

Lawrence Berkeley National Laboratory

LBL Publications

Title

Real-Time, Adaptive Radiological Anomaly Detection and Isotope Identification Using Non-Negative Matrix Factorization

Permalink

<https://escholarship.org/uc/item/4s1159jr>

Journal

IEEE Transactions on Nuclear Science, 72(11)

ISSN

0018-9499

Authors

Jones, AC
Bandstra, MS
Faaland, S
[et al.](#)

Publication Date

2025-11-01

DOI

10.1109/tns.2025.3617856

Copyright Information

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

Peer reviewed

Real-time, Adaptive Radiological Anomaly Detection and Isotope Identification Using Non-negative Matrix Factorization

A. C. Jones, M. S. Bandstra, S. Faaland, Y. Lai, N. Abgrall, S. Suchyta, and R. J. Cooper

Abstract—Spectroscopic anomaly detection and isotope identification algorithms are integral components in nuclear nonproliferation applications such as search operations. The task is especially challenging in the case of mobile detector systems due to the fact that the observed gamma-ray background changes more than for a static detector system, and a pretrained background model can easily find itself out of domain. The result is that algorithms may exceed their intended false alarm rate, or sacrifice detection sensitivity in order to maintain the desired false alarm rate. Non-negative matrix factorization (NMF) has been shown to be a powerful tool for spectral anomaly detection and identification, but, like many similar algorithms that rely on data-driven background models, in its conventional implementation it is unable to update in real time to account for environmental changes that affect the background spectroscopic signature. We have developed a novel NMF-based algorithm that periodically updates its background model to accommodate changing environmental conditions. The Adaptive NMF algorithm involves fewer assumptions about its environment, making it more generalizable than existing NMF-based methods while maintaining or exceeding detection performance on simulated and real-world datasets.

I. INTRODUCTION

Spectroscopic anomaly detection and isotope identification algorithms are integral components in nuclear nonproliferation applications such as detection of radioactive materials that are lost, stolen, or otherwise outside of regulatory control. Such algorithms have been used in applications including homeland security [1]–[4]; contamination mapping, monitoring, and remediation [5], [6]; and emergency response (e.g., recovery of lost sources) [7], [8].

Numerous methods have been developed to detect and identify anomalous sources of gamma radiation in different systems, including for networked arrays of detectors [9]–[18], radiation portal monitors [19]–[23], aerial detector systems [24]–[27], and ground-based mobile detector systems [28]–[33]. Each of these data collection methods presents unique challenges in algorithm development. The main challenges of search tasks performed using mobile detector systems (the use

case on which this paper will focus), particularly in complex environments such as cities, are the following:

- 1) **Rapidly changing background gamma-ray radiation signatures:** Variations of concentrations of naturally occurring radioactive material (NORM) in building materials, density of buildings in a given area, visibility of the sky, and static and dynamic occlusion from objects in the environment are just a few factors that influence the temporally and spatially varying background spectrum. The ability to accurately account for changes in the background is key to achieving high detection sensitivity and controlling false positive rates [25], [34]–[36].
- 2) **Weak/low signal-to-noise ratio (SNR) sources:** Effective algorithms often need to be able to detect weak anomalies, as shielding and other environmental attenuation factors will reduce the observable activity. The ability to detect weak anomalies is closely tied to maintaining an accurate background model.
- 3) **Maintaining an operationally relevant false-alarm rate (FAR):** From an operational standpoint, it is essential that detection and identification algorithms maintain a FAR at or below some predetermined level (e.g., 1 per 8 hours). False alarms waste time and resources as an operator will have to investigate the alarm. In order to maintain a low FAR, it is helpful for the alarm metric distribution to be stable and well-understood to make threshold setting simple.

To date, many approaches which use varying degrees of spectral information have been developed. Some algorithms use only the gross count rate of gamma rays registered in the detector [37], [38], disregarding their energies. Others rely on truncated spectral regions of interest [39] or ratios between coarse spectral bins [26], [35], making them better suited to detect weak anomalies (as in point 2, above) than algorithms using gross count rate alone. Algorithms that perform full-spectrum analysis [20], [23], [40]–[43] make use of all available spectroscopic information, which generally results in better detection performance. In recent years, various classical [22], [40], [44] and deep learning [18], [41], [42], [45], [46] approaches to background estimation and anomaly detection have also been developed.

Though, in general, more complex algorithms seem to obtain better performance, recent research in machine learning and artificial intelligence has shown that even for sophisticated models there is an inevitable drift over time away from optimal

A.C. Jones, M.S. Bandstra, S. Faaland, Y. Lai, N. Abgrall, and R.J. Cooper, are with the Nuclear Science Division at Lawrence Berkeley National Laboratory, Berkeley, CA 94720 USA, e-mail: acjones@lbl.gov. S. Suchyta is with Nevada National Security Sites, Remote Sensing Laboratory - Andrews, Joint Base Andrews, MD 20762.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory (LBNL) under Contract DE-AC02-05CH11231. This work was done by Mission Support and Test Services, LLC, under Contract No. DE-NA0003624 with the U.S. Department of Energy and the National Nuclear Security Administration's Office of Defense Nuclear Nonproliferation DOE/NV/03624--2197.

performance due to the widening gap between the data used to train the model and the current data on which it is being evaluated [47]–[49]. Model drift is closely related to point 1 above, and effective anomaly detection algorithms should have methods in place to combat it.

Non-negative matrix factorization (NMF) [50] has been shown to be a useful framework for modeling gamma-ray background signatures [17], [27], [51], [52] when using “global” models trained on large datasets containing data representative of background signatures present in test cases. However, the existing Berkeley Lab-developed NMF-based approach [51] lacks the flexibility required for deployment on mobile detector systems in scenarios where a pre-trained model is unable to capture the complexity of unseen data. In Section II, we will discuss why NMF-based algorithms are inherently well suited for detection and identification tasks and how the constraints of the conventional approach lead to clear limitations of performance on mobile detection tasks. We then present Adaptive NMF in Section III, an update to the conventional approach that seeks to address the issue of temporal drift and extend the NMF framework to unseen environments.

Finally, Section IV we test the detection performance of Adaptive NMF against a suite of baseline algorithms including gross counts k-sigma [37], [38], multiplexed censored energy window (mCEW) [39], the “conventional” NMF approach [51], and Nuisance-Rejection Spectral Comparison Ratio (NSCRAD) [26], [35].¹ We show that Adaptive NMF achieves anomaly detection performance superior to the benchmarks while maintaining a low FAR under rapidly changing environmental conditions.

The algorithm was developed and tested using a combination of synthetic and experimental data from mobile detection systems. The benchmark algorithms against which Adaptive NMF is tested are implementations based on their descriptions in the literature from the Radiological Anomaly Detection and Identification (RADAI) project [54].

II. NON-NEGATIVE MATRIX FACTORIZATION

The notation used in the following presentation of NMF closely follows [51] and [17].

A. Background Modeling with NMF

NMF is a dimensionality reduction technique which takes a data matrix $\mathbf{X} \in \mathbb{R}_{\geq 0}^{n \times d}$ and decomposes it into a product of two low-rank matrices

$$\mathbf{X} \approx \hat{\mathbf{X}} = \mathbf{A}\mathbf{V}, \quad \mathbf{A} \in \mathbb{R}_{\geq 0}^{n \times k}, \quad \mathbf{V} \in \mathbb{R}_{\geq 0}^{k \times d}, \quad (1)$$

where d is the dimension of each of the vectors, n is the number of vectors in the dataset, and k is the rank of the resulting approximation matrix, $\hat{\mathbf{X}}$. The columns of the matrix

¹These algorithms were selected as benchmark comparisons because they are commonly fielded in real-world applications, there exist open-source implementations, and were recommended by the Detecting Radiation Algorithms Group [53]. This is by no means an exhaustive list of algorithms used for anomaly detection tasks, and there exist methods used by entities in private industry. However, due to the proprietary nature of such methods, we were unable to test against them.

$\mathbf{V}$ are a set of basis vectors for a lower-dimensional space that contains the Poisson expectation values of the original training data points². In practice, the basis vectors can be interpreted as background spectral templates developed from the training process. The rows of $\mathbf{A}$ function as the coefficients for linear combinations of the basis vectors that approximate the members of the original dataset. By decomposing a set of gamma-ray spectra in this way, the columns of $\mathbf{V}$ can be interpreted as the constituent components of the local gamma-ray background signature. In order to force the NMF components to be consistent with Poisson statistics, the background model is fit by optimizing the Poisson negative log likelihood loss function:

$$\mathcal{L} = -\log L(\mathbf{X}|\mathbf{A}, \mathbf{V}) \equiv \sum (\mathbf{A}\mathbf{V} - \mathbf{X} \odot \log(\mathbf{A}\mathbf{V})) \quad (2)$$

where $\odot$ is the element-wise matrix product and the sum is taken over the entire matrix. In equation 2, the matrix $\mathbf{X}$ is now also constrained to have nonnegative integer-valued entries, representing the counts in each gamma-ray spectral bin.

The non-negativity constraint of NMF admits background models that are more consistent with photon physics than other dimensionality reduction methods such as singular value decomposition (SVD). Consequently, NMF models lend themselves well to physical interpretability and potential correlation with real features of the environment without the addition of explicit information about nuclear physics [17], [27]. While increased interpretability is often observed in practice, NMF does not guarantee interpretability for the following reasons:

- **NMF has no closed form solution:** NMF is computed by minimizing a loss function via maximum likelihood expectation maximization.
- **NMF solutions are not unique:** NMF requires an initial guess, which is often randomly initialized, so there is no guarantee that two solutions trained on the same dataset will admit the same decomposition if they are instantiated differently. Furthermore, one could construct a positive-definite linear transformation $\mathbf{W}$, apply it to $\mathbf{A}$ and its inverse to $\mathbf{V}$, and be left with a distinct and equally valid decomposition, as: $\mathbf{X} \approx \hat{\mathbf{X}} = \mathbf{A}\mathbf{V} = \mathbf{A}(\mathbf{W}\mathbf{W}^{-1})\mathbf{V} = (\mathbf{A}\mathbf{W})(\mathbf{W}^{-1}\mathbf{V}) \equiv \mathbf{A}'\mathbf{V}'$

Another important difference between NMF and methods like SVD is that the rank, k , of the approximation must be set prior to calculation. There exist many methods by which to choose an appropriate number of NMF components for a given training dataset. In this work, multiple models are trained with different values of k and the Akaike Information Criterion [55] is used as a heuristic for choosing the best model.

B. Anomaly Detection and Isotope Identification

In prior work, NMF has been used to perform anomaly detection and isotope identification using the following protocol:

- 1) Reconstruct an incident spectrum with the NMF background model and call its loss value $\mathcal{L}_B$.

²The space is not a true vector space because of the non-negative constraints. In particular, the set more closely resembles a conical hull with the additional constraints that all vectors in the set be non-negative.

- 2) For each source template in a template library t_i , add the template to the background model and reconstruct again. Call its negative Poisson log likelihood loss $\mathcal{L}_{t_i}$.
- 3) Perform a likelihood ratio test (LRT) for each of the sources in the library. This test statistic will function as the alarm metric $\ell_{t_i} \equiv 2(\mathcal{L}_B - \mathcal{L}_{t_i})$.
- 4) If any of the ℓ_{t_i} is above a precomputed threshold, raise an alarm. Take the argument maximum over all of the ℓ_{t_i} to be the isotope identification inference.

The distribution of the LRT statistic should roughly follow a χ^2 distribution with a single degree of freedom, but because the statistic is clipped at zero, the distribution will on average have about half of its values at zero. Thresholds can be set to yield a desired FAR by calculating the value where the integral of the χ^2 tail distribution is twice the desired probability of false alarm (PFA).

Algorithm 1 Conventional NMF Analysis

```

1: Require:  $\mathbf{x}$ , incident spectrum
2: Require:  $\mathcal{T}$ , detection threshold
3:  $\hat{\mathbf{x}}_B \leftarrow \text{fit}_B(\mathbf{x})$ 
4:  $\mathcal{L}_B \leftarrow \text{NPLL}(\mathbf{x}, \hat{\mathbf{x}}_B)$   $\triangleright$  Neg. Poisson log likelihood
5: for  $t_i \in \{t_0, t_1, \dots, t_m\}$  do  $\triangleright$  Iterate through templates
6:    $\hat{\mathbf{x}}_{t_i} \leftarrow \text{fit}_{B+t_i}(\mathbf{x})$ 
7:    $\mathcal{L}_{t_i} \leftarrow \text{NPLL}(\mathbf{x}, \hat{\mathbf{x}}_{t_i})$ 
8:    $\ell_{t_i} \leftarrow \max(2(\mathcal{L}_B - \mathcal{L}_{t_i}), 0)$   $\triangleright$  Likelihood Ratio Test
9: end for
10:  $is\_alarm \leftarrow \max(\ell_{t_i} | \forall i \in \{0, \dots, m\}) > \mathcal{T}$ 
11:  $id\_inference \leftarrow \arg \max(\ell_{t_i} | \forall i \in \{0, \dots, m\})$ 

```

^{137}Cs template is shown in Fig. 1, adapted from [51]. In this example, the addition of the source template improves the quality of the reconstruction and yields an alarm metric value $\ell_{^{137}\text{Cs}} = 36.8$ significantly above the detection threshold $\mathcal{T} = 21.6$.

Conventional NMF has been shown to perform well on detection and identification tasks when trained on a large amount of data representative of the testing environment. However, if this requirement is not met, as is often the case on mobile detector systems, the background model will no longer effectively represent the current environment as the detector enters unfamiliar territory, and detection sensitivity will suffer. Here we introduce ‘Adaptive NMF’, a more generalized approach to NMF-based anomaly detection and isotope identification whose purpose is to overcome the limitations of Conventional NMF.

III. ADAPTIVE NMF

The object of the Adaptive NMF algorithm is to maintain or exceed the detection sensitivity offered by the conventional NMF-based approach, but with the ability to adapt to changing environments and deployment modes without the need for a system operator to manually tune hyperparameter values, retrain the background model, or restart the algorithm altogether. In order to do so, the algorithm continuously collects background gamma-ray spectra and uses them to periodically refit its background model to remain consistent with local background conditions. The core functionality (and by extension the computational performance) of the algorithm is identical to Algorithm 1 with the following modification to the logic in the analysis step:

- Analyze the current spectrum with the most up-to-date background model
- If the algorithm does not flag the spectrum as anomalous, add it to a rolling buffer of background data to be used for retraining
- Refit the background model periodically.

The following sections describe additional features implemented in the Adaptive NMF framework to improve its detection accuracy and sensitivity.

A. Automatic Refitting

The adaptive NMF algorithm must be able to retrain its background model while maintaining real-time performance




Fig. 1. A ^{137}Cs -containing spectrum with the NMF background model fit and individual components (top), and the background+source template fit with the ^{137}Cs template (bottom). NMF components and ^{137}Cs template have been scaled by their corresponding reconstruction weights.

The algorithm set forth above (also shown in pseudocode in Algorithm 1) will be referred to henceforth as ‘Conventional NMF.’ This method of analysis runs at a rate of approximately 60 iterations per second on a 10-core Apple M1 chip. Similar performance was observed on an Intel® Core™ i7-1370P 14-core Processor. An example of a reconstruction of a spectrum containing ^{137}Cs using NMF background components and a




Fig. 2. Adaptive NMF alarm metric values and inferred encounter windows (shown in red) using the temporal windowing method.

(typically operating on incident spectra at 1–2 Hz). The model retraining step was implemented as a parallel process which waits idly until it receives a signal to begin. The background buffer on which the NMF model is trained contains the most recent hour of spectra (3600 spectra at 1 Hz) that did not raise alarms, and retraining takes place every 10 minutes. Both the buffer length and the retraining frequency are adjustable parameters that can be tuned based on a user’s specific practical or computational needs.

When collecting spectra from large-volume (e.g., 2”×4”×16” NaI) detectors in mobile operations, we found that three NMF components are typically sufficient to describe the background. Such models take about 20 seconds on average to retrain. In practice, this translates to approximately 3% of spectra being analyzed with the previously trained background model. No case in which this behavior leads to a degradation of detection accuracy or sensitivity has been observed in practice.

B. Tikhonov Regularization

Gamma ray signatures of NORM sources are spectroscopically similar to the background which can lead to a mathematical degeneracy during analysis. In order to suppress the effects of this degeneracy — namely, false alarms arising from NORM anomalies — a Tikhonov regularization penalty was applied to the alarm metric value during analysis. This regularizer is functionally equivalent to applying an L_2 penalty only to the source template component’s coefficient. This forces the algorithm to preferentially use the background model components over the source templates during reconstruction. As a result, the more distinct a source template’s spectral shape is from the background components, the smaller the regularizer’s effect on the template’s coefficient will be, and vice versa.

The addition of this regularizer greatly improves the consistency of the desired and observed false-alarm rates. However, one drawback of this method is that the algorithm loses sensitivity to NORM in the environment. In practical applications, NORM is often not regarded as a threat source. Therefore, in the context of an urban search task, the benefits of this method outweigh the risks associated with decreased sensitivity. The strength of the regularizer is also an adjustable parameter that can be tuned based on a user’s specific needs.

C. Temporal Windowing

Because the Adaptive NMF algorithm is allowed to construct its own training datasets using uncontrolled, unlabeled data, there is a danger that signatures associated with weak sources (i.e., those that did not trigger an alarm) may be inadvertently incorporated into the background model. This would reduce the algorithm’s sensitivity to similar source signatures in the future, as they would be more easily described by the background model. Therefore, mitigating false negatives is critical in maintaining detection performance. This is addressed by implementing a subalgorithm which constructs temporal exclusion windows around NMF alarms. In practice, we expect there to be periods during which a source is present but the signal strength is too weak to be detected. The temporal windowing method aims to prevent these low-SNR source-containing spectra found at the beginnings and ends of encounters from entering the background buffer.

When a spectrum raises an alarm, a window containing only the alarming spectrum is constructed. The algorithm then iterates forwards and backwards in time, integrating the window as it goes, and analyzes the resulting multi-second spectrum. This process continues until the alarm metric of the integrated window falls below a secondary threshold, set as a fraction of the primary detection threshold. This fraction is a tunable parameter, with higher values corresponding to shorter windows, and vice versa.

In addition to preventing algorithm sensitivity loss, the temporal windowing mechanism also provides users with an estimate of the bounds of source encounters. It handles asymmetric encounter windows and encounters of different lengths, as shown in Fig. 2. The blue curve shows the algorithm alarm metric value over two separate simulated source encounters, the black curve shows the actual value of SNR (calculated as $\frac{s}{\sqrt{s+b}}$ with s , the number of photons originating from the source and b , the number of photons coming from the background) at each second, and the red boxes are the exclusion windows determined by the temporal windowing logic. The subalgorithm successfully excludes multiple spectra with nonzero SNR from the background buffer that did not raise alarms.

Another potential approach to mitigating cascading sensitivity loss due to the incorporation of source-containing spectra into the background buffer would be the addition

of a secondary detection threshold, lower than the primary detection threshold. One could then exclude spectra from the background buffer whose maximum alarm metric exceeds this threshold even if the primary detection threshold is not exceeded. We did not pursue this approach due to its potential to filter out true background spectra with high variability. This would impede the background buffer from capturing some background fluctuations that are important for developing the NMF background model. The implementation of such an approach and comparing its efficacy to the temporal windowing method described here is a potential avenue for future work.




Fig. 3. Number of NMF components being used at each second of analysis (green), gross count rate (black), and false alarms (red). Running Adaptive NMF with the model switching logic applied to analysis (bottom) yields far fewer false alarms than without (top).

D. Background Complexity Detection

The ability of an algorithm to detect abrupt changes in the background and adapt accordingly is crucial for mobile search applications. Consider a scenario in which a detector system is at rest for an extended period of time (e.g., longer than the algorithm background buffer). Because gamma-ray backgrounds are less variable when observed on a static detector, the background buffer is populated with low-variability spectra, and Adaptive NMF will select a simple model. If the detector suddenly began to move, Adaptive NMF would likely raise many false alarms due to the complexity of the current (static) background model being insufficient to describe the dynamic and changing environment.

In order for the algorithm to detect an abrupt change and adjust its background model accordingly, an additional step was added to the analysis method. Namely, previous background models of varying complexity are kept in memory,

and the algorithm performs an additional likelihood ratio test between the current k -component background model and the most recent $(k + 1)$ -component model. If this LRT statistic is the highest one (i.e., a more complex background is the most likely explanation for the incident spectrum), the background model is replaced with the more complex one. This process is repeated until the background-background LRT statistic, $\ell_{B_{k+1}} \equiv 2(\mathcal{L}_{B_k} - \mathcal{L}_{B_{k+1}})$, is no longer the highest value.

Fig. 3 shows the efficacy of this approach in suppressing false alarms that would otherwise have been raised by the algorithm. The detector is static for approximately one hour, from approximately $t = 22000 - 25600$, after which point the gross count rate, shown in black, becomes highly variable. The detector remains mobile for the remainder of the experiment, with the exception of a short period at $t \approx 29000$. The number of components Adaptive NMF uses for its background model at each second is shown in green. When the algorithm is allowed to switch its background to a more complex previously-used model upon encountering a sharp change in the background, it is able to avoid becoming stuck in a state of false alarm.

E. Multiple Integration Times

Because the optimal integration window to maximize the SNR of a given source encounter is unknown when running in arbitrary environments, multiple Adaptive NMF instances are run in parallel with varying integration times. We found that using three integration timescales (short: 1–3 s, medium: 4–6 s, and long: 7–10 s) provided sufficient flexibility for the algorithm to detect source encounters of different lengths. Alarm metric values from each subprocess are concatenated, and the argument maximum overall is reported as the isotope identification inference label.

It was empirically observed that models derived from 1-second integrated spectra could be used to perform inference on multi-second integrated spectra with no noticeable loss in performance. For this reason, all background models used by the individual subprocesses are trained on 1-second integrated spectra. Each subprocess also performs model switching independently of the others, providing additional flexibility. Analysis time increases approximately linearly with the number of subprocesses employed.

F. Automatic Energy Calibration

In order to account for gain shifts and other detector effects, an automatic energy calibration routine [56], [57] was applied to each incoming spectrum before being sent to Adaptive NMF for analysis. The algorithm fits background spectra to full-spectral templates to perform continuous calibration. The method was tested on a NaI detector of a similar volume to those used throughout this work that experienced temperatures between -40°C and $+50^\circ\text{C}$, during which time less than 0.4% peak variability was observed at 122 keV and less than 0.2% at 1460 keV. Fig. 4 shows the logical flow of the full Adaptive NMF algorithm with all of the features described in the preceding subsections.




Fig. 4. The complete system architecture diagram of the Adaptive NMF algorithm.

TABLE I

Simulated Sources in RADAI Dataset			
NORM	Industrial	Medical	(S)NM
⁴⁰ K	¹⁹² Ir	¹³¹ I	NatU
²³² Th	⁶⁰ Co	¹⁸ F	HEU
²²⁶ Ra	¹³³ Ba	^{99m} Tc	LEU
	¹³⁷ Cs	⁵⁷ Co	Refined U
		²⁰¹ Tl	²⁴¹ Am
		⁹⁰ Sr	FGPu
		⁶⁷ Cu	WGPu

(HEU: High-Enriched Uranium, LEU: Low-Enriched Uranium, NatU: Natural Uranium, FGPu: Fuel-Grade Plutonium, WGPu: Weapons-Grade Plutonium)

IV. RESULTS

Adaptive NMF was tested against gross counts k-sigma [37], [38], mCEW [39], Conventional NMF [51], and NSCRAD [26], [35]. The implementations of the benchmark algorithms come from the RADAI algorithm library [54].

Analysis was performed by running sequential one-second integrated gamma-ray spectra through each algorithm. If there was any alarm within the ground-truth bounds of each encounter, it was counted as a detected encounter. Any alarm outside the bounds of any encounter was counted as a false alarm. Because only a subset of the benchmark methods are capable of performing isotope identification, only anomaly detection performance will be reported in this section.

A. Simulated Datasets

RADAI provides a dataset [58], [59] that simulates a 2"×4"×16" NaI detector moving through an urban environment. The purpose of the RADAI dataset is to provide users with simulated source encounters in environments where data collection would be prohibitively expensive or otherwise infeasible. It includes sources one might typically encounter in an urban environment like those used in medical and industrial applications, threat sources like highly enriched Uranium and weapons-grade Plutonium, and NORM. All sources also have varying shielding configurations common for their individual applications. Single-source encounters with all 21 of the sources listed in Table I are included in the dataset. The data were prepared by integrating list-mode events into a series of one-second spectra in 129 nonlinear energy bins from 45–3000 keV. The width of these bins is proportional to the square root of the energy in order to ensure that the ratio of the bin widths to the detector energy resolution is approximately constant throughout the entire spectrum. The template library used by each of the benchmark algorithms that make use of spectroscopic information (i.e., all but k-sigma) was also generated from RADAI source data.

For these tests, Adaptive NMF was instantiated with a background buffer length of 10000 1-second integrated spectra; a refit period of 600 seconds (10 minutes); integration times of 1, 2, 3, 4, and 5 seconds running in parallel; and a Tikhonov regularization factor of 10^{-4} . The Conventional NMF instance




Fig. 5. A set of ROC curves comparing the detection performance as a function of false alarm rate for each of the benchmark algorithms and Adaptive NMF. Adaptive NMF performs the best at functional false alarm rates. Note the log scale on the x-axis.

was trained on 10 hours' worth of background data (36000 spectra).

Fig. 5 shows a set of Receiver Operating Characteristic (ROC) curves comparing the detection performance of Adaptive NMF against the suite of RADAI benchmark algorithms tested on all 21 simulated sources. The x-axis of the ROC curve is the PFA per second, and the y-axis gives the probability of detection (PD) per source encounter. Points on the curve are calculated by varying each algorithm's detection threshold value and observing the corresponding PD and PFA. All other algorithm hyperparameter values are fixed. At functionally relevant false alarm rates (e.g., 1 per 8 hours, corresponding to a false alarm probability of about 3.5×10^{-5}), Adaptive NMF outperforms all benchmark methods.

Fig. 6 shows the performance (i.e., PD as a function of SNR at a fixed PFA) of Adaptive NMF against Conventional NMF and NSCRAD on two single-isotope detection tasks. Similarly, Adaptive NMF demonstrates superior detection performance over the other algorithms, achieving 50% probability of detection at an SNR of 3.53 on ^{60}Co compared to 4.62 for Conventional NMF and 7.85 for NSCRAD. On ^{241}Am encounters, Adaptive NMF achieves a probability of detection of 50% at an SNR of 5.47 compared to 7.70 for Conventional NMF and 9.00 for NSCRAD.

B. Field Tests

The following results come from tests performed on data collected in urban and rural locations with vehicle-borne $2'' \times 4'' \times 16''$ NaI detectors. We performed these tests with an

instance of Adaptive NMF equipped with a background buffer of 3600 1-second integrated spectra; a refit period of 600 seconds (10 minutes); integration times of 1, 2, 4, and 8 seconds running in parallel; and a Tikhonov regularization factor of 5×10^{-4} . Adaptive NMF was tested against Conventional NMF (trained on a dataset of 8742 spectra representative of the testing environment) and an implementation of NSCRAD that has been used in field applications. These tests consisted of encounters with four different sources at six different distances. Eight passes were performed at each distance.

Test results are shown in Fig. 7. Adaptive NMF is consistently able to detect non-NORM sources at larger standoff distances than NSCRAD and Conventional NMF at a fixed FAR of 1 per hour. In the case of NORM sources like ^{226}Ra , the detection performance of Adaptive NMF is poorer than that of the benchmark algorithms because such sources closely resemble background, and the Tikhonov penalty imposed during analysis (described in Section III-B) suppresses their alarm rates. As such, encounters with NORM sources are largely rolled into the background buffer and used for subsequent model refits, further exacerbating the difficulty in detecting them. However, while natural uranium, for example, is difficult to detect due to its high similarity with background spectral features, the ability of Adaptive NMF to detect uranium increases as the isotopic composition of the uranium is moved further away from secular equilibrium. Fig. 7 shows that Adaptive NMF remains the most sensitive of the three algorithms on encounters with depleted uranium at large standoff distances.

C. Nuisance Alarm Suppression

The Adaptive NMF algorithm is able to effectively adapt to changes in background environment and maintain a low and consistent false alarm rate compared to the conventional approach, as illustrated in Fig. 8. Both methods were trained on a dataset of spectra collected in Washington D.C. and then tested on a dataset of spectra collected in rural New Mexico. Conventional NMF (bottom) raises numerous false alarms due to the inconsistency of its background model with the environment on which it is being tested. In contrast, Adaptive NMF (top) begins analysis with a high false alarm rate as it gathers new background data, but is able to effectively capture the spectral signature in its background model after the first model refit.

Adaptive NMF has a clear advantage over the other benchmark methods — particularly Conventional NMF — during rain events. As it begins to rain, there is a gradual increase in radon progeny in the background [60]. If a Conventional NMF instance faced such a scenario without having been previously trained on spectra collected during rain, it would likely raise false alarms as the concentration of radon exceeds the expected level. Without operator intervention, the false alarms would continue for hours after the rain stops until the radionuclides from the rain decay to a point where they no longer contribute to a significant component of the background.

Fig. 9 shows the Conventional and Adaptive NMF alarm metric values over a 25-hour period which included a rain




Fig. 6. Logistic regression-fit sigmoids representing the probability of NSCRAD (green), Conventional NMF (orange), and Adaptive NMF (blue) as a function of SNR at a fixed FAR (1 per hour) for ^{60}Co (left) and ^{241}Am (right) encounters. The further a curve is to the left, the better the algorithm is at detecting weak sources. Adaptive NMF achieves a 50% probability of detection on both sources at a lower SNR than either of the other two algorithms. Error bars show one-sigma uncertainty in both the x and y axes.




Fig. 7. The fraction of real-world source encounters that were detected by Adaptive NMF (blue), Conventional NMF (orange), and NSCRAD (green) at each standoff distance at a fixed FAR of 1 per hour.

event. These data were collected from a static detector node deployed as part of the PANDAWN project [61]. Adaptive NMF (right) was able to incorporate the background change into its model and raise far fewer false alarms than Conventional NMF (left) during the rain event.

V. DISCUSSION

We have presented an NMF-based spectroscopic anomaly detection and isotope identification algorithm that generalizes more easily to unseen environments than existing comparable methods. In particular, Adaptive NMF met or exceeded the detection performance of Conventional NMF in most test cases. This is an encouraging result because Adaptive NMF is generally trained on far less data than Conventional NMF at

any given point and makes fewer *a priori* assumptions about the underlying environmental conditions. By demonstrating that a local model is sufficient for robust anomaly detection, we have extended the existing NMF framework to rapidly changing background conditions.

We note the following potential weakness of the Adaptive NMF algorithm: As depicted in Fig. 7, our method has reduced sensitivity to NORM anomalies. The Tikhonov parameter was implemented to suppress the effects of the mathematical degeneracy arising from the spectroscopic similarity of NORM signatures with the background and to reduce false alarms arising from NORM anomalies. However, as a result, the values of the NORM template alarm metrics are suppressed, which makes it even more challenging to distinguish true




Fig. 8. Conventional (bottom) and Adaptive (top) NMF alarm metric values when initially trained on a dataset of gamma-ray spectra collected in Washington D.C. and subsequently tested on spectra collected in rural New Mexico.




Fig. 9. Conventional (left) and Adaptive (right) NMF alarm metric values over a rain event from the PANDAWN project. The blue curves are the maximum alarm metric values at each second, the black lines show the instantaneous accumulation value from rain sensor on the PANDAWN node, and the red dotted lines denote times at which the alarm metric exceeded the detection threshold, $\mathcal{T} = 21.6$, calculated as in Section II-B.

signatures from typical Poisson fluctuations. The strength of the regularizer can be tuned or set to 0, giving users the flexibility to increase the algorithm's sensitivity to NORM, though it comes at the expense of the FAR.

Perhaps unsurprisingly, the performance of the methods tested against the RADAI dataset, as shown in Fig. 5, follows the same hierarchy as the amount of information being leveraged in each method. K-sigma uses only gross count rate information and had the poorest performance, which was to be expected in the presence of a rapidly changing background. mCEW, which constructs spectral regions of

interest from its training data to perform detection, had greater sensitivity than k-sigma but still suffered when faced with the highly variable RADAI background. NSCRAD, too, constructs spectral windows to detect broad classes of anomalies, but it also incorporates information about its environment via a continually updating background model. This enables it to outperform mCEW and k-sigma, but it is susceptible to becoming stuck in a state of false alarm when abrupt background shifts occur. Conventional NMF achieved greater sensitivity by preserving the most spectral information compared to the other methods described and by making use of a broad library of

templates of potential sources and a background model trained on many hours of data. The fact that Adaptive NMF was able to, in most test cases, match or improve upon the detection sensitivity of Conventional NMF is particularly encouraging due to the tradeoff in the information being utilized by each method. Adaptive NMF builds its background models from much smaller datasets — on the order of one hour’s worth of data — while Conventional NMF commonly uses many hours of data to build a complex model that is capable of capturing the richness of an environment over time in a single model. We have shown that in the case of highly variable background compositions, a flexible local model is not only sufficient but oftentimes superior to a highly complex static model.

As we have demonstrated with our method, the NMF framework lends itself particularly well to the integration of additional features that further improve its detection performance. As such, there are many directions for future work that could further improve NMF-based methods. For instance, integrating an anomaly clustering method similar to that presented in [17] could allow Adaptive NMF to update its template library with observed anomalies, making it more sensitive to the specific shapes of source signatures in its environment. The inclusion of temporal information in the analysis pipeline could also increase detection sensitivity by allowing previous results to inform the algorithm’s beliefs about future observations. Additionally, a continuously updating background model would generally be preferable to one that refits periodically to ensure that it does not become out-of-date between refits. A cursory exploration of such a method was performed, but it proved to be computationally intense and yielded no discernible difference in performance. However, it is possible that further investigation could uncover a feasible way to achieve a continuously updating model.

ACKNOWLEDGMENTS

The authors wish to thank Jayson Vavrek and Brian Quiter for their helpful comments and suggestions for this manuscript.

CREDIT STATEMENT

A.C.J.: Writing – original draft, Formal analysis, Visualization, Software, Investigation (lead), Methodology. **M.S.B.:** Writing – review & editing, Supervision, Methodology, Formal analysis. **S.F.:** Investigation (supporting). **Y.L.:** Investigation (supporting). **N.A.:** Data curation. **S.S.:** Validation, Formal analysis, Data curation. **R.J.C.:** Writing – review & editing, Funding acquisition, Project administration, Supervision.

REFERENCES

- [1] R. Kouzes, “Detecting Illicit Nuclear Materials,” *American Scientist*, vol. 93, no. 5, p. 422, 2005. [Online]. Available: <http://www.americanscientist.org/issues/page2/detecting-illicit-nuclear-materials>
- [2] A. J. P. Robert C. Runkle, L. Eric Smith, “The photon haystack and emerging radiation detection technology,” *J. Appl. Phys.*, 2009. [Online]. Available: <https://doi.org/10.1063/1.3207769>
- [3] B. Geelhood, J. Ely, R. Hansen, R. Kouzes, J. Schweppe, and R. Warner, “Overview of portal monitoring at border crossings,” in *2003 IEEE Nuclear Science Symposium. Conference Record (IEEE Cat. No.03CH37515)*, vol. 1, 2003, pp. 513–517 Vol.1.
- [4] L. Wein, “Preventing the importation of illicit nuclear materials in shipping containers,” *Risk Analysis*, vol. 26, pp. 1377–1393, 2006.
- [5] R. Engelbrecht, “Chapter 1 - Environmental Radioactivity Monitoring,” in *Handbook of Radioactivity Analysis*, 4th ed., M. F. L’Annunziata, Ed. Academic Press, 2020, vol. 2, pp. 1–40. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780128143957000015>
- [6] S. Hardie and I. McKinley, “Fukushima remediation: status and overview of future plans,” *Journal of Environmental Radioactivity*, vol. 133, pp. 75–85, 2014, environmental Radioactivity: Implications for Human and Environmental Health - International Symposium at Plymouth University, UK. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0265931X13001847>
- [7] Nuclear Newswire, “Cs-137 sealed source found in Western Australia,” *American Nuclear Society*, 2023.
- [8] D. C. W. Sanderson and J. D. Allyson, “An aerial gamma ray search for a missing 137Cs source in the Niger Delta,” *Scottish Universities Res. React. Centre, East Kilbride, Scotland, Tech. Rep.*, 1991. [Online]. Available: <https://eprints.gla.ac.uk/58144/>
- [9] M. Chandy, C. Pilotto, and R. McLean, “Networked sensing systems for detecting people carrying radioactive material,” in *2008 5th International Conference on Networked Sensing Systems*, Jun. 2008, pp. 148–155.
- [10] D. A. Cooper, R. J. Ledoux, K. Kamieniecki, S. E. Korbly, J. Thompson, J. Batcheler, S. Chowdhury, N. Roza, J. Costales, and V. Aiyawar, “Intelligent Radiation Sensor System (IRSS) Advanced Technology Demonstrator (ATD),” in *2012 IEEE Conference on Technologies for Homeland Security (HST)*, Nov. 2012, pp. 511–516.
- [11] R. Cooper, N. Abgrall, G. Aversano, M. Bandstra, D. Hellfeld, T. Joshi, V. Negut, B. Quiter, E. Rofors, M. Salathe, K. Vetter, P. Beckman, C. Catlett, N. Ferrier, Y. Kim, R. Sankaran, S. Shahkarami, S. Amitkumar, E. Ayton, J. Kim, and S. Volkova, “Networked sensing for radiation detection, localization, and tracking,” *J. Phys.: Conf. Ser.* 2586, 2023.
- [12] R. Nemzek, J. Dreicer, D. Torney, and T. Warnock, “Distributed sensor networks for detection of mobile radioactive sources,” *IEEE Transactions on Nuclear Science*, vol. 51, no. 4, pp. 1693–1700, Aug. 2004, conference Name: IEEE Transactions on Nuclear Science.
- [13] N. Hoteling, E. T. Moore, W. P. Ford, T. McCullough, and L. McLean, “An analysis of gamma-ray data collected at traffic intersections in Northern Virginia,” *arXiv:2104.04137 [physics]*, Mar. 2021, arXiv: 2104.04137. [Online]. Available: <http://arxiv.org/abs/2104.04137>
- [14] S. Brennan, A. Mielke, D. Torney, and A. Maccabe, “Radiation detection with distributed sensor networks,” *Computer*, vol. 37, no. 8, pp. 57–59, 2004.
- [15] B. Deb, “Iterative Estimation of Location and Trajectory of Radioactive Sources With a Networked System of Detectors,” *IEEE Transactions on Nuclear Science*, vol. 60, no. 2, pp. 1315–1326, Apr. 2013.
- [16] N. S. V. Rao, S. Sen, N. J. Prins, D. A. Cooper, R. J. Ledoux, J. B. Costales, K. Kamieniecki, S. E. Korbly, J. K. Thompson, J. Batcheler, R. R. Brooks, and C. Q. Wu, “Network algorithms for detection of radiation sources,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 784, pp. 326–331, Jun. 2015. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900215000686>
- [17] M. S. Bandstra, N. Abgrall, R. J. Cooper, D. Hellfeld, T. H. Y. Joshi, V. Negut, B. J. Quiter, M. Salathe, R. Sankaran, Y. Kim, and S. Shahkarami, “Background and anomaly learning methods for static gamma-ray detectors,” 2023.
- [18] M. D. Verlie, D. T. Murnane, A. Butko, R. Cooper, M. Kiran, E. Rofors, and K. Vetter, “Graph neural networks for radiological/nuclear detection with static detector networks,” in *2023 IEEE Nuclear Science Symposium, Medical Imaging Conference and International Symposium on Room-Temperature Semiconductor Detectors (NSS MIC RTS)*, 2023.
- [19] 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 Transactions on Nuclear Science*, vol. 53, no. 3, pp. 1418–1423, Jun. 2006.
- [20] D. Boardman, M. Reinhard, and A. Flynn, “Principal Component Analysis of Gamma-Ray Spectra for Radiation Portal Monitors,” *IEEE Transactions on Nuclear Science*, vol. 59, no. 1, pp. 154–160, Feb. 2012.
- [21] M. Ghelman, V. Berdichevsky, E. Vax, A. Osovitzky, E. Marcus, D. Ginzburg, E. Gonen, T. Mazor, and Y. Cohen, “RPMs performances enhancement - Using matched filter algorithm and GLRT to estimate source activity and location,” in *2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC)*, Oct. 2013, pp. 1–4, ISSN: 1082-3654.

- [22] Y. Altmann, A. Di Fulvio, and M. Paff, "Expectation-propagation for weak radionuclide identification at radiation portal monitors," *Sci Rep* 10, 6811, 2020.
- [23] B. R. Cosofret, K. N. Shokhirev, P. A. Mulhall, D. Payne, B. Harris, E. Arsenaault, and R. Moro, "Utilization of advanced clutter suppression algorithms for improved spectroscopic portal capability against radionuclide threats," in *2013 IEEE International Conference on Technologies for Homeland Security (HST)*, 2013, pp. 618–622.
- [24] A. J. Cresswell, D. C. W. Sanderson, and D. C. White, "¹³⁷Cs measurement uncertainties and detection limits for airborne gamma spectrometry (AGS) data analysed using a spectral windows method," *Applied Radiation and Isotopes*, vol. 64, no. 2, pp. 247–253, Feb. 2006. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0969804305002654>
- [25] P. Kock, C. Rääf, and C. Samuelsson, "On background radiation gradients – the use of airborne surveys when searching for orphan sources using mobile gamma-ray spectrometry," *Journal of Environmental Radioactivity*, vol. 128, pp. 84–90, Feb. 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0265931X13002373>
- [26] R. S. Detwiler, D. M. Pfund, M. J. Myjak, J. A. Kulisek, and C. E. Seifert, "Spectral anomaly methods for aerial detection using KUT nuisance rejection," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 784, pp. 339–345, Jun. 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0168900215000716>
- [27] 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 Transactions on Nuclear Science*, vol. 67, no. 5, p. 777–790, May 2020. [Online]. Available: <http://dx.doi.org/10.1109/TNS.2020.2978798>
- [28] L. Mitchell, B. F. Philips, W. Johnson, E. A. Wulf, A. Hutcheson, C. Lister, K. Bynum, B. Leas, and G. Guadagno, "Mobile Imaging and Spectroscopic Threat Identification (MISTI): System Overview," in *2009 IEEE Nuclear Science Symposium Conference Record (NSS/MIC)*, Oct. 2009, pp. 110–118. [Online]. Available: <https://ieeexplore.ieee.org/document/5401849>
- [29] R. D. Penny, W. E. Hood, R. M. Polichar, F. H. Cardone, L. G. Chavez, S. G. Grubbs, B. P. Huntley, R. A. Kuharski, R. T. Shyffer, L. Fabris, K. P. Ziocok, S. E. Labov, and K. Nelson, "A dual-sided coded-aperture radiation detection system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 652, no. 1, pp. 578–581, Oct. 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0168900211002634>
- [30] S. Zelakiewicz, R. Hctor, A. Ivan, W. Ross, E. Nieters, W. Smith, D. McDevitt, M. Wittbrodt, and B. Milbrath, "SORIS—A standoff radiation imaging system," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 652, no. 1, pp. 5–9, Oct. 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0168900211004244>
- [31] 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. Zakhov, R. Zhang, and K. Vetter, "RadMAP: the Radiological Multi-sensor Analysis Platform," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 840, pp. 59–68, 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900216309780>
- [32] 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," *Information Systems*, vol. 57, pp. 195–206, Apr. 2016. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306437915001866>
- [33] K. Miller, J. H. Good, I. Fawaz, D. Howarth, and A. Dubrawski, "Gamma-Ray Source Detection Under Occlusions and Position Errors in Cluttered Urban Scenes," *IEEE Transactions on Nuclear Science*, vol. 67, no. 6, pp. 1185–1194, Jun. 2020.
- [34] T. Aucott, M. Bandstra, V. Negut, J. Curtis, D. Chivers, and K. Vetter, "Effects of Background on Gamma-Ray Detection for Mobile Spectroscopy and Imaging Systems," *IEEE Transactions on Nuclear Science*, vol. 61, no. 2, pp. 985–991, Apr. 2014. [Online]. Available: <https://ieeexplore.ieee.org/document/6786359>
- [35] D. M. Pfund, K. K. Anderson, R. S. Detwiler, K. D. Jarman, B. S. McDonald, B. D. Milbrath, M. J. Myjak, N. C. Paradis, S. M. Robinson, and M. L. Woodring, "Improvements in the method of radiation anomaly detection by spectral comparison ratios," *Applied Radiation and Isotopes*, vol. 110, pp. 174–182, Apr. 2016. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0969804315304024>
- [36] I. R. Stewart, A. D. Nicholson, D. E. Archer, M. J. Willis, M. W. Swinney, I. Garishvili, and W. R. Ray, "Understanding and quantifying the systematic effects of clutter within a radiation detection scene," *Journal of Radioanalytical and Nuclear Chemistry*, Aug. 2018. [Online]. Available: <https://doi.org/10.1007/s10967-018-6159-8>
- [37] 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," *Applied Radiation and Isotopes*, vol. 66, no. 3, pp. 362–371, 2014-09-19. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0969804307002886>
- [38] D. E. Archer, D. E. Hornback, J. O. Johnson, A. D. Nicholson, B. W. Patton, D. E. Peplow, T. M. Miller, and B. Ayaz-Maierhafer, "Systematic Assessment of Neutron and Gamma Backgrounds Relevant to Operational Modeling and Detection Technology Implementation," Oak Ridge National Lab. (ORNL), Oak Ridge, TN (United States), Tech. Rep. ORNL/TM-2014/687, Jan. 2015. [Online]. Available: <https://www.osti.gov/biblio/1185844>
- [39] E. Lei, "Robust Detection of Radiation Threat," in *IEEE Nuclear Science Symposium and Medical Imaging Conference*, 2017, p. 20.
- [40] J. Lee, R. J. Cooper, T. H. Joshi, K. J. Bilton, D. M. Raji, M. S. Bandstra, and K. Vetter, "An ensemble approach to computationally efficient radiological anomaly detection and isotope identification," *IEEE Transactions on Nuclear Science*, vol. 69, no. 10, pp. 2168–2178, 2022.
- [41] P. Lalor, "Enhancing radioisotope identification in gamma spectra with transfer learning," 2024. [Online]. Available: <https://arxiv.org/abs/2412.07069>
- [42] J. M. Ghawaly, A. D. Nicholson, D. E. Archer, M. J. Willis, I. Garishvili, B. Longmire, A. J. Rowe, I. R. Stewart, and M. T. Cook, "Characterization of the Autoencoder Radiation Anomaly Detection (ARAD) model," *Engineering Applications of Artificial Intelligence*, vol. 111, p. 104761, 2022. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0952197622000550>
- [43] C. Sullivan, M. Martinez, and S. Garner, "Wavelet analysis of sodium iodide spectra," in *IEEE Nuclear Science Symposium Conference Record*, 2005, vol. 1, 2005, pp. 302–306.
- [44] Y. Altmann, S. McLaughlin, and A. Hero, "Robust linear spectral unmixing using anomaly detection," *IEEE Transactions on Computational Imaging*, vol. 1, no. 2, pp. 74–85, 2015.
- [45] M. Alamaniotis, "Mode-driven explainable artificial intelligence approach for estimating background radiation spectrum in a measurement applicable to nuclear security," *Annals of Nuclear Energy*, vol. 209, p. 110849, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306454924005127>
- [46] M. Kamuda, J. Zhao, and K. Huff, "A comparison of machine learning methods for automated gamma-ray spectroscopy," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 954, p. 161385, Feb. 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218313779>
- [47] F. Bayram, B. S. Ahmed, and A. Kassler, "From concept drift to model degradation: An overview on performance-aware drift detectors," 2022. [Online]. Available: <https://arxiv.org/abs/2203.11070>
- [48] K. Nelson, G. Corbin, M. Anania, M. Kovacs, J. Tobias, and M. Blowers, "Evaluating model drift in machine learning algorithms," in *2015 IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, 2015, pp. 1–8.
- [49] D. Vela, A. Sharp, R. Zhang, T. Nguyen, A. Hoang, and O. S. Pianykh, "Temporal quality degradation in AI models," *Scientific Reports*, vol. 12, no. 1, p. 11654, Jul 2022. [Online]. Available: <https://doi.org/10.1038/s41598-022-15245-z>
- [50] 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: <http://www.nature.com/nature/journal/v401/n6755/abs/401788a0.html>
- [51] 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 Transactions on Nuclear Science*, vol. 66, no. 5, pp. 827–837, May 2019. [Online]. Available: <https://ieeexplore.ieee.org/abstract/document/8673769>
- [52] M. S. Bandstra, B. J. Quiter, M. Salathe, K. J. Bilton, J. C. Curtis, S. Goldenberg, and T. H. Y. Joshi, "Correlations Between Panoramic Imagery and Gamma-Ray Background in an Urban Area," *IEEE Transactions on Nuclear Science*, vol. 68, no. 12, pp. 2818–2834, Dec. 2021, conference Name: IEEE Transactions on Nuclear Science.

- [53] M. S. Bandstra, C. Britt, J. Ghawaly, T. Grimes, T. Haard, P. Heimberg, T. H. Joshi, H. Komkov, S. Labov, N. McFerran *et al.*, “Metrics and methods for radiation detection algorithm characterization for nuclear/radiological source search,” Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), Tech. Rep., 12 2023. [Online]. Available: <https://www.osti.gov/biblio/2439899>
- [54] T. Joshi, “RADAI: Radiological Anomaly Detection And Identification,” <https://gitlab.com/lbl-anp/radai/radai>, 2020.
- [55] H. Akaike, “A new look at the statistical model identification,” *IEEE Transactions on Automatic Control*, vol. 19, no. 6, pp. 716–723, 1974.
- [56] M. Salathe, T. Joshi, M. Bandstra, D. Hellfeld, B. Quiter, and R. Cooper, “Real-time energy calibration and gain stabilization in statically deployed nai detectors,” in *2020 IEEE Nuclear Science Symposium*, 2020.
- [57] M. S. et al., “Automatic calibration of gamma-ray detectors deployed in urban environments,” *arXiv*, 2025.
- [58] D. Peplow, D. Archer, J. Ghawaly, T. H. Joshi, M. S. Bandstra, B. Quiter, and A. C. Nachtsheim, “Detection algorithm virtual testbed for urban search with SCALE.” Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), 09 2022. [Online]. Available: <https://www.osti.gov/biblio/1897851>
- [59] “RADAI Interactive Datasets,” 2024, <https://bdc.lbl.gov/wiki/public/radai-interactive-datasets/>, Last accessed on 2025-06-30. [Online]. Available: <https://bdc.lbl.gov/wiki/public/radai-interactive-datasets/>
- [60] R. J. Livesay, C. S. Blessinger, T. F. Guzzardo, and P. A. Hausladen, “Rain-induced increase in background radiation detected by Radiation Portal Monitors,” *Journal of Environmental Radioactivity*, vol. 137, pp. 137–141, Nov. 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0265931X1400215X>
- [61] R. Cooper, N. Abgrall, G. Aversano, M. Bandstra, D. Hellfeld, T. Joshi, V. Negut, B. Quiter, E. Rofors, M. Salathe, K. Vetter, P. Beckman, C. Catlett, N. Ferrier, Y. Kim, R. Sankaran, S. Shahkarami, S. Amitkumar, E. Ayton, J. Kim, and S. Volkova, “Networked sensing for radiation detection, localization, and tracking,” *Journal of Physics: Conference Series*, vol. 2586, no. 1, p. 012125, sep 2023. [Online]. Available: <https://dx.doi.org/10.1088/1742-6596/2586/1/012125>