







Determining urban material activities with a vehicle-based multi-sensor system

M. Salathe ^{*}, B. J. Quiter , M. S. Bandstra , J. C. Curtis , R. Meyer, and C. H. Chow 
Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

 (Received 21 December 2020; revised 28 December 2020; accepted 24 March 2021; published 23 April 2021)

Integration of contextual sensors into vehicle-borne mobile radiation detection systems delivers a rich description of the environment to inform estimates of the complex and variable gamma-ray signals observed in urban areas. Models based on these data streams could provide realistic inputs to urban radiological search algorithms and potentially improve the system's sensitivity to detect illicit radiological and nuclear materials. In this work, LiDAR and inertial data are combined using simultaneous localization and mapping techniques to create a three-dimensional (3D) representation of the surrounding scenery. Semantic segmentation of concurrently collected video imagery enables the division of the 3D model into distinct material categories. The radioactive flux of surfaces associated with these categories are inferred through maximum likelihood estimation maximization and the activity of the three most common isotopes (K-40, U-238 series, and Th-232 series) in the respective materials is predicted. The results, found to be in agreement with ground truth measurements performed at the facility, suggest that it is possible to quickly infer the composition of naturally occurring materials in structures that comprise a radiological scene. Such a capability could be used to inform radiological search algorithms and enable data-driven modeling of radiological search problems, which could facilitate system testing and operator training activities.

DOI: [10.1103/PhysRevResearch.3.023070](https://doi.org/10.1103/PhysRevResearch.3.023070)

I. INTRODUCTION

Mobile gamma-ray detection systems are often fielded in uncontrolled environments such as ports, borders, and urban areas for homeland security purposes. Understanding and predicting the measurements that are generated by such fielded detection systems is nontrivial primarily due to the unknown and continuously changing distribution of Naturally Occurring Radioactive Material (NORM) in the vicinity of the fielded detection systems [1]. This NORM comprises an irreducible background atop of which “nuisance” radiological sources, such as medical isotopes and sources in use for industrial purposes, are also present. In addition to these background contributions to the radiological environment, there may be sources present for more nefarious reasons (i.e., threats), whose detection motivates the fielding of such gamma-ray detection systems in the first place. The complexity associated with the background signals motivated the development of complex algorithms, designed to better distinguish radiological threats from background [2–6]. This complexity results in overall probabilities of detecting threat radiological material that are lower than those predicted by a simple model that compares the strength of a hypothesized gamma-ray signature to the associated instantaneous

background [7]. Also, operators of gamma-ray detection equipment must frequently participate in training activities to keep abreast as to how these radiologically complex environments impact the behavior of their systems. Methods that provide better understanding of instantaneous gamma-ray backgrounds typically encountered by mobile detection systems or that can infer radiological activities of materials as they are encountered are therefore potentially valuable for multiple purposes. These methods can facilitate the development of more capable radiological threat detection algorithms and enable high-fidelity inference of the radiological composition of materials within a measured area, thereby enabling realistic high-fidelity radiological models for the purposes of system development, testing and evaluation, and operator training.

The primary sources of naturally occurring radioactivity are primordial potassium-40, uranium-238, thorium-232, and the radioactive progeny of the latter two. These are referred to herein as KUT. KUT is present in nearly all material, including the majority of materials used in construction. Typical KUT concentrations are summarized in [8,9]. In addition to terrestrial gamma rays originating from the decays of KUT isotopes, gamma rays from airborne radon-222, cosmic rays, and skyshine [10] inevitably contribute to the background in urban environments.

A limited number of earlier studies have tried to establish a connection between contextual data and radiological backgrounds. One such study proposed to use a terrain database to formulate a Maximum-Likelihood Estimation Maximization (MLEM) model that predicts radiological background distributions for an airborne system [11] and another featured a panoramic image-based neural network to improve detec-

^{*}msalathe@lbl.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.




FIG. 1. The meshed 3D model of the MOUT facility. The blue line represents the trajectory of the truck (center of detector array). Colors defined by the labels below the image represent the different classification labels resulting from the transfer-learned DeepLabv3+ [18,19] semantic segmentation process. Pastel versions of these same colors indicate labels used for areas with no mesh coverage.

tion performance [12]. That work was based on a previous study that linked segmented panoramic images to radiological measurements [13] and a subsequent effort that expanded the model to three dimensions by leveraging LiDAR-based Simultaneous Localization And Mapping (SLAM) [14]. The current work builds upon these previous studies by leveraging physics-based modeling to better constrain the spectral characteristics of the KUT emissions from the facility and provides quantitative error estimates describing the fidelity of the model. It suggests that a vehicle-borne sensor system driving at typical speeds for an urban setting can infer the complete radiological environment, e.g., the emitted spectra from the constituent structures.

This series of inquiry has leveraged measurements performed by the Radiological Multi-sensor Analysis Platform (RadMAP) system [15] at a controlled facility that is similar in scale to an urban environment. RadMAP was developed to investigate how contextual sensors can enhance and complement radiological data of mobile systems. The components of the RadMAP vehicle of importance in this work are: 100 NaI(Tl) detectors in a coded mask array, a NovAtel SPAN-CPT GPS/INS receiver, two Velodyne HDL-32E LiDAR units, and two Point Grey Ladybug 3 spherical digital video cameras. When using the term *detector* within this publication, we refer to individual detector modules. When talking about the full array the expression *detector array* is used. *Detection systems* is used to refer to the full suite of sensors, i.e., radiological as well as contextual sensors.

The $4 \times 4 \times 2$ in. NaI(Tl) detectors are arranged in a 10×10 grid mounted in an upright fashion such that rows are aligned with the travel direction of the truck. The NaI(Tl) detectors are read out using photomultiplier tubes affixed to the port-side 4×4 in. faces of each detector. The starboard side of the detector array is covered at a distance of approximately 40 cm by a 1-in.-thick Pb half-filled coded aperture array that extends an additional four detector lengths beyond the detector array in the aft-fore direction (for a total of 18 grid elements) and an additional two detector lengths above the detector array. This arrangement was an update, as docu-

mented in [15], to the original design described in Ref. [16]. We do not use the coded mask geometry to perform coded mask imaging in this publication (which has been performed on other vehicle-based systems [17]), since we are interested in measuring the radiation from the entire surroundings rather than localizing point sources. The LiDAR units are installed on the front, port, and starboard corners of the truck and the two Ladybug camera units are mounted directly above the LiDAR units. In this configuration the camera and LiDAR units are located approximately 4.5 m in front of and 1 m exterior to the center of the detector array. The NovAtel GPS/INS unit is mounted on the fore edge of the roof, centered between the camera and LiDAR units. As part of the Multiagency Urban Search Experiment (MUSE) collaboration, RadMAP was brought to the Military Operations in Urban Terrain (MOUT) facility at the Fort Indiantown Gap (FtIG) National Guard training facility in 2016. A model of the facility is shown in Fig. 1. The concentration of NORM at this site was determined through a series of over 70 10- to 15-min measurements with a high-purity germanium detector enclosed in a box with Pb shielding on five sides to minimize sensitivity to the environment except in the direction of surface under investigation. The per-unit-mass activity ranges derived from measurements and the subsequent analysis are described in [20] and are listed in Table I as ground truth. The vehicle-borne data analyzed herein were collected over 165 s of continuous motion of the RadMAP system being driven in a single “figure eight” lap of the facility. The trajectory is shown as the blue line in Fig. 1. Vehicle speeds ranged from 2–8 m/s, where the vehicle driver slowed when cornering and accelerated to a nominal 15 mph when traveling along straight portions of the roadway.

We will describe the mathematical formulation used to determine NORM activities in constituent scenes in Sec. II. This formulation entails detailing the angular sensitivity model of detectors in Sec. II A, the effects of photon scattering and absorption in air in Sec. II B, and the use of semantic segmentation and SLAM to infer the content of the system’s field of view at varying times in Sec. II C. Additional constraints leveraging simulated spectra emitted from materials are de-

TABLE I. The activity attributed to potassium-40 (K), uranium-238 (U), and thorium-232 (T) concentrations in materials and the total fluxes emitted from labeled surfaces compared to other analyses and measurements.

Class	Activities (Bq/kg)						Flux ($\gamma/s/cm^2$)			
	Ground truth [Min, Max] [20]			This work (3D KUT)			Ground truth [20]	Video [13]	3D full [14]	This work (3D KUT)
	K	U	T	K	U	T				
Asphalt	[83.0, 101.0]	[22.0, 25.3]	[2.9, 5.0]	92 ± 12	36 ± 2	2 ± 1	0.56	0.48 ± 0.03	0.87	0.71 ± 0.03
Building brown	[67.0, 257.0]	[6.9, 14.3]	[4.9, 12.1]	154 ± 17	22 ± 2	19 ± 1	0.39	0.72 ± 0.02	0.9	0.90 ± 0.05
Building red	[98.5, 107.5]	[9.9, 10.5]	[5.45, 5.95]	207 ± 31	20 ± 4	21 ± 3	0.28	0.83 ± 0.02	0.98	0.98 ± 0.09
Building roof	N/A	N/A	N/A	162 ± 83	11 ± 1	0 ± 1	N/A	0.00 ± 0.21	0.8	0.36 ± 0.10
Building white	[71.0, 171.0]	[8.1, 12.9]	[4.6, 6.1]	28 ± 49	30 ± 7	21 ± 5	0.29	0.71 ± 0.04	0.98	0.91 ± 0.15
Concrete	[129.0, 260.0]	[14.0, 18.6]	[8.9, 11.9]	216 ± 10	27 ± 1	9 ± 1	0.69	0.73 ± 0.03	1.01	0.86 ± 0.03
Forest	N/A	N/A	N/A	19 ± 38	40 ± 5	18 ± 4	N/A	0.64 ± 0.03	1.19	1.01 ± 0.12
Grass	[337.0, 513.0]	[22.7, 30.9]	[30.8, 43.8]	530 ± 25	42 ± 3	54 ± 2	1.67	1.53 ± 0.03	2.6	2.35 ± 0.07
Gravel	[114.0, 190.0]	[14.8, 23.5]	[2.7, 6.3]	274 ± 47	23 ± 5	10 ± 2	0.56	0.73 ± 0.04	1.12	0.88 ± 0.10
Pole	N/A	N/A	N/A	0 ± 191	0 ± 8	0 ± 8	N/A	N/A	0.47	0.00 ± 0.31
Sky	N/A	N/A	N/A	N/A	N/A	N/A	N/A	0.37 ± 0.02	0.31	0.51 ± 0.02
Vehicle	N/A	N/A	N/A	237 ± 100	40 ± 10	0 ± 1	N/A	0.90 ± 0.10	1.0	0.91 ± 0.19

scribed in Sec. IID, followed by the methods of evaluating the formulated system response in Sec. IIE, before discussing how uncertainties are estimated in Sec. IIF. The results of this formulation when used to process the measured data are presented in Sec. III, followed by a discussion of the fidelity of the model results and implications in Sec. IV before concluding.

II. METHODOLOGY

We consider the case of a vehicle-borne radiation detection system operating in an environment such as one typically observed from an urban or suburban roadway. Such a scene is typically composed of static components such as ground cover, vegetation, buildings, and roadway. Furthermore, there are also moving components such as pedestrian and cars. In general we cannot assume that the KUT composition of these different components are uniform and identical across different entities, but as a simplification we assume that the exterior

surfaces of those components that are semantically labeled identically have identical emission spectra. This allows the overall complexity of the scene to be limited by manifesting the scenes as a limited number of components, each with a distinct gamma-ray signature.

To model the gamma-ray transport between the surfaces of these components and the detectors, the field of view of each detector is divided into small elements (pixels). Each pixel (denoted by k) was divided into equally sized bins in azimuth and elevation. Together, the pixels comprise an image of the full 4π steradians of solid angle surrounding the detector. Pixels near the horizon cover more solid angle area than pixels near the poles, however at a constant elevation all pixels have the same size. The type of structures covered by each pixel and the distance to those structures evolves as RadMAP drives through the facility.

We formulate a model for the average rate of photon-induced events λ_γ in each detector d at energy E'' and time t as given by

$$\lambda_\gamma(d, t, E'') = \sum_{i, E} \alpha_i(E) \underbrace{\sum_{E', k} A_k(E''|E', d) S_{\text{air}}(E'|E, r_k(d, t)) \frac{\delta_{i, l_k(d, t)} \Delta \Omega_k}{\pi}}_{R_i^{\text{3D}}(d, t, E'', E)}. \quad (1)$$

The energy and time dependence of the different terms in Eq. (1) is assumed to be discrete and thus the various terms are approximated with tensors, matrices, and vectors. If they would be continuous the respective sums would need to be replaced with integrals. The numerical derivation of each of the terms will be described in the following subsections, their meaning in the next paragraph.

The first term, the effective area A_k is effectively a product of detector efficiency and geometric area. It describes the probability of energy E'' being deposited into each detector when a photon of energy E' is incident upon the truck, and

factors in the effects of scattering and absorption due to the other materials contained within the truck. The effective area is calculated for a range of photon energies E' and angles of incidence upon the truck. The latter are referenced through k indexing a specific pixel covering a fraction of the full 4π steradian solid angle. The angles that describe a solid angle element are with respect to the detector's reference frame and the term does not depend on time. The second term S_{air} incorporates scattering and absorption in air, linking the energy of the gamma ray emitted from a surface E to the energy E' it retains after traveling a specific distance through air. The term

is defined for all possible distances, but it is only evaluated at distances described by the first geometric component r_k . This function describes the distance from detector d to a surface element present in the scene, covered by pixel k in that detector's field of view at time t . Similarly, the second geometric term, the function l_k , expresses which class label is covered by solid angle pixel k as seen by detector d at time t . This implies that detected gamma rays must originate from surfaces visible from the detector. While this simplification could be lifted by handling gamma-ray transport through intermediate surfaces in this publication occluded surfaces are neglected. The factor of π , included in Eq. (1), originates from Lambert's cosine law, which predicts that the surface emission of gamma rays from a material with uniformly distributed sources follows a cosine distribution. Integrating this flux over 2π steradians solid angles leads to $\alpha^{\max}(E) = \alpha(E)/\pi$, expressing the maximum flux observed perpendicular to the surface as a function of the integrated flux, introducing the factor of π . As in [13], the cosine and r^2 dependence of the flux can be absorbed into the solid angle element $\Delta\Omega_k$. It describes the solid angle coverage of each pixel, which for an angular bin defined by the azimuthal angle range ϕ_{\min} to ϕ_{\max} and elevation angle range θ_{\min} to θ_{\max} is given by

$$\Delta\Omega_k = (\phi_{\max} - \phi_{\min})(\sin\theta_{\max} - \sin\theta_{\min}). \quad (2)$$

The three-term product in Eq. (1), R_i^{3D} , is referred to as the system response.

A. Effective area

The effective area has been derived through a series of Monte Carlo simulations conducted to characterize the RadMAP NaI(Tl) detectors' response to gamma radiation approaching from various directions. The simulations were performed by modeling the densest materials in the truck and transporting photons using the GEANT4-based [21] code, MEGALib [22]. For each of the 520 different incident directions (described by elevation and azimuth angles), photons were emitted in a limited cone directed towards the vehicle. The cone apex points were located in vacuum, 10 m from the center of the detector array. The truck was surrounded by a vacuum and no photon interactions outside the truck were simulated. Thirteen discrete energies were simulated (see [15] for more information) and the energy deposited by the simulation of each gamma particle in each of the 100 detectors was used as input to create a histogram with 128 bins covering the energy range from 0 to 3072 keV. The histogram was interpolated between simulated energy, handling the full energy peak, escape peak and the 511 keV separately from the Compton continuum. The energies were convolved with an energy dependent Gaussian distribution of width $\sigma = 0.15 \times E''[\text{keV}]^{0.77}$. This only mirrors the average resolution observed in RadMAP and the energy resolutions of each individual detector were not handled separately, as highlighted in Sec. IV. Some of the detectors had an energy resolution considerably larger than was used for creating the effective area. Thus, 13 out of 100 detectors were not used in this analysis. The result, a square matrix, holds elements for the effective area as a function of the simulated deposited energy in detectors on one axis and the energy at which

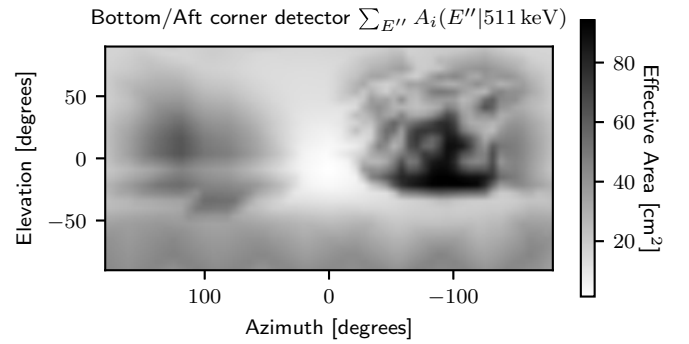


FIG. 2. The effective area of the NaI(Tl) detector in the bottom-aft corner of the array. Displayed is the total sensitivity (sum over all detected energies) for an impending gamma ray of 511 keV.

gamma particles were approaching the truck on the other. The full 4π field of view of each detector was divided into 600×300 pixels. For each pixel the effective area was reconstructed through linear interpolation of the respective elements of the four closest simulated coordinates. The resulting tensor is of size $600 \times 300 \times 128 \times 128$ and has dimensions representing azimuth, elevation, detected energy, and incoming energy, in that order. The tensor was calculated separately for each of the 100 detectors. Index k in Eq. (1) refers to the individual pixel for which the effective area is calculated and the effective area is a function of deposited energy E'' given a certain incoming energy E' . Figure 2 shows a result of this simulation campaign, in the form of total effective area for a single detector and for a gamma ray of $E' = 511$ keV. The zero azimuth is pointing towards the fore of the truck. The coded mask pattern is visible in the positive elevation and negative azimuth area of the figure. It has high effective area where the detector is not covered by Pb and nearby low effective area where the Pb mask is present. Negative elevations are towards the bottom of the truck, where the detector array frame, surrounding Pb, and the structure of the box truck reduce the effective area relative to the upper elevations. The positive azimuth region of the effective area shows some modulation due to the presence of the photomultiplier tube on the back (port) side of the detector.

B. Scattering/absorption in air

The function $S_{\text{air}}(E'|E, r)$ was also determined by Monte Carlo simulations. However, these simulations leveraged more rudimentary physics and geometry engines than GEANT4, similar to that in [23]. The goal of these simulations were to account for the effects of down-scatter as the photons traversed from the identified surfaces in the surroundings to the exterior of the RadMAP truck. The highest fidelity simulations would include specific geometries for each pair of surface segments and detector array stances and would include adjacent solid surfaces and the Lambertian angular emission profile. However, this would be computationally very expensive both in terms of the number of simulations required and in terms of the resulting size of the S_{air} tensor, so two simplifications were made: The effect of differing viewing angles of a series of Lambertian profiles was simplified as a single isotropic profile, thereby removing an angular dependence; and the presence of other solid material was ignored, resulting

in only a volume of air between the emitting surface and RadMAP. Therefore, the simulations comprised the transport of gamma rays through a slab of dry air with a square cross section. The thickness of the slab was varied between 0 and 80 m in 31.25 cm increments and the side length of the slab was set to be large (600 m) relative to the thickness. Gamma rays of a fixed energy were emitted from a point source placed at the center of one face of the slab. The energy of the particles reaching a square area of side length 0.3 times the thickness of the slab was histogrammed. The energy binning of the histogram was the same as that for A_k . The energy at which gamma rays were emitted was also varied over the same range, which results in a matrix with 128×128 elements, for each of the 256 distances. The resulting tensor $S_{\text{air}}(E'|E, r)$ was normalized such that the probability to detect a 3 MeV gamma ray, at the emitted energy or lower, after traveling through any of the simulated distances of air was 1. This normalization implies that the probability that a gamma ray is absorbed at 3 MeV is negligible, however at lower energies the probability of absorption is still larger than zero. Thus, the probability to detect a gamma ray including any down-scattering to lower energies is larger than predicted by the attenuation coefficient [24], but the probability of observing a gamma ray, at the energy it was emitted at, is lower.

C. Detection field of view

The missing quantities in the system response—the two functions l_k and r_k mapping detector, time, and solid angle element (pixel) to labels and to distances—was derived from video, LiDAR, and inertial measurement unit (IMU) data collected with RadMAP while driving the “figure eight” along the main streets of the facility. The corresponding analysis pipeline is summarized in Fig. 3. Concurrently collected images from the two Ladybug camera arrays were combined as described in [13] to produce a single panoramic image. The portion of those panoramic images that contained RadMAP were unchanged frame-to-frame and were manually removed. The images do show slight parallax effects at azimuth angles 0 and 180 deg where images from the port and starboard cameras are stitched together. An example resulting stitched panoramic image is shown as the top panel in Fig. 4. A set of 46 panoramic images were hand-labeled with 11 classes: *asphalt, building roof, building red, building brown, building white, concrete, forest, grass, gravel, sky, and vehicle*. A semantic image segmentation neural network, DeepLabv3+ [18], was transformed to be sensitive to these 11 labels through transfer learning. The existing weights, pretrained on the Cityscapes data set [19], were frozen, except for the last fully connected layer, where new weights were found by retraining the network with the hand-labeled images. Some areas in those images were not covered by any label and were marked to be ignored during transfer learning. After retraining, the converted network was used to segment a fraction of the images recorded during the 165 s period, sampled at 3 Hz. The same set of images were also analyzed with the unaltered version of DeepLabv3+ and areas segmented as *pole* were added in as such in the transfer learned segmented images. This increases the number of labels distinguished in this analysis to 12. The labeling results from the unaltered version

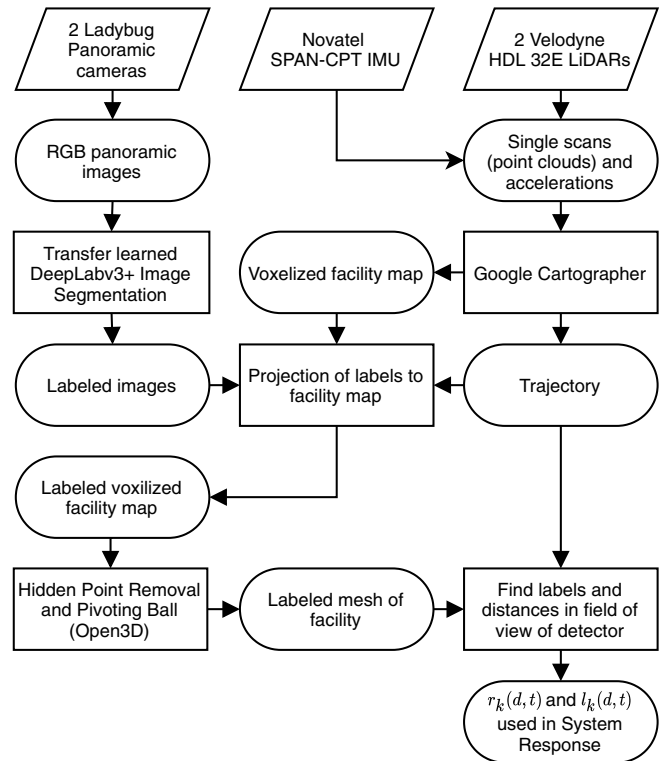


FIG. 3. A diagram of the analysis procedure used to find the two functions $r_k(d, t)$ and $l_k(d, t)$ in Eq. (1). Parallelepipeds represent methods by which data are generated, round boxes represent intermediate data products, and rectangles represent analysis steps.

of DeepLabv3+ and after application of transfer learning are shown in the middle and bottom panels of Fig. 4, respectively.

Range data from the LiDAR and acceleration data from the IMU were used to perform SLAM using the Cartographer code package [25]. The two LiDAR units were configured to generate data packets comprising a 360° view (or “scans”) at 10 Hz. SLAM algorithms [26,27] operate by matching these scans to an internal representation based on the data collected up to this point in time by minimizing a cost function creating an estimation of the truck’s position and orientation within the facility at each time step. The localization result is then used to project the LiDAR measurement coordinates collected at that location into space to recreate a representation of the facility in the form of a point cloud. As the name implies SLAM creates both a trajectory of the truck in the facility and a map of the facility.

Using the concurrent trajectories and facility map, the results of the semantic segmentation are assigned to point cloud constituents. Each pixel in every segmented image, starting from the location of the truck at the time the image was recorded, is projected outwards in increments of 0.75 m. If points from the point cloud are found to be within the wedge formed by the pixel’s solid angle and the radial segment, those points are marked with the class label of the pixel and the search for that pixel is terminated. If the stepping algorithm reaches 80 m without encountering any points the stepping is halted and continued with the next pixel in the image. After processing all the images, the class label that has been




FIG. 4. Semantic image segmentation: The top panel shows the stitched panoramic image with the hood of the RadMAP truck filtered out in the black region in the bottom center of the image. The second panel shows segmentation results with DeepLabv3+ [18] pre-trained with the Cityscapes data set [19]. The bottom panel shows the same image segmented with the network obtained from applying transfer learning with hand-labeled images of the facility to the DeepLabv3+ pre-trained model.

assigned to a point most often is selected to represent that point. Any point that ends up without a label is classified with the label of the nearest point with a valid label.

Next, the resulting point cloud was converted into a triangular mesh. This was done by first estimating the surface normals for each point based on the location of neighboring points within a radius of 0.5 m. The sense of each normal is randomly assigned (inward vs outward) initially. The hidden point removal algorithm [28] then updates the senses of normals to point towards the trajectory for points that are visible from the truck during its journey through the facility. The resulting point cloud with oriented normals is converted into a mesh with the ball pivoting algorithm [29] before being sparsified through vertex clustering [30]. All of these methods were based on algorithms implemented in Open3D [30]; the final mesh is displayed in Fig. 1.

To find the mapping of l_k and r_k in Eq. (1), the truck's locations are sampled at 10 Hz resulting in 1650 locations. From each of the detector's positions at these sampled locations a ray is cast in the direction of every solid angle element's (pixel) center. The intersection between the outgoing rays and

the triangular faces of the mesh are calculated with a ray-triangle intersection algorithm [31]. For each outgoing ray, the three vertices with the smallest intersection-detector distance are selected. The corresponding pixel is then marked by the label of the vertex that is closest to the intersection. In this process we keep track of the distance between intersection and detector, not only to create a tensor representing r_k , but also to decide if a given intersection is closer to the detector and thus must replace the previously assigned label. The resulting matrices construe a rendered version of the mesh visible from a detector at a given point in time and the distances to the visible triangular faces. Some pixels will not have a valid label assigned at this point, because the field of view of the LiDAR—used to create the mesh—and the one of detectors are different. Furthermore, only vertices within 80 m of the detector are used in this process leaving a unlabeled band near the horizon. These pixels are labeled by searching for the intersection with a plane that was aligned with the lowest laying vertices in the mesh. The nearest vertex to the intersection is used as a substitute for that pixel's label. The plane colored by the nearest valid label is shown in Fig. 1 with pastel versions of the label colors. Any pixel still unlabeled at this point was considered to be *sky* at 80 m distance. This process results in two tensors of dimension $1650 \times 300 \times 600$ for each detector. The first tensor assigns to each pixel in the detector's field of view a label that covers that pixel at time t , thus representing l_k in Eq. (1). The second assign a distance to each pixel as a function of time, thus representing r_k .

D. Spectral constraints on the flux

There are two important restrictions to the flux $\alpha_i(E)$ that were imposed here as well as in [14]: A lower energy threshold at 216 keV and an energy bin width of 24 keV. The simulations are unreliable at low energies, where the detector efficiency drops off sharply to zero and can no longer be neglected. The energy threshold was introduced to not be sensitive to that range. The bin width of 24 keV was chosen to limit the total number of free parameters in the model, influencing the same choice in the effective area and air scattering tensor. These two choices restrict $\alpha_i(E)$ to a vector of size 119, which, multiplied by 12 labels, results in




FIG. 5. Simulated emission spectra from a 1-cm² surface area based on 1 Bq/kg activity for the three radioisotopes considered as the the main source of environmental backgrounds.

a set of equations with 1428 free parameters. A solution that does not restrict this parameter set any further is discussed in [14].

This paper focuses on a different method that conveys a more detailed understanding of the radioisotope concentrations of various materials and is based on modeling the energy dependence of the terrestrial labels with simulated KUT templates. The emission spectrum radiated from a slab was calculated using the same simplified Monte Carlo gamma-ray transport code that was used in Sec. II B. Homogeneous source distributions of potassium-40, uranium-238 series, and thorium-232 series gamma-ray energies throughout the slab were simulated with a density of the slab of 1.52 g/cm^3 and its composition as reported for U.S. average earth in [32].

$$\lambda_\gamma(d, t, E'') = \sum_{(i,m) \in \mathcal{I}} \alpha_{im} \underbrace{\sum_E [S_m(E) \delta_{m \in \text{KUT}} + \delta_{m=E}] R_i^{3\text{D}}(d, t, E'', E)}_{\hat{R}_{im}^{3\text{D}}(d, t, E'')} \quad (3)$$

The index m refers to KUT components and the 119 energy bins E considered in this analysis simultaneously. The sum running both over m and i can be restricted to the set of parameters defined as

$$\mathcal{I} = \{(i, m) : i \neq \text{sky}, m \in \{K, U, T\} \text{ or } i = \text{sky}, m \in \{E_1, \dots, E_N\}\} \quad (4)$$

This set includes only 155 elements; three that reference KUT in each of the 11 terrestrial labels and 119 that describe the *sky* spectrum. $S_m(E)$ is the simulated spectral shape emitted from soil by KUT as introduced earlier in this section. Together with the sum over all emitted energies E , it can be absorbed into the definition of the system response, forming $\hat{R}_{im}^{3\text{D}}(d, t, E'')$.

The simulation was set up to have an arbitrary 1 Bq/kg of radioactivity so that α_{im} reflects directly the activity of a given isotope in a material and is no longer its surface flux. However, the resulting emission spectra from a surface can be normalized to unity so that it expresses the gamma-ray flux in photons/s cm^{-2} . This conversion was used in Table I to convert activity to flux, before adding the KUT components together. *Sky* was assumed to be at the maximum distance of 80 m from any detector all the time and the respective amplitude represent the gamma flux per unit surface at that distance. With these restrictions in place, we are left with solving the system of equations defined in Eq. (3).

E. Maximum-likelihood estimation maximization

Given a set of detected gamma rays characterized by the detector module d_n in which they were detected, the time of detection t_n and the detected energy E_n'' , MLEM can be used to predict α_{im} . We focus on a method that minimizes the negative Poisson log-likelihood because the observation in detectors

We note that the simulated thicknesses of materials was 2 m, which is essentially infinite and is likely a good approximation for soil (*grass*) and the roadway materials, but most likely is not accurate for structures of limited thickness such as buildings and *vehicle*. The spectra emitted from a 1 cm^2 area of surface predicted by this model is shown in Fig. 5 for each KUT component. Consequently, $\alpha_i(E)$ in Eq. (1) can now be expressed with only three unknown parameters per label, the activities of the KUT components, respectively. This approach is not valid for the *sky* component, as it has important contribution that do not originate from KUT [10]. Instead all 119 energy bins of the *sky* component's spectral distribution were only constrained to be positive. Hence, Eq. (1) can be expressed as

are discrete events:

$$\alpha_{im}^{j+1} = \frac{\alpha_{im}^j}{\sum_{d, E''} \int_0^T \hat{R}_{im}^{3\text{D}}(d, t, E'') dt} \times \sum_n \frac{\hat{R}_{im}^{3\text{D}}(d_n, t_n, E_n'')}{\sum_{(\tilde{i}, \tilde{m}) \in \mathcal{I}} \alpha_{\tilde{i}\tilde{m}}^j \hat{R}_{\tilde{i}\tilde{m}}^{3\text{D}}(d_n, t_n, E_n'')} \quad (5)$$

The above equation does not require any time binning and is often referred to as list-mode MLEM in literature [33]. In this publication, $\hat{R}_{im}^{3\text{D}}$ at (d_n, t_n, E_n'') stands for selecting the respective element from the multidimensional tensor. This simplification was considered sufficient for this analysis, however, in a more advanced investigation this could be substituted with an interpolation between nearby elements.

$R_{im}^{3\text{D}}$ was evaluated for every one of the 2 million events observed during the 165 s measurement duration. Simultaneously, the denominator in the first fractions, often referred to as the sensitivity, was precalculated, as it does not change during the update procedure. We used equal amplitude as an initial guess and processed the update rule for 100 000 iterations, using CuPy [34] to run the algorithm on the GPU for faster execution. It was observed that a large number of iterations is necessary to have convergence of the result.

F. Uncertainty estimation

Neither negative activities in a material nor negative values in the *sky* spectrum are physical. This non-negativity constraint is automatically satisfied by the MLEM approach that we chose in Eq. (5). But as discussed in Sec. III, some of the values in the solution are zero or very small. Thus, our solution will not be distributed according to Gaussian statistics and the Fisher information and the Cramér-Rao bound are unreliable as an uncertainty estimate. An alternative approach to understand the spread of such an experiment, and thus estimate uncertainties, is Monte Carlo simulations. The simulation was conducted by replacing α_{im} in Eq. (1) with the solution found through MLEM and calculating $\lambda_\gamma(t, E'')$.

Running these simulations for histogram mode is simpler than for list-mode MLEM because it is possible to Poisson sample the average gamma-ray rate expressed in Eq. (1) directly, without the need to also sample the temporal distribution, as would be necessary to resample the list-mode model. Unfortunately, the full expression of the system response for all 87 detectors used in this analysis (≈ 1 GB/detector) will not fit into the random-access memory (RAM) available on the computer system we used. Also the computations would be much more demanding than in the list-mode case. Thus we simplified the linear system by accumulating all counts from a 0.1 s interval into a single detector placed in the center of the array. The effective areas of the individual detectors were summed to reflect the effective area of the entire detector array. The simplified histogram version of the MLEM algorithm thus can be formulated as

$$\alpha_{im}^{j+1} = \frac{\alpha_{im}^j}{\sum_{t,E''} \hat{R}_{im}^{3D}(t, E'')} \sum_{t,E''} \frac{n(t, E'') \hat{R}_{im}^{3D}(t, E'')}{\sum_{(\tilde{i}, \tilde{m}) \in \mathcal{I}} \alpha_{\tilde{i}\tilde{m}}^j \hat{R}_{\tilde{i}\tilde{m}}^{3D}(t, E'')}, \quad (6)$$

where $n(t, E'')$ is a two-dimensional histogram of the measured events with time and energy on the respective axes.

These two approximations additionally reduce the number of iterations required for the MLEM algorithm to converge and make it possible to run MLEM sufficiently fast for use in a Monte Carlo approach. The average rate of the optimized histogram solution was Poisson sampled and the respective amplitudes for every sample of photon-induced events solved with histogram MLEM. We simulated 5000 realizations each with 3000 MLEM iterations. The uncertainties used throughout this publication were approximated with the sample variance of the simulated distributions. Furthermore, the sample covariance was used to calculate the correlation matrix, discussed in the next section.

III. RESULTS

Moving on to the results of this approach. The results are presented in Table I and Fig. 6 by juxtaposing the activity ascribed to different NORM constituents with the range of activities obtained through ground truth measurements conducted in [20]. The range of measurements from that publication is expressed by the extrema in Table I and represented as error bars in Fig. 6. Activities of the materials classified as light and dark tan concrete masonry units, and white and gray concrete masonry units in [20] were averaged to form label *building brown* and *building white*, respectively. The fluxes in Table I that were taken from other publications were adjusted to reflect only gamma rays emitted between 216 and 3072 keV. Furthermore, the flux in [13] was divided by a factor of 2 missing in that publication. The sky fluxes describe the emission off a surface at 80 m distance from the detector. Therefore, there is considerable attenuation included in the sky results presented in the two last columns of Table I. The results from [13] did not consider any attenuation. Most of our activities, represented as black dots in Fig. 6, are within a factor of 2 of the ground truth average value, marked with a central gray bar in Fig. 6. Only *building red* and *building white* have absolute deviations up to a factor of 5. Thus, indeed some meaningful information about the radioisotope

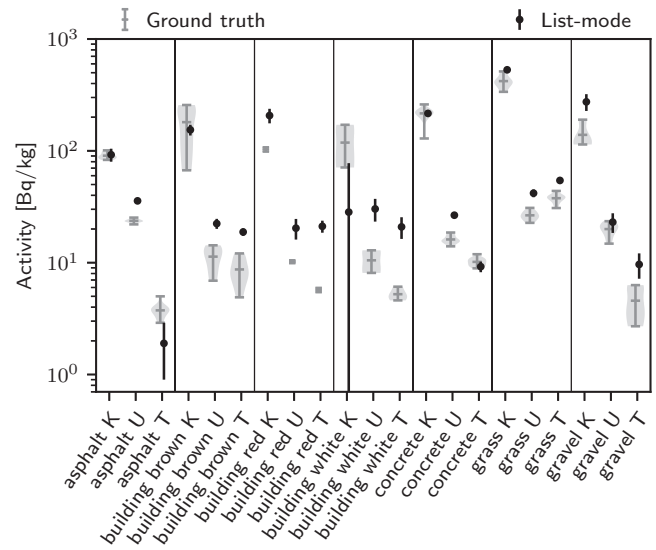


FIG. 6. Our predictions for the activities based on list-mode MLEM (listed in Table I) are represented with black dots. The gray error bars indicate the range of measured values of the same material class but different locations and the central gray bar marks the location of the average.

concentration of materials in the facility was extracted with only 165 s of data. The ground truth error bars in Fig. 6 cover the full range of measured values of the same material class but different locations observed in [20] and are quoted in Table I by the minimum and maximum observed value. Figure 7 studies the sigma level separating our result from the closest value reported in [20]. About 1/3 of all activities are within the range covered by the ground truth measurements or within one sigma of an extremum and only three activities are beyond 5σ . The thorium and uranium predictions are in bigger conflict than the potassium activities. Their activity is in general about 10 times lower and both are spectrally more difficult to resolve, having more Compton continuum and only the 2614 keV line as a clear separating feature. However, it is not entirely clear why the two uranium activities for *asphalt* and *concrete* are discrepant, but possible reasons will be highlighted throughout the following discussion. A first possible explanation could be that large scale structures, such




FIG. 7. The number of sigmas between the nearest value measured in [20] and our values. A zero indicates the value is within the range of observed ground truth values.

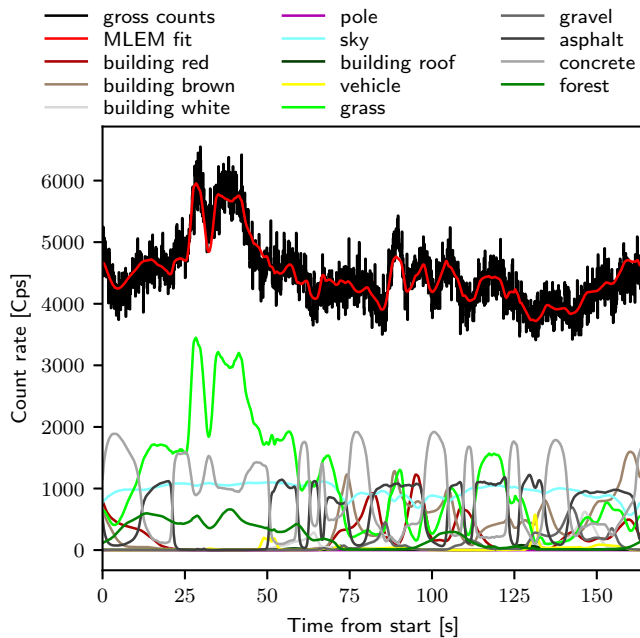


FIG. 8. Count rate as a function of time: The black line describes the count rate observed in the detector array, the red line the rate predicted by the list-mode MLEM result. The remaining colored lines describe what fraction of the total rate is attributed to each material label at a given time. The figure is in analogous to the results shown in [14], but are reconstructed here through adding the contributions associated with each of the three radioisotopes.

as *grass* and *asphalt/concrete*, might not be uniform across the entire facility. The measurements in [20] were conducted mostly along the street segment running through the center of the facility and the left loop of the “figure eight” displayed in Fig. 1. Our data set averages the activities observed in materials in the entire facility, including regions with little coverage in the ground truth data set. The observation that our averages are not in agreement with their ranges could be related to their measurement not capturing the variability of a given material class completed. It is interesting to note that the simplified histogram-based analysis does not show these discrepancies and it merely could be that the sample variance, based on that result, might underestimate the uncertainties of the list-mode approach for these activities. Figure 8 shows the fraction of the total count rate measured across the detector array that was attributed to each label, and helps to expand our understanding of why some labels are in better agreement with ground truth measurements than others. The labels that continuously make up a substantial fraction of the total count rate, i.e., *asphalt*, *concrete*, *grass*, and *gravel*, are better when comparing activities and thus are well placed inside the range of values observed in [20]. The building labels, particular *building red* and *building white*, are only visible from the truck for a limited amount of time and therefore fail to substantially contribute to the total count rate, which results in less agreement with the ground truth measurement range. *Building roof*, *vehicle*, and *pole*, are most likely also not very well confined, but there is no ground truth measurement available to compare these labels to. It should be noted that overall, the results presented in Fig. 8 are close to the full spectral analysis performed

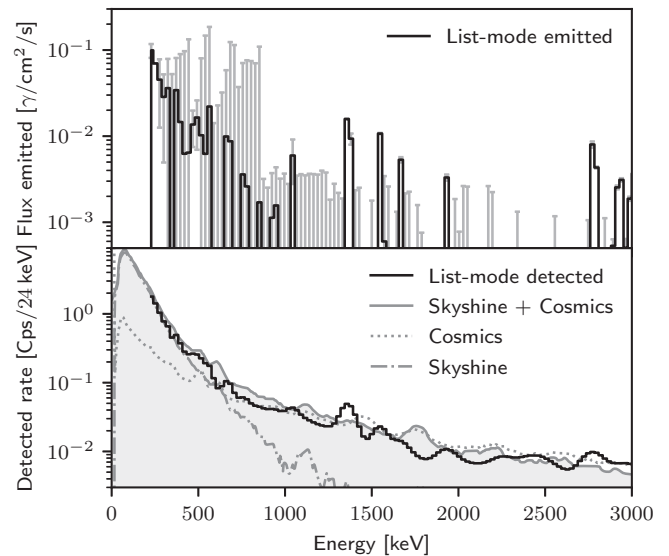


FIG. 9. The top panel shows the emission spectra for the *sky* component as a black histogram with gray error bars. It is based on *sky* being a surface at 80 m distance from the detector. The bottom panel shows the detected spectra, i.e., the spectrum shown in the upper panel folded with the system response, as a black histogram. Also shown are results from [10] for skyshine, cosmics, and the sum of both with different types of gray lines.

in our previous analysis [14]. Particularly, the reduction in free parameters by almost an order of magnitude seems to only have a marginal effect on the agreement between the fit, represented by the red line, and the total gross counts, shown in black. One important difference of the solution presented in Fig. 8 to the analysis in [14] is that our *sky* allotment is larger. This difference is reflected in Table I, which compares the total flux per label to the earlier analyses of [13,14] and the measurements from [20]. The higher apportionment to *sky* results in a smaller quota to other labels, moving the respective fluxes closer to the ground truth value than was observed in [14]. We note that the result still overestimates the ground truth result and that the camera-only result, which was corrected from the results in [13] by removing an erroneous factor of 2, still is in better agreement with the ground truth numbers. The analysis performed in [13] ignored scattering and attenuation in air. Including air attenuation here and the analysis from [14] forces the predicted fluxes to be larger in magnitude. This suggest that the better agreement of [13] could be accidental. All three studies are likely affected by the same systematic uncertainties. One likely source of systematic errors is the simulations of the effective area, discussed in Sec. II A and used for all analyses alike. These simulations might underestimate some of the housing of the truck or neglect some of the structures present in the trucks interior and thus result in a overestimation of the fluxes. The parameters associated with activity incident from *sky* are shown in the top panel of Fig. 9 in black in the form of an energy spectrum with 24 keV wide bins. The respective distribution falls off quickly with energy, with the most significant contribution being below 600 keV. The only bins with large amplitudes above 600 keV are found on both sides of the 1460

and 2615 keV transitions. The energy resolution used in the MLEM model may not sufficiently describe the measurement resulting in nonphysical features in the only unconstrained portion of the reconstructed parameter space that is able to accommodate the difference. The absence of other large contributions implies that three KUT emission components and the formulation of the response matrix otherwise adequately captures the spectral features observed in the detectors. The detected *sky* component, corresponding to the emitted component folded with the system response, is shown in the bottom panel of Fig. 9. Its energy spectrum follows closely, both in shape and amplitude, measurements conducted in [10] for a skyshine with cosmics component. All data in the second panel of Fig. 9 were normalized by volume to represent the spectrum observed in a $4 \times 4 \times 2$ in. detector. All spectra were transformed to have 24 keV wide bins. The bin widths used in [10] were not stated, but were inferred to be 6 and 4 keV for “skyshine + cosmics” and “cosmics”, respectively, resulting in self-consistent spectral scaling for those spectra. Our *sky* component seems to underpredict the 1.8 MeV region, where strong uranium/radon lines are present. This indicates that we might attribute some of radon and terrestrial back scattering erroneously to other labels, particularly *asphalt* and *concrete*, which could be another possible explanation for the unexplained excess of uranium in these two labels. Beyond this region, only the two peak regions pointed out earlier differ considerably between the two measurements. It is noteworthy that neither the cosmics nor the skyshine component described in [10] by themselves are a good fit for the measured spectral distribution. This is in agreement with the expected emission of areas classified as *sky* being partially cosmics and partially terrestrial components scattered in the air back towards the ground. This implies that our *sky* component realistically models the emission expected from a surface element covering the sky.

The activities and *sky* spectrum of the simplified histogram based approach does not substantially differ from the list-mode results presented here. The differences is covered by the uncertainties estimated through simulations. However, the total allotment to *sky* is about 50% larger for the histogram-based result, placing the activities mostly below the list-mode values. A possible explanation for this observation is that the detectors in the top rows of the detector array are particularly sensitive to the *sky* component, and some of that sensitivity is lost when consolidating the detectors to a single module. The similarity between the two analyses legitimizes using the uncertainties obtained from Monte Carlo simulations for the list-mode result as well, but as expressed earlier might not be an equally accurate substitute for all activities.

Now we turn our attention to the Monte Carlo results themselves. There was no systematic bias found between the mean of the sample distributions and the sampled parameters. It is important to note that, as expected, the distributions for values close to zero were not distributed symmetrically around the mean and the sample variance overestimates the lower bound and underestimates the upper bound of these variables. In particular, the *sky* component has many elements with zero being included within the bounds of the uncertainties. The correlation matrix, displayed in Fig. 10, is obtained by normalizing the covariance matrix across the diagonal. The

information about the absolute uncertainty is thus lost. While the correlation matrix reveals a lot of interesting relationships, those based on elements with large uncertainties might not be always especially meaningful. The correlation matrix in Fig. 10 represents strongly anticorrelated variables with blue and strongly correlated variables with red, while uncorrelated variables are marked white. The intralabel KUT relations are represented as a 3×3 block diagonal matrix in the upper left quadrant. They have large off-diagonal intralabel anticorrelations, indicating that the total flux is better constrained than the individual allocations to isotopes. The strongest intralabel anticorrelations are present between potassium and uranium, and between uranium and thorium. These patterns suggest that the 5σ deviations observed in some labels for uranium indeed might be caused by intralabel confusions of isotopes. No intralabel anticorrelations are apparent for the labels *building roof* and *pole*, which in general are poorly constrained because of little coverage by the detector array.

Next, advancing to the off-diagonal matrix elements in the upper left quadrant of the correlation matrix, where the label based block matrix diagonal elements describe interlabel correlations between identical isotopes. The most noticeable features are found in labels that are geometrically linked to each other. Both *asphalt* and *concrete* are part of the streets; one of the *building brown* labeled buildings is located at the edge of the facility near *forest*; the *building red* labeled buildings are positioned towards the center of the facility near large patches of *grass*. These pairs of labels are encountered with a similar coverage of the detectors field of view but rarely observed simultaneously, explaining the interlabel correlations between identical isotopes. Contrastingly, the KUT components in *gravel* are anticorrelated with *building white*, two labels encountered in close proximity at the same time. *Concrete* is found around *building brown* and *building red*, thus also observed simultaneously and the KUT components anticorrelate between labels as well. Last but not least, *grass* and *forest* simultaneously fill an important fraction of the field of view near the horizon for most of the measurements duration and thus are strongly anticorrelated. Figure 8 can be inspected to better gauge which labels appear concurrently and which do not.

The most distinct feature visible in the off-diagonal quadrants of the correlation matrix are centered around the two gamma-ray transitions at 1460 and 2615 keV. Particularly, the potassium activity of *asphalt* and *concrete* correlate with the *sky* spectrum, while the uranium activity anticorrelations around these transitions. The anticorrelation pattern of uranium-238 with the *sky* component might be related to a radon-222 trace in the *sky* spectrum. Radon-222 is part of the uranium-238 decay chain and there might be a balance between what fraction of the observed events is assigned to terrestrial NORM and sky-borne radon. By attributing more uranium to the two street labels and reducing the amount of radon in *sky*, the model furthermore compensates for the total flux being constrained by reducing thorium and potassium.

The inverse correlation patterns are observed for buildings. If the truck passes a building the presence of the building covers up a considerable amount of the sky. Thus, again we see the opposite behavior of labels seen at different times than those seen together. Noteworthy is that there is little correla-




FIG. 10. The correlation matrix of the KUT activities for each material label and *sky* component. The color scale goes from blue, highlighting strongly anticorrelated components (−1), to white, indicating no correlations at all, to red, marking strong positive correlations (1).

tion between *forest/grass* and *sky*, which are seen during most of the measurement.

The most prominent features in the bottom right quadrant are strong anticorrelation features close to the diagonal and some far off-diagonal correlations around 1.8 MeV. The former are found near lines in the KUT spectra, and probably related to the energy resolution of the detector not accurately being modeled by the effective area. As mentioned earlier, the uncertainties in the *sky* component are large and these patterns near the diagonal probably not especially meaningful. The latter is most likely due to an interplay between a radon-222

component and terrestrial uranium-238. In the spectrum of both of these isotopes strong lines are present around 1.8 MeV, arising from the progeny bismuth-214.

IV. DISCUSSION

We have shown that realistic NORM activities present in a urban area can be measured in very brief measurements, here with less than 3 min of continuous motion of a mobile system. Thus, realistic models of urban environments can be built

from contextual data that captures the radiological complexity and variability of materials in proximity to the mobile system.

The construction of the segmented model (simultaneous localization and mapping runs in real time, semantic segmentation took ≈ 2 s per image, and meshing requires about 6 h on a Intel Core i7-5930K CPU), the computation of the system response (≈ 5 h/detector on a Intel Core i7-5930K CPU) and finding the maximum likelihood estimation maximization solution (≈ 1.5 h/100 000 iteration on a GeForce GTX 1080 Ti GPU) was performed offline. The analysis in its current form does not translate easily into a real-time framework for long extensive search missions. However, it could be used offline together with large data sets to populate a database of structures that is used for urban radiological searches. Novel real-time algorithms could leverage this information to infer radiological backgrounds based on Eq. (1) when passing by a known structure or a structure that resembles a member of the database. With advances in machine learning and with more powerful computing resources such an analysis could potentially be performed in real time by inferring a current background prediction from a model based on data collected in the recent past. A real-time implementation of the presented algorithm would leverage a localized approach and not be based on the creation of a full scene. As such, local meshing based on single scans from LiDARs has been shown to be possible in real time [35]. Furthermore, computing the system response on the GPU instead of the CPU could lead to a considerable speed-up and allow for close to real-time realization of the full analysis.

Although this method was sufficient to approximate the radioactive composition of materials present in the scene, it also highlights areas with room for improvement. First, the model that was used to conduct the effective area simulations simplifies aspects of the truck and is partially based on assumptions. For example, altering the thickness of the metal sheet used as floor will increase or reduce the detector arrays sensitivity to labels typically encountered at low elevation such as *asphalt* and *concrete* and thus change the partition of the total count rate to labels. Unfortunately, the truck is no longer available to cross-check the simulation result with actual source measurements, but when applying the method to different mobile system it is crucial to properly understand the effective area used in the calculations. Second, the effective area also assumed that the energy resolution is uniform across the different detectors, but in reality there are considerable deviations. The energy resolution of each detector could be characterized and modeled independently. Third, a more complex scattering and attenuation model, properly taking into account the detector to surface orientations, could increase the accuracy of the method. To understand some of the systematic errors arising from scattering in air we investigated the effect of using a point source with a Lambert's cosine emission profile in place of the isotropic emission used to calculate the air scattering and absorption tensor in Sec. II B. The differences were found to be within the uncertainties, except for potassium in *asphalt* and uranium in *building roof*, where the Lambert's cosine emission analysis is lower by 20% and 40%, respectively. Fourth, we limited our sensitivity to objects in direct line of sight from the detectors. However, at 1 MeV about 20% of all gamma rays will penetrate through a




FIG. 11. Comparing the simulated emission spectra from [20] (gray filled histogram) to the distributions predicted in this work (black histogram).

concrete wall of 10 cm thickness. It would be more accurate to also include the effects of obstructed structures by properly accounting for attenuation. This will be particularly important when operating in a more realistic urban scene with vegetation obstructing structures. Fifth, we assumed identical emission spectra for all structures in the scene modeled with a slab of 2 m thick soil. Figure 11 compares the spectral distribution emitted from a given surface predicted by our calculation and those used in [20]. Whereas most spectra agree quite well, some difference is apparent at the highest energies for *gravel* and the simulated building spectra disagree in both magnitude and in shape. The summed flux for the building is overestimated by roughly a factor of 3 and our simulation predicts a smaller peak-to-Compton ratio. When using the ground truth activities listed in Table I to scale the simulated emission spectra for KUT, the summed flux in the “3D KUT” portion of Table I remains overestimated by a factor of ≈ 1.5 . For buildings (walls), instead of using an infinite thick slab [20], used a brick wall with air and a flooring of soil and asphalt behind to approximate the flux from the surface. This demonstrates that the prediction obtained with the presented method could be improved by using individually tailored simulations for the different labels. Lastly, this study was solely based on radioactive events measured with the NaI(Tl) array. However,

RadMAP also features 14 mechanically cooled high purity germanium detectors. Albeit, the number of recorded events is much lower, the improved resolution of these detectors would enable a better separation of the thorium and uranium series spectra and thus might have the potential to more confidently reconstruct the activity of these two components.

Some aspects of the observed result, in particular the inconsistency of our result with average values measured in [20], suggest that there might be some natural limit to the applicability of such an approach. The large spread in activities measured in [20] suggest that even within homogeneous structures, such as those present in this urban mock facility, the isotopic compositions in materials could be quite variable. While it still should be possible to predict average activities, even for materials with large dispersion in activity, these predictions could be of little use for modeling the backgrounds in the detectors that arise as a result of NORM activities.

V. CONCLUSION

A gamma-ray transport and sensing model of real-world measurements has been created by combining the contextual data collected during the measurement and processed via computer-vision techniques with Monte Carlo-based physics simulations of the detectors and radiation transport processes. This model mimics the physical world sufficiently well to enable approximate predictions of the activities of naturally occurring radioactive materials (NORM) present in the scene by performing MLEM on the gamma-ray data measured in the detector array during a 165-s-long vehicle-borne measurement. This agreement between model and reality

demonstrates substantial understanding of the radiological environment that has previously only been achieved through labor-intensive ground truth measurement campaigns, which are only practicable in controlled environments. The focus was placed on obtaining the potassium-40, uranium-238, and thorium-232 content within various classes of materials. However, this procedure can also be applied in reverse to predict detector count rates, based on known NORM in the vicinity of a mobile system. This is demonstrated best through Fig. 8 that shows how the total gamma-ray count rate was apportioned to the different classes, deduced through semantic image segmentation. Obviously if an estimate of the activities is known, Eq. (1) can be evaluated and the count rate in the detectors approximated by tallying the gamma-ray emission from the various classes. Such a prediction could be used in algorithms, for example as a Bayesian prior, to improve the sensitivity of the system to detecting and separating “nuisance” radiological sources, such as medical isotopes and sources in use for industrial purposes from illicit radiological and nuclear materials.

ACKNOWLEDGMENTS

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231. The project was funded by the U.S. Department of Energy, National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development (DNN R&D). This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.

-
- [1] T. J. Aucott, M. S. Bandstra, V. Negut, D. H. Chivers, R. J. Cooper, and K. Vetter, Routine surveys for gamma-ray background characterization, *IEEE Trans. Nucl. Sci.* **60**, 1147 (2013).
 - [2] B. R. Cosofret, K. Shokhirev, P. Mulhall, D. Payne, and B. Harris, Utilization of advanced clutter suppression algorithms for improved standoff detection and identification of radionuclide threats, in *Chemical, Biological, Radiological, Nuclear, and Explosives (CBRNE) Sensing XV*, edited by A. W. F. III, International Society for Optics and Photonics (SPIE, Bellingham, WA, 2014), Vol. 9073, pp. 253–265.
 - [3] P. Tandon, P. Huggins, R. Maclachlan, A. Dubrawski, K. Nelson, and S. Labov, Detection of radioactive sources in urban scenes using Bayesian aggregation of data from mobile spectrometers, *Inf. Syst.* **57**, 195 (2016).
 - [4] D. Pfund, K. Anderson, R. Detwiler, K. Jarman, B. McDonald, B. Milbrath, M. Myjak, N. Paradis, S. Robinson, and M. Woodring, Improvements in the method of radiation anomaly detection by spectral comparison ratios, *Appl. Radiat. Isot.* **110**, 174 (2016).
 - [5] T. H. Joshi, R. J. Cooper, J. Curtis, M. Bandstra, B. R. Cosofret, K. Shokhirev, and D. Konno, A comparison of the detection sensitivity of the Poisson clutter split and region of interest algorithms on the RadMAP mobile system, *IEEE Trans. Nucl. Sci.* **63**, 1218 (2016).
 - [6] K. J. Bilton, T. H. Joshi, M. S. Bandstra, J. C. Curtis, B. J. Quiter, R. J. Cooper, and K. Vetter, Non-negative matrix factorization of gamma-ray spectra for background modeling, detection, and source identification, *IEEE Trans. Nucl. Sci.* **66**, 827 (2019).
 - [7] K.-P. Ziocok and K. E. Nelson, Maximum detector sizes required for orphan source detection, *Nucl. Instrum. Methods Phys. Res. Sect. A* **579**, 357 (2007), Proceedings of the 11th Symposium on Radiation Measurements and Applications.
 - [8] R. Sonkawade, K. Kant, S. Muralithar, R. Kumar, and R. Ramola, Natural radioactivity in common building construction and radiation shielding materials, *Atmos. Environ.* **42**, 2254 (2008).
 - [9] Directorate-General for Environment, Radiation protection: Radiological protection principles concerning the natural radioactivity of building materials, Tech. Rep. 112 (European Commission, Office for Official Publications of the European Communities, Luxembourg, 2000).
 - [10] G. A. Sandness, J. E. Schweppe, W. K. Hensley, J. D. Borgardt, and A. L. Mitchell, Accurate modeling of the terrestrial gamma-ray background for homeland security applications, in

- 2009 *IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, Orlando, FL, USA (IEEE, Piscataway, NJ, 2009), pp. 126–133.
- [11] R. D. Penny, T. M. Crowley, B. M. Gardner, M. J. Mandell, Y. Guo, E. B. Haas, D. J. Knize, R. A. Kuharski, D. Ranta, R. Shyffer, S. Labov, K. Nelson, B. Seilhan, and J. D. Valentine, Improved radiological/nuclear source localization in variable NORM background: An MLEM approach with segmentation data, *Nucl. Instrum. Methods Phys. Res. Sect. A* **784**, 319 (2015), Symposium on Radiation Measurements and Applications 2014 (SORMA XV).
- [12] C. Kaffine, B. Pires, A. Laddha, D. Bayani, K. Miller, M. Hebert, and A. Dubrawski, Background spectrum estimation from panoramic images (2017), 2017 IEEE Nuclear Science Symposium (unpublished).
- [13] M. S. Bandstra, B. J. Quiter, J. C. Curtis, K. J. Bilton, T. H. Joshi, R. Meyer, V. Negut, K. Vetter, D. E. Archer, D. E. Hornback, D. E. Peplow, C. E. Romano, M. W. Swinney, T. L. McCullough, and M. L. McLean, Attribution of gamma-ray background collected by a mobile detector system to its surroundings using panoramic video, *Nucl. Instrum. Methods Phys. Res. Sect. A* **954**, 161126 (2020).
- [14] M. Salathe, M. S. Bandstra, B. J. Quiter, and J. C. Curtis, Using 3d-scene data from a mobile detector system to model gamma-ray backgrounds, in *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Manchester, UK (IEEE, Piscataway, NJ, 2019), pp. 1–4.
- [15] M. S. Bandstra, T. J. Aucott, E. Brubaker, D. H. Chivers, R. J. Cooper, J. C. Curtis, J. R. Davis, T. H. Joshi, J. Kua, R. Meyer, V. Negut, M. Quinlan, B. J. Quiter, S. Srinivasan, A. Zakhor, R. Zhang, and K. Vetter, RadMAP: The radiological multi-sensor analysis platform, *Nucl. Instrum. Methods Phys. Res. Sect. A* **840**, 59 (2016).
- [16] E. A. Wulf, B. F. Philips, W. Neil Johnson, B. Leas, and L. J. Mitchell, MISTI imaging and source localization, in *2008 IEEE Nuclear Science Symposium Conference Record, Dresden, Germany* (IEEE, Piscataway, NJ, 2008), pp. 2413–2417.
- [17] K. P. Ziocck, J. W. Collins, W. W. Craig, L. Fabris, R. C. Lanza, S. Gallagher, B. K. P. Horn, N. W. Madden, E. Smith, and M. L. Woodring, Source-search sensitivity of a large-area, coded-aperture, gamma-ray imager, in *IEEE Symposium Conference Record Nuclear Science 2004, Rome, Italy*, Vol. 1 (IEEE, Piscataway, NJ, 2004), pp. 322–326.
- [18] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, Encoder-decoder with atrous separable convolution for semantic image segmentation, in *European Conference on Computer Vision - ECCV 2018*, Lecture Notes in Computer Science, Vol. 11211, edited by V. Ferrari, M. Hebert, C. Sminchisescu, and Y. Weiss (Springer, Cham, 2018), pp. 833–851.
- [19] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, The Cityscapes dataset for semantic urban scene understanding, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA (IEEE, Piscataway, NJ, 2016), pp. 3213–3223.
- [20] M. W. Swinney, D. E. Peplow, B. W. Patton, A. D. Nicholson, D. E. Archer, and M. J. Willis, A methodology for determining the concentration of naturally occurring radioactive materials in an urban environment, *Nucl. Technol.* **203**, 325 (2018).
- [21] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytráček, G. Cooperman *et al.*, Geant4—A simulation toolkit, *Nucl. Instrum. Methods Phys. Res. Sect. A* **506**, 250 (2003).
- [22] A. Zoglauer, R. Andritschke, and F. Schopper, MEGAlib—The medium energy gamma-ray astronomy library, *New Astron. Rev.* **50**, 629 (2006), Astronomy with Radioactivities. V.
- [23] F. Arqueros and G. D. Montesinos, A simple algorithm for the transport of gamma rays in a medium, *Am. J. Phys.* **71**, 38 (2003).
- [24] M. Berger, J. Hubbell, S. Seltzer, J. Chang, J. Coursey, R. Sukumar, D. Zucker, and K. Olsen, XCOM: Photon Cross Sections Database, Tech. Rep. (NIST Standard Reference Database 8, 2010).
- [25] W. Hess, D. Kohler, H. Rapp, and D. Andor, Real-time loop closure in 2D LiDAR SLAM, in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, Stockholm, Sweden (IEEE, Piscataway, NJ, 2016), pp. 1271–1278.
- [26] H. Durrant-Whyte and T. Bailey, Simultaneous localization and mapping: Part I, *IEEE Robotics Automation Mag.* **13**, 99 (2006).
- [27] T. Bailey and H. Durrant-Whyte, Simultaneous localization and mapping (SLAM): Part II, *IEEE Robotics Automation Mag.* **13**, 108 (2006).
- [28] S. Katz, A. Tal, and R. Basri, Direct visibility of point sets, *ACM Trans. Graph.* **26**, 24-es (2007).
- [29] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin, The ball-pivoting algorithm for surface reconstruction, *IEEE Trans. Visualization Comput. Graph.* **5**, 349 (1999).
- [30] Q.-Y. Zhou, J. Park, and V. Koltun, Open3D: A modern library for 3D data processing, *arXiv:1801.09847*.
- [31] T. Möller and B. Trumbore, Fast, minimum storage ray-triangle intersection, *J. Graph. Tools* **2**, 21 (1997).
- [32] R. McConn, Jr., C. Gesh, R. Pagh, R. Rucker, and R. Williams III, Compendium of Material Composition Data for Radiation Transport Modeling, Tech. Rep. PIET-43741-TM-963, PNNL-15870 Rev. 1 (Pacific Northwest National Laboratory, Richland, WA, 2011).
- [33] L. Parra and H. H. Barrett, List-mode likelihood: EM algorithm and image quality estimation demonstrated on 2-D PET, *IEEE Trans. Med. Imaging* **17**, 228 (1998).
- [34] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, CuPy: A NumPy-compatible library for NVIDIA GPU calculations, in *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017)*, Long Beach, CA, USA (learningsys.org, 2017).
- [35] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, Voxblox: Incremental 3D euclidean signed distance fields for on-board MAV planning, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) Vancouver, BC, Canada* (IEEE, Piscataway, NJ, 2017), pp. 1366–1373.