The use of statistical mixture models to reduce noise in SPAD images of fog-obscured environments

Joyce Mau^a, Vladimyros Devrelis^b, Geoffrey Day^a, Jochen Trumpf^c, and Dennis Delic^a

^aDefence Science Technology Group, Edinburgh, Australia ^bBallistic Systems Pty Ltd, Edinburgh, Australia ^cSoftware Innovation Institute, Australian National University, Canberra, Australia

ABSTRACT

Navigating through fog plays a vital part in many remote sensing tasks. In this paper, we propose an Expectation-Maximization (EM) algorithm for fitting a mixture of lognormal and Gaussian distributions to the probability distributions of photon returns for each pixel of a 32x32 Single Photon Avalanche Diode (SPAD) array image. The distance range of the target can be determined from the probability distribution of photon returns by modeling the peak produced due to fog scattering with a lognormal distribution while the peak produced by the target is modeled by a Gaussian distribution. In order to validate the algorithm, 32x32 SPAD array images of simple shapes (triangle, circle and square) are imaged at visibilities down to 50.8m through the fog in an indoor tunnel. Several aspects of the algorithm performance are then assessed. It is found that the algorithm can reconstruct and distinguish different shapes for all of our experimental fog conditions. Classification of shapes using only the total area of the shape is found to be 100% accurate for our tested fog conditions. However, it is found that the accuracy of the distance range of the target using the estimated model is poor. Therefore, future work will investigate a better statistical mixture model and estimation method.

Keywords: SPAD, statistical mixture models, LiDAR, direct time-of-flight imaging, classification, obscurant, fog, Expectation-Maximization

1. INTRODUCTION

Reconnaissance using remote imaging can often be hindered by environmental factors. It is not uncommon for the imaging path to be obscured by fog, causing low visibility when navigating unfamiliar territory and observing adversary targets. Fog often presents an issue for Single Photon Avalanche Diode (SPAD) LiDAR imaging systems because of its light scattering nature which suppresses the return signal from the target. While there is considerable work on fog modeling and its removal for normal cameras¹⁻³ and airborne laser scanners,^{4,5} the literature for SPAD imaging systems is sparse. For some of the methods reported in the literature for SPAD, it is common to examine the probability distribution of photon returns.

SPAD imaging systems can be divided into FLASH and scanning systems. For FLASH systems, a two dimensional array of SPAD pixels is used. In 6, Satat et al. modeled the peak of the probability distribution of photon returns from the fog as a Gamma distribution and the peak from the target as a Gaussian distribution. He accurately recovered depth for up to an optical thickness (OT) of up to 2.1, where OT is calculated as $-log(\frac{P_t}{P_0})$. Here, P_t is the power measured at the opposite side of the chamber to the laser at time t and P_0 is the power when there is no fog. In his experimental setup, he uses a 32x32 SPAD camera paired with a 580nm pulsing laser. His indoor fog chamber is $0.5x0.5x1 m^3$. Similarly, in 7, the same imaging system was used to research the classification of different poses of a small wooden mannequin behind different thickness of fog without extracting a clear image of the model. However, both of these pieces of work^{6,7} by the same lead author perform fog imaging in a small space and did not investigate the imaging of longer distance targets.

For scanning systems, longer distance imaging of targets through fog has been addressed. In 8 and 9, the authors tested different algorithms for reconstructing scanned SPAD images of a polystyrene head in a 26m indoor tunnel filled with water-based fog. One of them is the Multidimensional Nonlocal Reconstruction of 3D images method (M-NR3D), which is proposed by 9. This algorithm solves for the distance range and reflectivity of the object by minimizing a cost function that assumes the photon return data histogram can be modeled

by a Poisson probability distribution.^{8,9} The cost function involves the log-likelihood of the data as well as regularization terms that account for several assumptions involving correlations between pixels as well as multi-temporal 3D images.^{8,9} The cost function is solved by using a variant of the alternating direction method of multipliers (ADMM) algorithm.^{8,9}

Another algorithm is the unmixing algorithm¹⁰ which focuses first on de-noising the measured photon returns and then determines the distance and reflectivity of the pixel by solving a constrained maximum likelihood problem.¹⁰ It assumes the de-noised photon returns can be modeled by a mixture of Poisson distributions. For de-noising, it uses a combination of a windowing method and a variant of superpixels.¹⁰ The windowing method uses a threshold to determine if there is a minimum number of photon returns within a window. This distinguishes target pixels from background pixels by assuming that the probability of having a certain number of photon returns in a window of time is much smaller if the pixel does not correspond to a target. The superpixel variant increases the number of photon returns for a pixel by combining its returns with those of its neighboring pixels. This assumes that natural scenes are generally smooth except at object boundaries and that edges in reflectivity and object boundaries in depth tend to be co-located.¹⁰

The authors in 8 found that the M-N3RD and unmixing algorithms perform best for removing fog in scanned SPAD images for up to 3.8 attenuation lengths in water-based fog at 1550nm. In 9, the author found that the M-NR3D algorithm performs better at a very low signal-to-background ratio (SBR) of 0.2, but has comparable performance with the unmixing algorithm for SBR of 1.7 or higher. However, the SPAD used in 8 and 9 is operated with a 1550nm laser, whilst in this paper, we are interested in the performance of the SPAD with a 532nm laser. Furthermore, both the M-NR3D and unmixing algorithms include the use of spatial correlation between pixels for processing.

In this paper, we are focused on imaging a single static target obscured by fog using a 32x32 direct timeof-flight (dToF) SPAD LiDAR system that operates at 532 nm. We investigate the use of statistical mixture modeling of the probability distribution of the photon returns to extract the target's location. Three simple shapes (circle, triangle and square) are imaged individually at 40m away from the sensor, with the fog spreading across the tunnel and up to the sensor. Target extraction is conducted by modeling the probability of photon return time values for each pixel as a mixture of lognormal and Gaussian distributions, where the target distance is defined as the mean of the Gaussian component furthest away from the imaging system. The Expectation Maximization (EM) algorithm¹¹ is chosen because it is a popular method for fitting statistical mixture models. The EM algorithm has been used in estimating depth from single-waveform multispectral lidar (SW-MSL) using a single scanning SPAD.¹² We are interested in the classification accuracy of the target from the resulting image. Classification is achieved by determining a set of rules on metrics involving the area of the targets. This work can be applied to long-range imaging using SPAD LiDAR systems which has potential for detection and classification of land, air, and sea-surface targets.

Preliminary work on 3D image collection and classification of various targets in the presence of fog has been presented in 13. This paper uses a different algorithm to extract the target and different types of target are used (i.e. simple shapes instead of frigates). Also, the fog distribution in this work is different as it spreads further to the SPAD LiDAR system than 13 where the fog stayed in a much shorter span of the tunnel. Furthermore, range-gating is not used in this work during imaging to filter out scattering from the fog like it was in 13. This is because this work is interested in the probability distribution of photon returns caused by the fog.

2. SPAD LIDAR IMAGING SYSTEM

The SPAD LiDAR imaging system is shown in Fig. 1 where it is placed at the start of a dark tunnel for imaging. The 32x32 SPAD microchip was designed as a planar device in a standard CMOS 130nm process. Its active area is 20 μ m in diameter and uses a passive quenching front-end circuit which is on the same die. Micro-lensing was used to improve its 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 were used. The bias voltage is at 29.1V, with the dark count rate being 2000-6000 cps. The digital and analogue circuitry operated on 1.8V and 3.3V. The laser had a wavelength of 532 nm, a pulse energy of 35 mJ and a pulse length of 5 ns which resulted in a range resolution of 0.75 m and a pulse repetition frequency of 20 Hz. 532nm was chosen because it was easy to



Figure 1. DST built SPAD LiDAR system in lab.

obtain lasers in this wavelength with high powers and it is close to the SPAD's peak photon detection efficiency of 11%. A diffuser was placed in front of the laser to make it eve safe and this resulted in a laser beam divergence of approximately 115 degrees.¹³ These parameters remained the same for all image collections. Furthermore, as the SPAD camera was placed close to the laser, the range-gate was set to 5 clock cycles (equivalent to 10 meters) to avoid scattering directly from the laser, while still being able to image as much of the fog in the tunnel as possible. When discussing distance in this work, clock cycles are used instead of meters as it is an equivalent measure. One clock cycle is equivalent to half a meter which is based on this system's parameters.



3. FOG TUNNEL SETUP

Figure 2. Setup of tunnel for measurements through fog.

23 meters

48 meters

Lidar

system

The imaging was conducted in a 54 m long tunnel laboratory. The setup is shown in Figure 2. White cardboard cut-outs of three different shapes (circle, triangle and square) were used as targets in this work. Their sizes are such that a 610 mm square could circumscribe around them. The shapes were installed onto three planks of long timber and each of them could be pulled up from the tunnel entrance via a long string individually for imaging. This is to ensure minimum interruption of the flow of the fog in the tunnel so the fog condition stays the same for all three shapes. Additional fog was released only after all three shapes were imaged separately. The fog generator used is the Rave AF1214 Fog Machine and was placed 25m away from the LiDAR system. The fog liquid is a Rave heavy fog water-based liquid. A fan was placed facing the end of the tunnel to keep the fog around the target. However, the fog during this imaging process was observed to finish at 6 m from the back of the tunnel while the start of the fog varied between the location of the fog generator and the start of the tunnel, as indicated by Figure 2. This is because air conditioning, located at the start of the tunnel, was used during imaging so the fog was able to spread past the fan.



Figure 3. For in the tunnel and the circle target is at the end of the tunnel.

The average visibility of the fog is determined to vary between 185 to 50.8m. This measurement describes how far the naked eye can see from the start of the fog, which is assumed to be at the start of the tunnel due to the varying nature of the spread of the fog. These values are calculated from measuring the transmittance of a continuous 532 nm laser beam over 35.5m, which is the average span of the fog. An average span is used since it is difficult to measure the exact distribution of fog for each imaging condition and the equations assume uniform distribution of fog. The attenuation coefficient of the laser is calculated using Beer-Lambert's Law¹⁴ and then the coefficient value is used to calculate the visibility value.¹⁵ It is assumed that there is negligible laser power dissipation in the area without fog.¹³

4. FOG MODELING ALGORITHM

The aim of the algorithm is to determine the distance of the target at each of the pixels where the target is present. Only a portion of photons are returned by the target, while others are returned by fog. This results in a probability distribution of photon return counts being a mixture of returns from both the fog and the target. This algorithm is inspired by the method in 6 where a mixture of statistical distributions are used to represent the background noise (fog) and target separately to extract the target's image through fog. The probability distribution of photon returns is created using the kernel density estimator (KDE) method,^{16,17} which follows the same method as 6, where Gaussian kernels are used. However, the difference in this work is that the bandwidth being used is approximately 320ps instead of 80ps because it is found to provide a smoother probability distribution.

For modeling the distribution of photon returns, up to 2 lognormal distributions are used in this work instead of a Gamma distribution to model the portion of the probability distribution that represents photon returns from the fog. A single lognormal distribution is used for scenarios where the visibility is below 72.9m while two lognormal distributions are used for visibilities of 57.2m and 50.8m. Similar to 6, we use one Gaussian distribution to model the portion of the distribution that represents the photon returns from the target. Altogether, this means the probability distribution of photon returns can be modeled by a mixture of lognormal and Gaussian distributions. The Gaussian is assumed to be located after the lognormal distributions because the experiment setup is such that the target is at the furthest point. An example of this is shown in Figure 4. The distance of the target is determined as the mean of the furthest Gaussian. There is no peak after the target to represent the photon returns from the back of the tunnel because they are range-gated out in post-processing of the measurements.

Another major difference to 6 is that the Expectation-Maximization (EM) algorithm is used to fit the mixture of lognormal and Gaussian distributions to the photon returns' probability distribution. This is possible by



Figure 4. A probability distribution of photon returns of pixel (15,15) for a square image at fog visibility of 94.5m fitted with 1 lognormal and 1 Gaussian distributions

adapting the EM algorithm for Gaussian Mixture Models (GMM) to a combination of lognormal and Gaussian distributions. In particular, to derive the new update formula in the EM algorithm, we use change of variables to relate a Gaussian model of variable in the form of log(x) to a lognormal model of the variable x. Here, x represents the probability distribution of photon returns. To initialize the parameters for this algorithm, the k-means++ algorithm,¹⁸ a variant to the k-means algorithm,¹⁸ is used.



5. CLASSIFICATION AND FURTHER IMAGE PROCESSING

Figure 5. Before and after images for the processing steps used before classification.

The classification algorithm in this work uses the total number of pixels on the target to make a distinction between the type of shape in the image. This is possible by running a few steps beforehand. Figure 5 shows the images before and after these processing steps for an image of a square in fog visibility of 50.8m. Firstly the image is turned into a binary image. The threshold used for binarization of each processed image is the 60th percentile pixel value of a running median of 200 raw frames. The median is used to reduce the effect of noise and the 60th percentile is chosen because it is observed that the shape's distance values make up the top 40% of the distance values in the image. After the image is converted to a binary format, holes are filled and scattered pixels are removed from it. This is because as the fog becomes more thick, there are less photon returns from the target which decrease the number of pixels representing the target. This causes more holes as well as noise in the binary image. Matlab's regionprops function is used to validate that the shape is the only detected target in the image. This is the case for all the images used for classification in this work. This series of processing steps is similar to the steps used for classification of shapes in 19. The difference is no selection criteria is used to identify the shape because there is only one detected target in this work.

6. RESULTS AND DISCUSSION

The performance of this algorithm can be discussed in terms of several different aspects. First, the performance of shape reconstruction is discussed in Section 6.1, then an examination of how well the shapes can be distinguished between each other is presented in the form of classification accuracy in Section 6.2. Finally, a discussion of the algorithm's localization accuracy of the target will be discussed in Section 6.3. In particular, an examination on the fit of the statistical mixture model on the probability distribution of photon returns is presented.



6.1 Shape Reconstruction

Figure 6. Percentage of pixels at the target's distance. The target distance is defined as the distance with the most pixels as the target occupies most of each image.

The percentage of pixels at the shape's distance is calculated to understand the shape reconstruction success. As the distance is different for each image, the distance of the target for each image is defined as the one clock cycle range that contains the most pixels in image. This is determined using a histogram of the pixel values for each image. The difference in the distances of targets will be discussed in Section 6.3 as part of understanding the localization accuracy of the target's distance range. Figure 6 shows images for different shapes in each row at different fog conditions in each column. The percentage under each image is the average of 10 different processed images made from a running of 500 raw frames of the same shape at the same fog condition. The percentage value for no fog is used as the ground truth for each shape. By comparing the percentages in different fog conditions with the ground truth, it provides an indication of the relative reconstruction quality. The percentage for the circle at visibility of 185m is higher than the percentage with no fog. One possible reason for this is there are a lot of noise pixels at the same distance. For the remaining images, there is a general trend of the percentages decreasing with increasing visibility. This is as expected because the fog creates difficulty in identifying the

peak at the distance of the target in the probability distribution of photon returns. There are some occasional instances where the percentage does not follow the trend. This is because the model does not fit the probability distribution perfectly, an issue that will be explored in future work. These discrepancies will be discussed in Section 6.3.

6.2 Shape Distinction

Classification is performed on 10 different processed images for each shape and each experimental fog condition. The area of the target is used as there was a clear threshold value between all the shapes for all visibility conditions. The average accuracies for all experimental fog conditions are found to be 100% for all shapes. This is because the distance of the shape from the SPAD sensor is still within the visibility distance for all conditions. Further work will be conducted to investigate the classification accuracy for lower visibility conditions, in particular where the visibility is shorter than the distance of the target.

6.3 Localization Accuracy

Table 1. Percentage of pixels within 3 clock cycles of their corresponding pixel's value in the ground truth image. The ground truth is defined as the image taken in no fog. The percentage is averaged over 10 different processed images for each shape in each visibility condition.

Visibility (m)	185	104.7	101.7	94.5	72.9	57.2	50.8
Circle	47.70%	7.36%	42.02%	7.64%	12.42%	34.26%	19.05%
Triangle	14.22%	10.10%	9.36%	8.13%	7.71%	10.65%	11.90%
Square	7.29%	12.58%	4.90%	10.52%	5.10%	9.48%	13.81%

In order to understand localization accuracy of the algorithm proposed in this work, the distance of the target for each fog condition is compared with the distance of the same target in no fog. As referenced in Section 6.1, the distance range of the target is the clock cycle with the most pixels in the image. Table 1 contains percentages of pixels within 3 clock cycles of the corresponding pixel in the ground truth, which is defined as the image of the shape with no fog. It shows the performance of all of the shapes' images at all the various fog conditions used in this paper. Each percentage is an average over the percentage values calculated from the same set of processed images used in Section 6.1, which contains 10 processed images made from 500 running raw frames from each shape at each fog condition. As shown in Table 1, there is no consistency in the values and the percentages are below 50%. Hence, this shows that this algorithm is not suitable for estimating the distance range for the localization of the target. This is the case because the estimated statistical mixture model does not fit the probability distribution of photon returns perfectly. This can be shown by measuring the L2 norm between the statistical mixture model and the probability distribution.

Table 2. Average L2 norm distance over the target's pixels between the predicted model and the probability distribution of the photon returns. The pixel is defined to belong to the target if it is within the clock cycle of the target's distance, which is defined at the clock cycles with the highest number of pixels in the image.

Visibility (m)	185	104.7	101.7	94.5	72.9	57.2	50.8
L2 norm	430.14e-07	268.68e-07	106.91e-07	76.415e-07	54.536e-07	29.423 e- 07	64.191e-07

The average L2 norm values over 10 processed images of the square target for each fog condition are presented in Table 2. Note as referenced in Section 4, the statistical mixture model used for visibilities between 185m to 72.9m is 1 lognormal and 1 Gaussian distributions while for visibilities 57.2m and 50.8m, 2 instead of 1 lognormal distributions are used. There is an overall decreasing trend in the L2 norm distance as the visibility through the fog decreases. One possible explanation is that the statistical mixture model provides a better fit when the peak from the fog noise is comparable to the peak of the target, as shown in the left plot of Figure 7. This is



Figure 7. A single lognormal and single Gaussian mixture model fit and residual plot of the probability distribution of photon returns of pixel (15,15) for fog visibility of 94.5m (left) and 104.7m (right)

also supported by the right plot of Figure 7 where the fit of the estimated model is poor when the peak of the fog (which is known to be at low clock cycles) is smaller than the peak of the target (which is known to be at a higher clock cycle).



Figure 8. Comparing processed images using 1 lognormal and 1 Gaussian (top row) versus 2 lognormals and 1 Gaussian (bottom row) in terms of the L2 norm distance for the estimated model's fit and the shape reconstruction in the image.

As mentioned previously, 2 lognormals and 1 Gaussian provides a better shape reconstruction than 1 lognormal and 1 Gaussian for visibilities of 57.2m and 50.8m. This can be seen in Figure 8 where the processed images as a result of 1 and 2 lognormal distributions are shown. However in the same figure, the average L2 norm (using the same calculation as used in Table 2) is shown to be worse for the case of using 2 lognormal distributions. This is because the L2 norm increases significantly with any large discrepancies, even if it occurs at only one point of the model. From Figure 9, it can be deduced that the maximum absolute value from the residual plot



Figure 9. Mixture model plot and residual plot for visibility of 57.2m at pixel (15,15) using 1 Gaussian and 1 lognormal (left) or 2 lognormals (right)

for 2 lognormals is higher than the absolute values from the residual plot for 1 lognormal. Therefore, a different metric needs to be used as a measure for the fit of the estimated model to prevent the L2 norm value from hindering a good geometric fit for the target's peak. It is clear from Figure 9 that using 2 lognormals enable the small peak of the target to be detected, hence providing a better distance range estimation to the target and providing better shape reconstruction. Furthermore, Figure 9 highlights that the EM algorithm aims to estimate the whole probability distribution of photon returns, but fails to model the target's peak when only using 1 lognormal because of the unknown mid-peak at approximately 60 clock cycles. Therefore, a more robust fitting criteria for the estimated statistical mixture model will be investigated in future work. Furthermore, a more suitable statistical probability distribution of photon returns will be investigated as well.

7. CONCLUSION

The proposed algorithm of fitting a mixture of lognormal and Gaussian distributions using the EM algorithm is found to be very good for shape reconstruction and classification. In the future, this algorithm will be tested on images with lower visibility through fog to understand its limitations. Furthermore, a more robust fitting criteria will be used to ensure a fair evaluation of the algorithm's performance. The L2 norm distance did not provide insight to the geometric fit of the estimated model to the probability distribution of photon returns. In addition, a different mixture model and estimation method will be investigated to provide a more accurate estimate of the range distance of the target, allowing for spatial and temporal correlations between SPAD pixels to be included as part of the algorithm.

REFERENCES

- Tripathi, A. K. and Mukhopadhyay, S., "Efficient fog removal from video," Signal, Image and Video Processing 8, 1431–1439 (2014).
- [2] Consani, C., Druml, N., Dielacher, M., and Baumgart, M., "Fog Effects on Time-of-Flight Imaging Investigated by Ray-Tracing Simulations," *Proceedings* 2, 859 (2018).
- [3] Stankovic, L. and Dakovic, M., "Reconstruction of Randomly Sampled Sparse Signals Using an Adaptive Gradient Algorithm," *arXiv* (2014).
- [4] Soederman, U., Persson, A., Toepel, J., and Ahlberg, S., "On analysis and visualization of full-waveform airborne laser scanner data," Proc. SPIE 5791, 184 (2005).
- [5] Chauve, A., Mallet, C., Bretar, F., Durrieu, S., Pierrot-Deseilligny, M., and Puech, W., "Processing Full-Waveform Lidar Data: Modelling Raw Signals," Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, 102–107.

- [6] Satat, G., Tancik, M., and Raskar, R., "Towards photography through realistic fog," IEEE Int. Conf. on Comput. Photography (ICCP), 1–10 (2018).
- [7] Satat, G., Tancik, M., Gupta, O., Heshmat, B., and Raskar, R., "Object classification through scattering media with deep learning on time resolved measurement," *Optics Express* 25(15), 17466 (2017).
- [8] Tobin, R., Halimi, A., McCarthy, A., Laurenzis, M., Christnacher, F., and Buller, G. S., "Three-dimensional single-photon imaging through obscurants," *Optics Express* 27(4), 4590 (2019).
- [9] 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 Comp. Imag.* 6, 138–152 (2020).
- [10] Rapp, J. and Goyal, V. K., "A Few Photons Among Many: Unmixing Signal and Noise for Photon-Efficient Active Imaging," *IEEE Transactions on Computational Imaging* 3(3), 445–459 (2017).
- [11] McLachlan, G. J. and Krishnan, T., [*The EM algorithm and extensions*], vol. 382, John Wiley & Sons, New Jersey (2007).
- [12] Legros, Q., Meignen, S., McLaughlin, S., and Altmann, Y., "Expectation-Maximization Based Approach to 3D Reconstruction From Single-Waveform Multispectral Lidar Data," *IEEE Trans. on Comput. Imag.* 6, 1033–1043 (2020).
- [13] Mau, J., Devrelis, V., Day, G., Nash, G., Trumpf, J., and Delic, D., "Through Thick and Thin: Imaging Through Obscurant using SPAD array," *IEEE Sensors* (2020, in press).
- [14] Smith, F. G., Accetta, J. S., and Shumaker, D. L., [*Transmission along a Slant Path*], ERIM and Infrared Information Analysis Center, Michigan (1993).
- [15] Smith, F. G., Accetta, J. S., and Shumaker, D. L., [Visibility Range, vol. 2], ERIM and Infrared Information Analysis Center, Michigan (1993).
- [16] Parzen, E., "On estimation of a probability density function and mode," The annals of mathematical statistics 33(3), 1065–1076 (1962).
- [17] Davis, R. A., Lii, K.-S., Politis, D. N., and Rosenblatt, M., "Remarks on some nonparametric estimates of a density function," in [Selected Works of Murray Rosenblatt], 95–100, Springer, New York (2011).
- [18] Arthur, D. and Vassilvitskii, S., "K-Means++: The Advantages of Careful Seeding," Proc. of the Annu. ACM-SIAM Symp. on Discrete Algorithms 8, 1027–1035 (2007).
- [19] Mau, J., Devrelis, V., Day, G., Trumpf, J., and Delic, D., "Impact of water quality on Single Photon Avalanche Diode direct time-of-flight imaging," OCEANS 2020 MTS/IEEE (2020, in press).