

Robust Spectral Unmixing of Sparse Multispectral Lidar Waveforms Using Gamma Markov Random Fields

Yoann Altmann ¹, Member, IEEE, Aurora Maccarone, Aongus McCarthy ², Member, IEEE, Gregory Newstadt, Gerald S. Buller, Steve McLaughlin, Fellow, IEEE, and Alfred Hero, Fellow, IEEE

Abstract—This paper presents a new Bayesian spectral unmixing algorithm to analyze remote scenes sensed via sparse multispectral Lidar measurements. To a first approximation, in the presence of a target, each Lidar waveform consists of a main peak, whose position depends on the target distance and whose amplitude depends on the wavelength of the laser source considered (i.e., on the target reflectivity). Besides, these temporal responses are usually assumed to be corrupted by Poisson noise in the low photon count regime. When considering multiple wavelengths, it becomes possible to use spectral information in order to identify and quantify the main materials in the scene, in addition to estimation of the Lidar-based range profiles. Due to its anomaly detection capability, the proposed hierarchical Bayesian model, coupled with an efficient Markov chain Monte Carlo algorithm, allows robust estimation of depth images together with abundance and outlier maps associated with the observed three-dimensional scene. The proposed methodology is illustrated via experiments conducted with real multispectral Lidar data acquired in a controlled environment. The results demonstrate the possibility to unmix spectral responses constructed from extremely sparse photon counts (less than 10 photons per pixel and band).

Index Terms—Anomaly detection, depth imaging, Markov chain Monte Carlo, multispectral Lidar, robust spectral unmixing.

I. INTRODUCTION

LASER altimetry (or Lidar) is an acknowledged tool for extracting spatial structure from three-dimensional (3D) scenes. Using time-of-flight to create a distance profile, signal analysis can recover, for instance, tree and canopy heights, leaf

area indices and ground slope by analyzing the reflected photons from a target. Conversely, passive multispectral (MSI) and hyperspectral images (HSI) are widely used to extract spectral information about the scene which can provide useful parameters about the composition and health condition of the canopy. The most natural method to extract spatial and spectral information from sensed scenes is to couple Lidar data and multi/hyperspectral images [1], [2]. Although the fusion of Lidar data and HSIs can improve scene characterization, data synchronization issues in space (alignment, resolution) and time (dynamic scene, change of observation conditions, etc.) make this problematic and these are still open issues. For these reasons, multispectral Lidar (MSL) has recently received attention from the remote sensing community for its ability to extract both spatial and spectral information from 3D scenes. The key advantage of MSL is the ability to potentially provide information on the full 3D distribution of materials, especially for scenes including semi-transparent objects (e.g., vegetation or fences). When the Lidar return signal is sufficiently strong the received light field will exhibit easily separable spatial and spectral peaks corresponding to the different surfaces and material properties. In this case, classical methods for 3D reconstruction can be applied [3]–[6]. For instance, in [4], for each waveform, a series of peak positions and reflectivity parameters is estimated iteratively by identifying and subtracting sequentially the peak with the highest amplitude, until a pre-defined threshold is reached. The pixels and spectral bands are processed independently, leading to one point cloud per spectral band. In [3], [5], a Bayesian approach is adopted to first estimate the number and positions of the peaks (from a single band) and these parameters are then used to estimate the reflectivity parameters associated with the remaining spectral bands. Another motivation for MSL is that HSIs, even when fully synchronized, can only integrate the spectral response along the path of each optical ray, and cannot measure the spectral response as a function of distance, e.g., depth into a forest canopy.

In [5]–[7], spectral unmixing techniques were developed to analyze 3D scenes composed of multi-layered objects, assuming that the spectral signatures of the materials composing the scenes were known and assuming linear mixing processes. In contrast to [5], [7], the unmixing method proposed in [6], which consists of a Metropolis-within-Gibbs sampler, assumes that there is either a single peak (whose position is estimated) or

Manuscript received October 13, 2016; revised February 21, 2017 and April 19, 2017; accepted April 19, 2017. Date of publication May 10, 2017; date of current version November 6, 2017. This work was supported in part by EPSRC via Grants EP/J015180/1, EP/K015338/1, and EP/M01326X/1 and in part by US DOE Consortium for Verification Technology DE-NA0002534 and US Army Grant W911NF-14-1-0479. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Brendt Wohlberg. (Corresponding author: Yoann Altmann.)

Y. Altmann, A. Maccarone, A. McCarthy, G. S. Buller, and S. McLaughlin are with the School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, U.K. (e-mail: Y.Altmann@hw.ac.uk; am827@hw.ac.uk; A.McCarthy@hw.ac.uk; G.S.Buller@hw.ac.uk; S.McLaughlin@hw.ac.uk).

G. Newstadt is with the Google Inc., Pittsburgh, PA 15206 USA (e-mail: newstage37@gmail.com).

A. Hero is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: hero@eecs.umich.edu).

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/TCI.2017.2703144

multiple peaks whose positions are known. This method has been extended in [8] to account for and identify possible deviations from the classical linear mixing model (LMM) used to estimate the amount/abundances of each endmember (assumed known) present in the scene. In this work, we further improve the robust unmixing method of [8] in order to enhance abundance and range profile estimation when there are extremely low photon counts per pixel, which is relevant to situations where the acquisition time is very restricted, e.g., in extreme low light imaging. Specifically, we propose a new abundance prior model which promotes smooth abundance maps and we propose a depth prior model that promotes piece-wise homogeneous depth profiles. As will be shown in Section VI of this paper, by accounting for the intrinsic spatial organization of natural images, our algorithm provides significantly better ranging and unmixing results when compared to pixel-wise estimation algorithms, in particular when considering Lidar with extremely low photon counts.

Single-photon Lidar MSL systems usually record, for each pixel/region of the scene, a histogram of time delays between emitted laser pulses and the detected photon arrivals. Due to the low number of photon arrivals detected, Poisson noise models are more appropriate for single-photon MSL data than Gaussian noise models that are typically used for HSIs in the high photon density regime. Such models are particularly relevant for challenging scenarios where the recorded waveforms consists of very few photons, i.e., less than 10 on average across the image pixels for each wavelength, which occurs when reducing the overall acquisition time. In this paper, we demonstrate the ability of robust Bayesian Poisson unmixing methods to simultaneously estimate endmember fractions, extract depth information, and detect anomalous regions that are poorly represented by the assumed nominal LMM. The performance is illustrated on an experimental testbed with clay objects of different colours and an MSL imaging system (33 wavelengths ranging from 500 nm to 820 nm), under favourable observation conditions.

The benefit of Bayesian approaches is that prior distributions can be chosen for the unknown parameters of the model, effectively smoothing the unmixing solution through regularization. For example, a total-variation (TV) regularization is proposed in [9], [10] to regularize the depth estimation problem and gamma Markov random fields (MRFs) are used to model the spatial dependencies affecting the unknown abundance maps. As illustrated in [6], such Markovian models can be used to promote local spatial smoothness of the estimated abundances while providing enough flexibility to handle sharp transitions. Gamma-MRFs have the nice property of leading to conjugate prior models under Poisson noise assumption.

To allow for moderate deviations from the linear mixing model we propose a sparse anomaly model within a hierarchical Bayesian framework. Such anomalies can occur in the presence of scarcely represented materials or when there are local variations in the main spectral signatures of the scene. To capture these anomalies, we use a 3D Ising model for sparse deviations from the standard LMM. Although the proposed method is able to detect deviations from the LMM that are not necessarily sparse (such as endmember variability or mis-specification

and nonlinear mixtures), such deviations are likely to yield locally biased abundance estimates. Nonetheless, the method that we propose can potentially be used to identify pixels for which more complex mixing models might be used.

The proposed prior models are incorporated into a hierarchical Bayesian model and the joint posterior distribution of these parameters is derived using Bayes' theorem. A simulation-based method is then developed to estimate the unknown parameters. More precisely, we construct a stochastic gradient Markov chain Monte Carlo (SG-MCMC) algorithm to jointly generate samples according to the posterior of interest and adjust the MRF hyperparameters during the burn-in period of the sampler. This strategy has several important advantages in the context of estimating MRF hyperparameters whose conditional distributions are highly complex and have no closed-form expressions. Firstly, it allows for the automatic adjustment of the hyperparameters for each dataset and, secondly it has a computational cost that is several times lower than that of competing approaches, such as those that include the hyperparameters in the Bayesian model through hierarchical priors. The proposed Bayesian approach produces a predictor of optimal estimator performance (through the derivation of posterior measures of uncertainty) while reducing potential convergence issues arising from the non-concavity of the log-posterior (due to the presence of anomaly terms). More sophisticated optimization techniques, e.g., variational methods, are worth additional study but are out-with the scope of this current paper.

The remainder of the paper is organized as follows. Section II introduces the observation model associated with MSL returns for a single-layered object to be analyzed. Section III presents the hierarchical Bayesian model associated with the robust spectral unmixing problem considered and the associated posterior distribution. Section IV describes the SG-MCMC method used to sample from the posterior of interest and subsequently approximate appropriate Bayesian estimators. The relation between the proposed model and the Poisson Factor Analysis is addressed in Section V. Results of experiments conducted on real MSL data are shown and discussed in Section VI and conclusions are reported in Section VII.

II. PROBLEM FORMULATION

This section introduces the statistical model associated with MSL returns for a single-surface reflecting object which will be used in Section III for robust spectral unmixing of MSL data. We consider a 4-D array \mathbf{Y} of Lidar waveforms of dimension $N_{\text{row}} \times N_{\text{col}} \times L \times T$, where N_{row} and N_{col} stands for the number of rows and columns of the regular spatial sampling grid (in the transverse plane), L is the number of spectral bands or wavelengths used to reconstruct the scene and T is the number of temporal (corresponding to range) bins. Let $\mathbf{y}_{i,j,\ell} = [\mathbf{Y}]_{i,j,\ell,t} = [y_{i,j,\ell,1}, \dots, y_{i,j,\ell,T}]^T$ be the Lidar waveform obtained in the pixel (i, j) using the ℓ th wavelength. The element $y_{i,j,\ell,t}$ is the photon count within the t th bin of the ℓ th spectral band considered. Due to the design of the proposed experiments (performed indoors in this work), and to simplify the estimation problem, we further assume that ambient noise

counts (e.g., from additional illumination sources or dark counts from the detector) can be neglected. Thus, for each pixel, the detected photons (originally emitted by the laser sources) only result from direct path reflections onto the surface of the object of interest. Moreover, we assume that the laser beam (for each pixel) encounters a single surface which is assumed to be locally orthogonal to the beam direction. This is typically the case for short to mid-range (up to dozens of meters) depth imaging where the divergence of the laser source(s) can be neglected. Let $d_{i,j}$ be the position of an object surface at a given range from the sensor, whose mean spectral signature (observed at L wavelengths) is denoted as $\lambda_{i,j} = [\lambda_{i,j,1}, \dots, \lambda_{i,j,L}]^T$. According to [11], [12], each photon count $y_{i,j,\ell,t}$ is assumed to be drawn from the following Poisson distribution

$$y_{i,j,\ell,t} | \lambda_{i,j,\ell}, t_{i,j} \sim \mathcal{P}(\lambda_{i,j,\ell} g_{i,j,\ell}(t - t_{i,j})) \quad (1)$$

where $g_{i,j,\ell}(\cdot)$ is the photon impulse response, evaluated at discrete time positions as discussed in Section III-B and whose shape can differ between wavelength channels and pixel locations. In (1), $t_{i,j}$ is the characteristic time-of-flight of photons emitted by a pulsed laser source and reaching the detector after being reflected by a target at range $d_{i,j}$ ($d_{i,j}$ and $t_{i,j}$ are linearly related in free-space propagation). Moreover, the impulse responses $\{g_{i,j,\ell}(\cdot)\}$ are assumed to be known, as occurs when they can be accurately estimated during imaging system calibration. We further assume that the spectral signatures of the scene surfaces can be decomposed as linear mixtures of R known spectral signatures \mathbf{m}_r (also referred to as endmembers and gathered in the $L \times R$ matrix $\mathbf{M} = [\mathbf{m}_1, \dots, \mathbf{m}_R]$) possibly corrupted by sparse anomalies (or deviations from the linear mixture), that is

$$\lambda_{i,j} = \mathbf{M}\mathbf{a}_{i,j} + \mathbf{r}_{i,j}, \quad \forall i, j, \quad (2)$$

where $\mathbf{a}_{i,j} = [a_{i,j,1}, \dots, a_{i,j,R}]^T$ contains the abundances of the R endmembers in the pixel (i, j) and $\mathbf{r}_{i,j} \in \mathbb{R}_+^L$ is a sparse vector that captures anomalies that do not fit the LMM $\lambda_{i,j} = \mathbf{M}\mathbf{a}_{i,j}$. As explained in [13], [14], these anomalies in the resulting robust LMM can be due to actual outliers/corrupted data, nonlinear spectral mixtures or intrinsic endmember variability over the spectral bands. Note that due to physical considerations the unknown abundance vectors $\{\mathbf{a}_{i,j}\}_{i,j}$ can be assumed to have positive entries. It is important to recall that, in this work, we consider applications where the observed objects consist of a single visible surface per pixel. We do not consider cases where the photons can penetrate through objects (e.g., semi-transparent materials for which we would like to infer the internal composition) or be reflected from multiple surfaces. This assumption allows the spectral unmixing problem to be reduced to two spatial dimensions, which will be extended to distributed targets in future work. Moreover, we consider a single spectral signature per material, which does not depend on the object orientation with respect to the imaging system. Should the material signatures change significantly (due to the surface orientation or its intrinsic spectral variability), the changes are expected to be captured by the vectors $\mathbf{r}_{i,j}$. The problem addressed is to jointly estimate the range of the targets (for all the image pixels) and to solve the robust spectral unmixing

problem (e.g., estimating the abundance vectors and identifying the pixels corrupted by anomalies). The next section introduces the Bayesian model for this problem.

III. BAYESIAN MODEL

A. Likelihood

Assuming that the MSL waveforms $\mathbf{y}_{i,j} = \{y_{i,j,\ell,t}\}_{\ell,t}$ associated with a given pixel (i, j) result from photon reflection from a single surface associated with the spectrum $\lambda_{i,j}$ and according to (1), the likelihood associated with the pixel (i, j) can be expressed as

$$f(\mathbf{y}_{i,j} | \lambda_{i,j}, t_{i,j}) = \prod_{\ell,t} f_{\mathcal{P}}(y_{i,j,\ell,t}; \lambda_{i,j,\ell} g_{i,j,\ell}(t - t_{i,j})), \quad (3)$$

when it is assumed that the detected photon counts/noise realizations, conditioned on their mean in all channels/spectral bands, are conditionally independent. Note that in (3), $f_{\mathcal{P}}(\cdot; \lambda)$ denotes the probability mass function of the Poisson distribution with mean λ . Considering that the noise realizations in the different pixels are also conditionally independent, the joint likelihood can be expressed as

$$f(\mathbf{Y} | \Lambda, \mathbf{T}) = \prod_{i,j} f(\mathbf{y}_{i,j} | \lambda_{i,j}, t_{i,j}), \quad (4)$$

where $\Lambda = \{\lambda_{i,j}\}_{i,j}$ and \mathbf{T} is a matrix gathering the target ranges.

B. Prior distributions

1) *Range parameters*: Each target position is a discrete variable defined on $\mathbb{T} = \{t_{\min}, \dots, t_{\max}\}$, such that $1 \leq t_{\min} \leq t_{\max} \leq T$. In this paper we set $(t_{\min}, t_{\max}) = (301, T - 300)$ and the temporal resolution of the grid is set to the resolution of the single-photon detection (i.e., 2 picoseconds (ps) in Section VI). As in [12], to account for the spatial correlations between the neighbouring pixels, we propose to use a Markov random field as a prior distribution for $t_{i,j}$ given its neighbours $\mathbf{T}_{\mathcal{V}(i,j)}$, i.e., $f(t_{i,j} | \mathbf{T}_{\mathcal{V}(i,j)}) = f(t_{i,j} | \mathbf{T}_{\mathcal{V}(i,j)})$ where $\mathcal{V}(i, j)$ is the depth neighbourhood of the pixel (i, j) , $\mathbf{T}_{\mathcal{V}(i,j)} = \{t_{i',j'}\}_{(i',j') \in \mathcal{V}(i,j)}$ and $\mathbf{T}_{\setminus(i,j)} = \{t_{i',j'}\}_{(i',j') \neq (i,j)}$. More precisely, we propose the following discrete MRF

$$f(\mathbf{T} | \epsilon) = \frac{1}{G(\epsilon)} \exp[-\epsilon \phi(\mathbf{T})] \quad (5)$$

where $\epsilon \geq 0$ is a parameter tuning the amount of correlation between pixels, $G(\epsilon)$ is a normalization (or partition) constant and where $\phi(\cdot)$ is an arbitrary cost function modeling correlation between neighbours. In this work we propose to use the following cost function

$$\phi(\mathbf{T}) = \sum_{i,j} \sum_{(i',j') \in \mathcal{V}(i,j)} |t_{i,j} - t_{i',j'}|, \quad (6)$$

which corresponds to a total-variation regularization [9], [10] promoting piecewise constant depth image. Moreover, the higher the value of ϵ , the more correlated the neighboring pixels. Several neighborhood structures can be employed to define

$\mathcal{V}(i, j)$; here, a four pixel structure (1-order neighbourhood) will be considered in the rest of the paper for the depth parameters.

To illustrate the effect of this spatial regularization of the depth profile, we also consider an alternative depth prior model constructed from independent uniform priors

$$p(t_{i,j} = t) = \frac{1}{T'}, \quad t \in \mathbb{T}, \quad (7)$$

where $T' = \text{card}(\mathbb{T})$.

2) *Material Abundances*: It follows from (3) that gamma distributions are conjugate priors for the unknown abundances $a_{i,j,r}$ (the resulting conditional distributions are mixtures of gamma distributions, see [12] for details). Consequently, we propose such gamma priors for these parameters. Note that such priors also ensure the positivity of the proposed estimator of abundances. In a similar manner to [12], we assign $a_{i,j,r}$ the following gamma prior

$$a_{i,j,r} \sim \mathcal{G}\left(c_r, \frac{\bar{a}_{i,j,r}}{c_r}\right) \quad (8)$$

where $\bar{a}_{i,j,r} > 0$ is a local parameter related to the prior mean of $a_{i,j,r}$ and $c_r > 0$ is a global parameter (i.e., one per endmember) which controls the shape of the distribution tails and thus the prior deviation of $a_{i,j,r}$ from $\bar{a}_{i,j,r}$. Hierarchical Bayesian models generally allow the construction of elaborate prior models in which parameters can be related through the introduction of additional parameters which generally belong to higher levels in the Bayesian hierarchical model. For instance, setting $\bar{a}_{i,j,r} = \bar{a}_r$ in (8) reduces to choosing the same prior, characterized by (c_r, \bar{a}_r) , for all abundances associated with the r th endmember. Conversely, here we specify (8) to reflect the prior belief that abundances exhibit spatial correlations. In particular, due to the spatial organization of images, we expect the values of $a_{i,j,r}$ to vary smoothly from one pixel to another. In order to model this behaviour, we specify $\bar{a}_{i,j,r}$ such that the resulting prior for the abundances associated with each endmember $\mathbf{A}_r = \{a_{i,j,r}\}_{i,j}$ is a hidden gamma-MRF [15].

More precisely, we introduce R auxiliary matrices $\mathbf{\Gamma}_r$ of size $(N_{\text{row}} + 1) \times (N_{\text{col}} + 1)$, with elements $\gamma_{i,j,r} \in \mathbb{R}^+$ and define bipartite conditional independence graphs between \mathbf{A}_r and $\mathbf{\Gamma}_r$ such that each $a_{i,j,r}$ is connected to four neighbour elements of $\mathbf{\Gamma}_r$ and vice-versa. This 1st order neighbourhood structure is depicted in Fig. 1, where we point out that any given $a_{i,j,r}$ and $a_{i+1,j,r}$ are 2nd order neighbours via $\gamma_{i+1,j,r}$ and $\gamma_{i+1,j+1,r}$. Following the general gamma-MRF model proposed in [15], and specified here by the neighbourhood structure depicted in Fig. 1, we assign a gamma-MRF prior for each $(\mathbf{A}_r, \mathbf{\Gamma}_r)$, and obtain the following joint priors for $(\mathbf{A}_r, \mathbf{\Gamma}_r)$

$$\begin{aligned} f(\mathbf{A}_r, \mathbf{\Gamma}_r | c_r) &= \frac{1}{G(c_r)} \prod_{(i,j) \in \mathcal{V}_{\mathbf{A}_r}} a_{i,j,r}^{(c_r-1)} \\ &\times \prod_{(i',j') \in \mathcal{V}_{\mathbf{\Gamma}_r}} (\gamma_{i',j',r})^{-(c_r+1)} \\ &\times \prod_{((i,j),(i',j')) \in \mathcal{E}} \exp\left(\frac{-c_r a_{i,j,r}}{4\gamma_{i',j',r}}\right), \quad (9) \end{aligned}$$

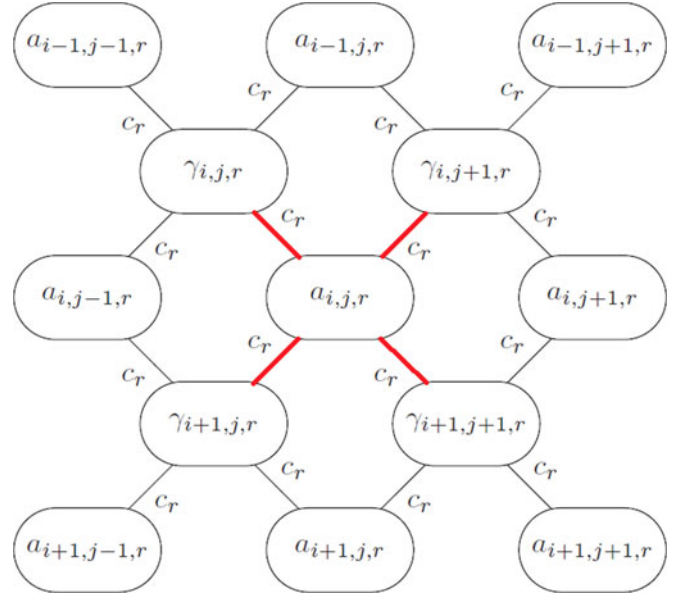


Fig. 1. Proposed 1st order gamma-MRF neighborhood structure for the abundances $\{a_{i,j,r}\}, \forall (i,j) \in \mathcal{V}_{\mathbf{A}_r}$ and $\forall r \in \{1, \dots, R\}$. We set $a_{i,j,r} = 0.01, \forall (i,j) \notin \mathcal{V}_{\mathbf{A}_r}$.

where $\mathcal{V}_{\mathbf{A}_r} = \{1, \dots, N_{\text{row}}\} \times \{1, \dots, N_{\text{col}}\}$, $\mathcal{V}_{\mathbf{\Gamma}_r} = \{1, \dots, N_{\text{row}} + 1\} \times \{1, \dots, N_{\text{col}} + 1\}$, and the edge set \mathcal{E} consists of pairs $((i,j), (i',j'))$ representing the connection between $a_{i,j,r}$ and $\gamma_{i',j',r}$. It can be seen from (9) that

$$a_{i,j,r} | \mathbf{\Gamma}_r, c_r \sim \mathcal{G}\left(c_r, \frac{\bar{a}_{i,j,r}(\mathbf{\Gamma}_r)}{c_r}\right) \quad (10a)$$

$$\gamma_{i,j,r} | \mathbf{A}_r, c_r \sim \mathcal{IG}(c_r, c_r \beta_{i,j,r}(\mathbf{R})) \quad (10b)$$

where

$$\begin{aligned} \bar{a}_{i,j,r}(\mathbf{\Gamma}) &= 4(\gamma_{i,j,r}^{-1} + \gamma_{i-1,j,r}^{-1} + \gamma_{i,j-1,r}^{-1} + \gamma_{i-1,j-1,r}^{-1})^{-1} \\ \beta_{i,j,r}(\mathbf{R}) &= (a_{i,j,r} + a_{i+1,j,r} + a_{i,j+1,r} + a_{i+1,j+1,r})/4. \end{aligned}$$

Notice that we denote explicitly the dependence of the gamma-MRFs on the value of $\{c_r\}_r > 0$, which act as regularization parameters that control the amount of abundance spatial smoothness enforced by each gamma-MRF (and which can differ among endmembers). Following an empirical Bayesian approach, the value of each c_r will remain unspecified and will be adjusted automatically during the inference procedure using maximum marginal likelihood estimation (see [12] for details).

3) *Anomaly model*: As in [13], [14], the outliers are assumed to be spatially and spectrally sparse, i.e., for most of the pixels and spectral bands there are no outliers. To model outlier sparsity, we factor each outlier vector as

$$\mathbf{r}_{i,j} = \mathbf{z}_{i,j} \odot \mathbf{x}_{i,j}, \quad (11)$$

where $\mathbf{z}_{i,j} = [z_{i,j,1}, \dots, z_{i,j,L}]^T \in \{0, 1\}^L$ is a binary label vector, $\mathbf{x}_{i,j} \in \mathbb{R}^L$ and \odot denotes the Hadamard (term-wise) product. This decomposition allows one to decouple the location of the sparse components from their values. More precisely, $z_{i,j,n} = 1$ if an outlier is present in the ℓ th spectral band of the pixel (i,j) with value equal to $r_{i,j,\ell} = x_{i,j,\ell}$.

Assuming that the potential anomalies *a priori* share the same statistical properties, we consider the following independent conjugate gamma priors

$$x_{i,j,\ell} | \alpha, \nu \sim \mathcal{G}(x_{i,j,\ell}; \alpha, \nu), \quad \forall i, j, \ell, \quad (12)$$

where (α, ν) are arbitrarily fixed parameters. In general it is difficult to empirically estimate outlier hyperparameters due to the fact that in most applications, outliers are by definition rare events, occurring in only a few pixels, and the observed waveforms are have unknown sparsity levels. Setting (α, ν) so that $E[x_{i,j} | (\alpha, \nu)] = \alpha\nu$ is too high might lead to poor detection performance (outliers not detected), in particular in the presence of low amplitude outliers. Conversely, setting (α, ν) so that $E[x_{i,j} | (\alpha, \nu)] = \alpha\nu$ is too small might lead to high probabilities of false alarm. However, in our experiments (See Section VI), we did not observe significant performance degradation when varying (α, ν) over the range considered in Section VI.

For many applications, the locations of outliers are likely to be spectrally correlated (e.g., water absorption bands) and/or spatially correlated (weakly represented components, local non-linear mixtures,...). An effective way to take correlated outliers/nonlinear effects into account is to use a Markov random field (MRF) as a joint prior for the anomaly labels in $\mathbf{Z} = \{z_{i,j,\ell}\}_{i,j,\ell}$. In this paper, we use the Ising model proposed in [14] for robust linear unmixing of HSIs to define the prior model for \mathbf{Z} . MRFs have the property that the conditional distribution of a label $z_{i,j,\ell}$ given the other labels of the image equals the conditional distribution of this label vector given only its neighbors, i.e., $P(z_{i,j,\ell} | \mathbf{Z}_{\setminus z_{i,j,\ell}}) = P(z_{i,j,\ell} | \mathbf{Z}_{\mathcal{V}_{i,j,\ell}})$, where $\mathcal{V}_{i,j,\ell}$ is the index set of the neighbors of $z_{i,j,\ell}$, $\mathbf{Z}_{\setminus z_{i,j,\ell}}$ denotes the matrix \mathbf{Z} whose element $z_{i,j,\ell}$ has been removed and $\mathbf{Z}_{\mathcal{V}_{i,j,\ell}}$ is the subset of \mathbf{Z} composed of the elements whose indexes belong to $\mathcal{V}_{i,j,\ell}$. In this study, we consider that the spatial and spectral correlations can be different and thus consider two different neighborhoods. We decompose the neighborhood $\mathcal{V}_{i,j,\ell}$ as $\mathcal{V}_{i,j,\ell} = \mathcal{V}_{i,j,\ell}^L \cup \mathcal{V}_{i,j,\ell}^N$ where $\mathcal{V}_{i,j,\ell}^N$ (resp. $\mathcal{V}_{i,j,\ell}^L$) denotes the spatial (resp. spectral) neighborhood of $z_{i,j,\ell}$. Specifically, we consider an Ising model that can be expressed as

$$P(\mathbf{Z} | \beta') = \frac{1}{B(\beta')} \exp [\beta^T \phi(\mathbf{Z}) + \phi_0(\mathbf{Z}, \beta_0)] \quad (13)$$

where $\beta = [\beta_N, \beta_L]^T$, $\beta' = [\beta^T, \beta_0]^T$ and

$$\begin{cases} \phi_L(\mathbf{Z}) &= \sum_{i,j,\ell} \sum_{z_{i,j,\ell'} \in \mathcal{V}_{i,j,\ell}^L} \delta(z_{i,j,\ell} - z_{i,j,\ell'}), \\ \phi_N(\mathbf{Z}) &= \sum_{i,j,\ell} \sum_{z_{i',j',\ell} \in \mathcal{V}_{i,j,\ell}^N} \delta(z_{i,j,\ell} - z_{i',j',\ell}), \\ \phi(\mathbf{Z}) &= [\phi_L(\mathbf{Z}), \phi_N(\mathbf{Z})]^T, \\ \phi_0(\mathbf{Z}, \beta_0) &= \beta_0 \sum_{i,j,\ell} (1 - z_{i,j,\ell}) + (1 - \beta_0) \sum_{i,j,\ell} z_{i,j,\ell}, \end{cases}$$

and $\delta(\cdot)$ denotes the Kronecker delta function. Moreover, $\beta_N > 0$ and $\beta_L > 0$ are hyperparameters that control the spatial and spectral granularity of the MRF and $0 \leq \beta_0 \leq 1$ is an additional parameter that models the probability of having outliers in the image. Specifically, the higher the value of β_0 , the lower is the probability of outliers in the data. In a similar fashion to the gamma-MRFs parameters $\{c_r\}$ in (10), the Ising model hyperparameters will be adjusted with a reduced computational cost




Fig. 2. Directed acyclic graph representing the proposed hierarchical Bayesian model. Fixed quantities appear in solid line boxes and quantities adjusted via maximum marginal likelihood estimation appear in dashed line boxes.

via maximum marginal likelihood estimation. Different spectral and spatial neighbourhoods can be used in (13). In this paper, we consider a 4-neighbour 2D structure to account for the spatial correlation and a 2-neighbour 1D structure for the spectral dimension.

C. Joint Posterior Distribution

From the joint likelihood and prior model specified in Sections III-A and III-B, we can now derive the joint posterior distribution for \mathbf{T} , $\mathbf{A} = \{\mathbf{A}_r\}$, $\mathbf{\Gamma} = \{\mathbf{\Gamma}_r\}$, \mathbf{Z} , and $\mathbf{X} = \{x_{i,j}\}$, given the observed waveforms \mathbf{Y} , and the value of the two sets of hyperparameters $\Phi = (\alpha, \nu)$ and $\theta = (\beta', \epsilon, c_1, \dots, c_R)$. Note that the hyperparameters are organized into two groups: Φ are the fixed variables and θ are those that will be adjusted. Using Bayes' theorem, and assuming prior independence between \mathbf{T} , $(\mathbf{A}, \mathbf{\Gamma})$, \mathbf{X} and \mathbf{Z} , the joint posterior distribution associated with the proposed Bayesian model is given by

$$\begin{aligned} & f(\mathbf{T}, \mathbf{A}, \mathbf{\Gamma}, \mathbf{Z}, \mathbf{X}, \theta | \mathbf{Y}, \Phi, \theta) \\ & \propto f(\mathbf{Y} | \mathbf{T}, \mathbf{A}, \mathbf{Z}, \mathbf{X}) f(\mathbf{A}, \mathbf{\Gamma} | \mathbf{c}) f(\mathbf{T} | \epsilon) f(\mathbf{X} | \alpha, \nu) f(\mathbf{Z} | \beta') \end{aligned} \quad (14)$$

with $\mathbf{c} = [c_1, \dots, c_R]^T$. The directed acyclic graph (DAG) summarizing the structure of the proposed Bayesian model is depicted in Fig. 2.

IV. ESTIMATION STRATEGY

The posterior distribution (14) captures all of the information about the unknowns given the observed data and the priors on the unknown parameters. To perform joint depth estimation and spectral unmixing of the MSL data, we use the following four different Bayesian estimators: 1) the minimum mean square error estimator (MMSE) of the abundances, 2) the marginal maximum a posteriori (MMAP) estimator of the anomaly labels, 3) conditioned on the estimated outliers location, the MMSE estimator of the anomaly values (in a similar fashion to [16]) and 4) conditioned on the estimated outliers and abundances, the MMAP estimator of the target ranges. Note that we use the MMAP estimators for the target ranges and labels, as this Bayesian estimator is particularly well adapted for estimation of discrete parameters.

In order to approximate these estimators, we propose a Markov chain Monte Carlo method to generate samples

according to the joint posterior

$$f(\mathbf{T}, \mathbf{A}, \Gamma, \mathbf{Z}, \mathbf{X} | \mathbf{Y}, \Phi, \hat{\theta}), \quad (15)$$

where $\hat{\theta}$ denotes the maximum marginal likelihood estimator of the Ising and gamma-MRFs regularisation hyperparameter vector θ given the observed data \mathbf{Y} , i.e.,

$$\hat{\theta} = \operatorname{argmax}_{\theta \in \Theta} f(\mathbf{Y} | \Phi, \theta), \quad (16)$$

where Θ is the admissible set for θ . This can be viewed as an empirical Bayes approach for specifying θ where hyperparameters with unknown values are replaced by point estimates computed from observed data (as opposed to being fixed a priori or integrated out of the model by marginalisation) [17]. This strategy has several important advantages for variables such as θ whose conditional distributions are intractable. In particular, it has significantly lower computational cost compared to that of competing approaches, such as ones that marginalize the posterior (15) over θ during the inference procedure [18].

To sample from the posterior (15), we use a Metropolis-within-Gibbs sampler that includes Hamiltonian Monte Carlo (HMC) updates [19]. This sampling procedure generates sequentially the unknown parameters from their estimated conditional distributions. The corresponding update steps are specified in the remainder of this section.

A. Sampling the Depth Parameters

Sampling the target ranges can be achieved by sampling sequentially each depth from its conditional distribution

$$f(t_{i,j} | \mathbf{y}_{i,j}, \lambda_{i,j}, \mathbf{T}_{\mathcal{V}(i,j)} \theta) \quad (17)$$

i.e., by drawing randomly from discrete (with finite support \mathbb{T}) distributions. In a similar fashion to [12], we use a Gibbs sampler implemented using a colouring scheme such that many depths can be updated in parallel (2 steps required when considering a 1-order neighborhood structure). Note that if the prior model in (7) is used instead of (5), it can be seen for (14) that

$$f(\mathbf{T} | \mathbf{Y}, \mathbf{A}, \Gamma, \mathbf{Z}, \mathbf{X}, \Phi, \theta) = \prod_{i,j} f(t_{i,j} | \mathbf{y}_{i,j}, \lambda_{i,j}) \quad (18)$$

with $f(t_{i,j} | \mathbf{y}_{i,j}, \lambda_{i,j}) \propto f(\mathbf{y}_{i,j} | t_{i,j}, \lambda_{i,j}) / T'$. Thus, all the depth parameters can be updated independently and simultaneously.

B. Sampling the Anomaly Labels

From (14), we obtain

$$f(z_{i,j,\ell} = k | \mathbf{Y}, \mathbf{A}, \mathbf{X}, \mathbf{Z}_{\setminus z_{i,j,\ell}}, \beta') \propto \tilde{\pi}_{i,j,\ell}^{(k)}, \quad \forall (i, j, \ell), \quad (19)$$

where $k \in \{0, 1\}$,

$$\begin{aligned} \log \left(\tilde{\pi}_{i,j,\ell}^{(k)} \right) &= \sum_{t=1}^T \log \left(f(y_{i,j,\ell,t} | \lambda_{i,j,\ell}^{(k)}, t_{i,j}) \right) \\ &\quad - \beta^T \phi(\mathbf{Z}) - \phi_0(\mathbf{Z}, \beta_0), \end{aligned} \quad (20)$$

and $\lambda_{i,j,\ell}^{(k)} = \mathbf{m}_{\ell,:} \mathbf{a}_{i,j} + kx_{i,j,\ell}$ with $\mathbf{m}_{\ell,:}$ the ℓ th row of \mathbf{M} . Consequently, the label $z_{i,j,\ell}$ can be drawn from its conditional

distribution by drawing randomly from $\{0, 1\}$ with probabilities given by

$$f(z_{i,j,\ell} = k | \mathbf{Y}, \mathbf{A}, \mathbf{X}, \mathbf{Z}_{\setminus z_{i,j,\ell}}, \beta') = \frac{\tilde{\pi}_{i,j,\ell}^{(k)}}{\tilde{\pi}_{i,j,\ell}^{(0)} + \tilde{\pi}_{i,j,\ell}^{(1)}}. \quad (21)$$

In our experiments we incorporated a Gibbs sampler implemented using a colouring scheme such that labels which are not direct neighbours can be updated in parallel. It is important to mention here that the main computation cost associated with the label updates arises from the computation of the sum over T in (20), especially from large values of T (e.g., $T = 3000$ in the experiments presented in Section VI). Fortunately, this sum can be written (up to an addition constant that does not depend on $z_{i,j,\ell}$) as

$$\tilde{y}_{i,j,\ell} \log(\mathbf{m}_{\ell,:} \mathbf{a}_{i,j} + kx_{i,j,\ell}) - (\mathbf{m}_{\ell,:} \mathbf{a}_{i,j} + kx_{i,j,\ell}) \tilde{g}_{i,j,\ell},$$

where $\tilde{y}_{i,j,\ell} = \sum_{t=1}^T y_{i,j,\ell,t}$ and $\tilde{g}_{i,j,\ell} = \sum_{t=1}^T g_{i,j,\ell}(t - t_{i,j})$ are quantities that only need to be computed once (before running the sampler) and stored in a look-up table, as they do not depend on the unknown model parameters ($\tilde{g}_{i,j,\ell}$ can be evaluated in advance for any value of $t_{i,j} \in \mathbb{T}$).

C. Sampling the Abundances

It is easy to show that

$$\begin{aligned} f(\mathbf{A} | \mathbf{Y}, \mathbf{T}, \Gamma, \mathbf{Z}, \mathbf{X}, \Phi, \theta) \\ = \prod_{i,j} f(\mathbf{a}_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, \Gamma, \mathbf{X}, \mathbf{Z}, \mathbf{c}), \end{aligned} \quad (22)$$

i.e., the abundances of the $N_{\text{row}}N_{\text{col}}$ pixels can be updated independently. Here we update the elements of each vector $\mathbf{a}_{i,j}$ simultaneously using a constrained Hamiltonian Monte Carlo update [16], [19] for several reasons. First, the intrinsic correlation between the spectral signatures in \mathbf{M} (especially when considering materials spectrally similar) imposes strong correlations between the elements of $\mathbf{a}_{i,j}$, which slows down the convergence of algorithms updating the abundances sequentially. Second, by accounting for the local curvature of $f(\mathbf{a}_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, \Gamma, \mathbf{X}, \mathbf{Z}, \mathbf{c})$, such an approach usually yields better convergence and mixing properties than standard random walks. Note that it can be shown that when $c_r > 1, \forall r$, the conditional distribution $f(\mathbf{a}_{i,j} | \mathbf{y}_{i,j}, t_{i,j}, \Gamma, \mathbf{X}, \mathbf{Z}, \mathbf{c})$ is strictly log-concave, which makes the use of gradient-based updating schemes particularly well adapted. These constraints on the regularization parameters c_r can easily be incorporated within the estimation process by specifying Θ in (16) accordingly and projecting the updated parameters onto Θ at the end of the update step (this step is not detailed here but the interested reader is invited to consult [12], [14], [17] for further details).

D. Sampling the Anomaly Values \mathbf{X}

As the spatial/spectral organization of the anomalies is only encoded via the Potts model (13), it can be shown that

$$f(\mathbf{X} | \mathbf{Y}, \mathbf{A}, \mathbf{T}, \Gamma, \mathbf{Z}, \Phi, \theta) = \prod_{i,j,\ell} f(x_{i,j,\ell} | \mathbf{y}_{i,j}, t_{i,j}, \mathbf{a}_{i,j}, \alpha, \nu, z_{i,j,\ell}). \quad (23)$$

Algorithm 1: RSU-MSL algorithm.

```

1: Fixed input parameters: Endmember matrix  $\mathbf{M}$ ,  $(\alpha, \nu)$ 
   number of burn-in iterations  $N_{\text{bi}}$ , total number of
   iterations  $N_{\text{MC}}$ 
2: Initialization ( $u = 0$ )
   • Set  $\mathbf{T}^{(0)}$ ,  $\mathbf{A}^{(0)}$ ,  $\mathbf{\Gamma}^{(0)}$ ,  $\mathbf{X}^{(0)}$ ,  $\mathbf{Z}^{(0)}$ ,  $\boldsymbol{\theta}^{(0)}$ 
3: Iterations ( $1 \leq u \leq N_{\text{MC}}$ )
4: Sample  $\mathbf{A}^{(u)}$  from (22)
5: Sample  $\mathbf{\Gamma}^{(u)}$  from (10)
6: for  $i = 1 : N_{\text{row}}$  do
7:   for  $j = 1 : N_{\text{col}}$  do
8:     Sample  $t_{i,j}^{(u)}$  from (17) or (18)
9:     for  $\ell = 1 : L$  do
10:      Sample  $z_{i,j,\ell}^{(u)}$  from (21)
11:     end for
12:   end for
13: end for
14: Sample  $\mathbf{X}^{(u)}$  from (23)
15: if  $u < N_{\text{bi}}$  then
16:   Update  $\boldsymbol{\theta}^{(u)}$  using [17]
17: else
18:   Set  $\boldsymbol{\theta}^{(u)} = \boldsymbol{\theta}^{(u-1)}$ 
19: end if
20: Set  $u = u + 1$ .

```

when $z_{i,j,\ell} = 1$, the conditional distribution $f(x_{i,j,\ell} | \mathbf{y}_{i,j}, t_{i,j}, \mathbf{a}_{i,j}, \alpha, \nu, z_{i,j,\ell})$ is a non-standard distribution (although log-concave when $\alpha \geq 1$, similarly to the conditional distribution of the abundances), while it reduces to (12) when $z_{i,j,\ell} = 0$. Here, to avoid additional algorithmic complexity associated with rejection sampling and random walk update schemes, we use a standard Metropolis-Hastings update for which the proposal distribution is the prior distribution (12). When $z_{i,j,\ell} = 0$, which is satisfied for most of the elements of \mathbf{X} when the anomalies are indeed sparse, this update allows the candidates to be automatically accepted. In practice, we observed that this updating approach leads to reasonable acceptance rates when $z_{i,j,\ell} = 1$.

E. Sampling the Auxiliary Variables $\mathbf{\Gamma}$

Simply, as $\mathbf{\Gamma}$ does not appear in (4), sampling the elements of $\mathbf{\Gamma}$ reduces to sampling from inverse-gamma distributions (10). Note that due to the structure of the gamma-MRFs considered (the elements of each matrix $\mathbf{\Gamma}_r$ are not directly connected), these updates can be performed in a parallel manner.

The resulting sampler is summarized in Algorithm 1. Its output is a set of N_{MC} samples which are then used to approximate the Bayesian estimators of interest after sampler burn-in period which discards the N_{bi} first samples. Here, the length of the burn-in period is determined from preliminary runs by visual inspection of the chains. Although there appear to be nested loops in Algorithm 1, the labels and target ranges can usually be grouped such that, within a group of labels or range parameters, the variables are conditionally independent and can thus be updated independently and in a parallel manner. With the

neighbourhood structure used in (6) and (13), only 2 sequential steps are required to update the binary labels, and the same applies for the range parameters, thus improving the convergence speed of the sampler and reducing computation. Moreover, the Hamiltonian Monte Carlo updates used to sample the abundances further improve the convergence speed of the sampler, when compared to standard Metropolis-Hastings updates.

V. RELATION TO POISSON FACTOR ANALYSIS

It is interesting to note from (1) that

$$\tilde{y}_{i,j,\ell} | \lambda_{i,j,\ell}, t_{i,j} \sim \mathcal{P}(\lambda_{i,j,\ell} \tilde{g}_{i,j,\ell}). \quad (24)$$

In other words, each integrated waveform (summed over the time bins) for each pixel and wavelength, follows a Poisson distribution whose mean depends only on the spectral parameters $\lambda_{i,j,\ell} = \mathbf{m}_{\ell,:} \mathbf{a}_{i,j} + r_{i,j,\ell}$, scaled by $\tilde{g}_{i,j,\ell}$ which contains only information about the range of the target. If we assume that $\tilde{g}_{i,j,\ell} = \sum_{t=1}^T g_{i,j,\ell}(t - t_{i,j})$ is constant for all possible values of $t_{i,j} \in \mathbb{T}$ (which occurs in practice when there are no truncations of the impulse responses at the boundaries of (t_{\min}, t_{\max})), the problem can be reduced to a more standard spectral unmixing problem. This unmixing problem is over two spatial and one spectral dimension. If the spatial distortion of the instrument responses can be neglected, the observation model can be rewritten

$$\tilde{\mathbf{Y}} \sim \mathcal{P}(\tilde{\mathbf{M}}\tilde{\mathbf{A}} + \tilde{\mathbf{R}}) \quad (25)$$

where $\tilde{\mathbf{Y}}$ is an $L \times N_{\text{row}}N_{\text{col}}$ matrix with elements $\tilde{y}_{i,j,\ell}$, $\tilde{\mathbf{M}}$ corresponds to the endmember matrix whose columns have been scaled by $\tilde{g}_{i,j,\ell}$, $\tilde{\mathbf{A}}$ is the reshaped abundance matrix and $\tilde{\mathbf{R}}$ is a sparse matrix representing the anomalies. Equation (25) is closely related to the Poisson Factor analysis (PFA) model [20] and to the robust non-negative matrix factorization model under a Poisson noise assumption [21]. Unlike in factor analysis, we assume the endmember matrix $\tilde{\mathbf{M}}$ is known. In addition to the temporal aspect of the problem considered, the proposed model extends the PFA model from [20] to account for spatial correlation and the presence of outliers/anomalies. In contrast to the simulation based method developed in [20] to sample the posterior and which introduces discrete latent variables to simplify the sampling procedure (most updates are standard Gibbs steps), we used accept-reject procedures based on constrained HMC moves. In contrast to [20], where the columns of $\tilde{\mathbf{M}}$ were assumed to be uncorrelated, these columns can be correlated in the spectral unmixing problem considered in this paper.

VI. EXPERIMENTS

We evaluate the performance of the proposed method applied to estimation of the depth and the spectral profiles of a 5×5 cm scene (see Fig. 3(a)) composed of different objects made of opaque polymer clay and mounted on a dark-grey backboard at a distance of approximately 1.8 m from a time-of-flight scanning sensor. The sensor detects photons using a time-correlated single photon counting (TCSPC) technique. The transceiver system and data acquisition hardware used for this work are similar




Fig. 3. (a): Standard RGB image of the first scene, composed of different coloured clays fixed on a dark-grey backboard. (b) Reference depth/range image in millimeter obtained using the algorithm proposed in [22].

to that described in [23]–[28], developed at Heriot-Watt University. The scanning system used was a monostatic transceiver that utilises a galvanometer mirror pair to define the field position and to direct the return signal to an individual silicon-based single-photon detector with picoseconds timing resolution. The measurements have been acquired indoors, under dark conditions to limit the influence of ambient illumination. The scene has been scanned using a regular spatial grid of 190×190 pixels and $L = 33$ regularly spaced wavelengths ranging from 500 nm to 820 nm. Given the spatial structure of the scene considered in Fig. 3(a) and the short distance between the detector and the target (constrained by indoor measurements), most of the pixels can be associated with a single material. In the sequel we use an $R = 15$ endmember library/dictionary formed from the spectral response of each material and hence, for this experiment, almost all pixels in the image represent pure pixels without mixing. However, when the laser spot size on the target is broadened due to beam divergence then the assumption of a single material in each pixel is unlikely to hold. Such a situation would apply to longer range targets and close range targets whose material properties vary significantly over the spot size. Although classification methods such as [29] can effectively be used to classify pixels in this particular scene, our aim is to demonstrate the ability of our method to discriminate the materials and estimate depth even in the very low photon count regime. The histograms consisted of $T = 3000$ bins of 2 ps, which represents a depth resolution of $300 \mu\text{m}$ per bin. The optical power of the supercontinuum laser system was adjusted using data from preliminary runs, with a per-pixel acquisition time of 10 ms for each wavelength, in order to obtain accurate reference depth profiles (of approximately 1000 photons per pixel at each spectral band on average). This corresponds to an overall acquisition time of approximately 3 hours and 20 minutes using $L = 33$ sequential scans. The data format of timed events allows us to construct photon timing histograms associated with shorter acquisition times, as the system records the time of arrival of each detected photon (with respect to the previous synchronisation signal). Here, we evaluate our method for average photon counts of 1, 3 and 10 photons per pixel for each wavelength. Examples of distributions of detected photon counts, for the shortest acquisition time, are depicted in Fig. 4. The top sub-figures in Fig. 4 show that the number of detected photons varies depending both on




Fig. 4. Distributions of the number of detected photons (integrated over the $T = 3000$ bins) for the highest data sparsity level (on average 1 photon per pixel for each band). Top: Spatial distribution of the detected photons at 500 nm (left), 700 nm (middle) and total photon counts after integration over the $L = 33$ wavelengths (right). Bottom: Observed photon count distribution over all the pixels and bands. This figure shows that 45% of the pixels do not contain any detection.




Fig. 5. Depth RMSEs (in mm) obtained via maximum likelihood estimation when considering a single wavelength and using data constructed from 10 photons per pixel (on average) for each wavelength. Target ranges of empty pixels are obtained via nearest-neighbour interpolation. These results illustrate how the depth estimation highly rely on the wavelength considered, especially when considering sparse data. Averaging the $L = 33$ depth estimates generally reduces the RMSE and here, we obtained $\text{RMSE} = 2.08$ mm.

the object and wavelength. Consequently, the depth estimation performance, which is highly dependent on the number of detected photons, can be significantly affected by the considered wavelength(s). This observation is confirmed by the results in Fig. 5 which depicts the root mean squared error (RMSE) of the depth estimate, obtained using data constructed from 10 photons per pixel (on average) for each wavelength and using a single wavelength to estimate the target ranges via maximum likelihood (ML) estimation. In this work the RMSE is defined by

$$\text{RMSE} = \sqrt{\frac{\sum_{i,j} (t_{i,j} - \hat{t}_{i,j})^2}{N_{\text{row}} N_{\text{col}}}}, \quad (26)$$

where $\hat{t}_{i,j}$ is the estimated target range in the pixel (i, j) and $t_{i,j}$ is the reference target range. The reference depth profile shown in Fig. 3(b) has been obtained using the algorithm proposed in [22]. The algorithm provided for depth estimation and clustering the spectra in the MSL data. In a similar fashion to the method proposed in the paper, the algorithm considered in [22] uses all the wavelengths to estimate the target ranges. However, it does




Fig. 6. Examples of instrumental impulse responses measured with an acquisition of 100 s at different wavelengths (500, 550, 600, 650, 700, 750 and 800 nm).

not rely on a particular mixing model and is thus less sensitive to the presence of complex spectral mixtures. It was applied to the data with the longest acquisition time (10 ms) to reduce estimation errors and the resulting reference range profile is in very good agreement with the structure of the scene in Fig. 3(a) (the reference range being set to the range of the backboard).

The instrument impulse responses $g_{i,j,\ell}(\cdot)$ (partly depicted in Fig. 6 for an arbitrary central pixel of the field of view) were estimated from preliminary experiments by analyzing the distribution of photons reflected onto a Lambertian scatterer placed at a known distance with long acquisition time (100 s). Fig. 6 illustrates the fact that the response of the imaging system can change in amplitude and shape depending on the wavelength (full width at half maximum around 60 ps). This variability is due to the wavelength-dependent characteristics of the different elements in the imaging system, e.g., supercontinuum laser source, detector, lenses. Notice also the delays between the different peaks. These are mainly due to the wavelength-dependent response of the laser source and the different path lengths of the light in the imaging system. These delays can be compensated during the instrument calibration phase of the experiment and do not have a significant influence of the imaging performance. The spatial distortion of the responses $g_{i,j,\ell}(\cdot)$ due to potential limitations of the imaging system are taken into account during system calibration. If only a single wavelength was to be used to estimate the depth profile (e.g., as in Fig. 5), the variations of the peak shape could also make the choice of the most relevant wavelength difficult as the depth estimation accuracy mainly depends on the amplitude (reflectivity estimation) and width (depth estimation) of the peak.

Fig. 7 shows the spectral signatures of the backboard and the 14 kinds of clay used to create the objects. The $R = 15$ endmembers were obtained using the algorithm proposed in [22] for depth estimation and clustering on the spectra in the MSL data. To ensure high accuracy of the endmembers they were extracted from much cleaner data, i.e., images of the sample acquired with much longer acquisition time resulting in more than 1000 photons per pixel and per band on average. To be more explicit, the endmembers were obtained by averaging the estimated spectral responses over sub-regions (400 pixels per endmember) identified by the method proposed in [22].

For all results presented here, the proposed Bayesian Poisson unmixing algorithm has been applied with $N_{MC} = 5000$ sampler iterations, including $N_{bi} = 2000$ burn-in iterations. These

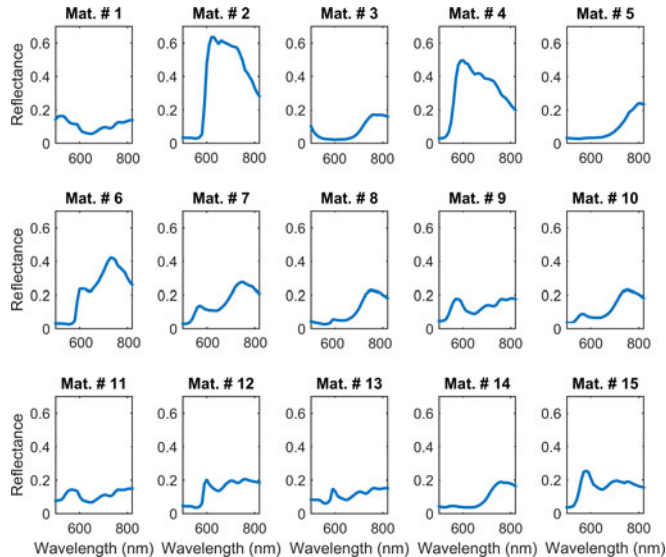


Fig. 7. Spectral signatures of the backboard and the 14 polymer clays used to create the objects in the scene shown in Fig. 3.

parameters were determined from preliminary runs by visual analysis of the chains and the variation of the results over independent runs. Moreover, we set $(\alpha, \nu) = (1, 0.05)$ for all the simulation results presented in this section. These parameters were selected based on the expectation that anomalies present low reflectivities. The results were not appreciably affected by small variation of (α, ν) about these values.

The depth profiles estimated by the proposed algorithm, referred to as R-PSU (for Robust Poisson Spectral Unmixing) are depicted in the second and third top rows of Fig. 8. R-PSU-TV denotes the algorithm with the TV-based depth regularization (and R-PSU without TV-based regularization). For completeness, the top row of Fig. 8, depicts the results obtained by the unmixing method proposed in [8], relying on the linearity of the mixtures. In contrast to the proposed method, the method in [8] unmixes the pixels independently using weakly informative abundance and depth prior models. Consequently, this method is denoted by ML in Fig. 8. As expected, the estimated depth profile becomes noisier as the average number of detected photons decreases. However, the TV-based depth prior model visually improves the estimated profiles, in particular for the highest data sparsity (1 photon on average per pixel and per band). The bottom rows of Fig. 8 depict the marginal posterior probabilities, for each pixel, of the actual object range to be within the estimated range bin (i.e., within a 0.3 mm interval centred around the estimated depth value). These images show that these probabilities increase with the number of detected photons, i.e., the joint posterior distribution becomes more concentrated around its modes and that the TV-based prior further concentrates this posterior distribution, thus reducing the *a posteriori* uncertainty. This observation is also illustrated in Fig. 9, which depicts the marginal posterior distributions of an arbitrary depth parameter (namely of the pixel (41, 61) chosen randomly), obtained using R-PSU and R-PSU-TV. Note that in the bottom rows of Fig. 8, the lowest probabilities are generally found at the object




Fig. 8. Top rows: Depth profiles estimated via pixel-wise ML estimation [8] and the proposed method, with and without spatial regularization of the depth profile. Bottom rows: Confidence map, which is the marginal posterior probability, for each pixel, that the actual object range is within the estimated range bin. The higher this probability, the more reliable the estimated depth.

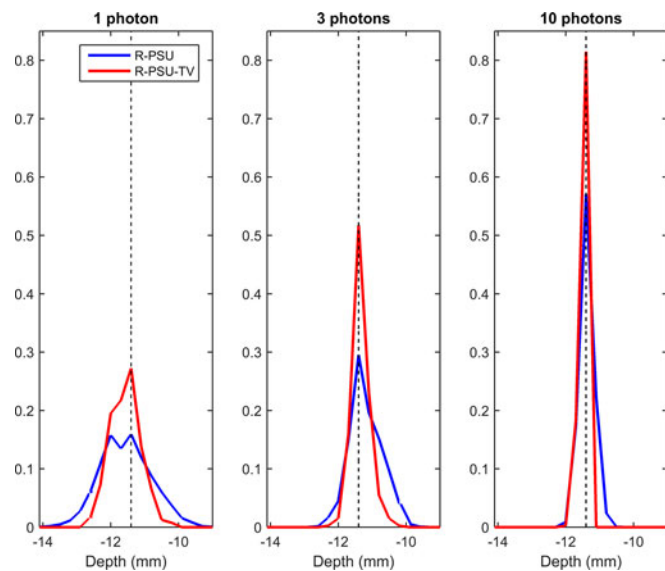


Fig. 9. Marginal posterior distributions of the object range estimated using R-PSU (blue lines) and R-PSU-TV (red lines), for the pixel (41, 61) (central region of the object #2). The black dashed lines represent the reference range estimated using the algorithm proposed in [22]. Total-variation (R-PSU-TV) provides a higher posterior confidence on the estimated depth.

TABLE I
DEPTH RMSES (IN MM) OBTAINED VIA PIXEL-WISE ML ESTIMATION [8] AND BY THE PROPOSED METHOD, WITH AND WITHOUT SPATIAL REGULARIZATION OF THE DEPTH PROFILES

	Average no. of photons per pixel		
	1	3	10
ML	3.66	1.09	0.65
R-PSU	1.91	1.05	0.65
R-PSU-TV	0.92	0.64	0.50

boundaries, where significant range changes occur and where the number of detected photons decreases due to the local orientation of the objects. These two effects generally increase the estimation uncertainty. Moreover, the RMSEs gathered in Table I confirm the performance degradation as the number of detected photons decreases and that the TV-based prior model for the depth parameters mitigates this degradation (the RMSE remains below 1 mm, even for the highest data sparsity). These results also illustrate the benefits of the proposed Bayesian approach for depth estimation since the RMSE obtained with the highest data sparsity (1 photons per pixel and per band) and R-PSU-TV (RMSE = 0.92 mm) is much lower than the RMSE obtained by the joint ML depth estimates proposed in [8], which processes the pixels independently (RMSE = 3.66 mm).

Fig. 10(a)–(c) compares the abundances estimated by the proposed method for the three photon sparsity levels considered (10, 3 and 1 photons per pixel for each band). The estimated abundances are generally in good agreement with the colour image in Fig. 3(a) as it is possible to identify the regions where the different clays are present. However, some unmixing errors are also visible, even when considering the lowest photon sparsity in Fig. 10(a). For instance, objects presenting different shades of green in Fig. 3(a) might not be perfectly unmixed due to their spectral similarity and their intrinsic spectral variability mainly caused by local orientation changes. When considering higher data sparsity levels (see Fig. 10(b) and (c)), the spectral variability of each object becomes negligible compared to the actual Poisson noise level and the proposed method identifies more accurately the different homogeneous regions thanks to the use of the gamma-MRFs in (9). To illustrate the benefits of the gamma-MRFs, we unmixed the data using the algorithm proposed in [8] and which does not promote spatially correlated abundances. For illustration, Fig. 10(d) presents the estimated abundance maps obtained using the highest data sparsity (1 photon per pixel on average). Without consideration of spatial correlation between abundance vectors, the estimated abundances maps become extremely noisy and it becomes impossible to clearly identify most of the homogeneous regions. In particular, the objects #8 and #14 are barely visible in Fig. 10(d) while more distinguishable in Fig. 10(c). Note that due to the spectral similarity between these objects and the board (#5), which is present in several regions on the scene, the objects #8 and #14 are also visible in the abundance map of the board in Fig. 10(c). Although the unmixing results in Fig. 10(a)–(c) present some errors, the




Fig. 10. (a)–(c): Abundance maps associated with the $R = 15$ main materials composing the scene and estimated by the proposed R-PSU-TV method. These results are obtained from data constructed from 10 (a), 3 (b) and 1 (c) photons per pixel (on average across the pixels) and for each spectral band. (d): Abundance maps estimated by the algorithm proposed in [8] based on the shortest acquisition time (1 photon per pixel). All images have the same dynamic, i.e., between 0 and 1.3.




Fig. 11. Anomaly maps ($\log(\|\mathbf{r}_{i,j}\|^2/L)$) associated with the region of interest for averages of 1, 3 and 10 detected photons per pixel and per band. The bottom row presents inset zooms of the squared red regions identified in the top row figures. The detected anomalies correspond to high reflectivity between 750 nm and 820 nm, and are due to the presence of residual glue used to fix the clay objects on the backboard.

estimates abundances maps are more accurate than those obtained without the abundance gamma-MRFs.

Fig. 11 shows the estimated anomaly maps obtained using the $R = 15$ endmembers of Fig. 7 for the three acquisition times. This figure illustrates significant local differences from the classical LMM, in particular for the longest exposures. In the central region of the scene, below and above the object #11, the two red strips correspond to high reflectivity between 750 nm and 820 nm, and are due to the presence of residual glue used to fix the clay objects on the backboard. When the exposure decreases,

the quality of the data decreases due to the statistical properties of the Poisson noise and it becomes more challenging to detect subtle spectral variations. In such a case, the proposed method, which promotes sparse deviations from the LMM, is no longer able to detect weak anomaly levels and only detects one glue strip (left sub-plots). It can be seen from Fig. 11 that deviations from the standard LMM are clustered spatially and only affect a reduced number of pixels. Thus, the abundance maps estimated while considering the LMM (i.e., by enforcing $\|\mathbf{r}_{i,j}\|^2 = 0$) are very similar, for most pixels, to those depicted in Fig. 10 and are not presented here for brevity.

VII. CONCLUSION

We have proposed a new Bayesian model and a joint depth estimation and robust spectral unmixing algorithm for 3D scene analysis from MSL data. Assuming the ambient illumination can be neglected, the spectra of the scene surfaces visible by the imaging system were decomposed into linear mixtures of known endmembers, potentially corrupted by sparse deviations/anomalies. Adopting a Bayesian approach, prior distributions were assigned to the unknown model parameters; in particular, a 3D Ising model was used to model the spatial organization of the anomalies and gamma Markov random fields were considered to promote spatially smooth abundances. Including ambient illumination and dark count levels in the observation model (as in [12], [28], [30], [31]) is the obvious next step from

a more general application (especially for long-range imaging applications) of the proposed method.

In this work, the experiments were performed indoor and the surface visible in each pixel was small ($\approx 0.01 \text{ mm}^2$) compared to the size of the objects. Applying the proposed method to large-standoff outdoor applications is under investigation. In such cases, the divergence of the laser sources can lead to changes in the shape of the returned pulses due to the local orientation of the observed objects with respect to the imaging system. Note that such distortion might also occur when there are semi-transparent objects. The divergence of the laser sources, which increases with the target range, results in reduced spatial resolution leading to increased mixing within a single pixel.

In analogy to passive hyperspectral image analysis, endmember and mixture characterization will become more challenging for large target ranges. Unsupervised or blind spectral unmixing approaches will need to be developed that can deal with unknown endmember spectral signatures. As explained in the introduction of this paper, our proposed anomaly model identifies “anomalous” pixels for which the linear mixing model is a poor fit. Such pixels identified by our model can be subsequently processed by more sophisticated, and yet to be developed, nonlinear mixing models. The development of such models is a worthy topic for future study.

Finally, for remote sensing applications, it will be crucial to account for the presence of distributed (multi-layered or semi-transparent) targets, which would require more complex models for the multiple returns in the MSL data. This could potentially lead to full 3D abundance profile estimation.

REFERENCES

- [1] M. Dalponte, L. Bruzzone, and D. Gianelle, “Fusion of hyperspectral and Lidar remote sensing data for classification of complex forest areas,” *IEEE Trans. Geosci. Remote Sens.*, vol. 46, no. 5, pp. 1416–1427, May 2008.
- [2] S. J. Buckley, T. H. Kurz, J. A. Howell, and D. Schneider, “Terrestrial Lidar and hyperspectral data fusion products for geological outcrop analysis,” *Comput. Geosci.*, vol. 54, pp. 249–258, 2013.
- [3] A. M. Wallace, C. J. Nichol, and I. H. Woodhouse, “Recovery of forest canopy parameters by inversion of multispectral Lidar data,” *Remote Sens.*, vol. 4, no. 2, pp. 509–531, 2012.
- [4] T. Hakala, J. Suomalainen, S. Kaasalainen, and Y. Chen, “Full waveform hyper-spectral Lidar for terrestrial laser scanning,” *Opt. Express*, vol. 20, no. 7, pp. 7119–7127, Mar. 2012.
- [5] A. M. Wallace *et al.*, “Design and evaluation of multispectral Lidar for the recovery of arboreal parameters,” *IEEE Trans. Geosci. Remote Sens.*, vol. 52, no. 8, pp. 4942–4954, Aug. 2014.
- [6] Y. Altmann, A. Wallace, and S. McLaughlin, “Spectral unmixing of multispectral Lidar signals,” *IEEE Trans. Signal Process.*, vol. 63, no. 20, pp. 5525–5534, Oct. 2015.
- [7] D. Ramirez *et al.*, “Developing hyperspectral Lidar for structural and biochemical analysis of forest data,” in *Proc. EARSEL Conf. Adv. Geosci.*, Mykonos, Greece, May 2012, pp. 1–11.
- [8] Y. Altmann *et al.*, “Robust spectral unmixing of multispectral Lidar waveforms,” in *Proc. IEEE GRSS Workshop Hyperspectral Image Signal Process.: Evol. Remote Sens.*, Los Angeles, CA, USA, Aug. 2016, 1–4.
- [9] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D*, vol. 60, no. 1–4, pp. 259–268, Nov. 1992.
- [10] A. Chambolle, “An algorithm for total variation minimization and applications,” *J. Math. Imag. Vis.*, vol. 20, no. 1/2, pp. 89–97, 2004.
- [11] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, “Bayesian analysis of Lidar signals with multiple returns,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 12, pp. 2170–2180, Dec. 2007.
- [12] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, “Lidar waveform based analysis of depth images constructed using sparse single-photon data,” *IEEE Trans. Image Process.*, vol. 25, no. 5, pp. 1935–1946, May 2016.
- [13] G. E. Newstadt, A. O. Hero, and J. Simmons, “Robust spectral unmixing for anomaly detection,” in *Proc. IEEE Workshop Stat. Signal Process.*, Gold Coast, Qld, Australia, Jul. 2014, pp. 109–112.
- [14] Y. Altmann, S. McLaughlin, and A. Hero, “Robust linear spectral unmixing using anomaly detection,” *IEEE Trans. Comput. Imag.*, vol. 1, no. 2, pp. 74–85, Jun. 2015.
- [15] O. Dikmen and A. Cemgil, “Gamma Markov random fields for audio source modeling,” *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 3, pp. 589–601, Mar. 2010.
- [16] Y. Altmann, N. Dobigeon, and J. Tourneret, “Unsupervised post-nonlinear unmixing of hyperspectral images using a hamiltonian Monte Carlo algorithm,” *IEEE Trans. Image Process.*, vol. 23, no. 6, pp. 2663–2675, Jun. 2014.
- [17] M. Pereyra, N. Whiteley, C. Andrieu, and J.-Y. Tourneret, “Maximum marginal likelihood estimation of the granularity coefficient of a Potts-Markov random field within an MCMC algorithm,” in *Proc. IEEE Workshop Stat. Signal Process.*, Gold Coast, Qld, Australia, Jul. 2014, pp. 121–124.
- [18] M. Pereyra, N. Dobigeon, H. Batatia, and J.-Y. Tourneret, “Estimating the granularity coefficient of a Potts-Markov random field within an MCMC algorithm,” *IEEE Trans. Image Process.*, vol. 22, no. 6, pp. 2385–2397, Jun. 2013.
- [19] S. Brooks, *Handbook of Markov Chain Monte Carlo* (Chapman & Hall/CRC Handbooks of Modern Statistical Methods). New York, NY, USA: Taylor & Francis, 2011.
- [20] M. Zhou, L. Hannah, D. B. Dunson, and L. Carin, “Beta-negative binomial process and Poisson factor analysis,” in *Proc. 15th Int. Conf. Artif. Intell. Stat.*, 2012, pp. 1462–1471.
- [21] C. Févotte and N. Dobigeon, “Nonlinear hyperspectral unmixing with robust nonnegative matrix factorization,” *IEEE Trans. Image Process.*, vol. 24, no. 12, pp. 4810–4819, Dec. 2015.
- [22] Y. Altmann, A. Maccarone, A. McCarthy, G. S. Buller, and S. McLaughlin, “Joint spectral clustering and range estimation for 3D scene reconstruction using multispectral Lidar waveforms,” in *Proc. Eur. Signal Process. Conf.*, Budapest, Hungary, Sep. 2016, pp. 513–517.
- [23] 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, Nov. 2009.
- [24] N. J. Krichel, A. McCarthy, and G. S. Buller, “Resolving range ambiguity in a photon counting depth imager operating at kilometer distances,” *Opt. Express*, vol. 18, no. 9, pp. 9192–9206, Apr. 2010.
- [25] A. M. Wallace, J. Ye, N. J. Krichel, A. McCarthy, R. J. Collins, and G. S. Buller, “Full waveform analysis for long-range 3D imaging laser radar,” *EURASIP J. Adv. Signal Process.*, vol. 2010, no. 1, 2010, Art. no. 896708.
- [26] A. McCarthy *et al.*, “Kilometer-range depth imaging at 1550 nm wavelength using an InGaAs/InP single-photon avalanche diode detector,” *Opt. Express*, vol. 21, no. 19, pp. 22 098–22 113, Sep. 2013.
- [27] A. Maccarone *et al.*, “Underwater depth imaging using time-correlated single-photon counting,” *Opt. Express*, vol. 23, no. 26, pp. 33 911–33 926, Dec. 2015.
- [28] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, “Target detection for depth imaging using sparse single-photon data,” in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Shanghai, China, 2016. [Online]. Available: <http://yoannaltmann.weebly.com/publications.html>
- [29] Y. Altmann, A. Maccarone, A. McCarthy, G. S. Buller, and S. McLaughlin, “Joint range estimation and spectral classification for 3D scene reconstruction using multispectral Lidar waveforms,” in *Proc. IEEE Workshop Stat. Signal Process.*, Jun. 2016, pp. 1–5.
- [30] A. Kirmani *et al.*, “First-photon imaging,” *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [31] D. Shin, A. Kirmani, V. Goyal, and J. Shapiro, “Computational 3D and reflectivity imaging with high photon efficiency,” in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2014, pp. 46–50.

Yoann Altmann (S'10–M'14) was born in Toulouse, France, in 1987. He received the Engineering degree in electrical engineering from Ecole Nationale Supérieure d'Electronique, d'Electrotechnique, d'Informatique, d'Hydraulique et des Télécommunications, Toulouse, France, and the M.Sc. degree in signal processing from the National Polytechnic Institute of Toulouse (INP Toulouse), Toulouse, France, both in 2010, and the Ph.D. degree from INP Toulouse in 2013. From 2014 until 2016, he was with the Heriot-Watt University, Edinburgh, as a Postdoctoral Researcher. Since 2017, he has been an Assistant Professor within the Institute of Sensors, Signals and Systems, School of Engineering and Physical Sciences. His current research interests include statistical signal and image processing, with a particular interest in Bayesian inverse problems with applications to remote sensing and biomedical imaging.

Aurora Maccarone received the B.Sc. degree in physical science in 2007 and the M.Sc. degree in physics in 2011, both from the Università degli Studi di Palermo, Palermo, Italy, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 2016. Since 2016, she has been a Postdoctoral Researcher in the Institute of Photonics and Quantum Sciences, Heriot-Watt University. Her research interests include underwater and multispectral time-of-flight depth imaging, single-photon counting technologies, free-space, and satellite quantum communication.

Aongus McCarthy (M'03) received the B.Sc. degree from the University College Galway, Galway, Ireland, in 1989, the Diploma degree in electronic engineering from the Institute of Technology, Carlow, Ireland, in 1990, the B.Sc. degree in physical optoelectronics from Essex University, Essex, U.K., in 1991, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 2002. He, then, worked in industry before completing the Ph.D. degree. He is currently a Research Fellow in the School of Engineering and Physical Sciences, Heriot-Watt University. His research interests include optical and optomechanical system design, time-of-flight depth imaging, single-photon counting technologies, and microscope systems. He is a member of the Optical Society of America and the IEEE Photonics Society.

Gregory Newstadt received the B.S. degrees (*summa cum laude*) in electrical engineering and engineering physics from Miami University, Oxford, OH, USA, in 2007. He received the M.S.E. degree in electrical engineering: systems in 2009, the M.A. degree in statistics in 2012, and the Ph.D. degree in electrical engineering: systems in 2013, all from the University of Michigan, Ann Arbor, MI, USA. He is currently a Software Engineer at Google, Inc., Pittsburgh, PA, USA. From 2013 to 2014, he was a Postdoctoral Researcher and a Lecturer at the University of Michigan of electrical engineering (systems). His research interests include detection, estimation theory, target tracking, sensor fusion, and statistical signal processing

Gerald S. Buller received the B.Sc. (Hons.) degree in natural philosophy from the University of Glasgow, Glasgow, U.K., in 1986, and the Ph.D. degree in physics from Heriot-Watt University, Edinburgh, U.K., in 1989. He is currently a Professor of physics at Heriot-Watt University. In 2002, he cofounded Helia Photonics, Ltd., based in Livingston, U.K. In 2015, he received an EPSRC Established Career Fellowship in quantum technology to research sparse photon imaging. His current research interests include single-photon detection methods, particularly at infrared wavelengths, and applications involving the use of single-photons. These applications include single-photon depth imaging and quantum communications protocols. He is a Fellow of the Royal Society of Edinburgh, a Fellow of the U.K. Institute of Physics, and a member of the Optical Society of America.

Steve McLaughlin (F'11) was born in Clydebank, U.K., in 1960. He received the B.Sc. degree in electronics and electrical engineering from the University of Glasgow, Glasgow, U.K., in 1981, and the Ph.D. degree from the University of Edinburgh, Edinburgh, U.K., in 1990. From 1981 to 1984, he was a Development Engineer in industry involved in the design and simulation of integrated thermal imaging and fire control systems. From 1984 to 1986, he worked on the design and development of high-frequency data communication systems. In 1986, he joined, as a Research Fellow, the Department of Electronics and Electrical Engineering, University of Edinburgh, where he studied the performance of linear adaptive algorithms in high noise and nonstationary environments. In 1988, he joined the academic staff at Edinburgh, and from 1991 until 2001, he held a Royal Society University Research Fellowship to study nonlinear signal processing techniques. In 2002, he received a personal Chair in electronic communication systems at the University of Edinburgh. In October 2011, he joined Heriot-Watt University as a Professor of signal processing and the Head of the School of Engineering and Physical Sciences. His research interests include the fields of adaptive signal processing and nonlinear dynamical systems theory and their applications to biomedical, energy, and communication systems. He is a Fellow of the Royal Academy of Engineering, Royal Society of Edinburgh, Institute of Engineering and Technology.

Alfred Hero (F'98) received the B.S. degree (*summa cum laude*) from Boston University, Boston, MA, USA, in 1980, and the Ph.D. degree from Princeton University, Princeton, NJ, USA, in 1984, both in electrical engineering. He is the John H. Holland Distinguished University Professor of electrical engineering and computer science and the R. Jamison and Betty Williams Professor of engineering at the University of Michigan, Ann Arbor, MI, USA. He is also the Co-Director of the University's Michigan Institute for Data Science, Ann Arbor, MI, USA. His primary appointment is in the Department of Electrical Engineering and Computer Science and he also has appointments, by courtesy, in the Department of Biomedical Engineering and the Department of Statistics. His recent research interests include high-dimensional spatio-temporal data, multimodal data integration, statistical signal processing, and machine learning. Of particular interest are applications to social networks, network security and forensics, computer vision, and personalized health. He was the President of the IEEE Signal Processing Society and a member of the IEEE Board of Directors. He has received numerous awards for his scientific research and service to the profession including the IEEE Signal Processing Society Technical Achievement Award in 2013 and the 2015 Society Award, which is the highest career award bestowed by the IEEE Signal Processing Society.