

COMPUTATIONAL 3D AND REFLECTIVITY IMAGING WITH HIGH PHOTON EFFICIENCY

Donggeek Shin* Ahmed Kirmani* Vivek K Goyal† Jeffrey H. Shapiro*

*Massachusetts Institute of Technology †Boston University

ABSTRACT

Capturing depth and reflectivity images at low light levels from active illumination of a scene has wide-ranging applications. Conventionally, even with single-photon detectors, hundreds of photon detections are needed at each pixel to mitigate Poisson noise. We introduce a robust method for estimating depth and reflectivity using on the order of 1 detected photon per pixel averaged over the scene. Our computational imager combines physically accurate single-photon counting statistics with exploitation of the spatial correlations present in real-world reflectivity and 3D structure. Experiments conducted in the presence of strong background light demonstrate that our computational imager is able to accurately recover scene depth and reflectivity, while traditional maximum likelihood-based imaging methods lead to estimates that are highly noisy. Our framework increases photon efficiency 100-fold over traditional processing and thus will be useful for rapid, low-power, and noise-tolerant active optical imaging.

Index Terms— Computational 3D imaging, low light-level imaging, time-of-flight, Poisson noise, convex optimization

1. INTRODUCTION

A light detection and ranging (LIDAR) system [1] builds a histogram of photon counts over time. The time delay and amplitude of the photon-count histogram, relative to the transmitted pulse’s temporal profile, contain object depth and reflectivity information, respectively, about the illuminated scene. LIDAR signal-acquisition time must be long enough to collect the 10^2 to 10^3 photons per pixel (ppp) needed to generate the finely-binned histogram required for accurate scene 3D and reflectivity images.

In this paper, we introduce an active optical imaging framework that recovers accurate reflectivity and 3D images simultaneously, using on the order of 1 detected ppp averaged over the scene (Fig. 1). Our computational imager combines physically accurate single-photon counting statistics with exploitation of the spatial correlations present in real-world reflectivity and 3D structures. Our method’s dramatically superior sensitivity enables accurate 3D imaging when very little backreflected light reaches the detector, as will be the case with low optical-power active imagers [2].

Prior work: Active 3D imaging systems differ in how they modulate their transmitted power. Temporal modulation enables distance measurement by the time-of-flight (TOF) principle. Examples of TOF methods are homodyne TOF [3], pulsed TOF cameras [4], and picosecond laser radar systems [5]. Spatial modulation techniques include structured light [6] and active stereo imaging [7]. These spatial-modulation techniques have low photon efficiencies because they use an always-on optical source, whereas pulsed-TOF systems

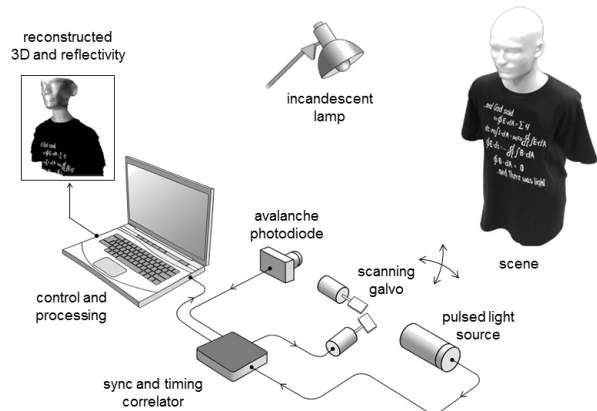


Fig. 1: Experimental imaging setup. A periodically pulsed light source illuminates the scene in a raster-scan pattern. The backscattered light is collected by a time-resolved single-photon detector. Each spatial location is illuminated with exactly N pulses (fixed dwell time). An incandescent lamp injects background light that corrupts the information-bearing signal. The photon arrival times and the total photon counts are recorded. This dataset is used to computationally reconstruct 3D structure and reflectivity. The setup is analogous to having a floodlight illumination source and an array of single-photon counting detectors operating at a fixed dwell time.

use sources that are on only for short intervals. The most photon-efficient imagers use single-photon avalanche diode (SPAD) detectors [8]. See [9] for a discussion of compressive methods [10–13] and earlier efforts in SPAD-based 3D imaging [14–16].

In low light-level imaging, a variety of optoelectronic techniques are employed for robust imaging. Active imagers use lasers with narrow spectral bandwidths and spectral filters to suppress background light and minimize the Poisson noise it creates. However, optical filtering alone cannot completely eliminate background light. Range-gated imaging [17] is another common technique, but it requires a priori knowledge of object location. Furthermore, a SPAD may be replaced with a superconducting nanowire single-photon detector (SNSPD) [18], which is much faster, has lower timing jitter, and has lower dark-count rate than a SPAD. However, SNSPDs have much smaller active areas and hence have narrower fields of view.

First-photon imaging (FPI) [19] is a method that allows accurate 3D and reflectivity reconstruction using only the first detected photon at every pixel in a raster-scanned scene. It combines first-photon detection statistics with the spatial correlations existing in natural scenes to achieve robust low light-level imaging. Its raster-scanning process to collect 1 ppp, however, makes the dwell time at each pixel a random variable. The dwell time is also a random variable when collecting any fixed number of photons per pixel, as in [15, 16] for 3D imaging separated from reflectivity. Thus, FPI does not extend naturally to operation using SPAD arrays—since simultaneous measurement implies equal dwell times—and precludes the dramatic speedup in image acquisition that such arrays enable.

This material is based upon work supported by the National Science Foundation under grant no. 1161413, a Samsung Scholarship, and a Microsoft PhD Fellowship.

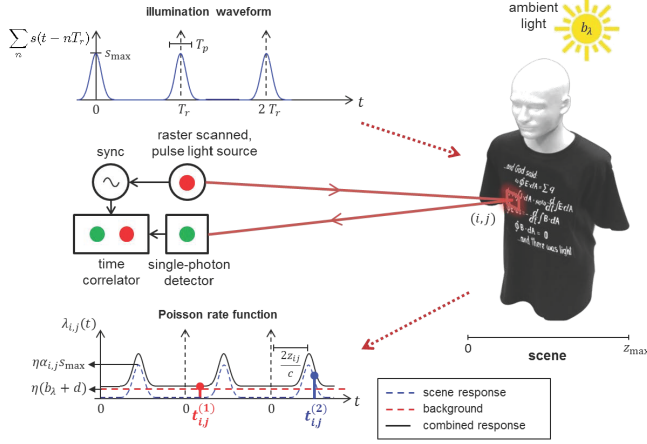


Fig. 2: Rate function of inhomogeneous Poisson process combining desired scene response and noise sources is shown. Here, $N = 3$ and $k_{ij} = 2$. A noise photon (red) was detected after the second transmitted pulse, and a signal photon (blue) was detected after the third transmitted pulse.

In this paper, we consider fixed dwell time at each pixel, thereby opening up the possibility of robust SPAD array-based imaging at low light-levels and with short dwell times.

Main contributions—Theoretical: We introduce a physically accurate model for the signal produced by a SPAD under low light-level conditions that incorporates arbitrary illumination pulse shapes and the inhomogeneous Poisson process characteristics (shot noise from the quantum nature of light) given a fixed acquisition time.

Main contributions—Experimental: We experimentally demonstrate that our proposed 3D imager’s photon efficiency, measured as recovery accuracy, is more than 100 times higher than that of traditional maximum likelihood (ML) estimation. We also show that our 3D imager achieves sub-pulse-width depth resolution under short acquisition times, in which 54% of the pixels have missing data, and at high background levels, when any given photon detection has approximately probability 0.5 of originating from ambient light.

2. IMAGING SETUP

Fig. 2 shows our imager’s signal-acquisition model. We aim to form reflectivity and depth images $\alpha, z \in \mathbb{R}^{n \times n}$ of the scene. We index the scene pixels as (i, j) , where $i, j = 1, \dots, n$. The distance to pixel (i, j) is $z_{ij} \geq 0$ and its reflectivity, $\alpha_{ij} \geq 0$, includes the effect of radial fall-off, view angle, and material properties.

Illumination: We use a periodically pulsed laser—pulse waveform $s(t)$ and repetition period T_r —to illuminate the scene in a raster-scanned fashion. Physically, $s(t)$ is the photon-flux waveform of the pulse emitted at $t = 0$ measured in counts/sec (cps). To avoid distance aliasing, we assume $T_r > 2z_{\max}/c$, where z_{\max} is the maximum scene range and c is the speed of light. With conventional processing, the root mean-square (RMS) pulse width T_p governs the achievable depth resolution in the absence of background light [20]. As is typically done in range imaging, we assume that $T_p \ll 2z_{\max}/c$.

Detection: A SPAD detector provides time-resolved single-photon detections [8]. Its quantum efficiency η is the fraction of photons passing through the pre-detection optical filter that are detected. Each detected photon is time stamped within a time bin of duration Δ , measuring a few picoseconds, that is much shorter than T_p .

Data acquisition: Each pixel (i, j) is illuminated with N laser

pulses. The total pixelwise acquisition time (dwell time) is thus $T_a = NT_r$. We record the total number of observed photon detections k_{ij} , along with their photon arrival times $\{t_{ij}^{(\ell)}\}_{\ell=1}^{k_{ij}}$, where the latter are measured relative to the immediately-preceding transmitted pulse. We also shine ambient light, with photon flux b_λ at the operating optical wavelength λ , onto the detector.

3. MEASUREMENT MODEL

Illuminating pixel (i, j) with the pulse $s(t)$ results in backreflected light with photon flux $r_{ij}(t) = \alpha_{ij}s(t - 2z_{ij}/c) + b_\lambda$ at the detector.

Poisson statistics: The photon detections produced by the SPAD in response to the backreflected light from transmission of $s(t)$ constitute an inhomogeneous Poisson process with time-varying rate function $\eta r_{ij}(t)$. To these photon detections we must add the detector dark counts, which come from an independent homogeneous Poisson process with rate d . Lumping the dark counts together with the background-generated counts yields the observation process at the SPAD’s output, viz., as shown in Fig. 2, an inhomogeneous Poisson process with rate function

$$\lambda_{ij}(t) = \eta r_{ij}(t) + d = \eta \alpha_{ij}s(t - 2z_{ij}/c) + (\eta b_\lambda + d), \quad (1)$$

when only a single pulse is transmitted. Fig. 2 shows the rate function $\lambda_{ij}(t)$ for the pulse-stream transmission.

Define $S = \int s(t) dt$ and $B = (\eta b_\lambda + d)T_r$ as the total signal and background count per pulse-repetition period, where, in what follows, we will include dark counts in what we refer to background. We assume that B is known, because it is straightforward to measure it before we begin data acquisition. The derivations to follow assume $\eta \alpha_{ij}S + B \ll 1$, meaning that the photon-flux per pixel per pulse-repetition period is much less than 1, as would be the case in low light-level imaging where an imager’s photon efficiency is paramount.

Signal vs. noise photons: A detection could originate from the backreflected signal or from background. The arrival statistics observed at the detector result from the merging of the Poisson processes corresponding to these sources. Using the theory of merged Poisson processes [21], we have that

$$\Pr[\text{detected photon at } (i, j) \text{ is signal}] = \frac{\eta \alpha_{ij}S}{\eta \alpha_{ij}S + B}.$$

3.1. Statistics of Number of Detected Photons

Using Poisson process statistics [21], we have that the probability of the SPAD detector’s *not* recording a detection at pixel (i, j) from one illumination trial is $P_0(\alpha_{ij}) = \exp[-(\eta \alpha_{ij}S + B)]$. Because we illuminate with a total of N pulses, and the low-flux condition ensures that multiple detections per repetition interval can be neglected, the number of detected photons K_{ij} is binomially distributed:

$$\Pr[K_{ij} = k_{ij}; \alpha_{ij}] = \binom{N}{k_{ij}} P_0(\alpha_{ij})^{N-k_{ij}} (1 - P_0(\alpha_{ij}))^{k_{ij}},$$

for $k_{ij} = 0, 1, \dots, N$.

Pixelwise maximum-likelihood reflectivity estimation: Given the total observed photon count k_{ij} at pixel (i, j) , the constrained ML reflectivity estimate is

$$\hat{\alpha}_{ij}^{\text{ML}} = \max\{\lceil \log(N/(N - k_{ij})) - B \rceil / \eta S, 0\},$$

where \log is the natural logarithm. Traditionally, the normalized photon-count value is used as the reflectivity estimate [22],

$$\tilde{\alpha}_{ij} = \frac{k_{ij}}{N \eta S}. \quad (2)$$

Note that the normalized count value estimate is equal to the ML estimate under the Poisson approximation to the binomial distribution when $B = 0$.

3.2. Statistics of Single-Photon Arrival Times

The distribution of the single-photon detection time for pixel (i, j) depends on whether the detection is due to signal or background. Under our low-flux assumption, the arrival time of a detected photon originating from backreflected signal is characterized by the normalized time-shifted pulse shape [21]. On the other hand, the photon arrival time from background is uniformly distributed over the T_r -sec pulse repetition period in which the detection occurred. The probability density function for the photon arrival time T_{ij} in the interval $[0, T_r)$ can thus be shown to satisfy

$$f_{T_{ij}}(t_{ij}; z_{ij}) = \frac{\eta\alpha_{ij}S}{\eta\alpha_{ij}S+B} \left(\frac{s(t_{ij}-2z_{ij}/c)}{S} \right) + \frac{B}{\eta\alpha_{ij}S+B} \left(\frac{1}{T_r} \right),$$

where infinitely fine time-binning has been assumed. This density is a mixture distribution, comprising pulse and uniform background distributions whose mixture weights depend on $\eta\alpha_{ij}S$ and B .

Pixelwise maximum-likelihood depth estimation: Using the photon arrival-time dataset $\{t_{ij}^{(\ell)}\}_{\ell=1}^{k_{ij}}$, the pixelwise constrained ML depth estimate is

$$\hat{z}_{ij}^{\text{ML}} = \underset{z_{ij} \in [0, cT_r/2)}{\operatorname{argmax}} \sum_{\ell=1}^{k_{ij}} \log \left[\eta\alpha_{ij}s(t_{ij}^{(\ell)} - 2z_{ij}/c) + B/T_r \right],$$

assuming that $k_{ij} \geq 1$. If $B > 0$, then the ML depth estimate is obtained by solving a non-convex optimization problem. Also, ML estimation requires the knowledge of the true reflectivity α_{ij} , which is not typically available. Thus, the log-matched filter [21] is instead traditionally used for estimating depth from k_{ij} photon detections:

$$\tilde{z}_{ij} = \underset{z_{ij} \in [0, cT_r/2)}{\operatorname{argmax}} \sum_{\ell=1}^{k_{ij}} \log \left[s(t_{ij}^{(\ell)} - 2z_{ij}/c) \right]. \quad (3)$$

We observe that the log-matched filter solution is equal to the constrained ML estimate when $B = 0$.

4. NOVEL IMAGE FORMATION

In the limit of large sample size or high signal-to-noise ratio (SNR), the ML estimate converges to the true parameter value [23]. However, when the data is limited or SNR is low—such as in our problem—pixelwise ML solutions yield inaccurate estimates. We compare our 3D imaging method with the baseline normalized-count reflectivity estimate $\tilde{\alpha}_{ij}$ and the log-matched filter depth estimate \tilde{z}_{ij} , which are ML estimates asymptotically. Along with using the single-photon detection statistics, we exploit the spatial correlations present in real-world scenes by regularizing the ML estimators. Our approach provides significant improvements over pixelwise ML estimators as well as traditional denoising techniques that may exploit scene sparsity but assume additive Gaussian noise. Our computational image construction proceeds in three steps.

1. Reflectivity estimation: The negative log-likelihood of scene reflectivity α_{ij} given count data k_{ij} is

$$\mathcal{L}_\alpha(\alpha_{ij}; k_{ij}) = (N - k_{ij})\eta S\alpha_{ij} - k_{ij} \log[1 - \exp\{-(\eta\alpha_{ij}S + B)\}]$$

after constants independent of α_{ij} are dropped. Since $\mathcal{L}_\alpha(\alpha_{ij}; k_{ij})$ is a strictly convex function in α_{ij} , it is amenable to global minimization using convex optimization, with or without the inclusion of sparsity-based regularization [24]. The penalized ML (PML) estimate for scene reflectivity image is obtained from noisy data $\{k_{ij}\}_{i,j}$ by solving the following convex program:

$$\hat{\alpha}^{\text{PML}} = \underset{\alpha: \alpha_{ij} \geq 0}{\operatorname{argmin}} \sum_{i=1}^n \sum_{j=1}^n \mathcal{L}_\alpha(\alpha_{ij}; k_{ij}) + \beta_\alpha \operatorname{pen}(\alpha),$$

where $\operatorname{pen}(\alpha)$ is a convex function that penalizes the non-smoothness of the estimate and β_α controls the degree of penalization.

2. Rejection of background detections: Direct application of a similar regularized-ML approach to depth estimation using time-of-arrival data is infeasible. This is because the background light contribution to the likelihood function creates a cost function with locally-optimal solutions that are far from the global optimum. Hence, before estimating depth, the second processing step attempts to censor the photon detections that are due to background light.

Detections from background light do not contain any scene-depth information. Their arrival times are mutually independent over spatial locations with variance $T_r^2/12$. In contrast, because light pulses have duration $T_p \ll T_r$ and depths z_{ij} are correlated over spatial locations, the detection times of signal photons have conditional variance, given data from neighboring positions, that is much lower than $T_r^2/12$. Based on this key observation, our method to censor a noisy detection at (i, j) is as follows:

1. Compute the rank-ordered mean (ROM) t_{ij}^{ROM} for each pixel, which is the median value of all the photon arrival times at the 8 neighboring pixels of (i, j) [25]. If t_{ij}^{ROM} cannot be computed due to missing data, then set $t_{ij}^{\text{ROM}} = +\infty$.
2. Estimate the set of uncensored signal photons as follows:

$$G_{ij} = \left\{ \ell : |t_{ij}^{(\ell)} - t_{ij}^{\text{ROM}}| < 2T_p \frac{B}{\eta\hat{\alpha}_{ij}^{\text{PML}}S+B}, \ell = 1, \dots, k_{ij} \right\}.$$

It is demonstrated in [25] that the method of rank-ordered means is effective in detecting pixels that are corrupted by high-variance uniform noise. Since photon detections from background light are uniformly distributed, we use the ROM method to reject such detections and only keep signal detections for further processing.

3. Depth estimation: With background detections rejected, the negative log-likelihood function of depth z_{ij} given data $\{t_{ij}^{(\ell)}\}_{\ell \in G_{ij}}$ is

$$\mathcal{L}_z(z_{ij}; \{t_{ij}^{(\ell)}\}_{\ell \in G_{ij}}) = - \sum_{\ell \in G_{ij}} \log \left[s(t_{ij}^{(\ell)} - 2z_{ij}/c) \right],$$

where $\mathcal{L}_z(z_{ij}; \{t_{ij}^{(\ell)}\}_{\ell \in G_{ij}}) = 0$ if $G_{ij} = \emptyset$. Our framework allows the use of arbitrary pulse shapes, but many practical pulse shapes are well approximated as $s(t) \propto \exp\{-v(t)\}$, where $v(t)$ is a convex function in t . Then, $\mathcal{L}_z(z_{ij}; \{t_{ij}^{(\ell)}\}_{\ell \in G_{ij}}) = \sum_{\ell \in G_{ij}} v(t_{ij}^{(\ell)} - 2z_{ij}/c)$ is a convex function in z_{ij} . Our penalized ML estimate for the scene depth image is thus obtained using uncensored data and solving the following convex optimization problem:

$$\hat{z}^{\text{PML}} = \underset{z: z_{ij} \in [0, cT_r/2)}{\operatorname{argmin}} \sum_{i=1}^n \sum_{j=1}^n \mathcal{L}_z(z_{ij}; \{t_{ij}^{(\ell)}\}_{\ell \in G_{ij}}) + \beta_z \operatorname{pen}(z).$$

5. EXPERIMENTAL RESULTS

To test our robust 3D structure and reflectivity imaging method, we used the dataset collected by D. Venkatraman for [19], which is available from [26]. The experimental setup used to collect data is shown in Fig. 1. A pulsed laser diode with pulse width $T_p = 270$ ps and repetition period $T_r = 100$ ns was used as the illumination

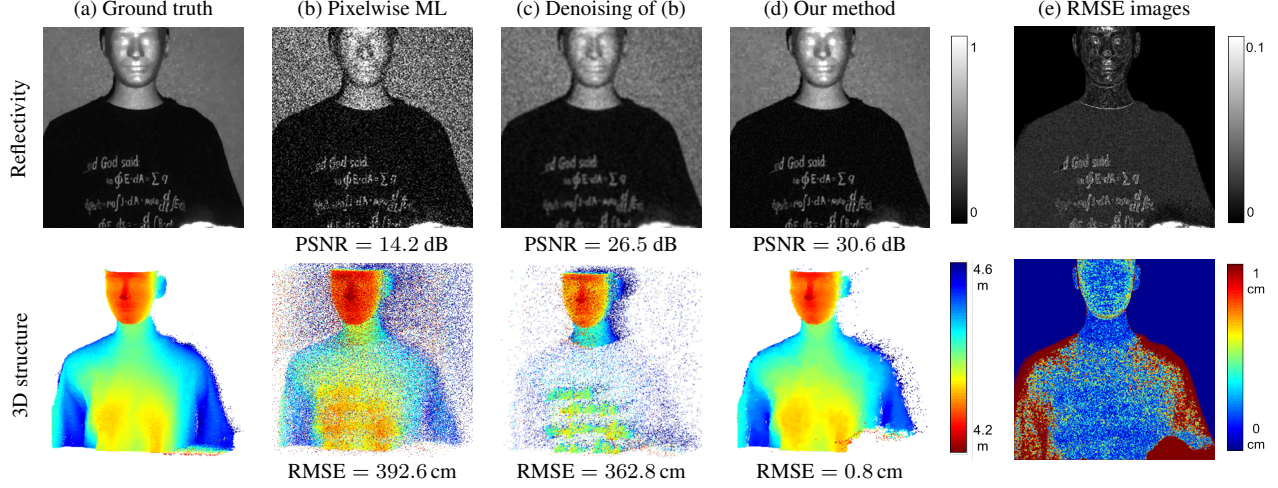


Fig. 3: Experimental results. We compare the 3D and reflectivity reconstruction performance of our proposed imager with pixelwise ML estimation methods (see Section 3). Pixelwise RMSE images of 3D and reflectivity reconstruction using our method were generated with 100 repeated trials of experiments.

source. A two-axis galvo was used to raster scan 1000×1000 pixels. A lensless SPAD detector with quantum efficiency $\eta = 0.35$ was used for detection. The background light level was set such that B equaled the scene-averaged value of $\eta\alpha_{ij}S$. Further details of the experimental setup are given in [19]. Because raster scanning with a fixed dwell time is equivalent to using a floodlight illumination source and a detector array, our experimental results are indicative of what can be accomplished in real-time imaging scenarios using SPAD arrays.

Fig. 3 shows the results of recovering depth and reflectivity of a life-size mannequin using traditional and our imaging methods. Ground truth images were generated using ML estimation from 200 photon detections at each pixel (Fig. 3(a)). The raw dataset for Figs. 3(b), (c) and (d), generated using dwell time $T_a = 100 \mu s$, has 1.21 detected photons per pixel averaged over the scene, with 54% of the pixels having zero detections. All reflectivity images were rescaled to have visual information in the range $[0, 1]$.

We see that the pixelwise ML approach (Eqs. (2), (3)) gives 3D and reflectivity estimates with high root mean-square error (RMSE) and low peak signal-to-noise ratio (PSNR) (Fig. 3(b)), due to signal and background light Poisson noise. Pixels with missing data were imputed with the average of their neighboring 8 pixelwise ML values. Denoising the ML reflectivity estimate using bilateral filtering [27] and the ML depth estimate using median filtering [28] improves their qualities (Fig. 3(c)). However, denoising the 3D structure of the mannequin shirt fails, because the region's very low reflectivity causes many pixels to have missing data. On the other hand, because our framework combines accurate photon-detection statistics with spatial prior information, it accurately reconstructs images with RMSE and PSNR values of 0.8 cm and 30.6 dB (Fig. 3(d)). We used the total variation semi-norm [29] as the penalty function in our method and the penalty parameters β_α and β_z were chosen to maximize PSNR for reflectivity imaging and minimize RMSE for 3D imaging. The pixelwise RMSE images in Fig. 3(e) show that

	Mean T_a	Mean k_{ij}	PSNR	RMSE
FPI	244 μs	1.0 ppp	29.5 dB	0.62 cm
Our method	244 μs	2.3 ppp	35.6 dB	0.30 cm

Table 1: Comparison between first-photon imaging (FPI) [19] and our imaging method using the mannequin dataset, when dwell time averaged over all pixels is equal.




Fig. 4: RMSE simulation results for 3D imaging. The RMSE plots were generated by simulating background levels on the ground truth mannequin dataset. Signal-to-background ratio (SBR) is defined as $\sum_{i,j=1}^n \eta\alpha_{ij}S/(n^2B)$. Note the differences in the colorbar scales.

our computational imager achieves sub-centimeter depth resolution and repeatedly recovers high reflectivity information. Fig. 4 demonstrates that the high photon efficiency of our 3D imager holds for a range of acquisition times and background strengths. Finally, Table 1 shows that our imager has somewhat better performance than the first-photon imager when imaging the mannequin for a fixed mean dwell time. See [9] for additional results and analysis.

6. CONCLUSION AND FUTURE WORK

We have demonstrated accurate 3D structure and reflectivity imaging using on the order of 1 detected ppp averaged over the scene, when any given photon detection has approximately probability 0.5 of originating from background light. Extending the FPI framework from [19], our computational imager is based on combining accurate photon-counting statistics with the spatial prior information in depth and reflectivity maps given a fixed acquisition time. Thus, our work motivates the development of accurate and low-power SPAD array-based 3D and reflectivity imagers.

Our framework can be used in a variety of low light-level imaging applications using photon-counting detectors, such as fluorescence lifetime imaging microscopy (FLIM) [22] and high-resolution LIDAR [8]. Future developments in optoelectronic methods can improve the accuracy of our 3D and reflectivity imager. In particular, our framework can benefit from improved background suppression techniques [2] and range-gating methods [17].

7. REFERENCES

- [1] B. Schwarz, "LIDAR: Mapping the world in 3d," *Nat. Phot.*, 2010.
- [2] A. McCarthy, R. J. Collins, N. J. Krichel, V. Fernández, A. M. Wallace, and G. S. Buller, "Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting," *Appl. Opt.*, vol. 48, no. 32, pp. 6241–6251, 2009.
- [3] S. B. Gokturk, H. Yalcin, and C. Bamji, "A time-of-flight depth sensor — system description, issues and solutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, 2004.
- [4] S. Lee, O. Choi, and R. Horaud, *Time-of-Flight Cameras: Principles, Methods and Applications*, Springer, 2013.
- [5] A. V. Jelalian, "Laser radar systems," in *EASCON'80; Electronics and Aerospace Systems Conference*, 1980, vol. 1, pp. 546–554.
- [6] Z. Zhang, "Microsoft Kinect sensor and its effect," *IEEE Multimedia*, vol. 19, no. 2, pp. 4–10, 2012.
- [7] D. A. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*, Prentice Hall, 2002.
- [8] B. F. Aull, A. H. Loomis, D. J. Young, R. M. Heinrichs, B. J. Felton, P. J. Daniels, and D. J. Landers, "Geiger-mode avalanche photodiodes for three-dimensional imaging," *Lincoln Lab. J.*, vol. 13, no. 2, pp. 335–349, 2002.
- [9] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3d and reflectivity imaging with single-photon detectors," arXiv:1406.1761 [stat.AP], June 2014.
- [10] A. Kirmani, A. Colaço, F. N. C. Wong, and V. K. Goyal, "Exploiting sparsity in time-of-flight range acquisition using a single time-resolved sensor," *Opt. Expr.*, vol. 19, no. 22, pp. 21485–21507, Oct. 2011.
- [11] G. A. Howland, P. B. Dixon, and J. C. Howell, "Photon-counting compressive sensing laser radar for 3d imaging," *Appl. Optics*, vol. 50, no. 31, pp. 5917–5920, Nov. 2011.
- [12] A. Colaço, A. Kirmani, G. A. Howland, J. C. Howell, and V. K. Goyal, "Compressive depth map acquisition using a single photon-counting detector: Parametric signal processing meets sparsity," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Providence, RI, June 2012, pp. 96–102.
- [13] G. A. Howland, D. J. Lum, M. R. Ware, and J. C. Howell, "Photon counting compressive depth mapping," *Opt. Expr.*, vol. 21, no. 20, pp. 23822–23837, Oct. 2013.
- [14] A. Kirmani, D. Venkatraman, A. Colaço, F. N. C. Wong, and V. K. Goyal, "High photon efficiency computational range imaging using spatio-temporal statistical regularization," in *Proc. CLEO*, San Jose, CA, June 2013, paper QF1B.2.
- [15] A. Kirmani, A. Colaço, D. Shin, and V. K. Goyal, "Spatio-temporal regularization for range imaging with high photon efficiency," in *SPIE Wavelets and Sparsity XV*, San Diego, CA, Aug. 2013, pp. 88581F–88581F.
- [16] D. Shin, A. Kirmani, A. Colaço, and V. K. Goyal, "Parametric Poisson process imaging," in *Proc. IEEE Global Conf. Signal Inform. Process.*, Austin, TX, Dec. 2013, pp. 1053–1056.
- [17] J. Busck and H. Heiselberg, "Gated viewing and high-accuracy three-dimensional laser radar," *Appl. Opt.*, vol. 43, no. 24, pp. 4705–4710, 2004.
- [18] G. Goltsman, O. Okunev, G. Chulkova, A. Lipatov, A. Semenov, K. Smirnov, B. Voronov, A. Dzardanov, C. Williams, and R. Sobolewski, "Picosecond superconducting single-photon optical detector," *Appl. Phys. Lett.*, vol. 79, no. 6, pp. 705–707, 2001.
- [19] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. C. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [20] B. I. Erkmén and B. Moision, "Maximum likelihood time-of-arrival estimation of optical pulses via photon-counting photodetectors," in *Proc. IEEE Int. Symp. Inform. Theory*, 2009, pp. 1909–1913.
- [21] D. L. Snyder, *Random Point Processes*, Wiley, New York, 1975.
- [22] Y. Chen, J. D. Müller, P. T. So, and E. Gratton, "The photon counting histogram in fluorescence fluctuation spectroscopy," *Biophys. J.*, vol. 77, no. 1, pp. 553–567, 1999.
- [23] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*, Prentice Hall PTR, 1998.
- [24] Z. T. Harmany, R. F. Marcia, and R. M. Willett, "This is SPIRAL-TAP: Sparse Poisson intensity reconstruction algorithms—theory and practice," *IEEE Trans. Image Process.*, vol. 21, no. 3, pp. 1084–1096, 2012.
- [25] E. Abreu, M. Lightstone, S. K. Mitra, and K. Arakawa, "A new efficient approach for the removal of impulse noise from highly corrupted images," *IEEE Trans. Image Process.*, vol. 5, no. 6, pp. 1012–1025, 1996.
- [26] "First-photon imaging project," <http://www.rle.mit.edu/first-photon-imaging/>.
- [27] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *Proc. 6th Int. Conf. Comput. Vis.*, 1998, pp. 839–846.
- [28] R. Jain, R. Kasturi, and B. G. Schunck, *Machine Vision*, vol. 5, McGraw-Hill New York, 1995.
- [29] S. Osher, A. Solé, and L. Vese, "Image decomposition and restoration using total variation minimization and the H^{-1} norm," *Multiscale Model. Simul.*, vol. 1, no. 3, pp. 349–370, 2003.