model for photon returns, using a mixture of different probability distributions instead.^{3,4} Different mixture components are used to model different parts of the empirical distribution that represent returns from the fog and the target, respectively. The fog can be modeled as either a Gamma distribution⁴ or a mixture of Gaussian and lognormal distributions³ while the target is usually modeled as a Gaussian distribution.^{3,4}

Whilst some of the existing algorithms use spatial correlation between pixels,^{5,6} there is very little literature that approaches reconstruction of SPAD images from a fog obscured environment as an image feature-based filtering problem. Only simulation work has been conducted in validating the use of deep learning techniques for classifying simulated 32×32 SPAD images of different handwritten digits from the MNIST (Modified National Institute of Standards and Technology) dataset that have been degraded by fog.¹² In experimental work, the most similar research would be using a Convolutional Neural Network for pose estimation of a mannequin imaged by a 32×32 SPAD LiDAR system behind a piece of paper, which acts as an obscurant.¹³

In this paper, we are interested in the reconstruction of SPAD images for localization of targets. An image feature-based approach for reconstructing SPAD LiDAR images and localizing a single static target in 3D space is proposed. Geometric characteristics of the target are used to differentiate between target and background photon returns. Combinations of different features such as Fourier shape descriptors and apparent target size are used to improve performance. Full image reconstruction is not necessary as we can rely on the geometric properties to inform us of the type of target in the image. From the knowledge of the type of target, we are able to deduce the depth of the target by inspecting the correct segment of the partially reconstructed image.

To validate the algorithm, a 32×32 silicon SPAD array Flash LiDAR system is used for collecting images through fog. A 532nm pulsed laser source emitting 5ns wide pulses at a rate of 25 Hz through a diffuser is used to flood-illuminate the scene with a laser power of 10mJ. Simple geometric shape targets (i.e. triangle, square, circle) are placed at a distance of 44.6m from the sensor indoors in a dark tunnel with fog decreasing the visibility in steps down to 12m (corresponding to a fog level of 3.1 attenuation lengths).

Preliminary results of the localization accuracy of the triangle target at a fog level of 1.4 attenuation lengths and visibility of 25.4m show promising performance. The algorithm can be extended to other target shapes and multiple targets, including at different depths, enabling an implementation of lookup tables of different target geometries for fast real-time classification on-board autonomous systems.

2. SPAD LIDAR SYSTEM

The SPAD LiDAR system used in this paper is the same system used in Ref. 3. The 32×32 SPAD microchip is designed as a planar device in a standard CMOS 130nm process. Its active area is 20 μm in diameter and uses a passive quenching front-end circuit on the same die. Micro-lensing is used to improve the fill-factor. Optically, a narrow 2nm band filter centered at 532nm (Edmund Optics part no. 68-970) and the Pentax TV Zoom lens 12.5 to 75 mm with F number of 1 to 1.8 are used. The bias voltage is at 29.1V, with the dark count rate being 2000-6000 cps. The digital and analog circuitry operate on 1.8V and 3.3V, respectively. The SPAD bin width is 3.33ns corresponding to a measured distance resolution of 1 m, or a target distance resolution of 0.5m after accounting for the round trip configuration where photons travel to and back from the target.

A pulsed laser operating at 532nm and at a flash rate of 25Hz is used to flood-illuminate the scene. A wavelength of 532nm is used because high power systems are more widely available commercially. This wavelength is also close to this particular SPAD sensor’s peak photon detection efficiency of 11%. A diffuser is placed in front of the laser to make the system eye safe and this results in a laser beam divergence of approximately 115°. The diffuser is a circle of tracing paper placed within a simple beam expander consisting of two negative lenses spaced 40mm apart. This gives a good even illumination at an energy level of 10mJ (measured over a 70 degree hemisphere). The laser pulse width is approximately 5ns long with approximately 1ns of jitter.

3. EXPERIMENT SETUP

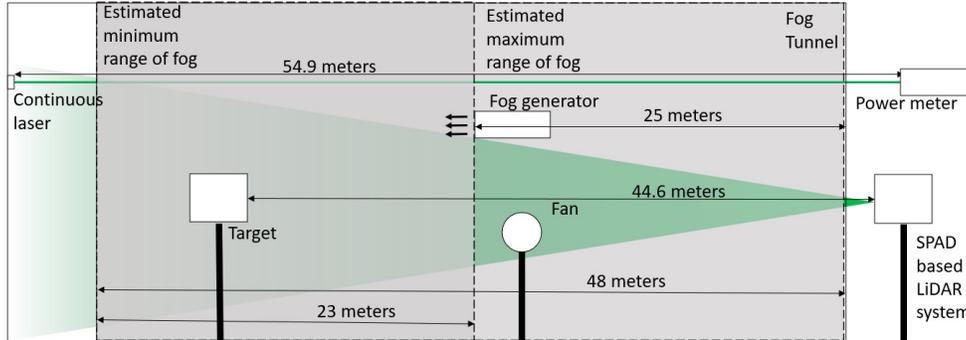


Figure 1. Schematic of experimental setup

The 32×32 silicon SPAD array Flash LiDAR system described in Sec. 2 is used to collect images of geometric targets, using the same experimental setup as Ref. 3, see Fig. 1. The imaging is performed in a 54m long dark tunnel that is filled with different levels of fog, from 185m down to 12m in average visibility (corresponding to a fog level of 3.1 attenuation lengths, cf. Sec. 4). Fig. 1 shows a schematic of the equipment layout in the dark tunnel. The SPAD LiDAR system and power meter are located in an approximately 3m x 5.5m room into which the start of the tunnel opens. The range-gate of the SPAD camera is set to 10 clock cycles (equivalent to a target distance of 5m) to avoid scattering directly from the laser, while still being able to image as much of the fog in the tunnel as possible. The SPAD system must be placed close to the laser to enable LiDAR imaging.

A pulley system is used to pull up several individual targets into the field of view of the SPAD LiDAR system allowing individual shapes to be imaged with minimum interruption to the flow of the fog in the tunnel. The targets are a white cardboard cut-outs of a right-angled triangle, square and circle and their sizes can be circumscribed by a 610 mm square. Additional fog is added only after all shapes are imaged separately. The fog generator, Rave AF1214 Fog Machine, is placed at 25m away from the LiDAR system. The fog liquid is a Rave heavy fog water-based liquid. A fan is placed facing the end of the tunnel to keep the fog confined to a region around the target. Visibility is calculated from the received power measurements of an additional 532nm 14mW continuous laser. The calculations for visibility and fog level in attenuation lengths are described in Sec. 4.

In this paper, SPAD data of a triangle in two different fog conditions are used to demonstrate the performance of the algorithm. The lower fog condition has a visibility of 43.8m and a fog level of 0.8 attenuation lengths while the higher fog condition has a visibility of 25.4m and a fog level of 1.4 attenuation lengths.

3.1 Timing Corrections

An offset of 4 clock cycles is added to the timing data before it is converted to a distance measurement. This offset is added because the FPGA controller for the LiDAR system delays the SPAD window opening. Throughout this paper, we report the timing data in terms of the raw clock cycles taken from the SPAD system, but add the offset when we talk about the equivalent distance measurement.

4. MEASURING FOG LEVELS

There is no standardized way of measuring fog levels. Two different measurements are used in this paper to attempt to provide a more comprehensive view of the level of fog used in the data collection. Both visibility and the number of attenuation lengths vary with wavelength making the measurement of fog level specific to a particular type of laser. Therefore, care must be taken when comparing results for different SPAD LiDAR systems operating at different wavelengths. In this paper, all measurements related to the visibility and attenuation lengths are for a 532nm system.

4.1 Visibility in meters

Visibility describes how far the naked eye can see from the start of the fog. Visibility is calculated by measuring the transmittance of a continuous 532nm laser beam. This is possible by using the formula for visibility:

$$V = \frac{1}{\alpha} \ln \left(\frac{C_0}{C_{th}} \right), \quad (1)$$

where α is the attenuation coefficient, C_0 is the target contrast ($C_0 = 1$ in this paper), C_{th} is the threshold contrast and \ln denotes the natural logarithm. In this paper, C_{th} is set to 0.05 according to the CIE's (International Commission on Illumination) definition of the human eye's minimal perceptible contrast.^{6,14} The attenuation coefficient, α , is calculated using the Beer-Lambert Law:¹⁵

$$\alpha = -\frac{1}{x} \ln \left(\frac{P}{P_0} \right), \quad (2)$$

where P_0 is the initial power of the continuous laser and P is the received power measured at the other end of the tunnel (see Fig. 1). x is the path length of the continuous laser beam. In this paper, the attenuation of the laser through air is considered negligible compared to attenuation by fog. Therefore, x is the distance traveled by the laser through fog to reach the power meter. The value used for x in this paper is 35.5m, which is the average span of the fog. An average span is used because the formula assumes a uniform distribution of fog and fog in nature varies in distribution. During the experiment, the span of the fog is observed to be between 23m to 48m as shown in Fig. 1. A more sophisticated experimental setup would be needed to stabilize this value.

4.2 Fog Level in Number of Attenuation Lengths

Number of attenuation lengths is a common measure of the level of fog when discussing SPAD LiDAR imaging through fog.^{2,5,6,16} The definition of one attenuation length is the distance at which the transmitted light power is reduced by a factor of $1/e$ of its initial value.² The attenuation length is thus calculated as:

$$N_{AL} = \alpha d = -\frac{d}{x} \ln \left(\frac{P}{P_0} \right), \quad (3)$$

where α is the attenuation coefficient and d is the distance between the SPAD sensor and the target. α is calculated using the Beer-Lambert Law as above.

5. IMAGE FEATURE-BASED APPROACH TO TARGET LOCALIZATION

5.1 Data structure

The algorithm analyzes the image features created by the empirical distribution of photon return timing values in each pixel on the 32×32 array. Specifically, it looks at the number of photon returns that each pixel has for a given timing value over a certain number of LiDAR pulse cycles. In this paper, one SPAD frame means the collection of photon returns measured by the SPAD array sensor during one LiDAR pulse cycle. The algorithm transforms batches of SPAD frames into a histogram stack. Fig. 2 shows an example of how four raw SPAD frames are transformed into a histogram stack. The histogram stack collates all the histograms from each pixel such that each slice of the histogram stack is a 32×32 array containing values of the histogram for each pixel at

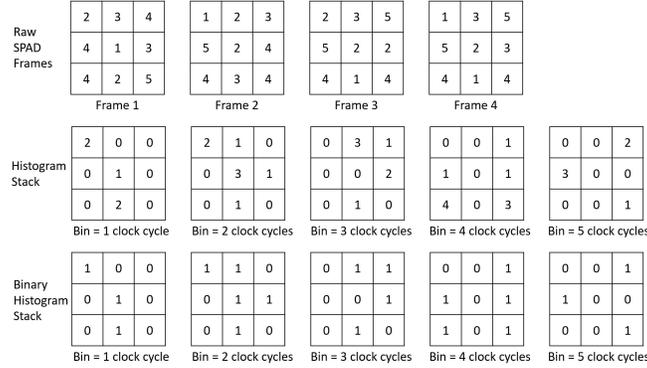


Figure 2. Example of turning frames into histogram stack and binary histogram stack

that slice’s timing bin value. This turns the histogram stack into a 3D array where each slice is 32×32 and the third dimension corresponds to the bin number. In this paper, the maximum bin value for the histogram stack is 100 clock cycles, corresponding to an approximate distance of 52m from the SPAD sensor to the back of the fog (see Fig. 1). The values along all the slices of the histogram stack for each pixel make up the complete histogram of timing values at that pixel. Collating histogram values from each pixel at each slice allows the algorithm to analyze the data spatially for each timing bin. The histogram stack is binarized to allow for easier and faster processing of the spatial information in the histogram slices. The binary histogram stack converts any histogram stack values larger than 0 into 1. An example is shown in Fig. 2 which shows the binary histogram stack values generated from the histogram stack.

The algorithm is targeted for implementation on an embedded board where it is presented with a set of raw SPAD frames as input data. The algorithm processes the SPAD frames in batches to enable real-time processing.

5.2 Algorithm

The purpose of this algorithm is to provide the location of the object that most resembles the target. In this paper, it is assumed that there is only one target imaged by the SPAD sensor and that the target has a depth smaller than the distance resolution. This means the target should only exist in one single slice of the binary histogram stack. Under this assumption of a single target at only one distance, the algorithm calculates only a maximum of one Fourier shape descriptor for each slice of the binary histogram stack. The location of the target is defined as the distance value of the slice with the Fourier shape descriptor closest to the target descriptor. Fourier shape descriptors are used as a frequency domain description of the contour of an object.¹⁷

For a slice of binary histogram stack to have a Fourier shape descriptor, an object in the image must be identified. In this paper, an object is defined to be a region with pixels of value 1 that are 8-connected. The algorithm first removes any noise regions and then fills in holes in a given slice. A region that contains pixels of value 1 is identified as a noise region if the total number of pixels of value 1 is smaller than a certain fraction of the number suggested by the apparent target size at the region’s location. Likewise, a region that contains pixels of value 0 is identified as a hole if the total number of pixels of value 0 is smaller than that same fraction. It has been determined empirically that a fraction of 0.02 works well for the lower fog condition data presented in this paper and a fraction of 0.035 works well for the higher fog condition. For the proof-of-concept implementation presented in this paper, MATLAB’s `bwareaopen` function is used. The apparent area of the target is calculated using the formula as follows:

$$\text{apparent area of target at the } k\text{-th bin} = \frac{\text{target's area at 85-th target bin}}{\left(\frac{k}{85}\right)^2} \quad (4)$$

As this is only a crude estimation, the area of a bounding box around the target is used instead of calculating the exact area of the target. This is mathematically sound as the area of a region is proportional to the area of the bounding box that circumscribes it.

Both the bounding box area and Fourier shape descriptor of the target are calculated beforehand. The target used in this paper is a right-angled triangle with three vertices at coordinates (6,5), (6,27) and (28,27). These vertices are determined from inspecting a non-obscured SPAD image of the target that has been averaged over 1000 SPAD frames. The target is defined to be at 85 clock cycles, given it is observed to be at 44.6m away as shown in Fig. 1 (see Sec. 3.1 for the correction made to convert clock cycle values to distance values).

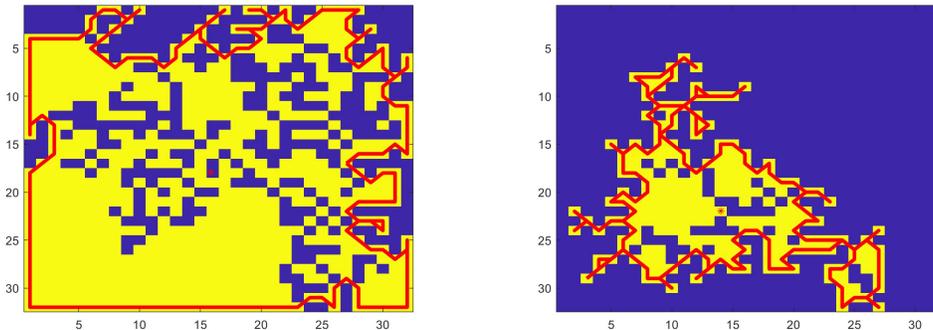


Figure 3. Example of boundary traces with the red asterisk representing the mean of the boundary trace coordinates. The slices are at bins 50 (left) and 85 (right) clock cycles. The fog condition of this data has visibility of 25.4m and the fog level is 1.4 attenuation lengths.

The calculation of Fourier shape descriptors in this paper is similar to the method described in Ref. 17 for shape recognition. A boundary is traced clockwise around a region of 8-connected pixels with a value of 1. In the proof-of-concept implementation used in this paper, this is performed using MATLAB's `bwtraceboundary` function. The resulting trace depends on the initial point and initial search direction that are used as input for the function. Fig. 3 shows two examples of boundary traces on slices generated from SPAD data used for this paper.

As the location of the target within each slice is unknown, it is not possible to set a pixel and direction as the initial point and search direction. Therefore, a set of 4000 unique and randomly generated combinations of initial pixels and search directions are used for boundary tracing. The initial pixels are chosen from all 32×32 pixels of the SPAD array and the initial search directions are any of the 8 possible directions (North, North-East, East, South-East, South, South-West, West and North-West). Not every initial boundary pixel and search direction will result in a boundary trace around an object but 4000 different initial conditions covers approximately half of the possible combinations of possible pixel and search directions. This makes it probable that an object is traced in each slice since opposite directions at the same pixel provide the same trace. For one binary histogram stack generated from one batch of SPAD frames, the same set of initial pixels and search directions is used. A different set is only generated when a new batch of SPAD frames is processed.

All generated boundary traces for each slice are further filtered by performing an area check. This is to check whether the object could be the target. The boundary trace is only kept if the object's area is within a tolerance of 50% of the apparent target size at the region's location as calculated in Eq. 4. The tolerance level has been determined empirically to work well with the data presented in this paper. The area of a bounding box that circumscribes the boundary trace is used instead of the exact area encompassed by the boundary trace to reduce computation time. This approach is mathematically sound as the area of a region is proportional to the area of the bounding box that circumscribes it. The area check is only a crude test that does not account for any differences in shapes between a detected object and the target. However, the algorithm subsequently uses a Fourier shape descriptor to correct for this issue.

The coordinates of each of the remaining boundary traces are converted to a radius signature. The radius signature contains the distances of consecutive boundary trace points to the mean point of the same trace. Examples of radius signatures are shown in Fig. 4. Once the radius signature is calculated, the Fast Fourier Transform (FFT) is performed to create the Fourier shape descriptor. Since the phase of the descriptor is dependent on the

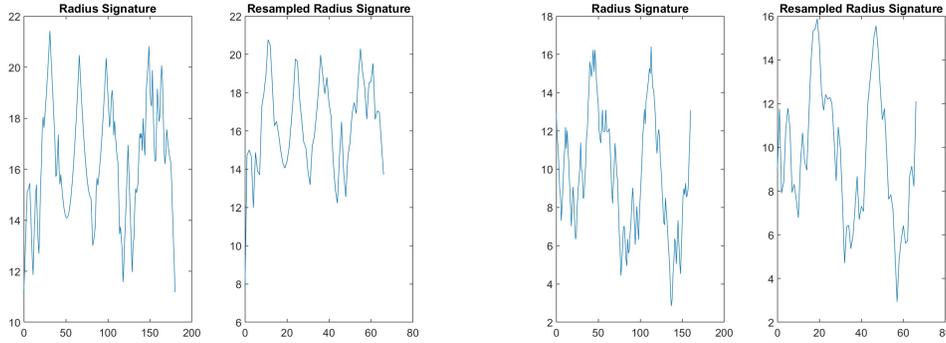


Figure 4. Example of boundary traces turned to radius signatures. The bin numbers are 50 (left) and 85 (right) clock cycles. The fog condition of this data has visibility of 25.4m and the fog level is 1.4 attenuation lengths.

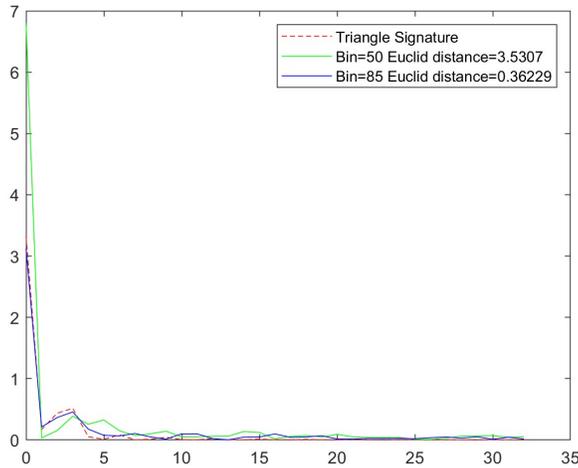


Figure 5. Example of Fourier shape descriptors of the radius signatures shown in Fig. 4. The fog condition of this data has visibility of 25.4m and the fog level is 1.4 attenuation lengths.

starting point of the boundary trace, only the magnitude of the descriptor is used for comparison. The Euclidean distance between the components of the object and target’s Fourier descriptors are calculated to determine if the object has a similar shape to the target. For comparison, the radius signature is resampled to the same length as the target’s radius signature to produce equivalent FFT components for the object’s descriptor. The resampled radius signatures are also shown in Fig. 4. All frequency components are used and normalized to create a size-invariant descriptor that allows comparison of detected shapes at different depths, hence different apparent sizes, to a fixed size pre-computed target descriptor. Fig. 5 shows two objects’ Fourier descriptors compared with the target’s Fourier descriptor.

6. RESULTS

The algorithm determines the location of the target as the arithmetic mean of the 100 location values generated from processing 100 batches of SPAD frames. Batches of frames are chosen such that they are evenly separated across the SPAD frames collected for this paper. Different parameters are used for the algorithm for different fog conditions. For the higher fog condition, batches of 1000 frames are used and a fraction of 0.035 is used for filling holes and removing noise. For the lower fog condition, 500 frames are used for each batch and a fraction of 0.025 is used instead since there is less noise in the data.

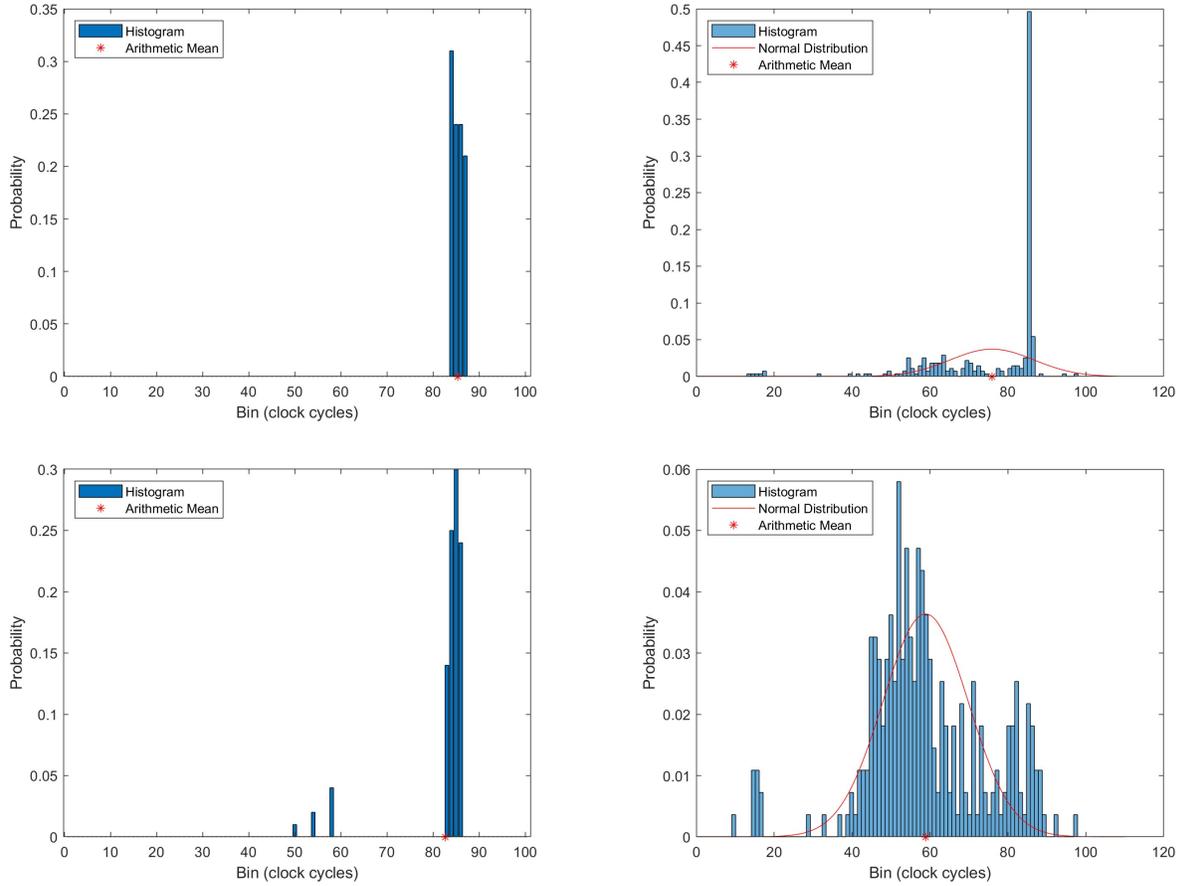


Figure 6. Comparison of the histogram and Gaussian generated from target location value results from the image-feature based algorithm (left column) and algorithm from Ref. 3 (right column). In the top row, the visibility is 43.8m and the fog level is 0.8 attenuation lengths. For the bottom row, the visibility is 25.4m and the fog level is 1.4 attenuation lengths. The red asterisk on the x-axis represents the mean of the Gaussian which represents the target location output from each algorithm. Its values are 85.35 (top left), 75.93 (top right), 82.66 (bottom left) and 58.86 (bottom right).

To provide an understanding of the performance of this image feature-based algorithm, the arithmetic mean of the target location values is compared with the arithmetic mean of the location values generated from applying the fog modeling algorithm proposed in Ref. 3 to the same SPAD dataset. The results of the location values deduced by both the algorithms and their arithmetic means are presented in Fig. 6. This shows that the image feature-based algorithm provides a significant improvement in the location estimation for the higher fog condition, and a better estimation for the lower fog condition. The aim of the fog modeling algorithm³ is to only locate the target along the timing bins for each pixel and not locate the target with respect to the 32×32 array. Therefore, the fog modeling algorithm³ is only applied to the pixels at the known location of the target (as discussed in Sec. 5.2) for a fair performance comparison. This implies that the fog modeling algorithm provides 276 location values for comparison.

Similar to the image feature-based algorithm in this paper, the fog modeling algorithm in Ref. 3 also processes SPAD data by transforming batches of SPAD frames into histograms. However, the fog modeling algorithm processes histogram values on a pixel-wise basis instead of processing the histogram value over the entire 32×32 array. In this paper, the fog modeling algorithm from Ref. 3 is applied to fit a mixture of two lognormal and one Gaussian probability distributions to the histogram of each pixel. The target corresponds to the Gaussian component. This implies that the arithmetic mean of the generated location values is normally distributed and

the variance is calculated from the individual pixel distributions. The resulting Gaussian distribution is plotted on top of the histogram in Fig. 6 to indicate that the fog modeling algorithm provides not only a point estimate, but an entire distribution of the target location.

7. CONCLUSION

We have demonstrated a new algorithm that utilizes image features such as Fourier shape descriptors and apparent target size to better localize a single target obscured by fog for SPAD LiDAR imaging. The algorithm shows better performance compared to the pixel-wise approach proposed in Ref. 3 for fog levels up to 1.4 attenuation lengths for a wavelength of 532nm. The use of batches will allow for real-time implementation on an embedded board. The use of Fourier shape descriptors enables fast execution of the algorithm for different target shapes by implementing a look-up table. In future work, this algorithm will be applied to other target shapes as well as in higher fog conditions. One potential improvement will be to introduce a better boundary tracing algorithm to enable better shape tracing for SPAD images with fog levels greater than 1.4 attenuation lengths. A better boundary trace will provide a more accurate Fourier shape descriptor, enabling better comparison with the target's descriptor.

REFERENCES

- [1] Pawlikowska, A. M., Halimi, A., Lamb, R. A., and Buller, G. S., "Single-photon three-dimensional imaging at up to 10 kilometers range," *Opt. Exp.* **25**(10), 11919–11931 (2017).
- [2] Halimi, A., Tobin, R., McCarthy, A., Bioucas-Dias, J., McLaughlin, S., and Buller, G. S., "Robust Restoration of Sparse Multidimensional Single-Photon LiDAR Images," *IEEE Trans. on Comput. Imag.* **6**, 138–152 (2020).
- [3] Mau, J., Devrelis, V., Day, G., Trumpf, J., and Delic, D. V., "The use of statistical mixture models to reduce noise in SPAD images of fog-obscured environments," *SPIE Future Sens. Technologies* **11525**, 115250P, SPIE, Online Only, Online Only (2020).
- [4] Satat, G., Tancik, M., and Raskar, R., "Towards photography through realistic fog," *IEEE Int. Conf. on Comput. Photography (ICCP)*, 1–10 (2018).
- [5] Tobin, R., Halimi, A., McCarthy, A., Soan, P. J., and Buller, G. S., "Robust real-time 3D imaging of moving scenes through atmospheric obscurant using single-photon LiDAR," *Scientific Reports* **11**, 11236 (2021).
- [6] Tobin, R., Halimi, A., McCarthy, A., Laurenzis, M., Christnacher, F., and Buller, G. S., "Three-dimensional single-photon imaging through obscurants," *Opt. Exp.* **27**(4), 4590–4611 (2019).
- [7] Sang, T.-H., Tsai, S., and Yu, T., "Mitigating effects of uniform fog on SPAD Lidars," *IEEE Sensors Letters* **4**(9), 1–4 (2020).
- [8] Tsai, S.-Y., Chang, Y.-C., and Sang, T.-H., "SPAD LiDARs: modeling and algorithms," *2018 14th IEEE Int. Conf. on Solid-State and Integr. Circuit Technol. (ICSICT)*, 1–4, IEEE, Qingdao (2018).
- [9] Rapp, J. and Goyal, V. K., "A Few Photons Among Many: Unmixing Signal and Noise for Photon-Efficient Active Imaging," *IEEE Trans. on Comput. Imag.* **3**(3), 445–459 (2017).
- [10] Chen, S., Halimi, A., Ren, X., McCarthy, A., Su, X., McLaughlin, S., and Buller, G. S., "Learning non-local spatial correlations to restore sparse 3D single-photon data," *IEEE Trans. on Image Process.* **29**, 3119–3131 (2020).
- [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. on Image Process.* **25**(5), 1935–1946 (2016).
- [12] Saxena, A., "A data-driven approach to object classification through fog," in [*Masters Thesis*], Department of Electrical engineering and Computer science, Massachusetts Institute of Technology (2018).
- [13] Satat, G., Tancik, M., Gupta, O., Heshmat, B., and Raskar, R., "Object classification through scattering media with deep learning on time resolved measurement," *Opt. Exp.* **25**(15), 17466–17479 (2017).
- [14] Favennec, P.-N. and de Fornel, F., [*Measurements Using Optic and RF Waves*], John Wiley & Sons, Great Britain (2013).

- [15] Smith, F. G., Accetta, J. S., and Shumaker, D. L., [*Transmission along a Slant Path*], ERIM and Infrared Information Analysis Center, Michigan (1993).
- [16] Tobin, R., Halimi, A., McCarthy, A., Laurenzis, M., Christnacher, F., and Buller, G. S., “Depth imaging through obscurants using time-correlated single-photon counting,” *Proc. SPIE Adv. Photon Count. Techniques XII* **10659**, 106590S, SPIE, Orlando, United States (2018).
- [17] Govender, N., Warrell, J., Torr, P., and Nicolls, F., “Probabilistic models for 2D active shape recognition using Fourier descriptors and mutual information,” *Proc. of the Scientific Cooperations Int. Workshops on Electr. and Comput. Eng. Subfields* , Koc University, Istanbul, Turkey (2014).