

Lawrence Berkeley National Laboratory

LBL Publications

Title

Correlations Between Panoramic Imagery and Gamma-Ray Background in an Urban Area

Permalink

<https://escholarship.org/uc/item/95p8p33j>

Journal

IEEE Transactions on Nuclear Science, 68(12)

ISSN

0018-9499

Authors

Bandstra, MS

Quiter, BJ

Salathe, M

et al.

Publication Date

2021-12-01

DOI

10.1109/tns.2021.3128472

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial License, available at <https://creativecommons.org/licenses/by-nc/4.0/>

Peer reviewed

Correlations Between Panoramic Imagery and Gamma-Ray Background in an Urban Area

M. S. Bandstra¹, B. J. Quiter², M. Salathe³, K. J. Bilton⁴, J. C. Curtis⁵,
S. Goldenberg⁶, and T. H. Y. Joshi⁷, *Member, IEEE*

Abstract—When searching for radiological sources in an urban area, a vehicle-borne detector system will often measure complex, varying backgrounds primarily from natural gamma-ray sources. Much work has been focused on developing spectral algorithms that retain sensitivity and minimize the false-positive rate even in the presence of such spectral and temporal variability. However, information about the environment surrounding the detector system might also provide useful clues about the expected background, which if incorporated into an algorithm, could improve performance. Recent work has focused on extensive measuring and modeling of urban areas with the goal of understanding how these complex backgrounds arise. This work presents an analysis of panoramic video images and gamma-ray background data collected in Oakland, California, by the radiological multisensor analysis platform (RadMAP) vehicle. Features were extracted from the panoramic images by semantically labeling the images and then convolving the labeled regions with the detector response. A linear model was used to relate the image-derived features to gamma-ray spectral features obtained using nonnegative matrix factorization (NMF) under different regularizations. We find some gamma-ray background features correlate strongly with image-derived features that measure the response-adjusted solid angle subtended by sky and buildings, and we discuss the implications for the development of future, contextually aware detection algorithms.

Index Terms—Environmental radiation effects, gamma-ray background, gamma-ray detectors, gamma-ray spectroscopy, non-negative matrix factorization (NMF), radiation environment characterization, security applications.

I. INTRODUCTION

Vehicle-borne gamma-ray detection systems have been developed and deployed for many years and play a key role in radiological and nuclear security missions, especially the search for sources outside of regulatory

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract DE-AC02-05CH11231, and was supported by the U.S. Department of Energy National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development through the Lawrence Berkeley National Laboratory under Contract DE-AC02-05CH11231.

M. S. Bandstra, B. J. Quiter, M. Salathe, J. C. Curtis, and T. H. Y. Joshi are with the Applied Nuclear Physics Program, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA (e-mail: msbandstra@lbl.gov).

K. J. Bilton is with the Department of Nuclear Engineering, University of California, Berkeley, CA 94720 USA.

S. Goldenberg is with the Computer Science Department, The College of William and Mary, Williamsburg, VA 23185 USA.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNS.2021.3128472>.

control [1]–[7]. These systems can be rapidly deployed and carry large volume detectors, giving them advantages in efficiency and the ability to cover wide areas relative to human-portable systems. However, with their higher efficiencies, these systems also suffer from the complex natural radiological backgrounds often found in urban areas [8], [9], whose spatial and temporal complexities can be exacerbated by the added mobility of vehicle-borne systems and limit the sensitivity of detection algorithms [3], [10], [11].

The natural backgrounds encountered by a vehicle-borne system include the three main terrestrial “KUT” sources (^{40}K , the ^{238}U decay series, and the ^{232}Th decay series), which are found in some quantity in most minerals, soils, asphalt, and building materials. The background also contains the progeny of ^{222}Rn , which, while a major portion of the U-238 series, can also be suspended in the air; and cosmic emission, which takes the form of a power-law continuum and 511 keV line emission. (The decay products of ^{220}Rn from the ^{232}Th series can also escape from the soil and be suspended in air, but due to the very short half-life of ^{220}Rn compared to ^{222}Rn , this decay chain is a negligible contribution to the background.) The reader is directed to [12] for a review of background sources and [13] for thorough measurements and modeling of the backgrounds encountered by a ground-based gamma-ray detector system. Because the compositions of KUT in building materials can vary by orders of magnitude [14] and the sizes of buildings and other structures can also vary widely in urban areas, urban radiological backgrounds can vary significantly, even over distances as small as several meters [8].

In order to improve their sensitivity to sources of interest, many recent detection algorithms focus on capturing the background complexity through analysis of the full gamma-ray spectrum instead of only a portion of the spectrum [15]–[19]. However, these algorithms are not yet close to the Poisson statistical limit [19], presumably due to the temporal variability of the background, which is difficult to compensate for. One way to potentially improve the performance of detection algorithms in urban settings may be to include some nonradiological contextual information. The most extreme hypothetical example would be an algorithm that is able to perfectly predict the Poisson mean of the current background spectrum through the use of contextual information. Spectroscopic algorithms in this case should therefore be able to achieve the Poisson limit of detection sensitivity. However, such a contextual algorithm does not exist, and instead we assert that finding correlations between contextual features and spectroscopic features, even if

they are weak, could provide useful information to algorithms that could allow them to improve their performance, e.g., by adapting their detection thresholds to the current background environment.

Some radiation portal monitor (RPM) research has focused on the use of contextual information to provide cues for detection algorithms. The advantage that portal monitors have is that they are stationary, so any rapid changes in the background must be due to shielding by the vehicles being monitored [20] or radon washout. Sensors to detect the presence of vehicles are typically used to estimate background suppression profiles as vehicles pass through the sensors [20], [21]. In some cases, cameras for monitoring the location of vehicles have been used, although to correlate particular vehicles with sources and not to identify background variations [22], [23]. Rainfall sensors have been used with RPMs to estimate the increases in the background due to radon washout [24].

For mobile detector systems, the problem is more complex; both the movement of the detector system itself and the movement of objects in the scene around the system can change the radiological conditions surrounding the detector. Previous research has shown that by being aware of the city [25] or the region within a city [8] that a mobile system is in can give some idea about the distribution of backgrounds encountered. Other research has focused on “clutter,” i.e., vehicles and people near the detector system, which temporarily shield some of the background emissions and depress measured count rates. Using contextual sensors such as cameras and light detection and ranging (LiDAR) units to detect nearby clutter can be used to identify when an algorithm threshold should be increased so as not to alarm on the changing rates [26]. One recent attempt was made to train a deep neural network on panoramic images to predict the measured spectrum, with some promising results [27].

This work is a part of the modeling urban scenarios and experiments (MUSE) collaboration [28], [29], where the radiological multisensor analysis platform (RadMAP) vehicle [30] was used to explore the connections between gamma-ray backgrounds and various contextual sensors. Previous work using RadMAP within MUSE consisted of analysis of panoramic imagery at a small mock urban area (the Military Operations in Urban Terrain or MOUT facility in Fort Indiantown Gap, PA, USA, or FtIG) [31], for which numerous gamma-ray ground truth measurements had been made [32]. This work investigates applying methods developed for analyzing data collected at the MOUT facility to data from a dense urban area for which no ground truth data exists but which offers realistic complexity and covers a much larger survey area. An earlier version of this analysis was presented in [33], and this version expands upon it by examining more spectral features, quantitatively comparing the results of the correlations between imagery and spectral features, discussing the results in more detail, and discussing future prospects for this type of analysis.

This article will discuss the preparation of the dataset (Section II), the linear model used to guide the analysis (Section III), the extraction of gamma-ray background features




Fig. 1. Cutaway view of the RadMAP vehicle (from [30]). The locations of cameras and the NaI(Tl) array and the relevant dimensions between the two sensor systems have been indicated.

(Section IV), the extraction of features from panoramic imagery (Section V), and then a search for correlations between the two feature sets (Section VI). Finally, the implications for detection algorithms and the development of improved models will be discussed in Section VII.

II. RADMAP DATASET

The evaluation dataset used in this analysis consisted of data from multiple sensors on board the RadMAP vehicle [30]. The RadMAP data offer a unique opportunity to explore correlations between gamma-ray backgrounds and panoramic imagery because of its large (1 m²) NaI(Tl) detector array and two panoramic video cameras, which are shown in Fig. 1. One long, continuous set of data was chosen for this analysis, during which RadMAP traversed much of downtown Oakland, California (Fig. 2). The data were taken on August 18, 2016, from 11:12:42 to 11:52:42 Pacific Daylight Time (PDT) (UTC-7), a total of 40 min. As in [31], images from the two Ladybug3 panoramic cameras were down-sampled from the maximum rate of 15 to 3 Hz so that the corresponding NaI(Tl) spectra had no fewer than approximately one count per bin on average, resulting in 7190 images from each camera being considered. The panoramic images from each camera were fused into a single panoramic image covering nearly the entire scene around the vehicle. The rest of the preparation of the contextual data followed the same procedure described in [31] except for an improvement in the method used to align radiation data with imagery, which is described below.

Because the panoramic cameras were located approximately 4.45 m forward on the vehicle from the NaI(Tl) array, spectral and image data that are coincident in time do not necessarily (and often do not) represent the same location in space, and even a few meters can lead to a significant difference in background [8]. The typical speed of the vehicle varied between 0 and 10 m/s, implying that there is an inherent time delay of at least 0.45 s between the imagery and spectra that must be compensated for if they are to represent measurements of the same location. Previously, this compensation was performed using data from RadMAP’s Inertial Navigation System (INS), specifically the Global Positioning System (GPS) coordinates

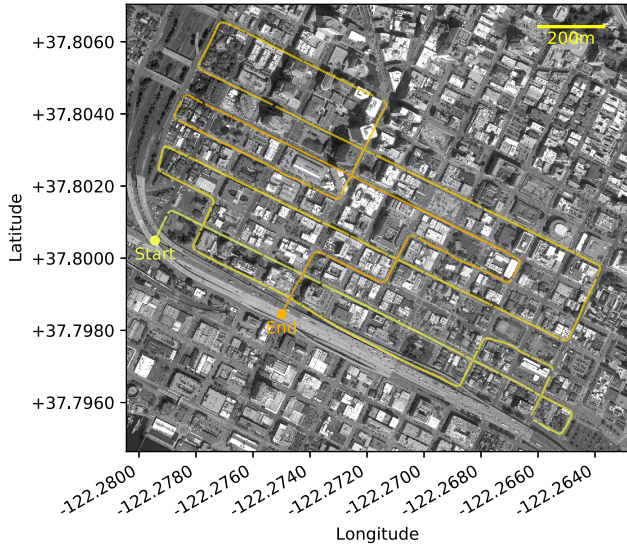


Fig. 2. Path taken by RadMAP during the evaluation dataset. (Map imagery: Google.)

and heading, whereas in this work, the spatial alignment was achieved by using RadMAP’s INS and two LiDAR units to obtain a simultaneous localization and mapping (SLAM) solution using Google Cartographer [34]. A SLAM solution consists of two simultaneously derived objects, the pose of the system (i.e., the 3-D position and orientation) at any timestamp during the measurement and a 3-D model of the environment, typically represented by a point cloud. The SLAM pose, rotated and translated to fit the highest quality GPS points with a least-squares optimization, is shown on the map in Fig. 2.

The spatial alignment was achieved by converting the timestamp of each panoramic image into a corresponding range of timestamps for the list-mode radiation data generated by the NaI(Tl) array. First, the pose solution was used to calculate the distance traveled by the vehicle as a function of time, denoted $s(t)$. Next, each image was assumed to “cover” the time range $(t_0^{\text{img}}, t_1^{\text{img}}) = t^{\text{img}} \pm (\delta t)/2$ where t^{img} is the timestamp of the image and δt is the inverse of the frame rate (1/3 s). The corresponding NaI(Tl) array timestamps were calculated by solving $s(t_0^{\text{det}}) - 4.45 \text{ m} = s(t_0^{\text{img}})$ and the analogous relationship for t_1^{img} . Note that the detector time bins will not necessarily be of equal dwell time.

In order to ensure that the NaI(Tl) measurement was taken before the scene had changed, image-spectra pairs were only accepted if the delay between the image time and the NaI(Tl) time was less than 2 s ($t_0^{\text{det}} - t_0^{\text{img}} \leq 2 \text{ s}$ and $t_1^{\text{det}} - t_1^{\text{img}} \leq 2 \text{ s}$), and if the duration of the NaI(Tl) measurement was less than 0.5 s ($t_1^{\text{det}} - t_0^{\text{det}} \leq 0.5 \text{ s}$). This cut reduced the number of images and spectra from 7190 to 4198, primarily due to times when RadMAP was stopped in traffic, often near intersections.

III. METHODOLOGY

The main analysis of this work will be to relate gamma-ray spectral features with features derived from the panoramic




Fig. 3. Schematic of the geometric definitions used in (1) and (2) relating measurement i to the emission from the surface element subtended by pixel k .

images. By “feature” we mean a scalar value that represents a more complex data structure through some encoding scheme to be determined. “Spectral features” (denoted y) will be one or more scalar values with units of counts per second that are derived from a single gamma-ray spectrum, and “image features” (denoted R) will be one or more scalar values with units of area derived from each image.

Spectral features and gamma-ray flux from the environment are assumed to be related through a simple model. The model is based on three assumptions: 1) all spectral features and gamma-ray emissions are nonnegative; 2) the intensity of a measured spectral feature is a linear combination of the intensity of emission from all visible material in the environment surrounding RadMAP; and 3) material that belongs to the same visual category has identical gamma-ray emission. Assumption 1) must be true because of the physical properties of gamma-ray emission, while assumption 2) is only an approximation to reality because of the effects of scattering between visible objects, shielding of emission from objects that are not visible, and downscatter in the air. The third assumption is the weakest, given the wide ranges of KUT concentrations in soil and building materials [14]; however, it is possible that within a single urban area the variations in KUT might not be that large. Since the model requires that each spectral feature be fully explained by a linear combination of nonnegative fluxes, as a consequence, the model has an intercept of zero. Here, we will explain the model in anticipation of the analysis later on in Section VI.

We start with the detector being surrounded by a series of solid surfaces and the sky, which we will simply model as a fictional solid surface at a distant radius. The panoramic images divide the world around the detector into the small segments subtended by each image pixel, and each image pixel k subtends some solid angle element $\Delta\Omega_k$. During measurement i , the pixel subtends some surface area element Δa_{ik} at distance r_{ik} and angle θ_{ik} from its normal. Fig. 3 shows the geometry we are considering and Table I summarizes the indices that will be used throughout.

We then assume gamma-ray emission of some type j is emitted from surface element Δa_{ik} with photon current ϕ_{ijk} (photons per second per area). An emission type may be comprised of a characteristic distribution of photon energies,




Fig. 4. Schematic of the process of extracting features from the spectra and images and then comparing the features through fitting a linear model with coefficients $\phi_{j\ell}$. The i th spectrum \mathbf{x}_i is a vector of counts in each energy bin, and Δt_i is its time duration. Spectral features y_{ij} have units of counts per second, and image features $R_{ij\ell}$ have units of area, leaving the photon currents $\phi_{j\ell}$ with units of counts per second per area.

TABLE I
DESCRIPTION OF THE INDICES USED IN THIS ARTICLE

Index symbol	Dimension
i	measurement
j	emission type / feature type
k	image pixel
ℓ	image label type

so any energy-dependent quantity will need to be indexed by j . We will also assume that surfaces are solid objects that are multiple photon scatter lengths thick (so the differential photon current in units of photons per second per area per solid angle is proportional to $\cos\theta_{ik}$). Letting \mathcal{A}_{jk} be the effective area (geometric area times efficiency) of the NaI array for emission type j in the direction of image pixel k , assuming the surfaces are in the far-field (i.e., $r_{ik}^2 \gg \Delta a_{ik}$ and $r_{ik}^2 \gg$ the area of the array), and neglecting attenuation and scattering from the air, then the measured count rate of emission type j is obtained by summing over all of the surface elements

$$y_{ij} \approx \sum_{k \in \text{pixels}} \mathcal{A}_{jk} \frac{\Delta a_{ik} \cos\theta_{ik}}{\pi r_{ik}^2} \phi_{ijk} \quad (1)$$

$$\approx \sum_{k \in \text{pixels}} \mathcal{A}_{jk} \frac{\Delta\Omega_k}{\pi} \phi_{ijk}. \quad (2)$$

Notably, the quantity $\Delta a_{ik} \cos\theta_{ik}/r_{ik}^2$ loses any dependence on the geometry and orientation of the surface element to become simply the solid angle element $\Delta\Omega_k$ subtended by each image pixel k . (Note that (1) in our previous work [31] contains an erroneous factor of 2π in the denominator that we have corrected to π here. This factor comes from the normalization for the differential photon current over the half unit sphere, which is $\iint \cos\theta d\Omega = \int_0^{2\pi} \int_0^{\pi/2} \cos\theta \sin\theta d\theta d\phi = \pi$.)

We will identify the quantities y_{ij} as gamma-ray spectral features with units of counts per second. These features are calculated independently of the panoramic imagery from the i th spectrum \mathbf{x}_i , a vector of counts in each energy bin, and its time duration Δt_i . For example, spectral feature extraction could be done using a linear model, e.g., $y_{ij} = \mathbf{U}_j^\top \mathbf{x}_i / \Delta t_i$ for spectrum \mathbf{x}_i and some matrix \mathbf{U}_j .

Letting the semantic label for each pixel of every image be denoted \mathcal{L}_{ik} , and if we assume that the photon currents ϕ_{ijk} are identical for identically labeled pixels across all the measurements (so that ϕ_{ijk} only needs to be indexed by emission type j and label ℓ), then this model further reduces to

$$y_{ij} = \sum_{\ell \in \text{labels}} \left[\sum_{k \in \text{pixels}} \mathcal{A}_{jk} \frac{\Delta\Omega_k}{\pi} \delta_{\mathcal{L}_{ik}, \ell} \right] \phi_{j\ell} \quad (3)$$

$$\equiv \sum_{\ell \in \text{labels}} R_{ij\ell} \phi_{j\ell}. \quad (4)$$

Once a scheme for extracting spectral features y_{ij} and a scheme for labeling the panoramic images have been chosen, the image feature tensor $R_{ij\ell}$, which describes the spectral response to each kind of image label, can be calculated and the linear model solved for the photon currents for each label $\phi_{j\ell}$. Fig. 4 presents a conceptual model of this methodology. The following sections will explain the calculation of y_{ij} (Section IV) and $R_{ij\ell}$ (Section V), and the fits of the linear model will be examined in Section VI.

IV. GAMMA-RAY SPECTRAL FEATURES FROM NMF

There are many choices one could make to extract features from measured gamma-ray background spectra. Simple examples of such features are the gross counts or counts in certain regions of each spectrum, while more complicated examples could use principal component analysis (PCA) [35] or convolutions with wavelets [36]. To preserve the physical basis of the linear model, we are interested in nonnegative features only. Previous work has shown nonnegative matrix factorization (NMF) [37], [38] to be a useful full-spectrum representation for gamma-ray background spectra [19], yielding physically relevant spectral features for vehicle- and aircraft-based systems [31], [39], so we will continue with that approach here.

Three different versions of NMF models will be discussed in more detail in Sections IV-A–IV-C, but first are a few general details on how we prepared our dataset for NMF. The NMF models generated in this section were trained using 3-Hz NaI(Tl) spectra from RadMAP measured in the same area of downtown Oakland as the evaluation dataset, on August 18 and 22, 2016, and part of the training data includes the route in the evaluation dataset. (Note that although there is temporal overlap with the evaluation dataset, the




Fig. 5. Diagram of the NMF decomposition [see (5)] describing the matrix dimensions and the interpretation of the columns of each of the matrices.

training dataset is larger and binned evenly in time.) Three nuisance sources were scrubbed from the training set by coarsely binning the data into seven bins (boundaries at 50, 80, 100, 150, 300, 700, 1200, and 3060 keV) and applying the spectral comparison ratio anomaly detection (SCRAD) method [10], [40] with an exponential weighting parameter of 0.01. All three nuisances found were encountered on 18 August but were outside of the time range of the evaluation dataset. The final training dataset consisted of 18 569 spectra totaling 6190 s.

A. NMF Decomposition With Poisson Loss (NR)

The list-mode gamma-ray event data from the NaI(Tl) array on RadMAP were histogrammed in time and energy. The energy bins were such that the widths were proportional to the square root of energy but broken into two spacing groups: 120 bins from 50 to 3000 keV and 10 bins from 3000 to 4500 keV. The spectra are arranged as the columns of an $m \times n$ matrix \mathbf{X} , with $m = 130$ the number of spectral bins and $n = 18\,569$ the number of time intervals.

A d -component NMF decomposition is a linear decomposition of the spectra of the form

$$\mathbf{X} \approx \mathbf{W}\mathbf{H} \quad (5)$$

where \mathbf{W} is an $m \times d$ matrix whose columns are the spectral components and \mathbf{H} is a $d \times n$ matrix whose rows are the component weights. Fig. 5 shows a diagram of the matrices in (5). Equation (5) is solved by minimizing the negative log likelihood of \mathbf{X} given $\hat{\mathbf{X}} \equiv \mathbf{W}\mathbf{H}$ assuming Poisson statistics

$$-\log L(\mathbf{X}|\mathbf{W}, \mathbf{H}) = \sum \left(\hat{\mathbf{X}} - \mathbf{X} \odot \log(\hat{\mathbf{X}}) + \log \mathbf{X}! \right) \quad (6)$$

where \odot denotes element-wise multiplication, the natural logarithm and factorial are applied element-wise, and the sum is over all matrix elements. To find an NMF solution, the multiplicative update rules from [38], [41] can be used

$$\mathbf{W} \leftarrow \mathbf{W} \odot \left(\frac{\left(\frac{\mathbf{X}}{\mathbf{W}\mathbf{H}} \right) \cdot \mathbf{H}^T}{\mathbf{1}_{m,n} \cdot \mathbf{H}^T} \right) \quad (7)$$

$$\mathbf{H} \leftarrow \mathbf{H} \odot \left(\frac{\mathbf{W}^T \cdot \left(\frac{\mathbf{X}}{\mathbf{W}\mathbf{H}} \right)}{\mathbf{W}^T \cdot \mathbf{1}_{m,n}} \right) \quad (8)$$

where $\mathbf{1}_{m,n}$ is an $m \times n$ matrix of ones. To preserve the normalization of the columns of \mathbf{W} , the following renormalization

is done at each step:

$$\mathbf{D} \equiv \text{diag}(\mathbf{W}^T \cdot \mathbf{1}_m) \quad (9)$$

$$\mathbf{W} \leftarrow \mathbf{W}\mathbf{D}^{-1} \quad (10)$$

where $\mathbf{1}_m$ is a length- m column vector of ones and diag creates a diagonal matrix from a column vector.

A series of NMF models were trained using the multiplicative update rules without any further regularization. Since the multiplicative update rules are only guaranteed to find a local, not necessarily global, optimum [41], the initialization of the model can influence the final result, and in the high-dimensionality cases considered here and in related work, the initialization appears to always affect the result. These models were initialized by setting the components to the mean spectral shape, with small random numbers between 0 and 10^{-6} added to break the degeneracy between them. In addition, since NMF models have no preferred order of the components, for ease of comparison between models, the components were sorted in order of increasing variance of their weights, i.e., by the row-wise variance of \mathbf{H} . This same NMF approach was taken in [39].

The models that result from this procedure we will denote NR- d (for *no regularization*, with d components). The first column of Fig. 6 shows the results of these unregularized models for $d = 2$ to 4 components when allowed to converge until the difference in $-(\log L)/n$ between subsequent iterations is less than the arbitrary level of 10^{-9} . Of note is that for all three models, exactly one component contains the cosmic continuum above 3 MeV, and each of these components also uniquely displays a slight 511 keV line, which is expected from cosmic emission. For all models, this component is always the component with the lowest variance of its weights (i.e., component 0). Also worth noting is that all of the components have different shapes for the low energy continuum “roll off” around 100 keV. For the two- and three-component models, all of the components display all of the prominent lines from the KUT background sources, whereas for NR-4 they do not—e.g., component 0 lacks prominent ^{238}U series lines, component 2 lacks the prominent ^{232}Th series line at 2614 keV, and component 1 lacks the ^{40}K line at 1460 keV. Additionally, particularly for NR-4, different ratios of the KUT background sources are clearly seen, with higher ^{40}K in component 3 and higher ^{238}U and ^{232}Th series in component 1. These observations imply that NMF is able

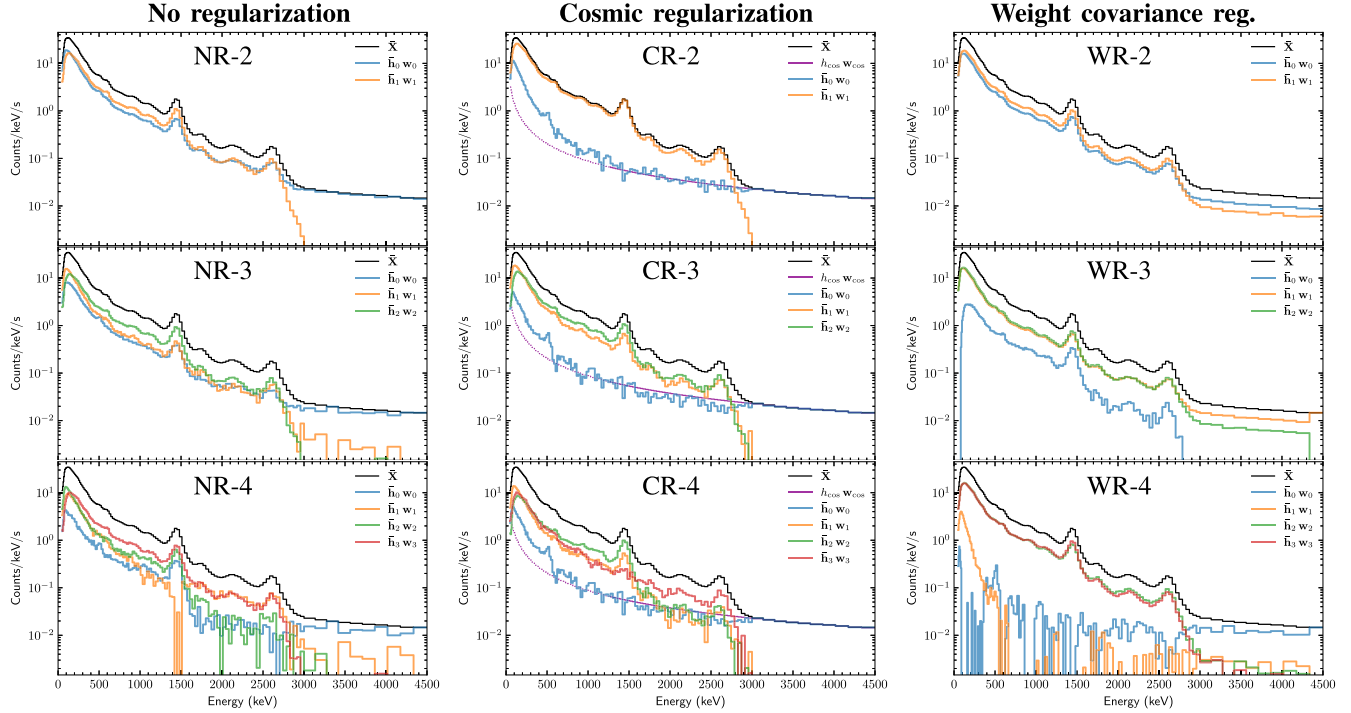


Fig. 6. Spectral components found by the various NMF approaches explored for two, three, and four components. In each plot, the average measured spectrum in the training dataset (black) is compared to the NMF components, each scaled by their mean weight. In addition, for the CR- d models, the power law used in the regularization function is shown.

to capture spectral features that arise from physics, but the various spectral features are not necessarily correlated across the different decompositions. For example, the high energy cosmic continuum is present in component 0 for all the NR models, but the component that has the highest mean count rate at 100 keV is different for each model (components 0, 1, and 2, respectively). Evidently, the dynamic range of KUT backgrounds encountered is not enough to more cleanly separate the KUT spectra from each other.

B. Cosmic Component Regularization

On its own, NMF does not contain any physics; it is a mathematical method for finding structure in the training data. However, physical priors and constraints, if expected to have relevance to the model, can be encoded in the training process in the form of regularization functions. Here we will present one such physical prior—namely, that one of the components should represent cosmic ray-induced gamma-ray emission, and the remainder only terrestrial emission.

Attempting to isolate a cosmic component may be worthwhile for the following reasons. First, we know that a cosmic component is present in the data because any detector exposed to the atmosphere will measure some background due to cosmic rays. Second, the shape of the cosmic emission can be partially estimated from the region above 3 MeV because there are negligible contributions in that spectral region from terrestrial emission. Third, repeatedly training with randomly initialized NMF models usually results in a model where the spectrum above 3 MeV is nearly entirely contained in only one

NMF component (e.g., the NR models in Fig. 6 all display this behavior).

Regularization to isolate the KUT components from each other was also attempted, but these results were less conclusive than for the cosmic component. Profiles for each of the KUT emission types were derived from simulations and used in a similar manner to what will be described for the cosmic regularization, however, these attempts have led to nonphysical shapes at low energies, e.g., components where the counts in the spectrum below 300 keV decrease with decreasing energy more rapidly than expected. Another issue with KUT regularization is that the training dataset itself appears to comprise an area with a limited range in the ratios of K, U, and T, so there is not enough data to support cleanly separating those background sources from each other.

Here we will initialize an NMF component meant to represent cosmic emission and describe the regularization and final results from training such models. The models that result from this treatment will be denoted CR- d .

1) *Initialization*: The cosmic component has been observed to consist of a continuum described by a power law with an index of ≈ 1.3 and 511 keV emission from pair production [13]. The power-law index can be estimated from the spectrum above 3 MeV because energy deposition events above the ^{208}Tl line at 2614 keV are almost entirely due to cosmic rays and cosmic-induced gamma rays. A weighted least-squares fit of the data in the ten spectral bins above 3 MeV was performed using a power-law model, resulting in a power-law index of 1.22. Although we do not expect the entire cosmic spectrum to follow this power-law shape




Fig. 7. Average background spectrum in the training data shown with the provisional cosmic and terrestrial components used to initialize the NMF model. The solid line denotes the portion of the cosmic power law used for the regularization function f_1 .

because changing detector efficiencies will cause a “roll off” at the lowest energies and we have not yet included the expected 511 keV emission, we nevertheless extrapolated this power-law fit to all spectral bins and normalized it to unity to generate the provisional cosmic component \mathbf{w}_{cos} . The first column of \mathbf{W} was initialized to \mathbf{w}_{cos} .

For the sake of initialization, an ansatz was made that the count rate of the cosmic component be approximately constant for all measurements. The constant weight h_{cos} was calculated as the average number of total counts above 3 MeV in \mathbf{X} divided by the sum of the portion of \mathbf{w}_{cos} above 3 MeV. The entire first row of \mathbf{H} was initialized to h_{cos} . The provisional cosmic spectrum \mathbf{w}_{cos} times the average cosmic counts h_{cos} is compared to the average spectrum in Fig. 7.

After the first column of \mathbf{W} was initialized to \mathbf{w}_{cos} , the following procedure was followed to initialize the remaining $d - 1$ components. First, the noncosmic spectrum (assumed to be largely terrestrial) was estimated

$$\mathbf{X}_{\text{terr}} = \mathbf{X} - (\mathbf{w}_{\text{cos}} \otimes \mathbf{1}_n)h_{\text{cos}} \quad (11)$$

where \otimes denotes the outer product. Any negative values encountered were clipped to zero. The terrestrial rates \mathbf{h}_{terr} were calculated as the sum over the columns of \mathbf{X}_{terr} , while the terrestrial component \mathbf{w}_{terr} was calculated by summing \mathbf{X}_{terr} over its rows and dividing by the sum of \mathbf{h}_{terr} . Fig. 7 shows the resulting terrestrial component \mathbf{w}_{terr} times the average terrestrial counts $\bar{\mathbf{h}}_{\text{terr}}$ compared to the average spectrum.

For all $d \geq 2$, the remaining $d - 1$ columns of \mathbf{W} were initialized with \mathbf{w}_{terr} , and the remaining $d - 1$ rows of \mathbf{H} were initialized with $\mathbf{h}_{\text{terr}}/(d - 1)$. Small random numbers between 0 and 10^{-6} were added to these rows and columns to avoid degeneracy between these components.

2) *Cosmic Regularization Functions*: With the provisional shapes of the cosmic and terrestrial components now estimated, we would like regularization functions that preserve the power-law shape of the cosmic component during training and keep the majority of events above 3 MeV in that component. Two regularization functions were added to the loss function

$$\Lambda(\mathbf{W}, \mathbf{H}) = -\log L(\mathbf{X}|\mathbf{W}, \mathbf{H}) + \alpha_1 n f_1(\mathbf{W}) + \alpha_2 n f_2(\mathbf{W}) \quad (12)$$

where α_1 and α_2 are dimensionless regularization parameters, n is used to account for the number of training spectra, and f_1 and f_2 are the additional penalty functions.

The f_1 regularization term is used to compare the shape of the first column of \mathbf{W} with \mathbf{w}_{cos} . We chose to use the symmetric Kullback–Leibler (KL) divergence [42] since it is a natural metric for comparing the shape of two normalized, nonnegative distributions. For example, this divergence has been used in the analysis of hyperspectral data to match NMF components with laboratory-measured spectra [43]. The symmetric KL divergence between normalized vectors \mathbf{p} and \mathbf{q} takes the form

$$D_S(\mathbf{p}, \mathbf{q}) = \sum_i (p_i - q_i) \log\left(\frac{p_i}{q_i}\right). \quad (13)$$

Since we desire that the cosmic component maintain its power-law shape at high energies, but we are uncertain what the shape will be at low energies (in fact, the NR models indicate we should anticipate a 511 keV peak feature), we only apply the penalty to the components of the first column of \mathbf{W} that are above 1250 keV. The exact form of f_1 and the positive and negative terms in its gradient (∇_1^+ and $-\nabla_1^-$) are given in the Appendix.

The f_2 regularization term is used to suppress the contributions above 3 MeV by any components except for the first column of \mathbf{W} . This regularization is needed as pressure to keep the cosmic contribution entirely in the first column of \mathbf{W} , since f_1 on its own does not guarantee that desired property. For f_2 we chose the sum of \mathbf{W} over the nominally terrestrial rows and spectral bins above 3 MeV. The exact form of f_2 and its gradient (∇_2^+) are given in the Appendix.

In the presence of regularization terms, the multiplicative update rules for NMF can be modified to include the gradients of the regularization terms (e.g., [44]). The multiplicative update rule for \mathbf{W} [see (7)] is changed to

$$\mathbf{W} \leftarrow \mathbf{W} \odot \left(\frac{\left(\frac{\mathbf{X}}{\mathbf{W}\mathbf{H}}\right) \cdot \mathbf{H}^T + \alpha_1 n \nabla_1^-}{\mathbf{1}_{m,n} \cdot \mathbf{H}^T + \alpha_1 n \nabla_1^+ + \alpha_2 n \nabla_2^+} \right) \quad (14)$$

where ∇^+ are the positive and ∇^- the negative parts of the gradients (see appendix). This change means that the NMF solution continues to approximately follow the negative gradient of Λ while maintaining nonnegativity, however, values of α_1 and α_2 that are too large can lead to an overprioritization of the regularization functions relative to the Poisson likelihood. To keep the regularizations from having an undue influence early on during the training, we delayed the application of the regularization terms until after 500 iterations of the multiplicative update rules so that gradients of the Poisson loss were first able to stabilize. We chose α values by starting with small values and increasing them until there was a small but noticeable effect on the shapes of the components relative to their unregularized counterparts after 10000 iterations. We found that $\alpha_1 = 10^{-2}$ and $\alpha_2 = 10^{+1}$ were suitable to meet this goal.

3) *Cosmic Component Training Results*: The final results of cosmic component initialization and regularization are shown in the middle column of Fig. 6 as CR- d for $d = 2$ to 4. These

models show that component 0 is consistent in shape even as the number of NMF components is increased. Even though the regularization is only applied above 1250 keV, the shape of component 0 below that energy is remarkably similar across the models, and all components show the expected 511 keV emission line from atmospheric positrons, which was not an engineered feature.

Also of note is that, similar to the NR models, as the number of components is increased, different KUT ratios are captured by the components, although now the KUT variations are all relegated to the noncosmic components. For example, in CR-4, component 3 has a high amount of ^{238}U and ^{232}Th series lines relative to the other components. Some of the non-cosmic components even resemble the NR components, such as CR-3 components 1 and 2, which look similar to NR-3 components 1 and 2, respectively.

C. Maximizing the Covariance of the NMF Weights (WR)

Finally, the third class of NMF models, which we shall refer to as weight covariance regularization (WR- d), was generated. Rather than using physics to inform the spectral shape (\mathbf{W}), this regularization applies a desired mathematical property to the temporal evolution of the weights (\mathbf{H}). For example, if we assume that the features arise from spatially distinct sources of radiation in the environment (e.g., soil versus buildings), we might expect them to evolve somewhat independently of one another over time.

To put this property in mathematical language, we calculated the row-wise mean of the weights \mathbf{H} , which is $(1/n)\mathbf{H}\mathbf{1}_{n,n}$, and then the row-wise covariance of \mathbf{H}

$$\text{var}[\mathbf{H}] = \frac{1}{n} \left(\mathbf{H} - \frac{1}{n} \mathbf{H}\mathbf{1}_{n,n} \right) \left(\mathbf{H} - \frac{1}{n} \mathbf{H}\mathbf{1}_{n,n} \right)^\top \quad (15)$$

$$= \frac{1}{n} \mathbf{H} \mathbf{C} \mathbf{C}^\top \mathbf{H}^\top \quad (16)$$

$$= \frac{1}{n} \mathbf{H} \mathbf{C} \mathbf{H}^\top \quad (17)$$

where $\mathbf{C} = \mathbf{I} - (1/n)\mathbf{1}_{n,n}$ and we have used the fact that \mathbf{C} is both symmetric and idempotent, meaning $\mathbf{C}\mathbf{C}^\top = \mathbf{C}^2 = \mathbf{C}$.

We want $\text{var}[\mathbf{H}]$ to be as diagonal as possible since that would mean the rows of \mathbf{H} vary independently. There are multiple ways to apply regularizations to achieve this goal. Following a similar constraint in local NMF (LNMF) [45], the trace of the covariance matrix was maximized. This choice allows the regularization function and its gradient (calculated symbolically using [46]) to take the forms

$$f_3(\mathbf{H}) = -\frac{1}{n} \text{tr}(\mathbf{H} \mathbf{C} \mathbf{H}^\top) \quad (18)$$

$$\frac{\partial f_3}{\partial \mathbf{H}} = -\frac{2}{n} \mathbf{H} \mathbf{C} \quad (19)$$

and the positive and negative parts of the gradient become $\nabla_3^+ = (2/n^2)\mathbf{H}\mathbf{1}_{n,n}$ and $\nabla_3^- = (2/n)\mathbf{H}$. As before, the NMF models were trained using a modified multiplicative update rule (the analogous version of (14) for \mathbf{H} , including weighting f_3 with a factor of n). This regularization was applied to the same training data with a coefficient $\alpha_3 = 10^{-7}$, arrived at using the same heuristic as previously. Once again, the

regularization was not applied until after 500 iterations of the multiplicative update rules.

The results of applying the WR to the training data for 2 to 4 components are shown in Fig. 6. Compared with the NR and component regularization (CR) models, the cosmic continuum above 3 MeV is shared more between multiple components, except for WR-4 where component 0 contains most of it. The two components of WR-2 look more similar to each other than those of the other two-component models (these components will be examined further in Section VI-C). Component 0 of WR-3 seems to contain mostly KUT emission, while the other two components of WR-3 are similar mixtures of KUT but differ in the low energy shape. Finally, WR-4 seems to break the spectra into cosmic emission (component 0), a sharp low energy continuum that resembles skyshine (i.e., emission downscattered in the air from distant sources, which has its largest contributions below 200 keV [13]) in component 1, and two similar KUT components (2 and 3) that seem to differ mostly in the relative amount of ^{40}K .

D. Extracting Spectral Features Using the NMF Models

Finally, all the NMF models \mathbf{W} ($\mathbf{W}_{\text{NR}-2}$, etc.) obtained in Sections IV-A–IV-C using the training data were applied to the evaluation dataset from August 18, 2016, to generate new NMF decompositions (e.g., $\mathbf{H}_{\text{NR}-2}$). These decompositions were performed by holding each model's component matrix \mathbf{W} fixed and minimizing (6) to obtain the weight matrix \mathbf{H} , i.e., by performing repeated applications of the multiplicative update rule for \mathbf{H} [see (8)]. The various regularization functions were only used during the respective training processes and not during this step. The final result for a d -component NMF model is that the $m \times n$ matrix of spectra ($m = 130$ and $n = 4198$) is reduced to the $d \times n$ matrix of weights \mathbf{H} , or in other words that the spectral data are reduced to d feature vectors of length n .

In addition, the NMF weights have measurement uncertainties due to Poisson statistics that will be needed to properly weight the different models later when being fit to the image features. To estimate the uncertainty of each element of \mathbf{H} , we use the Fisher information in the following manner. For each measurement \mathbf{x}_i (the i th column of \mathbf{X}), we will call the d corresponding best-fit weights \mathbf{h}_i , i.e., the i th column of \mathbf{H} . Since \mathbf{h}_i was determined using maximum likelihood, which is to say

$$\mathbf{h}_i = \arg \max_{\mathbf{h}} \log L(\mathbf{x}_i | \mathbf{W}\mathbf{h}) \quad (20)$$

then the covariance of \mathbf{h}_i can be approximated using the inverse of the observed Fisher information matrix

$$\mathbf{F}_i \equiv - \left. \frac{\partial^2 \log L}{\partial \mathbf{h} \partial \mathbf{h}} \right|_{\mathbf{h}_i} \quad (21)$$

$$= \mathbf{W}^\top \text{diag} \left(\frac{\mathbf{x}_i}{(\mathbf{W}\mathbf{h}_i)^2} \right) \mathbf{W}. \quad (22)$$

So, the covariance of \mathbf{h}_i is approximately

$$\text{var}[\mathbf{h}_i] \approx \mathbf{F}_i^{-1} \quad (23)$$




Fig. 8. Creation of semantically labeled panoramic images. The original stitched panoramic image (top) is labeled using the DeepLabv3+ model (middle). The classes that can be seen in this image are *sky* (light blue), *building* (brown), *road* (dark gray), *sidewalk* (light gray), *vegetation* (green), *car* (yellow), *person* (magenta), *pole* (purple), and *traffic sign* (light green). The foreground region is masked out in the middle image. The bottom image is the result of using the vehicle pose to transform nearby labeled images to estimate the missing foreground region and thus obtain a label for all 4π steradians around the vehicle.

and for simplicity the standard deviation of each element of \mathbf{h} was approximated as the square root of the diagonal elements

$$\sigma_{\mathbf{h}_i} \approx \sqrt{\text{diag}(\mathbf{F}_i^{-1})}. \quad (24)$$

Then call $\Sigma_{\mathbf{H}}$ the matrix with the same dimension as \mathbf{H} where column i is $\sigma_{\mathbf{h}_i}$. These uncertainties are used in Section VI.

V. FEATURES FROM PANORAMIC IMAGERY

In [31], the panoramic images were labeled according to known visual classes at the FtIG MOUT facility, such as *asphalt*, *concrete*, *gravel*, *red building*, *brown building*, etc. This class selection was motivated by the small variety of visual classes at the facility as well as the extensive ground truth measurements of those materials, which revealed relative uniformity within several of the visual classes [32].

For this dataset, since no ground truth exists but also because there are many possible visual classes, we decided to use other tools to label the imagery. To do so, we took advantage of recent advances in machine learning for

TABLE II
REDUCED SET OF EIGHT CLASSES MADE BY COMBINING
DEEPLABV3+ IMAGE CLASSES

Reduced classes	DeepLabv3+ classes
road	road
sidewalk	sidewalk
building	building, wall
vegetation	vegetation
terrain	terrain
sky	sky
vehicles and people	car, truck, bus, train, motorcycle, bicycle, person, rider
other	fence, pole, traffic light, traffic sign

the semantic segmentation of urban scenes by using a DeepLabv3+ model [47] trained on the Cityscapes dataset [48]. The model applies 19 separate labels: flat horizontal regions (*road*, *sidewalk*), humans (*person*, *rider*), vehicles (*car*, *truck*, *bus*, *train*, *motorcycle*, *bicycle*), vertical structures (*building*, *wall*, *fence*), small structures (*pole*, *traffic light*, *traffic sign*), nature (*vegetation*, *terrain*), and sky [48]. The results of the segmentation of one RadMAP panoramic image from the Oakland dataset are shown in Fig. 8. Although the model was trained on standard projection images, the results show that the model can still correctly identify many features of the panoramic images, so no retraining of the network was performed (although ideally such work should be performed in the future). The accuracy of the model observed elsewhere was 82.1% when measured as pixel intersection-over-union and averaged over all the classes in the evaluation [47]. The misidentifications present in Fig. 8 are fairly typical of the quality of labeling observed throughout the present work (e.g., sometimes clouds are mislabeled as *building* or *vegetation*, which may be due to the panoramic projection). As in [31], labeled images close in time and the relative motion of RadMAP were used to fill in the masked-out foreground region at the bottoms of all the images. Because some classes were rare, the number of classes was reduced from 19 to eight as summarized in Table II.

With labeled images providing \mathcal{L}_{ik} , the image feature tensor $R_{ij\ell}$ can be calculated using (4). As in [31], when calculating these features, the solid angle element of each pixel $\Delta\Omega_k$ was increased in the bottom half of the image to account for the vertical displacement of the cameras, which were appreciably higher than the array. In addition, the effective area \mathcal{A}_{jk} was derived from Monte Carlo simulations as described in [31], where the shapes of the NMF components were used to form a weighted sum of the effective areas calculated at a variety of discrete energies.

VI. CORRELATIONS BETWEEN IMAGE AND SPECTRAL FEATURES

This section will explain how weighted nonnegative least squares was used to fit the linear model in (4) to the evaluation

dataset, which figures of merit were used to analyze the results, and what those results might indicate about the connections between the spectra and contemporaneous imagery.

A. Fitting NMF Weights With the Linear Model

The linear model in (4) was fit in the following manner. For each NMF model, the linear model was fit to each of the rows of \mathbf{H} independently of one another. Specifically, for a given NMF model \mathbf{WH} (fit to the evaluation dataset according to Section IV-D), the following steps were performed for each feature j :

- 1) For feature j , choose the j th column of \mathbf{W} (\mathbf{w}_j) and the j th row of \mathbf{H} (\mathbf{h}_j).
- 2) Estimate the effective area \mathcal{A}_{jk} for spectral shape \mathbf{w}_j and the image feature tensor $R_{ij\ell}$ as described earlier. Since j is fixed, call the resulting 2-D matrix \mathbf{R}_j .
- 3) With j fixed, the 2-D spectral feature tensor y_{ij} in (4) is a 1-D feature vector \mathbf{y}_j . Identify the count rates $\mathbf{h}_j/\Delta t = \mathbf{y}_j$ as this feature vector.
- 4) Using Section IV-D, estimate the uncertainty of \mathbf{y}_j and call it $\sigma_j = \sigma_{\mathbf{h}_j}/\Delta t$, where $\sigma_{\mathbf{h}_j}$ is the j th column of $\Sigma_{\mathbf{H}}$.
- 5) Fit the linear model $\mathbf{y}_j = \mathbf{R}_j\phi_j$ using nonnegative least squares weighted by σ_j , finally obtaining the photon currents ϕ_j and the fit $\hat{\mathbf{y}}_j$.

The weighted nonnegative least-squares fit was performed using Lasso (least absolute shrinkage and selection operator) regression, which is implemented in `scikit-learn` [49]. Lasso performs nonnegative linear regression while trying to minimize the number of nonzero coefficients in the solution, i.e., enforcing sparsity on the solution. The sparsity parameter was empirically set to 0.1, which had the effect of forcing the fits to use approximately four out of the eight total image features.

B. Results of Linear Model Fits

In order to assess the quality and physical relevance of each fit, multiple statistics were considered.

First, in order to understand the prominence of individual image features within each fit, the fraction of the fit from each image feature was calculated. For measurement i , NMF feature j , and image feature ℓ , this fraction is

$$f_{j\ell} = \frac{\sum_i R_{ij\ell}\phi_{j\ell}}{\sum_i \hat{y}_{ij}}. \quad (25)$$

In addition to the fraction made up by each image feature, three goodness-of-fit metrics were considered to assess how well the fits represent the data. The reduced chi-squared statistic (χ_v^2) was chosen as a standard statistical measure of goodness of fit, however, a model with large weight uncertainties ($\sigma_{\mathbf{h}}$) can lead to an acceptable χ_v^2 value but an uninformative model. To also provide a sense of how well the model follows trends in the data, the Pearson's correlation coefficient was calculated between the image features and a smoothed version of the data. Finally, the usefulness of the model in predicting the count rates of each feature using

the imagery was calculated using the root mean squared deviation (RMSD).

The reduced chi-squared statistic is

$$(\chi_v^2)_j = \frac{1}{n} \sum_i \frac{(y_{ij} - \hat{y}_{ij})^2}{\sigma_{ij}^2} \quad (26)$$

and has an expected value of unity for a model that perfectly describes the data. Technically, the degrees of freedom are not n but could be as few as $n - 8$ depending on the sparsity of the linear model, but since n is so large this difference was ignored.

The second metric is Pearson's correlation coefficient r , to measure how closely the model count rate correlates with the count rate from the NMF weights

$$r_j = \frac{\sum_i (\tilde{y}_{ij} - \mu[\tilde{\mathbf{y}}_j])(\hat{y}_{ij} - \mu[\hat{\mathbf{y}}_j])}{\sqrt{\sum_i (\tilde{y}_{ij} - \mu[\tilde{\mathbf{y}}_j])^2} \sqrt{\sum_i (\hat{y}_{ij} - \mu[\hat{\mathbf{y}}_j])^2}} \quad (27)$$

where $\tilde{\mathbf{y}}_j$ is obtained by smoothing \mathbf{y}_j using a boxcar kernel of width 5 and $\mu[\cdot]$ is the sample mean over all n measurements. This filtering, which smooths the data over a timescale of at least 1.67 s, was done to reduce the influence of statistical fluctuations on r . The values r can take range from -1 (perfectly anticorrelated) to $+1$ (perfectly correlated).

The third metric is the RMSD, for a measure of each model's prediction error

$$\text{RMSD}_j = \sqrt{\frac{1}{n} \sum_i (\tilde{y}_{ij} - \hat{y}_{ij})^2} \quad (28)$$

where $\tilde{\mathbf{y}}_j$ is once again the smoothed \mathbf{y}_j .

The results of fitting all of the NMF models are shown in Fig. 9, which displays the fraction of the fits made up by each of the image features as well as the goodness-of-fit metrics. It is immediately apparent that the vast majority of the fits are not objectively good (i.e., the null hypothesis can be rejected with high confidence), since all but five of the χ_v^2 values are too large given the number of degrees of freedom, which, at approximately 4198, would require $\chi_v^2 > 1.05$ to reject the null hypothesis at the 10^{-2} level). In addition, the RMSD values, which range from 27 to over 2200 counts per second, indicate that the fits frequently under-predict and/or over-predict the spectral features. (For comparison, the mean gross count rate is 10 056 counts per second.) Even the five models with acceptable χ_v^2 values have some of the largest RMSD values. This comparison of RMSD and χ_v^2 values reveals that the NMF-WR models, probably as a side effect of the regularization itself, are creating inflated values of $\sigma_{\mathbf{h}}$ and thus appear to be good fits in a χ_v^2 sense but in fact have a high prediction error (RMSD value). From these metrics, we can conclude that the linear model approach of relating spectral and image features is not statistically correct, and a more complex model or different features would be needed for greater goodness of fit and lower prediction error.

On the other hand, the correlation coefficient r reveals that all models have positive correlations between the image features and the NMF-derived spectral features, ranging from weak (0.08) to somewhat strong (0.76). Of particular interest

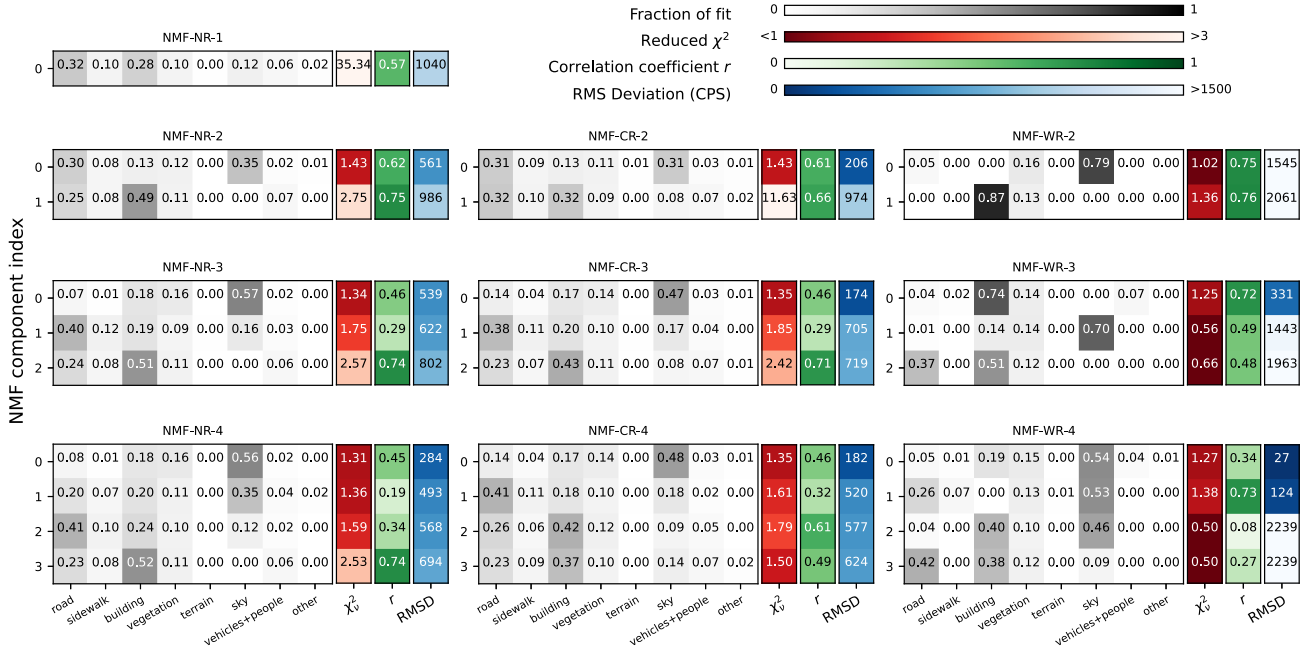


Fig. 9. Results of Lasso fits the evaluation RadMAP dataset for all of the NMF models considered. Shown in grayscale are the fractions of each fit made up by the image class, along with the goodness of fit as measured by the reduced χ^2 statistic (red), correlation coefficient (green), and root-mean-squared deviation (blue). Darker shades indicate better performance or stronger correlation, whichever is appropriate. The goodness-of-fit metrics represent how well the temporal evolution of the spectral features matches the linear model.

are those models indicative of strong correlations, and especially those with both the largest r values and the largest fractions of the fits made up of a single image feature $f_{j\ell}$. Though the overall fit may be statistically or predictively poor, if a large fraction of the fit is made up of a single image feature, that is an indication that the spectral data still correlate with that image feature in a meaningful sense. For example, each model except $d = 1$ has at least one component with a fit comprised of a significant fraction ($>30\%$) from *sky*. The component with the largest fraction from *sky* is always the component that contains the largest amount of emission over 3 MeV, as can be seen in Fig. 6, and it also tends to be the lowest-variance component (0). This finding comports with the prediction that atmospheric emission should have some connection to the *sky* feature, and that of the natural background sources, cosmic emission has the most stable flux and composition.

Further examining the fractions from *sky*, we can see that cosmic regularization, which was designed to produce an NMF component 0 that has a shape consistent with cosmic emission, does not enhance this fraction in the CR- d models when compared to the other models—component 0 in the equivalent NR- d and WR- d models counterintuitively have a similar or greater fraction from *sky*. The process of regularization, however, does force the component's spectral shape to be consistent across different values of d .

Besides *sky*, other features that have significant fractions are *road* and *building*. Of particular note is that *building* has some fraction values near and over 50% (e.g., NR-3 component 2 and WR-2 component 1). The components with large fractions of these image features tend to have

low fractions from *sky*, suggesting that these correlations are due to a difference between atmospheric and terrestrial emission.

The models exhibiting the largest fractions for both *sky* and *building* are the two- and three-component NMF models with WR (WR-2 and WR-3). We will examine WR-2 in further detail, since each of its two components has a large fraction associated with a single image feature. For this model, both NMF components have large fractions explained by image features—component 0 has a 79% fraction from *sky*, and component 1 has an 87% fraction from *building*. These fits and their breakdown by image feature are shown in Fig. 10, with a comparison to fit the simplest model, NMF-NR-1, which is just the gross count rate. Although there are significant departures between the data and the fits (which is the subject of Section VI-C), the general trends in both are strong.

C. Examination of WR-2

The WR-2 model will be further examined here, both for its spectral component shapes and when the linear model fit is poorest.

1) *Spectral Component Shapes*: The two components of this NMF model are shown in Fig. 6. Fig. 11, which shows the two components are rescaled to match in the region around 1460 keV, reveals that they have a similar shape between 500 and 2800 keV. The main contrasts between the components are that component 0 has relatively increased rates above 3 MeV, below 200 keV, and also near 511 keV. To draw these differences into sharper contrast, the two scaled components are subtracted from each other and shown in Fig. 11. The

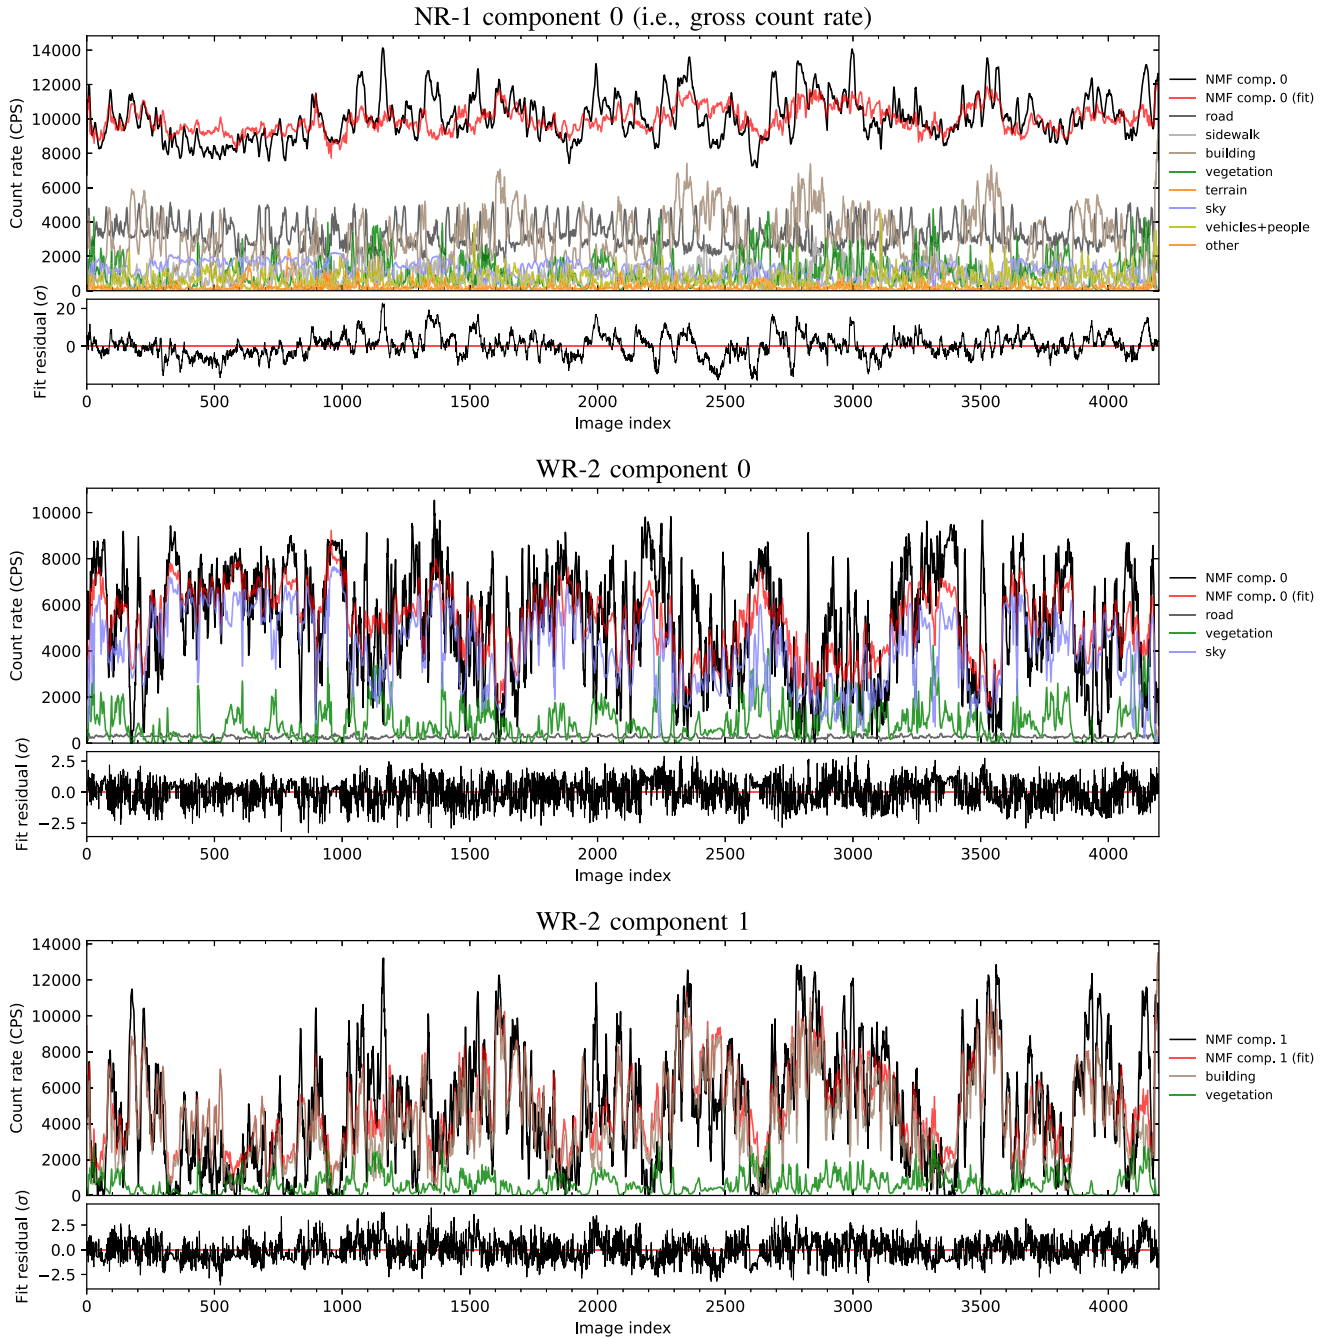


Fig. 10. Results of Lasso fits NMF component weights from model WR-2 (middle and bottom). For comparison, the Lasso fits the single component of NR-1, which are the gross counts, is also shown (top). The spectral features from NMF, shown in black, are smoothed with a boxcar filter of width 5 to improve clarity, but the residuals are not smoothed. Image features with coefficients of zero (due to Lasso’s sparsity regularization) are not shown.

shape of this difference spectrum is strikingly similar to the cosmic-regularized component (component 0) of the CR- d models in Fig. 6. This finding suggests that the difference between the two WR-2 components is that component 0 contains an additional amount of cosmic and skyshine emission in addition to the terrestrial, skyshine, and cosmic emission shared by both components. Therefore, component 0 may represent emission that has relatively more “distant” emission, and component 1 may represent emission from relatively more “nearby” emission, although both include a mixture of all emission types. Though different kinds of regularizations have

been applied to obtain many different NMF models, this cosmic and skyshine spectrum and its association with *sky* consistently emerge from this analysis.

2) *Discrepancies Between Fits and NMF Weights*: Examining the weights of WR-2, it is clear from Fig. 10 that there are certain regions where the fit and the data significantly diverge. The largest discrepancies were examined to see if there were consistent explanations for why the models may have failed in a manner consistent with their assumed (more distant versus more nearby) origins. The twelve regions containing the largest RMSD values were identified and the index of the




Fig. 11. Difference between the two components of NMF-WR-2, after scaling the components to be the same in the 1460-keV region. The resulting spectrum resembles the cosmic components of the NMF-CR models.




Fig. 12. Panoramic images at the times when the linear model under-predicts the spectral features from component 0 of WR-2. Intersections (top) and plazas or open fields (middle—the orange color is the *terrain* label) are common scenarios that occur when the model under-predicts. One large discrepancy was seen when a large tree obscured the sky behind it (bottom). DeepLabv3+ labeling is overlaid as transparency.

maximum was recorded. In many, but not all, cases, a plausible explanation could be found for the differences.

For component 0, the twelve largest discrepancies where the fit under-predicted the NMF weights were found and examined (these discrepancies are in raw count rate, not σ residual). In nine of these cases, the panoramic images showed that the vehicle was either in the center of an intersection or adjacent to a large open lot or plaza (see top and middle images in Fig. 12 for examples). In these scenarios, buildings




Fig. 13. Panoramic images at the times when the linear model under-predicts the spectral features from component 1 of WR-2. Trees (top) or large vehicles (middle) can visually obscure emission from buildings, as can errors in the labeling of buildings (bottom). DeepLabv3+ labeling is overlaid as transparency.

are farther away from the sides of the vehicle, which breaks up the typical “urban canyon” scenario found in the dataset, and so more distant, downscattered terrestrial emission such as skyshine might be expected due to greater exposure to the distant roads and terrain. At the same time, the portion of the image subtended by *sky* may increase somewhat when in these locations, but the increase might not fully capture what is actually a 3-D phenomenon that should include, e.g., scattering in the air from distant terrestrial emission (see [50] for the development of a 3-D emission model). In one of the other cases, a large tree obscured the open sky behind it, thus leading to a much lower estimated contribution from *sky* than if the tree were not there, presumably without commensurate attenuation by the tree (bottom images in Fig. 12). In two other cases, there were no clear explanations.

The 12 largest discrepancies where the fit over-predicted the component 0 weights were also examined. In six of these cases, the vehicle was directly adjacent to a large building on its right side, which is the side of the NaI array with a larger effective area. The other six cases offered no obvious common features, and, in general, it is not known what the main reasons for these over-predictions are.

For component 1, the same inspections were performed of the top 12 discrepancies caused by model underprediction.

There were five cases where a large foreground object obscured a building behind it (four were trees and one was a large truck). Two such examples are shown in the top and middle images of Fig. 13. It is likely that those particular objects do not fully attenuate the emission from the buildings that is passing through them, and since the model coefficients have low photon currents from *vegetation* and *vehicles and people*, there is a net deficit in the model prediction. These cases thus reveal another weakness in using 2-D imagery alone instead of a 3-D model (see [50]) that could account for material present behind the foreground objects. Two other cases involved encounters with two different facades of the same building, a possible indication that that particular building likely comprises material with higher concentrations of radioactive material than other buildings in the area. The last two cases involved the DeepLabv3+ model’s mislabeling of large parts of the buildings as *vegetation* or *vehicles and people*, thus decreasing the linear model’s prediction (e.g., the bottom image of Fig. 13).

Of the dozen cases where the model over-predicted component 1, five involved passing by wooden buildings, which were less common in the downtown area than concrete and brick buildings and possibly indicates that the lower emissivity of buildings constructed from wood results in poor fits in this model. Three cases involved the right side of the vehicle being very close to (nonwooden) building facades, which could be an indication that those buildings were lower in activity than the average building. The other scenarios offered no clear commonalities.

VII. DISCUSSION

The work presented here shows that features derived from measurements of the gamma-ray background in an urban area can be correlated with some fidelity to features derived from panoramic visual imagery of the same. In particular, the WR-2 model decomposition of the measured spectra resulted in one component that is strongly associated with the presence of buildings, and another that is strongly associated with the visibility to the sky. Although all of the model fits had high prediction error and only a few were statistically acceptable, the analysis presented here nevertheless finds connections between spectral and image features similar to what has previously been seen with the same mobile system but in a smaller and more radiologically uniform urban environment [31].

The shapes of the spectral components from NMF suggest there may be a reason for the associations since WR-2 component 0’s spectrum shows what is likely to be more cosmic and skyshine emission than WR-2 component 1, and visibility to the sky should be a rough proxy for exposure to cosmic and skyshine emission. In addition, nearby KUT emission, captured at a higher proportion by component 1, should increase when there are larger surfaces of KUT emission around the vehicle, a rough proxy for which is the size of nearby buildings. These same associations were noted in a separate RadMAP dataset at Fort Indiantown Gap, PA, USA [31], and a similar separation of gamma-ray emission

into nearby and distant emission has also been observed in airborne data [39]. The separation of urban gamma-ray backgrounds into nearby (*building-dominated*) and distant (*sky-dominated*) emission may therefore be a general phenomenon. The connection between cosmic emission and the *sky* feature is also reminiscent of the connection between measured fast neutron backgrounds and the fraction of the sky visible and not shielded by buildings [51]–[53]. The results of this work suggest an analogous result for gamma-ray backgrounds.

Another result of this work is that in all of the models, little to no emission is fit to clutter: the classes *vehicles and people* and *other*. The fact that they contribute so little to the fits comports with the notion their presence largely serves to attenuate background emission coming from the buildings or roads behind or beneath them, which is a known impact of such clutter on these systems [26]. Of course, the rare exception for clutter is when vehicle cargo is particularly radioactive or a pedestrian has undergone a nuclear medicine treatment, but such events have been screened for and are not in this dataset.

Obviously, models that tie gamma-ray backgrounds more strongly to surrounding imagery may exist than what is shown here. This work takes the brute-force approach of generating various sets of gamma-ray features, using NMF models that have been tailored by regularization approaches that were inspired by physics or statistical considerations. These gamma-ray features are then compared to image features that were derived using a single semantic segmentation approach with fixed visual categories. This approach results in our only examining a small parameter space within a high-dimensional dataset. Correlations were found to exist between some of these feature sets, but there could easily be other correlations that exist outside of the small set of features examined here. For example, different kinds of buildings (e.g., brick versus concrete versus wooden) might cluster together into different and predictable levels of KUT flux that the DeepLabv3+ model is currently blind to since it collapses all building types into one class. But even for the same type of material, the KUT activities can vary by orders of magnitude [14], and therefore any model that seeks to make a unique KUT activity prediction based on imagery alone is going to have difficulty. Also, the correlations might break down outside of this relatively small region of this particular city, if, for example, the flux from the roadbed or nearby soil is much larger in another location, or if the magnitude and composition of surface fluxes among buildings becomes more variable. Other radiological effects that might not be readily noticed in imagery are those due to weather, especially the increased background due to radon progeny during rainfall events, and while the data analyzed herein did not include any rainfall, we posit other contextual sensors could be incorporated into a model to address such complications. To be as robust and accurate as possible, any model like this would have to be trained and deployed in the same local area, but some phenomena such as sky-related features would likely persist even if a model were broadly trained.

One area for improvement over this work would be to extend the model beyond the simple linear model used here.

A trivial extension would be to include an intercept in the linear model, which would allow the model to capture spectral emission that does not change in time and may not be attributable to imagery. Some nonlinear models may also be useful.

Another area for improvement would be in the simultaneous engineering and correlation of spectral and image features. We attempted a small move in this direction via regularization with NMF. Training on the evaluation dataset, we used a regularization term that minimized the symmetric KL divergence between component 0's weights and the *sky* image feature, and also the symmetric KL divergence between component 1's weights and the *building* image feature. The resulting two-component model closely resembled NMF-WR-2 and had only slightly higher fractions f for the two image classes.

The simultaneous engineering and correlation of spectral and image features is perhaps a task better performed by deep convolutional neural networks (CNNs). Indeed, CNNs have already been used for this task with some initial success [27]. This work could provide guidance to the design of future CNNs as well as in aiding the interpretability of a trained CNN model. For example, a model that takes its inspiration from the DeepLabv3+ model for its initial layers, and with one or more fully connected layers at the end to predict the spectrum, might be a promising direction to pursue (one fully connected layer is equivalent to a linear model like NMF). Of course, one would want to explore such a model not only in Oakland, but in many other cities to understand the transferability of the knowledge, since different cities have different background distributions (e.g., [25]).

Given the complexity of urban scenes, a model that can generate a spectrum based only on a single image may be possible but may not be robust. Instead, another potential avenue of research is to explore machine learning models that update the current background estimate using both recently measured spectra and the changing context around the system. A reinforcement learning model, i.e., a model that is trained to make predictions based on the recent state and new data, could be appropriate here. Such a model could be more flexible in generic urban scenes by learning universal trends in spectral variability that correspond to noticeable trends in imagery, such as learning that a vehicle pulling up next to the system will depress the background count rate by a certain amount. Beyond these suggested machine learning approaches, there likely exist numerous other approaches that could leverage the observations made herein in effort to improve radiological anomaly detection algorithms' efficacy when confronted with the realistic spectral and temporal variability encountered when operating in urban environments.

A limitation of this current work is the small size of the dataset examined, which consisted of 40 min in the same neighborhood of a single city. To draw wider conclusions on the types of correlations present in urban backgrounds, more data from more neighborhoods and cities is needed. RadMAP data can provide some of this need, although its coverage is limited to certain portions of the San Francisco

Bay Area. Larger datasets will be crucial for the training of CNNs, which typically require enormous datasets due to their multitude of parameters. It is unknown whether a single model (of any type) could be useful for any urban area, or whether its domain would be limited to a single city or even neighborhood.

The ultimate application of the types of models discussed here would be to improve the performance of a detection algorithm by leveraging visual imagery. Whether this is achieved by creating an algorithm that directly predicts background, or whether algorithms simply leverage context to inform alarming behavior is unclear. However, it is arrived at, a model that is able to ingest contextual information such as panoramic imagery and improve algorithm performance beyond the state-of-the-art would potentially revolutionize the urban search problem.

APPENDIX COSMIC REGULARIZATION FUNCTIONS AND THEIR GRADIENTS

We only apply the penalty f_1 to the components of the first column of \mathbf{W} that are above 1250 keV, represented by spectrum index I_{1250} . To use the symmetric KL divergence, we generate normalized vectors for the cosmic component and first row of \mathbf{W} above index I_{1250}

$$p_i \equiv \frac{W_{i0}}{\sum_{i' > I_{1250}} W_{i'0}} \quad (29)$$

$$q_i \equiv \frac{w_{\text{cos},i}}{\sum_{i' > I_{1250}} w_{\text{cos},i'}} \quad (30)$$

$$f_1(\mathbf{W}) = \sum_{i > I_{1250}} (p_i - q_i) \log\left(\frac{p_i}{q_i}\right). \quad (31)$$

The gradient of f_1 is

$$\begin{aligned} \frac{\partial f_1}{\partial W_{ij}} &= \frac{[j = 0][i > I_{1250}]}{\sum_{i' > I_{1250}} W_{i'0}} \\ &\times \left\{ \left(1 - \log q_i - \sum_{i' > I_{1250}} p_{i'} \log p_{i'} \right) \right. \\ &\quad \left. - \left(\frac{q_i}{p_i} - \log p_i - \sum_{i' > I_{1250}} p_{i'} \log q_{i'} \right) \right\} \quad (32) \\ &\equiv (\nabla_1^+)_{ij} - (\nabla_1^-)_{ij} \quad (33) \end{aligned}$$

where $[\]$ is the Iverson bracket (1 if argument is true, 0 otherwise), and we have split the gradient into a sum of terms that are always positive (∇_1^+) and terms that are always negative ($-\nabla_1^-$).

The f_2 function is the sum of \mathbf{W} for the spectral bins above 3 MeV, represented by the spectral index I_{3000}

$$f_2(\mathbf{W}) = \sum_{i > I_{3000}} \sum_{j > 0} W_{ij}. \quad (34)$$

The gradient of f_2 is

$$\frac{\partial f_2}{\partial W_{ij}} = [i > I_{3000}][j > 0] \quad (35)$$

$$\equiv (\nabla_2^+)_{ij}. \quad (36)$$

REFERENCES

- [1] T. Hjerpe, R. R. Finck, and C. Samuelsson, "Statistical data evaluation in mobile gamma spectrometry: An optimization of on-line search strategies in the scenario of lost point sources," *Health Phys.*, vol. 80, no. 6, pp. 563–570, Jun. 2001. [Online]. Available: https://journals.lww.com/health-physics/Abstract/2001/06000/STATISTICAL%20DATA_EVALUATION_IN_MOBILE_GAMMA.6.aspx
- [2] H. K. Aage and U. Korsbech, "Search for lost or orphan radioactive sources based on NaI gamma spectrometry," *Appl. Radiat. Isotopes*, vol. 58, no. 1, pp. 103–113, Jan. 2003. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0969804302002221>
- [3] K. P. Ziock *et al.*, "Large area imaging detector for long-range, passive detection of fissile material," *IEEE Trans. Nucl. Sci.*, vol. 51, no. 5, pp. 2238–2244, Oct. 2004. [Online]. Available: <https://ieeexplore.ieee.org/document/1344316>
- [4] L. J. Mitchell *et al.*, "Mobile imaging and spectroscopic threat identification (MISTI): System overview," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec. (NSS/MIC)*, Oct. 2009, pp. 110–118. [Online]. Available: <https://ieeexplore.ieee.org/document/5401849>
- [5] R. D. Penny *et al.*, "A dual-sided coded-aperture radiation detection system," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 652, no. 1, pp. 578–581, Oct. 2011. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900211002634>
- [6] S. Zelakiewicz *et al.*, "SORIS—A standoff radiation imaging system," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 652, no. 1, pp. 5–9, Oct. 2011. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900211004244>
- [7] J. C. Curtis *et al.*, "Simulation and validation of the Mobile Urban Radiation Search (MURS) gamma-ray detector response," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 954, Feb. 2020, Art. no. 161128. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218310453>
- [8] 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.*, vol. 60, no. 2, pp. 1147–1150, Apr. 2013. [Online]. Available: <https://ieeexplore.ieee.org/document/6496304>
- [9] D. E. Archer *et al.*, "Systematic assessment of neutron and gamma backgrounds relevant to operational modeling and detection technology implementation," Oak Ridge National Lab. (ORNL), Oak Ridge, TN, USA, Tech. Rep. ORNL/TM-2014/687, Jan. 2015. [Online]. Available: <https://www.osti.gov/biblio/1185844>
- [10] K. D. Jarman, R. C. Runkle, K. K. Anderson, and D. M. Pfund, "A comparison of simple algorithms for gamma-ray spectrometers in radioactive source search applications," *Appl. Radiat. Isotopes*, vol. 66, no. 3, pp. 362–371, Mar. 2008. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969804307002886>
- [11] T. J. Aucott, M. S. Bandstra, V. Negut, J. C. Curtis, D. H. Chivers, and K. Vetter, "Effects of background on gamma-ray detection for mobile spectroscopy and imaging systems," *IEEE Trans. Nucl. Sci.*, vol. 61, no. 2, pp. 985–991, Apr. 2014. [Online]. Available: <https://ieeexplore.ieee.org/document/6786359>
- [12] R. C. Runkle, L. E. Smith, and A. J. Peurrung, "The photon haystack and emerging radiation detection technology," *J. Appl. Phys.*, vol. 106, no. 4, Aug. 2009, Art. no. 041101. [Online]. Available: <https://aip.scitation.org/doi/full/10.1063/1.3207769>
- [13] 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 *Proc. IEEE Nucl. Sci. Symp. Conf. Rec. (NSS/MIC)*, Oct. 2009, pp. 126–133. [Online]. Available: <https://ieeexplore.ieee.org/document/5401843>
- [14] R. Trevisi, S. Risica, M. D'Alessandro, D. Paradiso, and C. Nuccetelli, "Natural radioactivity in building materials in the European union: A database and an estimate of radiological significance," *J. Environ. Radioactivity*, vol. 105, pp. 11–20, Feb. 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0265931X11002402>
- [15] B. R. Cosofret, K. Shokhiev, P. Mulhall, D. Payne, and B. Harris, "Utilization of advanced clutter suppression algorithms for improved standoff detection and identification of radionuclide threats," *Proc. SPIE*, vol. 9073, May 2014, Art. no. 907316, doi: [10.1117/12.2049831](https://doi.org/10.1117/12.2049831).
- [16] D. M. Pfund *et al.*, "Improvements in the method of radiation anomaly detection by spectral comparison ratios," *Appl. Radiat. Isot.*, vol. 110, pp. 174–182, Apr. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969804315304024>
- [17] 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.*, vol. 57, pp. 195–206, Apr. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306437915001866>
- [18] K. Miller and A. Dubrawski, "Gamma-ray source detection with small sensors," *IEEE Trans. Nucl. Sci.*, vol. 65, no. 4, pp. 1047–1058, Apr. 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8305516>
- [19] K. J. Bilton *et al.*, "Non-negative matrix factorization of gamma-ray spectra for background modeling, detection, and source identification," *IEEE Trans. Nucl. Sci.*, vol. 66, no. 5, pp. 827–837, May 2019. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8673769>
- [20] C. A. L. Presti *et al.*, "Baseline suppression of vehicle portal monitor gamma count profiles: A characterization study," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 562, no. 1, pp. 281–297, Jun. 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900206004402>
- [21] T. Burr, J. R. Gattiker, K. Myers, and G. Tompkins, "Alarm criteria in radiation portal monitoring," *Appl. Radiat. Isot.*, vol. 65, no. 5, pp. 569–580, May 2007. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0969804307000048>
- [22] T. P. Karnowski *et al.*, "Design of dual-road transportable portal monitoring system for visible light and gamma-ray imaging," *Proc. SPIE*, vol. 7665, May 2010, Art. no. 76651J, doi: [10.1117/12.850191](https://doi.org/10.1117/12.850191).
- [23] K. P. Ziock *et al.*, "Performance of the roadside tracker portal-less portal monitor," *IEEE Trans. Nucl. Sci.*, vol. 60, no. 3, pp. 2237–2246, Jun. 2013. [Online]. Available: <https://ieeexplore.ieee.org/document/6530685>
- [24] R. J. Livesay, C. S. Blessinger, T. F. Guzzardo, and P. A. Hausladen, "Rain-induced increase in background radiation detected by radiation portal monitors," *J. Environ. Radioactivity*, vol. 137, pp. 137–141, Nov. 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0265931X1400215X>
- [25] L. J. Mitchell *et al.*, "Gamma-ray and neutron background comparison of U.S. metropolitan areas," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 784, pp. 311–318, Jun. 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900215000467>
- [26] I. R. Stewart *et al.*, "Understanding and quantifying the systematic effects of clutter within a radiation detection scene," *J. Radioanal. Nucl. Chem.*, vol. 318, no. 1, pp. 727–737, Aug. 2018, doi: [10.1007/s10967-018-6159-8](https://doi.org/10.1007/s10967-018-6159-8).
- [27] C. Kaffine *et al.*, "Background spectrum estimation from panoramic images," presented at the IEEE Nucl. Sci. Symp. Med. Imag. Conf., Atlanta, GA, USA, 2017.
- [28] A. D. Nicholson *et al.*, "Multiagency urban search experiment detector and algorithm test bed," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 7, pp. 1689–1695, Jul. 2017. [Online]. Available: <https://ieeexplore.ieee.org/document/7869400>
- [29] D. E. Archer *et al.*, "Modeling and urban search experiments: Fort indiantown gap data collections summary and analysis," Oak Ridge National Lab. (ORNL), Oak Ridge, TN, USA, Tech. Rep. ORNL/LTR-2017/371, Oct. 2017. [Online]. Available: <https://www.osti.gov/biblio/1410932>
- [30] M. S. Bandstra *et al.*, "RadMAP: The radiological multi-sensor analysis platform," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 840, pp. 59–68, Dec. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900216309780>
- [31] M. S. Bandstra *et al.*, "Attribution of gamma-ray background collected by a mobile detector system to its surroundings using panoramic video," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 954, Feb. 2020, Art. no. 161126. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016890021831043X>
- [32] 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.*, vol. 203, no. 3, pp. 325–335, May 2018, doi: [10.1080/00295450.2018.1458558](https://doi.org/10.1080/00295450.2018.1458558).
- [33] M. S. Bandstra *et al.*, "Correlations between panoramic imagery and gamma-ray background in an urban area," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf. (NSS/MIC)*, Oct. 2019, pp. 1–5.
- [34] W. Hess, D. Kohler, H. Rapp, and D. Andor, "Real-time loop closure in 2D LIDAR SLAM," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, May 2016, pp. 1271–1278.

- [35] R. C. Runkle, M. F. Tardiff, K. K. Anderson, D. K. Carlson, and L. E. Smith, "Analysis of spectroscopic radiation portal monitor data using principal components analysis," *IEEE Trans. Nucl. Sci.*, vol. 53, no. 3, pp. 1418–1423, Jun. 2006.
- [36] J. Stinnett, C. J. Sullivan, and H. Xiong, "Uncertainty analysis of wavelet-based feature extraction for isotope identification on NaI gamma-ray spectra," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 7, pp. 1670–1676, Jul. 2017.
- [37] P. Paatero and U. Tapper, "Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values," *Environmetrics*, vol. 5, no. 2, pp. 111–126, Jun. 1994, doi: [10.1002/env.3170050203](https://doi.org/10.1002/env.3170050203).
- [38] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, no. 6755, pp. 788–791, Oct. 1999. [Online]. Available: <https://www.nature.com/nature/journal/v401/n6755/abs/401788a0.html>
- [39] M. S. Bandstra, T. H. Y. Joshi, K. J. Bilton, A. Zoglauer, and B. J. Quiter, "Modeling aerial gamma-ray backgrounds using non-negative matrix factorization," *IEEE Trans. Nucl. Sci.*, vol. 67, no. 5, pp. 777–790, May 2020. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/9025259>
- [40] D. M. Pfund, R. C. Runkle, K. K. Anderson, and K. D. Jarman, "Examination of count-starved gamma spectra using the method of spectral comparison ratios," *IEEE Trans. Nucl. Sci.*, vol. 54, no. 4, pp. 1232–1238, Aug. 2007. [Online]. Available: <https://ieeexplore.ieee.org/document/4178949>
- [41] D. D. Lee and H. S. Seung, "Algorithms for non-negative matrix factorization," in *Advances in Neural Information Processing Systems*, T. K. Leen, T. G. Dietterich, and V. Tresp, Eds. Cambridge, MA, USA: MIT Press, 2001, pp. 556–562. [Online]. Available: <https://papers.nips.cc/paper/1861-algorithms-for-non-negative-matrix-factorization.pdf>
- [42] S. Kullback and R. A. Leibler, "On information and sufficiency," *Ann. Math. Statist.*, vol. 22, no. 1, pp. 79–86, 1951. [Online]. Available: <https://www.jstor.org/stable/2236703>
- [43] V. P. Pauca, J. Piper, and R. J. Plemmons, "Nonnegative matrix factorization for spectral data analysis," *Linear Algebra Appl.*, vol. 416, no. 1, pp. 29–47, Jul. 2006. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S002437950500340X>
- [44] J. M. Becker, M. Rohbeck, and C. Rohlfing, "Adaptive weights for NMF with additional priors," in *Proc. Int. Symp. Intell. Signal Process. Commun. Syst. (ISPACS)*, Nov. 2015, pp. 89–94.
- [45] S. Z. Li, X. W. Hou, H. J. Zhang, and Q. S. Cheng, "Learning spatially localized, parts-based representation," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)* vol. 1, Dec. 2001, pp. 207–212.
- [46] S. Laue, M. Mitterreiter, J. Giesen, and J. Mueller. (2020). *Matrix Calculus*. [Online]. Available: <https://www.matrixcalculus.org>
- [47] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 801–818. [Online]. Available: https://openaccess.thecvf.com/content_ECCV_2018/html/Liang-Chieh_Chen_E%ncoder-Decoder_with_Atrous_ECCV_2018_paper.html
- [48] M. Cordts *et al.*, "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.
- [49] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [50] M. Salathe, B. J. Quiter, M. S. Bandstra, J. C. Curtis, R. Meyer, and C. H. Chow, "Determining urban material activities with a vehicle-based multi-sensor system," *Phys. Rev. Res.*, vol. 3, no. 2, Apr. 2021, Art. no. 023070, doi: [10.1103/PhysRevResearch.3.023070](https://doi.org/10.1103/PhysRevResearch.3.023070).
- [51] A. Iyengar, M. Beach, R. J. Newby, L. Fabris, L. H. Heilbronn, and J. P. Hayward, "Systematic measurement of fast neutron background fluctuations in an urban area using a mobile detection system," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 773, pp. 27–32, Feb. 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900214012029>
- [52] J. R. Davis, E. Brubaker, and K. Vetter, "Fast neutron background characterization with the radiological multi-sensor analysis platform (RadMAP)," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 858, pp. 106–112, Jun. 2017. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900217303820>
- [53] A. Glick, E. Brubaker, B. Cabrera-Palmer, M. Gerling, B. J. Quiter, and K. Vetter, "Deployment of a double scatter system for directional detection of background neutron radiation," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 992, Mar. 2021, Art. no. 165029. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900221000139>