

Single-Photon Depth Imaging Using a Union-of-Subspaces Model

Donggeek Shin, *Student Member, IEEE*, Jeffrey H. Shapiro, *Life Fellow, IEEE*, and Vivek K Goyal, *Fellow, IEEE*

Abstract—Light detection and ranging systems reconstruct scene depth from time-of-flight measurements. For low light-level depth imaging applications, such as remote sensing and robot vision, these systems use single-photon detectors that resolve individual photon arrivals. Even so, they must detect a large number of photons to mitigate Poisson shot noise and reject anomalous photon detections from background light. We introduce a novel framework for accurate depth imaging using a small number of detected photons in the presence of an unknown amount of background light that may vary spatially. It employs a Poisson observation model for the photon detections plus a union-of-subspaces constraint on the discrete-time flux from the scene at any single pixel. Together, they enable a greedy signal-pursuit algorithm to rapidly and simultaneously converge on accurate estimates of scene depth and background flux, without any assumptions on spatial correlations of the depth or background flux. Using experimental single-photon data, we demonstrate that our proposed framework recovers depth features with 1.7 cm absolute error, using 15 photons per image pixel and an illumination pulse with 6.7-cm scaled root-mean-square length. We also show that our framework outperforms the conventional pixelwise log-matched filtering, which is a computationally-efficient approximation to the maximum-likelihood solution, by a factor of 6.1 in absolute depth error.

Index Terms—Computational imaging, greedy algorithms, LIDAR, single-photon imaging, union-of-subspaces.

I. INTRODUCTION

A CONVENTIONAL light detection and ranging (LIDAR) system, which uses a pulsed light source and a single-photon detector, forms a depth image pixelwise using the histograms of photon detection times. The acquisition times for such systems are made long enough to detect hundreds of photons per pixel for the finely binned histograms these systems require to do accurate depth estimation. Here, we introduce a framework for accurate depth imaging using only a small

number of photon detections per pixel, despite the presence of an unknown amount of spatially-varying background light in the scene. We use a Poisson observation model for the photon detections plus a union-of-subspaces constraint on the scene's discrete-time flux at any single pixel. Using a greedy signal-pursuit algorithm—a modification of CoSaMP [1]—we solve for accurate estimates of scene depth and background flux. Our method forms estimates pixelwise and thus avoids assumptions on transverse spatial correlations that may hinder the ability to resolve very small features. We experimentally demonstrate that our proposed depth imaging framework outperforms log-matched filtering, which is the maximum-likelihood (ML) depth estimator given zero background light.

The conventional estimation of depth using histograms of photon detections is accurate when the number of detections is high. In the low photon-count regime, the depth solution is noisy due to shot noise. It has been shown that image denoising methods, such as wavelet thresholding, can boost the performance of scene depth recovery in the presence of background noise [2]. Also, using an imaging model that incorporates occlusion constraints was proposed to recover an accurate depth map [3]. However, these denoising algorithms implicitly assume that the observations are Gaussian distributed. Thus, at low photon-counts, where depth estimates are highly non-Gaussian [4], their performance degrades significantly [5].

First-photon imaging (FPI) [6] is a framework that allows high-accuracy imaging using only the first detected photon at every pixel. It demonstrated that centimeter-accurate depth recovery is possible by combining the non-Gaussian statistics of first-photon detection with spatial correlations of natural scenes. The FPI framework uses an imaging setup that includes a raster-scanning light source and a lensless single-photon detector. More recently, photon-efficient imaging frameworks that use a detector array setup, in which every pixel has the same acquisition time, have also been proposed [5], [7], [8].

Prior methods have two common limitations that we avoid:

- **Over-smoothing:** Many methods assume spatial smoothness of the scene to mitigate the effect of shot noise. In some applications, it is important to capture features that only occupy a few image pixels. Methods that assume spatial correlations may yield erroneously over-smoothed images that wash out the scene's fine-scale features. In such scenarios, a robust *pixelwise* imager is preferable.
- **Calibration:** Many methods assume a calibration step to measure the amount of background flux existing in the environment. This calibration mitigates bias in the depth estimate caused by background-photon or dark-count detections, which have high temporal variance. In practical

Manuscript received July 24, 2015; revised August 27, 2015; accepted August 28, 2015. Date of publication September 01, 2015; date of current version September 04, 2015. This work was supported in part by a Samsung Scholarship, by the U.S. National Science Foundation under Grant 1422034, and by the MIT Lincoln Laboratory Advanced Concepts Committee. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Jie Liang.

D. Shin and J. H. Shapiro are with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

V. K. Goyal is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LSP.2015.2475274

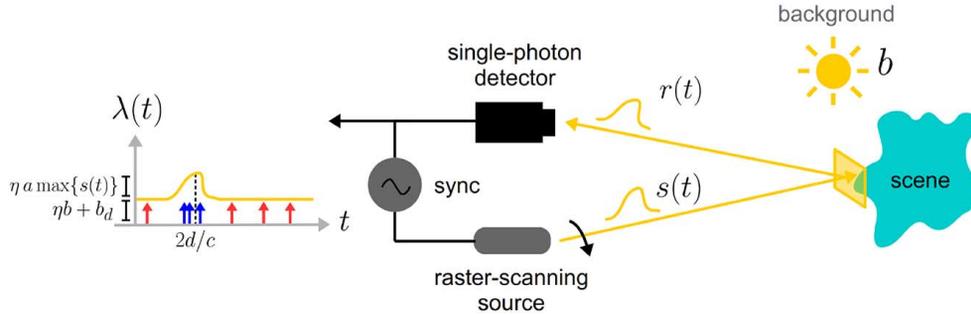


Fig. 1. An illustration of the single-photon imaging setup for one illumination pulse. A pulsed optical source illuminates a scene pixel with photon-flux waveform $s(t)$. The flux waveform $r(t)$ that is incident on the detector consists of the pixel return $as(t - 2d/c)$ —where a is the pixel reflectivity, d is the pixel depth, and c is light speed—plus the background-light flux b . The rate function $\lambda(t)$ driving the photodetection process equals the sum of the pixel return and background flux, scaled by the detector efficiency η , plus the detector’s dark-count rate b_d . The record of detection times from the pixel return (or background light plus dark counts) is shown as blue (or red) spikes, generated by the Poisson process driven by $\lambda(t)$.

imaging scenarios, however, the background response varies in time, and continuous calibration may not be practical. Furthermore, many methods assume background flux does not vary spatially. Thus, a *calibrationless* imager that performs simultaneous estimation of scene parameters and spatially-varying background flux from photon detections is useful.

Similar to [3], we use a union-of-subspaces constraint for modeling the scene parameters. Our union-of-subspaces constraint is defined for both the signal and background waveform parameters that generate photon detections; the framework in [3] assumes a system observing a noiseless signal waveform, not one corrupted by photon noise. We propose a greedy signal pursuit algorithm that accurately solves for the scene parameters at each pixel. We evaluate the photon efficiency of this framework using experimental single-photon data. In the presence of strong background light, we show that our pixelwise imager gives an absolute depth error that is 6.1 times lower than that of the pixelwise log-matched filter.

II. SINGLE-PHOTON IMAGING SETUP

Fig. 1 illustrates our imaging setup, for one illumination pulse, when the scene is illuminated in raster-scanning manner and a single-element photon detector is employed. (Alternatively, to reduce the time needed to acquire a depth map, our framework can be applied without modification when the scene is flood illuminated and a detector array is used.) A focused optical source, such as a laser, illuminates a pixel of the scene with the pulse waveform $s(t)$ that starts at time 0 and has root-mean-square pulsewidth T_p . This illumination is repeated every T_r seconds for a sequence of N_s pulses. The single-photon detector, in conjunction with a time correlator, is used to time stamp individual photon detections, relative to the time at which the immediately preceding pulse was transmitted. These detection times, which are observations of a time-inhomogeneous Poisson process, whose rate function combines contributions from pixel return, background light, and dark counts, are used to estimate scene depth for the illuminated pixel. This pixelwise acquisition process is repeated for $N_x \times N_y$ image pixels by raster scanning the light source in the transverse directions.

III. FORWARD IMAGING MODEL

For simplicity of exposition and notation, we focus on one pixel; this is repeated for each pixel of a raster-scanning or array-detection setup. Let a , d , and b be unknown scalar values that represent reflectivity, depth, and background flux at the given pixel. The reflectivity value includes the effects of radial fall-off, view angle, and material properties. After illuminating the scene pixel with a single pulse $s(t)$, the backreflected waveform that is incident at the single-photon detector is

$$r(t) = as(t - 2d/c) + b, \quad t \in [0, T_r]. \quad (1)$$

Using (1), we observe that the rate function that generates the photon detections is

$$\lambda(t) = \eta(as(t - 2d/c) + b) + b_d, \quad t \in [0, T_r], \quad (2)$$

where $\eta \in (0, 1]$ is the quantum efficiency of the detector and $b_d \geq 0$ is the dark-count rate of the single-photon detector.

Let Δ be the time bin duration of the single-photon detector. Then $M = \lceil T_r/\Delta \rceil$ is the total number of time bins that capture photon detections. Let \mathbf{y} be the vector of size $M \times 1$ that contains the number of photon detections at each time bin after we illuminate the pixel N_s times with pulse waveform $s(t)$. Then, from photodetection theory [9], we have that

$$\mathbf{y}_k \sim \text{Poisson} \left(N_s \int_{(k-1)\Delta}^{k\Delta} [\eta(as(t - 2d/c) + b) + b_d] dt \right), \quad (3)$$

for $k = 1, \dots, M$. We have assumed that our total pixelwise acquisition time $N_s T_r$ is short enough that b is constant during that period, the low-flux condition ensures that $\sum_{k=1}^M \mathbf{y}_k < N_s$, and the effect of reset time of the single-photon detector is negligible. We wish to reach an approximation in which the Poisson parameter of \mathbf{y}_k is given by the product of a known matrix and an unknown (and constrained) vector.

Choose $N \in \mathbb{Z}^+$ such that $\epsilon = T_r/N$ is adequate resolution for the estimated time of flight. (Our interest is in $N \geq M$ and hence $\epsilon \leq \Delta$.) Since $2d/c \in [0, T_r)$, $\mathbf{v} \in \mathbb{R}^N$ defined by

$$\mathbf{v}_j = \begin{cases} N_s \eta a, & \text{if } 2d/c \in [(j-1)\epsilon, j\epsilon); \quad j = 1, 2, \dots, N, \\ 0, & \text{otherwise,} \end{cases}$$

has exactly one nonzero entry. Using this vector,

$$N_s \eta a s(t - 2d/c) \approx \sum_{j=1}^N \mathbf{v}_j s\left(t - \left(j - \frac{1}{2}\right)\epsilon\right) \quad (4)$$

is a good approximation when ϵ is small enough; effectively, $2d/c$ has been quantized to an interval of length ϵ . Substituting (4) into the Poisson parameter expression in (3) gives

$$\sum_{j=1}^N \left(\int_{(j-1)\Delta}^{j\Delta} s\left(t - \left(j - \frac{1}{2}\right)\epsilon\right) dt \right) \mathbf{v}_j + N_s \Delta(\eta b + b_d).$$

Then, we can rewrite (3) as

$$\mathbf{y}_k \sim \text{Poisson}((\mathbf{S}\mathbf{v} + B\mathbf{1}_{M \times 1})_k), \quad (5)$$

for $k = 1, \dots, M$, where $\mathbf{1}_{M \times 1}$ is an $M \times 1$ vector of 1's, $\mathbf{S}_{i,j} = \int_{(i-1)\Delta}^{i\Delta} s(t - (j - \frac{1}{2})\epsilon) dt$, and $B = N_s \Delta(\eta b + b_d)$. Finally, defining $\mathbf{A} = [\mathbf{S}, \mathbf{1}_{M \times 1}]$ and $\mathbf{x} = [\mathbf{v}^T, B]^T$, we can further rewrite (5) as

$$\mathbf{y}_k \sim \text{Poisson}((\mathbf{A}\mathbf{x})_k). \quad (6)$$

Since \mathbf{v} has exactly one nonzero entry, \mathbf{x} lies in \mathcal{S}_N , the union of N subspaces defined as

$$\mathcal{S}_N = \bigcup_{k=1}^N \{\mathbf{x} \in \mathbb{R}^{N+1} : \mathbf{x}_{\{1,2,\dots,N\} \setminus \{k\}} = \mathbf{0}\}, \quad (7)$$

where each subspace is of dimension 2.

IV. SOLVING THE INVERSE PROBLEM

We have interpreted the problem of robust single-photon depth imaging as a noisy linear inverse problem, where the signal of interest \mathbf{x} lies in the union-of-subspaces \mathcal{S}_N . Using (6), the observed photon count histogram \mathbf{y} has the probability mass function

$$p_Y(\mathbf{y}; \mathbf{A}, \mathbf{x}) = \prod_{k=1}^M e^{-(\mathbf{A}\mathbf{x})_k} (\mathbf{A}\mathbf{x})_k^{\mathbf{y}_k} / \mathbf{y}_k!. \quad (8)$$

Thus, neglecting terms in the negative log-likelihood function that are dependent on \mathbf{y} but not on \mathbf{x} , we define

$$\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y}) = \sum_{k=1}^M [(\mathbf{A}\mathbf{x})_k - \mathbf{y}_k \log(\mathbf{A}\mathbf{x})_k]. \quad (9)$$

This objective function can be proved to be convex in \mathbf{x} .

We solve for \mathbf{x} by minimizing $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$ with the constraint that \mathbf{x} lies in the union-of-subspaces \mathcal{S}_N . Also, because photon flux is a non-negative quantity, the minimization results in a more accurate estimate when we include a non-negative signal constraint. Thus, we wish to solve

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} \quad \mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y}) \\ & \text{s.t.} \quad \mathbf{x} \in \mathcal{S}_N, \quad \mathbf{x}_i \geq 0, \quad i = 1, \dots, (N+1). \end{aligned} \quad (10)$$

We propose an algorithm that is inspired by compressive sampling matching pursuit (CoSaMP) [1], a greedy algorithm that finds a K -sparse approximate solution to an underdetermined linear inverse problem. CoSaMP iterates until it finds a solution that agrees with the observed data (according to some convergence metric), while the solution is a linear combination of K columns of the forward matrix \mathbf{A} . Unlike algorithms that only add to the solution support, never culling, CoSaMP has solution stability and accuracy properties that compete with globally-optimal ℓ_1 -based convex optimization methods for sparse approx-

imation [10]. Also, CoSaMP has been shown to be adaptable to applications in which the signal being estimated has a structured support [11], as is true for the union-of-subspaces model. Thus, we modified the CoSaMP algorithm to our specific use case, where we are interested in recovering a sparse solution in the union-of-subspaces \mathcal{S}_N using photon-noise corrupted data.

Our greedy algorithm is given in Algorithm 1. We define $\mathcal{T}_0(\mathbf{x})$ to be the thresholding operator setting all negative entries of \mathbf{x} to 0, $\text{supp}(\mathbf{x})$ to be the support of \mathbf{x} , and $\mathbf{x}_{[k]}$ to be the vector that approximates \mathbf{x} with its k largest terms. Also, we take \mathbf{A}_S to be a matrix with columns of \mathbf{A} chosen by the index set S . Finally, we use \mathbf{A}^T and \mathbf{A}^\dagger to denote the transpose and pseudo-inverse of matrix \mathbf{A} , respectively.

Algorithm 1 Depth imaging using a union-of-subspaces model

Input: $\mathbf{y}, \mathbf{A}, \delta$

Output: $\mathbf{x}^{(k)}$

Initialize $\mathbf{x}^{(0)} \leftarrow \vec{0}$, $\mathbf{u} \leftarrow \mathbf{y}$, $k \leftarrow 0$;

repeat

$k \leftarrow k + 1$;

$\hat{\mathbf{x}} \leftarrow \mathbf{A}^T \mathbf{u}$;

$\Omega \leftarrow \text{supp}((\hat{\mathbf{x}}_{1:N})_{[1]}) \cup \text{supp}(\mathbf{x}_{1:N}^{(k-1)}) \cup \{N+1\}$;

$\mathbf{b}|_\Omega \leftarrow \mathbf{A}_\Omega^\dagger \mathbf{y}$; $\mathbf{b}|_{\Omega^c} \leftarrow \mathbf{0}$;

$\mathbf{x}^{(k)} \leftarrow \mathcal{T}_0([\mathbf{b}_{1:N}]_{[1]}, \mathbf{b}_{N+1}]^T) \quad \triangleright$ Update solution

$\mathbf{u} \leftarrow \mathbf{y} - \mathbf{A}\mathbf{x}^{(k)}$

until $\|\mathbf{x}^{(k-1)} - \mathbf{x}^{(k)}\|_2^2 < \delta$

In Algorithm 1, for computational efficiency we have approximated $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$ with the ℓ_2 -loss $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$, which is the first-order Taylor expansion of $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$ up to a constant. Because CoSaMP also assumes an ℓ_2 -loss function, the only change from CoSaMP is then the update stage; instead of picking out the best k terms, we pick out the two terms from the intermediate solution based on the union-of-subspaces and non-negativity constraints. We iterate until the solution meets the convergence criterion: $\|\mathbf{x}^{(k-1)} - \mathbf{x}^{(k)}\|_2^2 < \delta$.

Many sparse pursuit algorithms, such as CoSaMP, are guaranteed to be successful when \mathbf{A} is incoherent. In our setup, however, \mathbf{A} is highly coherent due to ϵ being small and the pulse waveform $s(t)$ being smooth. Nevertheless, because the linear system's degree of underdetermination is extremely mild ($\mathbf{A} \in \mathbb{R}^{N \times (N+1)}$) and the sparsity level is fixed to a small number ($\dim(\mathcal{S}_N) = 2$) relative to the signal dimension (typically exceeding 100), our algorithm recovers the scene parameters of interest in a robust manner.

V. EXPERIMENTAL RESULTS

To validate our framework, we used a dataset collected by D. Venkatraman for the First-Photon Imaging project [6]; this dataset and others are available from [12]. The experimental setup uses a pulsed laser diode with pulsewidth $T_p = 270$ ps and repetition period $T_r = 100$ ns. A two-axis galvo was used to scan 350×350 pixels of a mannequin face at a distance of about 4 m. A lensless single-photon avalanche diode detector with quantum efficiency $\eta = 0.35$ was used for detection. The

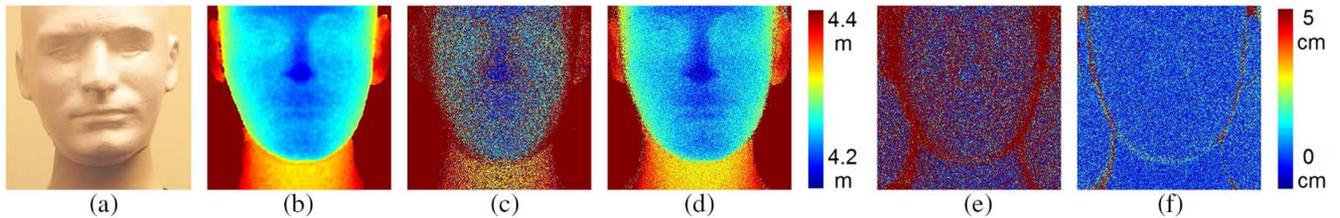


Fig. 2. Experimental pixelwise depth imaging results using single photon observations. The number of photon detections at every pixel was set to be 15. The figure shows the (a) photograph of imaged face, (b) ground-truth depth, (c) depth from log-matched filtering, which is approximately ML, and (d) depth using our method. The absolute depth-error maps for ML and our framework are shown in (e) and (f), respectively. (a) Photograph. (b) Truth. (c) Log-matched filter. (d) Proposed. (e) Error of (c). (f) Error of (d).

background light level was set using an incandescent lamp. The original mannequin data from [12] had the background count rate approximately equal to the signal count rate. Our experiment uses cropped data showing only the mannequin’s face, where the background count rate was approximately 0.1 of the average signal count rate. Although we used raster-scanning for our experiments, since our algorithm is applied pixelwise, it can be also used for imaging with a floodlight illumination source and a detector array.

We could compare our imaging method with the ML estimator for scene parameters $\{a, d, b\}$. Unfortunately, due to nonzero background flux, ML estimation requires minimizing a non-convex cost function, leading to a solution without convergence and accuracy guarantees. Thus, zero background is assumed conventionally such that the ML depth estimate reduces to the simple log-matched filter [13]:

$$\hat{d}_{\text{ML}} = \frac{1}{2} \epsilon \left(\arg \max_{i \in \{1, \dots, n\}} \log \mathbf{S}_i^T \mathbf{y} \right). \quad (11)$$

Note that this is equivalent to a one-step greedy algorithm (where a union-of-subspaces constraint is irrelevant) of minimizing $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$ for a 1-sparse solution. We use (11) as the baseline depth estimator that is compared with our proposed estimator using the union-of-subspaces model.

Fig. 2 shows the results of recovering depth of the mannequin face using single-photon observations. The kernel matrix \mathbf{S} was obtained by an offline measurement of the pulse shape. Note that this measurement depends only on the source, not on properties of the scene. The ground-truth depth, shown in Fig. 2(b), was generated separately by using background-calibrated ML estimation from 200 photons at each pixel.

In our depth imaging experiment, the number of photon detections at each pixel was set to 15. We observe that, due to extraneous background photon detections, the log-matched filter estimate in Fig. 2(c) (average absolute error = 10.3 cm) is corrupted with high-variance noise and the facial features of the mannequin are heavily obscured. On the other hand, our estimate, shown in Fig. 2(d), shows high-accuracy depth recovery (average absolute error = 1.7 cm). As shown by the error maps in Fig. 2(e), (f), both methods fail in depth recovery in the face boundary regions, where very little light is reflected back from the scene to the single-photon detector. This is because the signal-to-background ratio (SBR), which is the ratio of the probability of a detection coming from signal and the probability of a detection coming from background+dark counts, is very low in such regions. Also, we observe that our estimated average background level over all pixels was $\hat{B} = 1.4 \times 10^{-3}$,

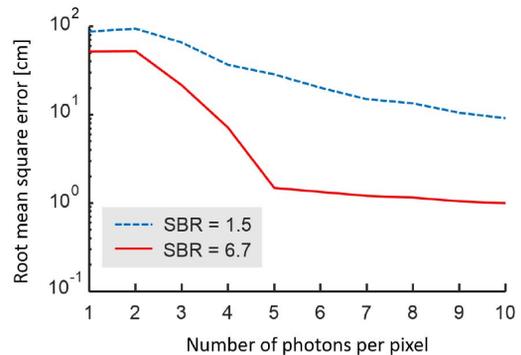


Fig. 3. Depth recovery performance of our algorithm at a face pixel (SBR = 6.7 and pixel coordinates (81,272)) and a depth-boundary pixel (SBR = 1.5 and pixel coordinates (237,278)) for varying numbers of photon detections.

which is very close to the calibrated true background level $B = 1.3 \times 10^{-3}$.

Fig. 3 shows how our depth reconstruction algorithm performs with varying numbers of photon detections for two different pixels, one in the facial region (with SBR 6.7) and one at the face boundary (with SBR 1.5). We observe that the algorithm performs better for higher SBR overall, and that the rate of decrease in depth error with increasing number of photon detections is faster for high SBR than for low SBR, especially at the very low-flux regime (2 to 5 detections). In this experiment, we had $M = N = 801$. Also, we set $\delta = 10^{-4}$ and the average number of iterations until convergence was measured to be 2.1 over all pixels. Code and data used to generate results can be downloaded from [14].

VI. CONCLUSIONS AND FUTURE WORK

We presented an imaging framework for calibrationless, pixelwise depth reconstruction using single-photon observations. Our method combined photon detection statistics with discrete-time flux constraints expressed using a union-of-subspaces model. We developed a greedy algorithm that recovers scene depth by solving a constrained optimization problem.

Our framework can be used in low light-level imaging applications, where the scene being imaged has fine features and filtering techniques that exploit patchwise smoothness can potentially wash out those details. For example, it can be useful in applications such as airborne remote sensing [15], where the aim is to recover finely-featured 3D terrain maps.

A straightforward generalization is to multiple-depth estimation, where more than one reflector may be present at each pixel. For K reflectors at a pixel, 1-sparsity must be changed to a K -sparsity when defining the union-of-subspaces.

REFERENCES

- [1] D. Needell and J. A. Tropp, "CoSaMP: Iterative signal recovery from incomplete and inaccurate samples," *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, 2009.
- [2] B.-Y. Sun, D.-S. Huang, and H.-T. Fang, "Lidar signal denoising using least-squares support vector machine," *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 101–104, 2005.
- [3] P. T. Boufounos, "Depth sensing using active coherent illumination," in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process*, 2012, pp. 5417–5420.
- [4] 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.
- [5] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Photon-efficient computational 3-d and reflectivity imaging with single-photon detectors," *IEEE Trans. Comput. Imag.*, vol. 1, no. 2, pp. 112–125, Jun. 2015.
- [6] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. Wong, J. H. Shapiro, and V. K. Goyal, "First-photon imaging," *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [7] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, "Computational 3d and reflectivity imaging with high photon efficiency," in *Proc. IEEE Int. Conf. Image Process*, Oct. 2014, pp. 46–50.
- [8] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, "Lidar waveform based analysis of depth images constructed using sparse single-photon data," *arXiv preprint arXiv:1507.02511*, 2015.
- [9] D. L. Snyder, *Random Point Processes*. New York, NY, USA: Wiley, 1975.
- [10] T. Goldstein and S. Osher, "The split Bregman method for L1-regularized problems," *SIAM J. Imag. Sci.*, vol. 2, no. 2, pp. 323–343, 2009.
- [11] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1982–2001, 2010.
- [12] GitHub repository for photon-efficient imaging [Online]. Available: <https://github.com/photon-efficient-imaging/sample-data/>
- [13] R. E. Blahut, *Principles and Practice of Information Theory*. Reading, MA, USA: Addison-Wesley, 1987.
- [14] GitHub repository for union-of-subspace imaging [Online]. Available: <https://github.com/photon-efficient-imaging/uos-imaging/>
- [15] M. Nilsson, "Estimation of tree heights and stand volume using an airborne lidar system," *Remote Sens. Environ.*, vol. 56, no. 1, pp. 1–7, 1996.