

An Ensemble Approach to Computationally Efficient Radiological Anomaly Detection and Isotope Identification

J. Lee¹, R. J. Cooper, *Member, IEEE*, T. H. Joshi², *Member, IEEE*, K. J. Bilton³, D. M. Raji⁴,
M. S. Bandstra⁵, and K. Vetter⁶, *Member, IEEE*

Abstract—Radiological source search is a challenging task involving detection and identification of weak sources in a constantly changing radiological background. As of now, many radiological source detection algorithms have been proposed; however, their computational complexity and, hence, reliance on power intensive processing units inhibit low-power applications of radiological source search systems. In this work, we introduce the anomaly filter (AF) algorithm; a computationally light, yet effective time-series source detection algorithm based on exponential weighted moving average (EWMA) and Poisson deviance statistics. Then, we demonstrate that the proposed algorithm can be used in ensemble with other more computationally intensive source detection and identification algorithms to achieve both increased detection performance and reduced power consumption. The proposed AF algorithm and the ensemble algorithms were thoroughly benchmarked against several existing source detection and identification algorithms. The results show that the AF algorithm outperforms existing conventional source detection algorithms, and the ensemble approach improves the overall performance of existing source detection and isotope identification algorithms. Furthermore, the AF algorithm and the non-negative matrix factorization approach-based source identification (NMF-ID) algorithm were combined and implemented on a single-board microcontroller, and the power consumption was measured. This ensemble algorithm reduced the power consumption of the NMF-ID algorithm almost by a factor of 100, while improving the detection performance of the overall system.

Index Terms—Anomaly detection, gamma-ray detection, low-power electronics, non-negative matrix factorization (NMF), radiation source search.

I. INTRODUCTION

RADIOLOGICAL source search, which encompasses the detection, identification, and localization of unknown

Manuscript received 30 March 2022; revised 1 August 2022; accepted 6 August 2022. Date of publication 17 August 2022; date of current version 18 October 2022. This work was supported by the Defense Threat Reduction Agency (DTRA) under Grant 13081-31571.

J. Lee and K. J. Bilton are with the Nuclear Engineering Department, University of California, Berkeley, Berkeley, CA 94720 USA (e-mail: jwonlee@berkeley.edu).

R. J. Cooper, T. H. Joshi, and M. S. Bandstra are with the Applied Nuclear Physics (ANP) Program, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA 94720 USA.

D. M. Raji is with the Nuclear Engineering Department, University of Tennessee, Knoxville, TN 37996 USA.

K. Vetter is with the Nuclear Engineering Department, University of California, Berkeley, Berkeley, CA 94720 USA, and also with the Applied Nuclear Physics (ANP) Program, Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA 94720 USA.

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/TNS.2022.3198906>.

Digital Object Identifier 10.1109/TNS.2022.3198906

radiological materials, continues to be of importance for homeland security [1], [2], [3]; emergency response [4]; environmental monitoring [5], [6]; and contamination remediation [7]. In particular, the goal of radiological source detection and identification is to find and characterize anomalies present in time-integrated radiation spectral histograms, obtained from mobile or stationary radiation detection systems. Thus, the source detection and identification problem can be classified as a multivariate time-series anomaly detection problem.

Time-series anomaly detection is a well-studied problem [8], yet there are several factors that make radiological source search particularly challenging. First, radiological backgrounds, which constantly evolve both spatially and temporally, are often not known *a priori*, making it difficult to differentiate anomalies arising from radiological sources from the background fluctuation and Poisson noise. Furthermore, operational constraints, such as the need to maintain low false positive rates (FPRs), further limit the sensitivity of the overall radiological source detection system.

As of now, to effectively deal with the aforementioned challenges, a variety of source detection and identification algorithms have been developed. For example, some algorithms estimate the background by means of dimensionality reduction techniques, such as principal component analysis (PCA) [9], [10] or non-negative matrix factorization (NMF) [11], [12]. Others make use of specific energy windows within which the source and background count contributions are estimated using the spectrum outside the windows [13]. Additionally, some recent works utilize modern machine learning techniques, such as deep neural networks (DNNs) [14], [15] and Gaussian processes (GPs) [16].

These source detection and identification algorithms have been successful to varying degrees; however, due to their computational cost, some of these algorithms primarily rely on processing units with high power consumption. This, in turn, inhibits deployment of radiation source search systems in circumstances where stable power sources are not readily available.

Although the necessity of low-power implementation of source detection and identification algorithms has become apparent and recognized, to the best of our knowledge, the literature on the topic is scarce. For example, Huang *et al.* [17] implemented a convolutional spiking neural network-based source detection algorithm on a field programmable gate array (FPGA) and achieved the overall power consumption of

75 mW [17]. However, the work lacks thorough benchmarking of the implemented algorithm, and the approach may not be easily applicable to other already existing algorithms. Hence, a more general approach is needed.

In this work, to achieve low-power implementation of the radiological source detection and identification algorithms, we consider the use of an ensemble approach where a computationally efficient algorithm is used in series with a computationally heavy, yet more effective algorithm. In addition, we demonstrate that the proposed ensemble approach can improve the overall performance of various source detection and identification algorithms. First, we introduce the anomaly filter (AF) algorithm—an effective, yet computationally light radiological source detection algorithm based on the exponential weighted moving average (EWMA) and Poisson deviance statistics. Then, we compare the performance of the AF algorithm with other existing source detection and identification algorithms to demonstrate the effectiveness of the AF algorithm. Furthermore, the performance gain of the ensemble algorithms, where the AF algorithm is used in conjunction with other existing algorithms, is explored through thorough benchmarking. Finally, we implement the ensemble of the AF algorithm and an NMF-based source identification algorithm on an off-the-shelf single-board microcontroller to demonstrate the low-power implementation of source detection and identification algorithms.

The rest of this article is organized as follows. Sections II and III introduce the AF algorithm and the ensemble approach for low-power implementation of source detection and identification algorithms. In Section IV, the datasets and existing source detection and identification algorithms, which were used for performance benchmark and power measurement, are introduced. The results of the benchmark and low-power implementation power measurement are introduced in Section V. Finally, in Section VI, we conclude the study, presenting the direction of the future work.

II. ANOMALY FILTER

A. Binning Scheme and Integration Time Selection

The AF algorithm aims to detect anomalies (i.e., non-background spectral features) from time-integrated energy spectra, and therefore, it is of importance to choose an appropriate integration time and binning scheme to optimize the performance of the algorithm. The choice of the optimum integration time depends primarily on the source-to-detector distance, the detector (or source) moving speed, and the detector dwell time. However, these parameters usually cannot be determined *a priori* in many source search scenarios. Thus, in this study, we use the integration time of 1 s for the AF algorithm and other algorithms, which are introduced later in this article.

An energy binning scheme is another important consideration, which can affect the performance of the algorithm. Too many energy bins result in fewer counts registered in each bin. This, in turn, may cause overdispersion of Poisson counts statistics, making it difficult to statistically interpret the results. On the other hand, too few bins may

lead to non-negligible quantization error, and therefore, some important spectral information may be lost. Also, due to the radiation energy dependency of the detector resolution [i.e., full-width at half maximum (FWHM)], the optimum bin width may differ depending on the energy and detector types.

For the AF algorithm and the other algorithms used for benchmarking, the time integrated counts were histogrammed with 128 energy bins ranging between 30 and 3000 keV. The widths of the bins were set to be proportional to the square root of the energy [$(E)^{1/2}$] to ensure that the FWHM of the detector resolution ranges over the same number of bins, independent of the energy. This number of bins was empirically determined to be the minimum number of bins required to ensure no spectral information loss. Note that the number of bins was chosen specifically for a medium resolution detector (i.e., NaI detector), and the optimum number of bins may vary depending on the type of detector used.

B. Background Tracking

Reliable tracking of constantly varying radiation backgrounds is crucial for all radiological source detection and identification algorithms, as background fluctuations may elicit false alarms, and thus, the overall performance of the source detection algorithm may be compromised. Suppose that the true time integrated mean background spectrum at time t is given as a vector form, $\boldsymbol{\mu}_t \in \mathbb{R}_{\geq 0}^d$, where d is the number of bins used for the spectral histogram. Then, a time integrated spectrum, $\mathbf{x}_t \in \mathbb{Z}_{\geq 0}^d$, is subject to Poisson noise as follows:

$$x_{t,i} \sim \text{Poisson}(\mu_{t,i}) \quad \forall i = 1, \dots, d. \quad (1)$$

However, as the true background $\boldsymbol{\mu}_t$ is typically not known *a priori* in most source search scenarios, most source detection and identification algorithms use a variety of techniques to estimate the true background spectrum $\boldsymbol{\mu}_t$. For the AF algorithm, this is done by continuously updating the background estimation using the EWMA. First, the time integrated spectrum \mathbf{x}_t is filtered to suppress the noise present in the spectrum. Since the energy spectra obtained from detectors are subject to Gaussian blur due to the limited energy resolution, it is reasonable to use a Gaussian filter of size proportional to $(E)^{1/2}$ to reduce the variance while preserving the prominent features (e.g., photopeaks). However, we have found that using a simple boxcar filter is not only computationally more efficient, but also just as good in noise suppression. Let $x_{t,i}$ be the counts registered in the i th bin at time t . Then, the moving average is applied as

$$\bar{x}_{t,i} = \frac{1}{w} \sum_{k=i-(w-1)/2}^{i+(w-1)/2} x_{t,k} \quad (2)$$

where w is the width of the boxcar filter, which is often chosen to be an odd number for the sake of symmetry.

Since the bin width is set to be proportional to $(E)^{1/2}$, the width of the filter does not need to vary to account for the energy dependency of detector resolution. For the type of the detector and the binning scheme considered in this study, it was found that the choice of $w = 3$ adequately

reduces the noise without much loss of spectral information. However, the optimal filter size may vary depending on the detection system setup.

With the found filtered spectra, the EWMA update rule for the mean background estimation $\hat{\boldsymbol{\mu}}_t \in \mathbb{R}_{\geq 0}^d$ is given as

$$\hat{\boldsymbol{\mu}}_t = (1 - \alpha)\hat{\boldsymbol{\mu}}_{t-1} + \alpha \bar{\mathbf{x}}_{t-1} / \|\bar{\mathbf{x}}_{t-1}\|_1 \quad (3)$$

where $\hat{\boldsymbol{\mu}}_t$ is the mean background spectrum estimation at time t , α is the EWMA weight, and $\|\cdot\|_1$ is l_1 -norm. The initial background spectrum estimation $\hat{\boldsymbol{\mu}}_0$ can be set using the mean background spectrum measured for a prolonged time. Note that $\hat{\boldsymbol{\mu}}_t$ is normalized, such that $\|\hat{\boldsymbol{\mu}}_t\|_1 = 1$, and thus, it only reflects the spectral shape, neglecting the gross counts. Then, the estimated mean background spectrum is

$$\boldsymbol{\mu}_t \approx \|\mathbf{x}_t\|_1 \hat{\boldsymbol{\mu}}_t. \quad (4)$$

In spite of its simplicity, it has been demonstrated that the EWMA performs as well or even better than more complex background estimation methods, such as Kalman filter [18].

The EWMA weight parameter α determines how much the estimated background reflects the most recently measured spectrum. Therefore, it is a hyperparameter that needs to be tuned differently depending on the background variability of the source search scenario. For example, in this work, $\alpha = 0.02$ was used as it was found to be effective in urban radiological source search scenarios. However, even lower values of α can be used in the case of static source detection systems. Although the EMWA background estimation performs well for slowly varying backgrounds, the estimation might fail when the count contribution from the source increases slowly. For example, when the source-to-detector distance is large, the source contribution may be updated together with the slowly changing background, yielding inaccurate background estimation. To prevent this from happening, a preset background update threshold, above which the background is not updated, is set as a fraction of the anomaly metric threshold. In this study, 80% of the anomaly metric threshold was used for background update threshold.

C. Anomaly Metric

Once the background is estimated, the spectral discrepancy between the filtered spectrum and the EWMA estimated mean background found from (4) is first quantified using unit Poisson deviance. Unit Poisson deviance spectrum measured at time t , \mathbf{d}_t is obtained as

$$\begin{aligned} \mathbf{d}_t &= 2(\log P(\bar{\mathbf{x}}_t; \bar{\mathbf{x}}_t)) - \log P(\bar{\mathbf{x}}_t; \boldsymbol{\mu}_t) \\ &\approx 2(\log P(\bar{\mathbf{x}}_t; \bar{\mathbf{x}}_t)) - \log P(\bar{\mathbf{x}}_t; \|\bar{\mathbf{x}}_t\|_1 \hat{\boldsymbol{\mu}}_t) \\ &= 2(\bar{\mathbf{x}}_t \odot \log(\bar{\mathbf{x}}_t \oslash \|\bar{\mathbf{x}}_t\|_1 \hat{\boldsymbol{\mu}}_t)) - (\bar{\mathbf{x}}_t - \|\bar{\mathbf{x}}_t\|_1 \hat{\boldsymbol{\mu}}_t) \end{aligned} \quad (5)$$

where \odot and \oslash are element-wise (Hadamard) product and division, respectively. Poisson deviance is a goodness-of-fit statistic often used for Poisson regression model selection. Typically, two different models (hypotheses) are compared to choose a model, which best fits the Poisson statistics of the data. In the context of AF algorithm, Poisson deviance is equivalent to likelihood ratio, where likelihoods of the null




Fig. 1. Example of an estimated background and measured spectrum with an $80\text{-}\mu\text{Ci } ^{133}\text{Ba}$ source present in the environment (top) and the unit Poisson deviances across all the spectral bins (bottom). The spectrum was integrated for 1 s, and 128 bins were used. Note that the peak at 356 keV is attributed to the ^{133}Ba source.

hypothesis and the alternative hypothesis are compared. Here, the null hypothesis means that the observed spectrum \mathbf{x}_t is a realized sample from the background spectrum $\boldsymbol{\mu}_t$, and the alternative hypothesis is that the observed \mathbf{x}_t is a sample from \mathbf{x}_t itself. Fig. 1 shows an example of a unit Poisson deviance spectrum, where the discrepancy between the EWMA estimated counts and measured counts in each bin is quantified using unit Poisson deviance. The measured spectrum contains counts from a ^{133}Ba source, which give rise to large unit Poisson deviances in bins around the ^{133}Ba photopeak energy (356 keV). Note that the measured spectrum in Fig. 1 was taken from the augmented dataset, which simulates gamma-ray source search runs in urban environment with a vehicle borne NaI(Tl) detector. More details of the dataset are described in Section IV.

For hypothesis testing, the total deviance is used to measure the discrepancy between two hypothesis

$$d_{\text{tot}} = \sum_{i=1}^n d_i. \quad (6)$$

Since the total deviance is a likelihood-ratio test, it provides a statistically interpretable means of determining anomalous spectra. According to Neyman–Pearson lemma, likelihood-ratio test is the most powerful hypothesis test when the probability distribution functions (PDFs) of the two hypothesis are known [19]. However, note that the alternative hypothesis \mathbf{x}_t does not account for some common attributes of the anomalies arising from radiological sources.

For example, if a measured spectrum contains counts from an unknown radiological source one wishes to find, the full energy peak(s) are the most prominent anomalous features, and thus, the Poisson unit deviances will be high in only a small number of spectral bins. On the other hand, if an anomalous




Fig. 2. Waterfall plots of the measured counts (top left), unit Poisson deviances (top middle), and weighted Poisson deviances (top right) from a simulated 180-s source search run with a ^{133}Ba source injected at 75 s. The run was taken from the augmented dataset described in Section IV. Marked as an ROI is the ^{133}Ba photopeak at 356 keV. The time-series plots of the gross counts (bottom left), total deviances (bottom middle), and l_2 -norm of the weighted Poisson deviance spectra (bottom right) are shown as well. Note that the weighted Poisson deviances can effectively suppress the noise arising from the background fluctuation, while preserving the prominent features from the ^{133}Ba source.

spectrum is registered due to background fluctuation and Poisson noise, the unit Poisson deviances are likely to be randomly distributed across spectral bins.

In addition, if high unit Poisson deviances arise from a radiological source in the environment, it is likely that the high deviances appear in the same bins in the next time-integrated spectrum. However, if the high unit Poisson deviances are from the background noise, they will rarely appear in those same bins again in the next time-integrated spectrum.

Taking into account the aforementioned characteristics of anomalies arising from extraneous radiological sources can greatly improve the sensitivity of the AF algorithm. Thus, instead of using total deviance as an anomaly metric, the obtained unit deviances are further processed, so that the algorithm is more robust to the background noise.

First, the unit Poisson deviances found from (5) are weighted as follows:

$$\mathbf{d}'_t = \mathbf{d}_t \odot \mathbf{w}_t \quad (7)$$

where \mathbf{w}_t is the weighting vector at time t . The weighting vector \mathbf{w}_t is obtained as

$$\mathbf{w}_t = (1 - \beta)\mathbf{w}_{t-1} + \beta \frac{n}{\|\mathbf{d}_t\|_1} \mathbf{d}_t. \quad (8)$$

In other words, the weighting vector is continuously updated with a newly measured spectrum using EWMA. Note that the weight is normalized, such that $\sum_{i=1}^n w_{t,i} = n$, and the initial weights $w_{0,i}$ can be initialized to ones. The hyperparameter β determines how fast the weights change as a new spectrum is registered, and depending on the expected source dwell time, it can be chosen differently. For the benchmark described in Section IV, $\beta = 0.5$ was used, as it was found to be suitable for urban radiological source search scenarios. Since the weight vector follows the most recent unit Poisson deviance spectrum, the weighted deviance becomes increasingly large

when the deviance is continuously high in certain bins. Since continuously high deviance values in certain bins are indicative of non-background radiological sources present in the scene, the deviance weighting scheme can greatly improve the overall performance of the algorithm.

Once the weighted deviance vector is found with (7), l_2 -norm of the weighted deviance vector is used as the anomaly metric as follows:

$$\text{AF}(\mathbf{x}_t) = \|\mathbf{d}'_t\|_2. \quad (9)$$

Total deviance is the sum of all unit deviances and, therefore, is equivalent to l_1 -norm of the deviance vector. Although the total deviance has favorable statistical interpretation as a likelihood-ratio test, using l_2 -norm improves the performance, since l_2 -norm yields a higher alarm metric when unit Poisson deviance spectrum is sparsely high in few bins.

To illustrate the impact of this approach, Fig. 2 compares the use of the following three anomaly metrics: gross counts, total Poisson deviance, and l_2 -norm of the weighted deviance (the AF anomaly metric). Shown in the top row are anomaly metric spectra of a source search run, where the different anomaly metrics are computed for each time and energy bin. Marked as region of interest (ROI) are count bins where the non-background ^{133}Ba source counts are registered. The plots in the bottom row show how the three different anomaly metrics change over the course of the source search run. The gross counts do not reflect any spectral information; therefore, it fails to yield a high anomaly metric when the ^{133}Ba source appears at around 75 s. The total Poisson deviance, on the other hand, gives a relatively high alarm metric when the source is nearby, but it is not robust to background fluctuations. As a result, it is impossible to detect the source without frequent false alarms. Finally, the l_2 -norm of the weighted deviance suppresses the

high Poisson deviances arising from the background fluctuations, yielding a distinctively high anomaly metric at 75 s.

Once the anomaly metric is found from (9), it is then compared with a preset threshold T_{AF} , and an alarm is registered when $AF(\mathbf{x}_t) \geq T_{AF}$. The preset threshold can be empirically determined from background measurements.

III. ENSEMBLE ALGORITHM

In this study, we investigate the use of the ensemble algorithms where the AF algorithm is used together with other existing source detection and identification algorithms. The purpose of the ensemble approach is to improve the performance of the existing algorithms while simultaneously reducing the computational cost and, therefore, power consumption, associated with their implementation.

A. Ensemble Algorithm

Suppose that T_{AF} and T_{alg} are the anomaly metric thresholds of the AF algorithm and a source detection algorithm of choice, respectively. First, an AF alarm metric $AF(\mathbf{x}_t)$ is obtained from a measured spectrum \mathbf{x}_t and compared with the AF threshold. If $AF(\mathbf{x}_t) \geq T_{AF}$, the spectrum \mathbf{x}_t is further processed by the source detection algorithm, and an alarm is registered if $alg(\mathbf{x}_t) \geq T_{alg}$. In other words, measured spectra are first “filtered” by the AF algorithm, with those spectra exceeding the AF thresholds then further inspected by the source detection algorithm.

Since the ensemble algorithm requires two different preset thresholds, T_{AF} and T_{alg} , theoretically, an infinite number of threshold combinations can yield the same overall FPR of the ensemble algorithm. Therefore, it is important to tune the thresholds on a training dataset, so that the desirable results [i.e., better computational efficiency or highest true positive rate (TPR)] can be achieved.

B. Low-Power Implementation and Performance Gain

Because of the low computational cost of the AF algorithm, the computation burden of the ensemble algorithms can be significantly lower than that of standalone source detection and identification algorithm. Furthermore, as demonstrated in Section V, this method not only reduces the overall computational cost, but also improves the detection performance.

Since the AF algorithm is computationally much more efficient than many other source detection and identification algorithms, the reduction in computational cost (the overall power consumption) ultimately depends on the FPR of the AF algorithm (i.e., the rate at which the more computationally intensive algorithm is called). However, setting the FPR too low may negatively impact the overall performance of the ensemble algorithm. Thus, the optimal threshold settings must be carefully chosen depending on the intended use of the algorithm. Heuristically, we have found good overall algorithm performance with the FPR of the AF between 10^{-2} and 10^{-1} s^{-1} .

IV. BENCHMARK

To demonstrate the performance of the AF algorithm and ensemble algorithms, a thorough benchmark was performed against other existing source detection and identification algorithms on two different datasets. In this section, we introduce the dataset used to train and test the algorithms and the existing source detection and identification algorithms used for the benchmarks.

A. Datasets

The first dataset comes from the radiological source search competition held in 2018 and is publicly available [20], [21]. Hence, it will be referred to as the competition dataset hereinafter. It contains Monte Carlo simulated list-mode data, mimicking source search runs with a vehicle-borne 2-in \times 4-in \times 16-in NaI detector. The background data were obtained from a simulated mid-sized U.S. street, populated with objects (e.g., streets, buildings, trees, sidewalk, and air) with different naturally occurring radiological material (NORM) compositions. In addition, the total of six different types of sources (highly enriched uranium (HEU), weapons-grade plutonium (WGPu), ^{99m}Tc , ^{131}I , ^{60}Co , and a combination of ^{99m}Tc and HEU) were injected in some of the runs with or without shielding, at a random location along the street. A more detailed description for the data generation process is available in [22]. For benchmarking, the list-mode data were time-integrated with the integration time of 1 s, and the time corresponding to the spectrum yielding the largest alarm metric was considered to be the time of closest approach to a source. The algorithms were then evaluated using the same scoring rubric used for the radiological source search competition. The rubric accounts for detection, identification, and source location prediction and yields a single score ranging from 0 (worst) to 100 (best). More details on the rubric are described in [23].

The second dataset, which will be referred to as the augmented dataset hereinafter, consists of 4900 source search runs and was prepared as follows; first, the background data were taken from the background only runs of the competition dataset. Then, each background run was modified by injecting one of the 21 different sources with various activities. The detector angular response of the injected source was obtained using Geant4, and the source signal from a 2” \times 4” \times 16” NaI detector was simulated accordingly. In all runs, the detector moving speed of 5 m/s and the distance of closest approach between the detector and a source of 10 m were assumed. Table I shows the list of source types and activities used for the dataset. The same dataset was also used for benchmarking by Bilton *et al.* [14], in which a more detailed data generation procedure is described. The algorithms were assessed on the augmented dataset in terms of receiver operating characteristic (ROC) curves and minimum detectable amount (MDA).

B. Algorithms

To benchmark the performance of the AF algorithm and the ensemble approach, the following source detection and

TABLE I

LIST OF RADIOLOGICAL ISOTOPES USED IN THE AUGMENTED DATASET

Isotope	Activity (μCi)	Isotope	Activity (μCi)	Isotope	Activity (μCi)
^{241}Am	10-210	^{137}Cs	10-85	^{177}Lu	140-520
^{198}Au	5-75	^{152}Eu	60-180	^{54}Mn	20-90
^{133}Ba	40-140	^{67}Ga	70-170	^{124}Sb	20-90
^{82}Br	20-100	^{123}I	25-95	^{46}Sc	15-80
^{57}Co	20-150	^{131}I	25-125	^{75}Se	15-90
^{60}Co	15-105	^{111}In	25-105	^{113}Sn	40-105
^{51}Cr	300-520	^{192}Ir	25-105	^{201}Tl	50-150

identification algorithms were implemented, trained, and tested on the aforementioned dataset: gross counts $K\sigma$, spectral anomaly detection (SAD), nuisance-rejection spectral comparison ratio anomaly detection (N-SCRAD) algorithm, censored energy window (CEW), and NMF-based source identification (NMF-ID) algorithm. The $K\sigma$, SAD, N-SCRAD, and CEW algorithms were implemented with the radiological anomaly detection and identification (RADAI) Python library [24], and the NMF-ID algorithm was implemented with the Berkeley anomaly detection (BAD) Python library [25].

1) *Gross Counts $K\sigma$ Algorithm*: Gross counts $K\sigma$ algorithm ($K\sigma$) is one of the rudimentary radiological anomaly detection algorithms, which do not account for any spectral information. The algorithm uses the z -score of the gross counts from a spectrum \mathbf{x} as an anomaly metric as follows:

$$K\sigma(\mathbf{x}) = \frac{\|\mathbf{x}\|_1 - \mu}{\sigma} \quad (10)$$

where μ and σ are the mean and standard deviation of the gross counts, respectively, and are found from background-only training data. Then, the measured z -score is compared with the preset threshold to identify anomalous spectra.

2) *Spectral Anomaly Detection*: SAD projects measured spectra onto a PCA-based orthonormal subspace and then reconstructs the spectra using the orthonormal components. Then, the anomaly metric is obtained by measuring the reconstruction error [26], [27]. Suppose the orthonormal subspace is denoted as \mathbf{U} . Then, the anomaly metric on a measured spectrum \mathbf{x} is

$$\text{SAD}(\mathbf{x}) = \|\mathbf{x} - \mathbf{U}\mathbf{U}^T\mathbf{x}\|_2. \quad (11)$$

The orthonormal subspace \mathbf{U} is often found using PCA on the background only training data.

3) *N-SCRAD Algorithm*: The N-SCRAD algorithm uses spectral comparison ratios (SCRs) defined for a set of coarse spectral bins [18], [28], [29], [30], [31]. The SCR defined for bin 1 and bin j at time index t , and α_{1j_t} is defined as follows:

$$\alpha_{1j_t} = x_{1_t} - \frac{\bar{B}_{1_t}}{\bar{B}_{j_t}} x_{j_t} \quad (12)$$

where x_{1_t} and x_{j_t} are the number of counts registered in bin 1 and j at time t , and \bar{B}_{1_t} and \bar{B}_{j_t} are reference counts in bin 1 and j . Assuming that $(\bar{B}_{1_t}/\bar{B}_{j_t})x_{j_t}$ is an unbiased estimator of x_{1_t} , the SCR α_{1j_t} follows a zero-mean normal distribution. Therefore, the SCR spectrum $\boldsymbol{\alpha}_t$, which consists of $N - 1$ SCRs, is a zero-mean multivariate normal distribution with a covariance matrix \mathbf{S}_t . Hence, any non-noise

spectral change in the count spectra will cause the SCR vector $\boldsymbol{\alpha}_t$ to deviate from the normal distribution, leading to a larger Mahalanobis distance $(\boldsymbol{\alpha}_t^T \mathbf{S}_t^{-1} \boldsymbol{\alpha}_t)^{(1/2)}$. To discriminate the deviation from benign sources and threat sources, the SCR vector $\boldsymbol{\alpha}_t$ is projected onto the subspace spanned by the SCRs of the nuisance sources, such as K-, U-, and Th-related naturally occurring radioactive materials (NORMs). The projection matrix \mathbf{P} is given by

$$\mathbf{P} = \bar{\mathbf{A}}(\bar{\mathbf{A}}^T \mathbf{S}_t^{-1} \bar{\mathbf{A}})^{-1} (\bar{\mathbf{A}}^T \mathbf{S}_t^{-1}) \quad (13)$$

where the matrix $\bar{\mathbf{A}}$ has SCRs of the nuisances as columns. The anomaly metric is then found as

$$\text{N-SCRAD}(\boldsymbol{\alpha}_t) = (\boldsymbol{\alpha}_t^T \mathbf{S}_t^{-1} (\mathbf{I} - \mathbf{P}) \boldsymbol{\alpha}_t)^{\frac{1}{2}}. \quad (14)$$

The time-dependent reference count spectrum $\bar{\mathbf{B}}_t$ and covariance matrix \mathbf{S}_t can be estimated using Kalman filter or EWMA. For more detailed description of the algorithm, readers are encouraged to refer to [18], [28], [29], [30], and [31]. For this study, the N-SCRAD algorithm implemented in the RADAI package was used, and it should be noted that this version of implementation may differ from the original implementation described in the literature. First, the implemented N-SCRAD algorithm uses a genetic optimization algorithm for the binning scheme optimization. Furthermore, since the nuisance spectra were not readily available, three background basis vectors learned from the NMF-based source identification algorithm were used as nuisances. For benchmarking, 15 bins were used, and the time-dependent parameters, $\bar{\mathbf{B}}_t$ and \mathbf{S}_t , were tracked using EWMA.

4) *Censored Energy Window*: CEW measures the estimated signal-to-noise ratio (SNR) within selected energy windows to identify anomalous spectra [13]. First, an indicator vector $\hat{\mathbf{w}} = \{0, 1\}^d$, which defines the energy windows within which the SNR is calculated, is obtained by solving the following optimization problem:

$$\hat{\mathbf{w}} = \underset{\mathbf{w}}{\text{argmax}} \frac{\mathbf{w}^T \mathbf{s}}{\sqrt{\mathbf{w}^T \mathbf{b}}} \quad (15)$$

where \mathbf{s} is a source template, and \mathbf{b} is a background only training spectrum. The anomaly metric is then obtained by calculating the SNR as follows:

$$\text{CEW}(\mathbf{x}) = \frac{(\mathbf{w} - \hat{\boldsymbol{\theta}}^T) \mathbf{x}}{\sqrt{\hat{\boldsymbol{\theta}}^T \mathbf{x}}}. \quad (16)$$

Here, $\hat{\boldsymbol{\theta}}$ is a regression vector used to estimate the background counts within the window defined by \mathbf{w} . The regression vector is typically learned from background only training data using ridge regression.

5) *NMF-ID Algorithm*: NMF-ID is a source identification algorithm based on the NMF dimensionality reduction technique. The algorithm first decomposes a data matrix with background only spectra, $\mathbf{X} \in \mathbb{R}^{n \times d}$, into non-negative components $\mathbf{V} \in \mathbb{R}^{n \times k}$ and weights $\mathbf{A} \in \mathbb{R}^{k \times d}$, such that $\mathbf{X} \approx \hat{\mathbf{X}} = \mathbf{A}\mathbf{V}$ [11], [12], [32]. Note that $k \ll d$, and therefore, $\hat{\mathbf{X}}$ is a low-rank approximation of \mathbf{X} . The decomposition is

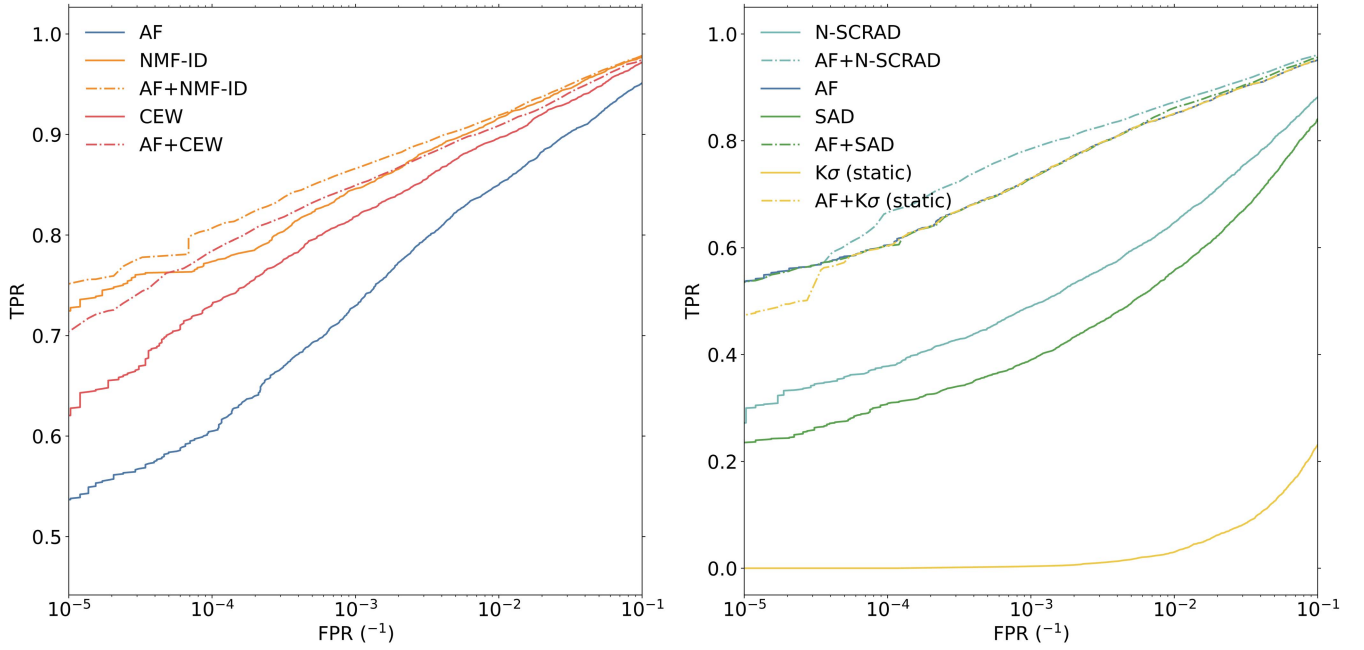


Fig. 3. ROC curves for the AF, the source identification algorithms, and ensemble algorithms (left). AF, the source detection algorithms, and ensemble algorithms (right). The dotted lines represent the ROC curves for the ensemble algorithms where the AF algorithm was used in conjunction with various detection and identification algorithms.

found by minimizing the following Poisson loss:

$$\ln P(\mathbf{X}|\hat{\mathbf{X}}) = \sum_{i=1}^n \sum_{j=1}^d -\hat{x}_{ij} + x_{ij} \ln \hat{x}_{ij} - \ln x_{ij}! \quad (17)$$

Once the non-negative components for background are found, the component matrix \mathbf{V} is augmented with a source template \mathbf{s} , yielding $\mathbf{V}_s = [\mathbf{V}^T | \mathbf{s}]^T$. Then, a given test spectrum is fit to both \mathbf{V} and \mathbf{V}_s , yielding fit spectra $\hat{\mathbf{x}}$ and $\hat{\mathbf{x}}_s$. Finally, the fit spectra are used to find the anomaly metric as follows:

$$\text{NMF-ID}(\mathbf{x}) = -2(\ln P(\mathbf{x}|\hat{\mathbf{x}}) - \ln P(\mathbf{x}|\hat{\mathbf{x}}_s)) \quad (18)$$

where $P(\mathbf{x}|\hat{\mathbf{x}})$ is a Poisson likelihood, and therefore, (18) is a likelihood-ratio test. The found anomaly metric is compared with a preset threshold to determine anomalous spectra.

V. RESULTS

In this section, first, we compare the performance of AF and ensemble algorithms on the datasets introduced in Section IV using ROC curves, MDA, and data competition scores. Then, the power measurement results from the ensemble algorithm implementations on NVIDIA Jetson Nano [33] and Particle Photon microcontroller [34] to demonstrate the low-power implementation capability of the ensemble approach.

A. ROC Curves

To compare the performance of the source detection and identification algorithms, all algorithms were trained and tested on the augmented dataset, and the ROC curves were produced, as illustrated in Fig. 3. For the anomaly detection algorithms (AF, CEW, N-SCRAD, $K\sigma$, and SAD), the ROC curves were

obtained by evaluating the TPR while varying the thresholds of the trained algorithms. For the source identification algorithms (NMF-ID and CEW), the results from 21 different models, each of which was trained on a single source, were aggregated to yield the overall FPR and TPR. In addition, each source detection and identification algorithm was used in ensemble with the AF algorithm, and the ROC curves were drawn with dotted lines to compare the performance of the ensemble algorithms with the standalone algorithms. As stated before, for ensemble algorithms, an infinite number of threshold combinations can yield the same FPR. Therefore, for a given FPR, the threshold combination resulting in the best TPR was chosen to draw the ROC curves.

As can be seen in Fig. 3, the NMF-ID algorithm performs the best, among the source identification algorithms, followed by the CEW algorithms. For the source detection algorithms, the AF algorithm achieved the best results, followed by the N-SCRAD, SAD, SPRT, and $K\sigma$ algorithms. Generally speaking, identification algorithms perform better than detection algorithms, as they are trained for specific radioisotopes used in the simulation. In addition, all ensemble combinations surpassed the performance of standalone algorithms, indicating the effectiveness of the ensemble approach.

B. MDA₉₅ Comparison

Although the ROC curves can well characterize the algorithm performance as a function of FPR, it is often helpful to use more physically interpretable metrics to compare them. Thus, in this study, we also use MDA₉₅ to compare the performance of different detection, identification, and ensemble algorithms. For the sake of brevity, we only estimate MDA₉₅

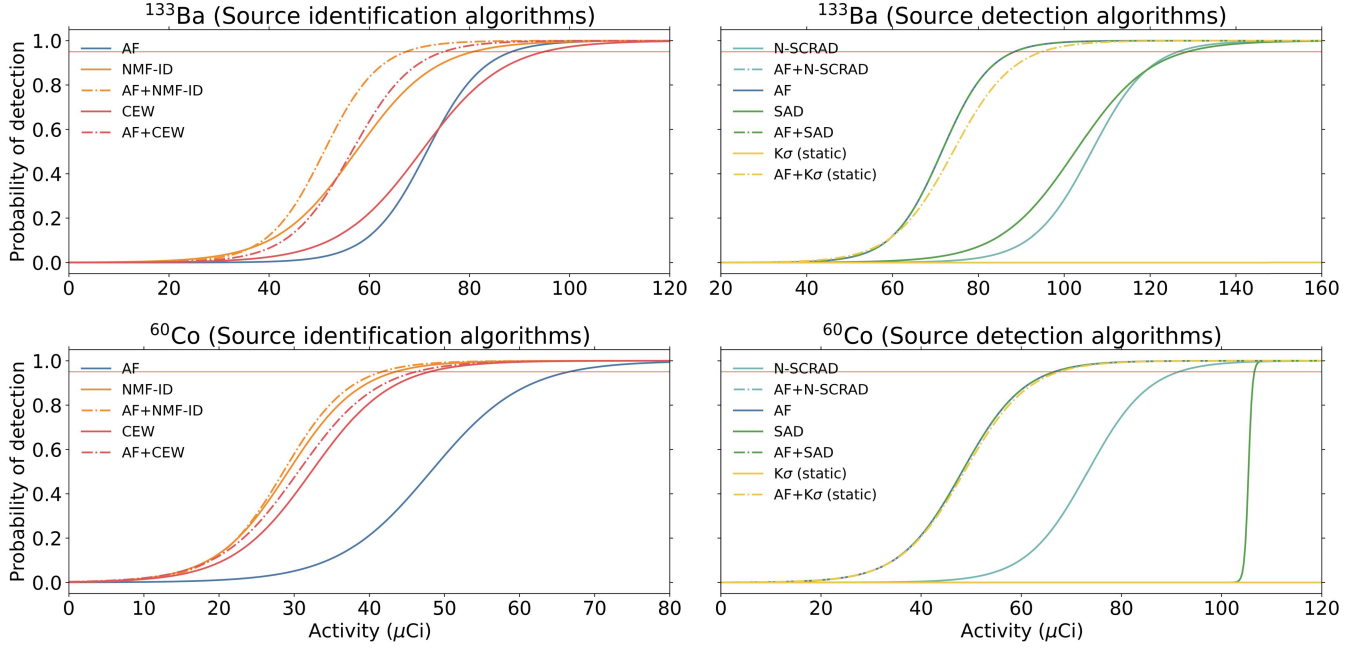


Fig. 4. Probability of detection curves for ^{133}Ba (top) and ^{60}Co (bottom). The 95% probability of detection line is drawn in red. Note that the curves for ensemble algorithms are drawn in dotted lines. Figures on the left are for source identification algorithms, whereas figures on the right are for source detection algorithms.

for ^{133}Ba and ^{60}Co sources used in the augmented dataset. The estimation of MDA_{95} was done as follows: first, using the background only dataset, the algorithm thresholds were set to yield the same FPR of $1/8 \text{ h}^{-1}$. Then, the source injected augmented dataset was used to get the probability of detection P_D at different source activities, and the obtained probability of detection was fit to a sigmoid function

$$P_D(x) = \frac{1}{1 + e^{-\frac{(x-\mu)}{w}}}. \quad (19)$$

Once the fit function is obtained, MDA_{95} was obtained by setting $P_D(x) = 0.95$ and solving (19) for x .

Fig. 4 illustrates the fit probability of detection curves, and Table II shows the MDA_{95} found for detection, identification, and ensemble algorithms. Generally speaking, as demonstrated with the ROC curves, identification algorithms outperform the source detection algorithms, while the AF algorithm performs the best among the source detection algorithms. Furthermore, all ensemble algorithms achieved lower MDA_{95} than the standalone algorithms, demonstrating the performance improvement of the ensemble approach.

C. Data Competition Score

To benchmark the source identification algorithms considered in this study, the algorithms were trained and tested on the competition dataset, and scored using the same scoring rubric used in the data competition. Since the scoring rubric is based on identification of the radiological sources, only the identification algorithms and the ensemble algorithms were used and compared. Fig. 5 shows the best scores obtained from a single identification algorithm, as well as the scores achieved

TABLE II
 MDA_{95} FOR BARIUM-133 AND COBALT-60

Algorithm	MDA_{95} (μCi)	
	^{133}Ba	^{60}Co
AF	88.4 ± 3.6	66.6 ± 3.3
NMF-ID	80.7 ± 2.8	43.4 ± 1.7
AF+NMF-ID	67.3 ± 1.8	41.6 ± 1.7
CEW	94.8 ± 3.6	48.0 ± 1.2
AF+CEW	74.8 ± 2.0	46.2 ± 1.5
N-SCRAD	126.4 ± 3.2	91.4 ± 1.6
AF+N-SCRAD	88.4 ± 3.6	66.6 ± 3.3
SAD	128.6 ± 4.0	105.0 ± 3.3
AF+SAD	88.3 ± 3.6	66.6 ± 3.3
$K\sigma$	N/A ^a	N/A ^a
AF+ $K\sigma$	94.8 ± 3.5	67.5 ± 3.4

^a The MDA_{95} was too large and couldn't be estimated reliably.

by ensemble algorithms. The score is ultimately dependent on the FPR settings of the algorithms, and therefore, in Fig. 5, the score of the ensemble algorithms is shown as contour plots. Note that the best score the ensemble algorithms can achieve is higher than the scores achieved with the standalone algorithms.

D. Low-Power Implementation

To demonstrate the power reduction capability of the ensemble approach, the AF and NMF-ID algorithms were implemented on both the Nvidia Jetson Nano single-board computer and the Particle Photon microcontroller, and the power consumption was measured. The NMF-ID algorithm was chosen for this measurement, as it exhibits the best performance among all algorithms we used for benchmark. However, it should be noted that the ensemble approach




Fig. 5. Data competition score of AF + NMF-ID (left) and AF + CEW (right) ensemble algorithms as a function of AF and algorithm FPRs. The best score achieved from each ensemble algorithm is marked with a red “X.” The red contour lines represent the best scores achieved from the standalone algorithms.




Fig. 6. Overall TPR of the AF + NMF-ID ensemble algorithm on the augmented dataset as a function of AF and NMF-ID FPR settings. The contour lines drawn in black represent the overall FPR of the ensemble algorithm. Note that even for the same overall FPR, the overall TPR varies depending on the individual algorithm FPR settings.

described in this article can be generally applied to any other source detection and identification algorithms.

As the overall computational burden of the ensemble algorithm is a function of the AF algorithm FPR (i.e., NMF-ID calling rate), it is crucial to first characterize the overall performance of the ensemble algorithm as a function of individual algorithm threshold settings. Fig. 6 shows the overall TPR and FPR of the ensemble algorithm on the augmented dataset as a function of AF and NMF-ID FPR settings. For the power measurement, the AF and NMF-ID FPRs were set to 1×10^{-2} and $8 \times 10^{-4} \text{ s}^{-1}$, as the FPR setting yields the best TPR for the overall FPR of $1/8 \text{ h}^{-1}$. To mimic the real-world

TABLE III

POWER CONSUMPTION OF THE AF, NMF-ID, AND ENSEMBLE ALGORITHM ON THE NVIDIA JETSON NANO

Algorithm	Overall power consumption (mW)	Algorithm power consumption (mW) ^a
Dummy ^b	225.70 ± 1.36	0
AF	226.13 ± 2.71	0.43 ± 0.13
NMF-ID	660.49 ± 3.20	434.79 ± 3.48
AF+NMF-ID	228.82 ± 2.90	3.11 ± 0.20

^a Dummy algorithm power consumption is subtracted from the power consumption of each algorithm to compare the power required to run each algorithm.

^b Dummy algorithm only reads the spectrum data and does not process it. Therefore it serves to measure the base power required to run the computing machine.

measurement scenario, the 1-s time integrated spectrum was read every second and processed by algorithms in real time.

Table III shows the results of the power consumption measurements on the Nvidia Jetson Nano with its Quad-core ARM Cortex-A57 MPCore processor. Note that the dummy algorithm only reads the data and does not process it and, therefore, represents the power required to run the necessary electronics for the system. Therefore, the “algorithm power consumption” in Table III represents the additional amount of power required to run the algorithms only. The power measurement results suggest that while using the AF algorithm alone does not require much power (0.43 mW), running the NMF-ID algorithm is computationally taxing and much more power consuming (434.79 mW). On the other hand, the ensemble approach could reduce the algorithm power consumption by two orders of magnitude (3.11 mW), while improving the overall performance.

Although the ensemble approach can significantly reduce the power consumption, the overall power consumption is not only determined by algorithms used, but also computing platforms and processors on which the algorithms run. Hence, to study the impact of changing the computing platform, another power consumption measurement was done by running the ensemble algorithm (AF + NMF-ID) on a Particle

TABLE IV

POWER CONSUMPTION OF THE AF, NMF-ID, AND ENSEMBLE ALGORITHM ON THE PARTICLE PHOTON MICROCONTROLLER

Algorithm	Average power consumption (mW)	Algorithm power consumption (mW) ^a
Dummy ^b	46.5±0.3	0
AF	46.7±0.2	0.2±0.1
NMF-ID	N/A ^c	N/A ^c
AF+NMF-ID	52.7±0.2	6.2±0.4

^a Dummy algorithm power consumption is subtracted from the power consumption of each algorithm to compare the power required to run each algorithm.

^b Dummy algorithm only reads the spectrum data and does not process it. Therefore it serves to measure the base power required to run the computing machine.

^c The standalone NMF-ID could not run in real time on Particle Photon due to the high computational cost.

Photon microcontroller. The measurement was performed on the competition dataset with the individual algorithm FPR settings being 9.1×10^{-2} and $2.3 \times 10^{-2} \text{ s}^{-1}$ for AF and NMF-ID algorithms, respectively. Note that the FPR settings were chosen to yield the best ensemble algorithm competition score, as illustrated in Fig. 5. Table IV shows the results of the measurement. Note that the base power consumption (i.e., dummy algorithm power consumption) of Particle Photon (46.5 mW) is much lower than that of NVIDIA Jetson Nano (225.70 mW). Also, while the NMF-ID power consumption could not be measured due to the computational cost of the NMF-ID algorithm, the ensemble algorithm could achieve the power consumption as low as 52.7 mW. Note that, theoretically, the power consumption could be even more reduced if the AF FPR is set lower, at the expense of slightly lower competition score.

VI. CONCLUSION

Accurate and rapid detection of abnormal radiological sources in a constantly fluctuating background is a difficult problem, and although a variety of algorithms have been proposed for the challenge, many of them are computationally so expensive that they cannot be easily deployed in power-limited detection scenarios.

In this study, we investigated the use of ensemble algorithms to achieve low-power implementation of radiological source detection and identification algorithms. First, we introduced the AF algorithm: a computationally light, yet effective source detection algorithm based on the EWMA and Poisson deviance statistics. Then, we proposed an ensemble approach where the AF algorithm is first used to filter some potentially anomalous spectra, and those filtered spectra are further processed by other existing source detection and identification algorithms to reduce the computational cost and, hence, power consumption of the overall algorithm.

We carried out two power measurement experiments on commercial off-the-shelf computing platforms where the AF algorithm was used together with NMF-ID algorithm to demonstrate the power consumption reduction capability of the ensemble approach. The results suggest that the power consumption of the AF algorithm alone is very low, and therefore, the overall power consumption ultimately depends on the FPR of the AF algorithm. In our experiments with Nvidia Jetson Nano, we have found that the ensemble approach

can significantly reduce the algorithm power consumption up to a factor of 100, while improving the sensitivity of the algorithms.

Through thorough benchmark with two different datasets and various algorithms, we have shown that the proposed AF algorithm can be used as a powerful standalone anomaly detection algorithm. Moreover, the proposed ensemble algorithm can not only achieve low-power implementation, but also attain improved detection performance, as demonstrated in Figs. 3–5.

We emphasize that the ensemble approach described in this study can be generically applied to any source detection or identification algorithms. In addition, the proposed approach to low-power implementation can be easily implemented on many off-the-shelf computing platforms, providing great flexibility in algorithm implementation. Moreover, comparing the Tables III and IV suggests that the overall power consumption can be further lowered once the ensemble algorithm is implemented on more power efficient computing platforms, such as FPGAs.

REFERENCES

- [1] 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: <http://aip.scitation.org/doi/10.1063/1.3207769>
- [2] B. Geelhood, J. Ely, R. Hansen, R. Kouzes, J. Schewpe, and R. Warner, "Overview of portal monitoring at border crossings," in *Proc. IEEE Nucl. Sci. Symp. Conf. Rec.*, vol. 1, Portland, OR, USA: IEEE, Oct. 2003, pp. 513–517, [Online]. Available: <http://ieeexplore.ieee.org/document/1352095/>
- [3] L. M. Wein, A. H. Wilkins, M. Baveja, and S. E. Flynn, "Preventing the importation of illicit nuclear materials in shipping containers," *Risk Anal.*, vol. 26, no. 5, pp. 1377–1393, Oct. 2006. [Online]. Available: <https://onlinelibrary.wiley.com/doi/10.1111/j.1539-6924.2006.00817.x>
- [4] International Atomic Energy Agency, *Accident Monitoring Systems for Nuclear Power Plants*. Vienna, Austria: International Atomic Energy Agency, 2015.
- [5] R. Engelbrecht, "Environmental radioactivity monitoring," in *Handbook of Radioactivity Analysis*, vol. 2. Amsterdam, The Netherlands: Elsevier, 2020, pp. 1–40. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/B9780128143957000015>
- [6] R. L. Kathren, *Radioactivity in the Environment: Sources, Distribution and Surveillance*. Chur, Switzerland: Academic, 1984.
- [7] S. M. L. Hardie and I. G. McKinley, "Fukushima remediation: Status and overview of future plans," *J. Environ. Radioactivity*, vol. 133, pp. 75–85, Jul. 2014. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0265931X13001847>
- [8] A. Blázquez-García, A. Conde, U. Mori, and J. A. Lozano, "A review on outlier/anomaly detection in time series data," *ACM Comput. Surv.*, vol. 54, no. 3, pp. 1–33, Apr. 2022. [Online]. Available: <https://dl.acm.org/doi/10.1145/3444690>
- [9] D. Boardman, M. Reinhard, and A. Flynn, "Principal component analysis of gamma-ray spectra for radiation portal monitors," *IEEE Trans. Nucl. Sci.*, vol. 59, no. 1, pp. 154–160, Feb. 2012. [Online]. Available: <http://ieeexplore.ieee.org/document/6138889/>
- [10] 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. [Online]. Available: <http://ieeexplore.ieee.org/document/1645055/>
- [11] 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/document/8673769/>
- [12] 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/document/9025259/>

- [13] K. Nelson and S. Labov, "Aggregation of mobile radiation data," Lawrence Livermore Nat. Lab., Livermore, CA, USA, Tech. Rep., 2010.
- [14] K. J. Bilton, T. H. Y. Joshi, M. S. Bandstra, J. C. Curtis, D. Hellfeld, and K. Vetter, "Neural network approaches for mobile spectroscopic gamma-ray source detection," *J. Nucl. Eng.*, vol. 2, no. 2, pp. 190–206, May 2021. [Online]. Available: <https://www.mdpi.com/2673-4362/2/2/18>
- [15] M. Kamuda, J. Stinnett, and C. Sullivan, "Automated isotope identification algorithm using artificial neural networks," *IEEE Trans. Nucl. Sci.*, vol. 64, no. 7, pp. 1858–1864, Jul. 2017. [Online]. Available: <http://ieeexplore.ieee.org/document/7898423/>
- [16] G. R. Romanchek, Z. Liu, and S. Abbaszadeh, "Kernel-based Gaussian process for anomaly detection in sparse gamma-ray data," *PLoS ONE*, vol. 15, no. 1, Jan. 2020, Art. no. e0228048. [Online]. Available: <https://dx.plos.org/10.1371/journal.pone.0228048>
- [17] X. Huang *et al.*, "An FPGA implementation of convolutional spiking neural networks for radioisotope identification," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, May 2021, pp. 1–5.
- [18] D. M. Pfund, K. D. Jarman, B. D. Milbrath, S. D. Kiff, and D. E. Sidor, "Low count anomaly detection at large standoff distances," *IEEE Trans. Nucl. Sci.*, vol. 57, no. 1, pp. 309–316, Feb. 2010. [Online]. Available: <http://ieeexplore.ieee.org/document/5410021/>
- [19] J. Neyman and E. S. Pearson, "IX. On the problem of the most efficient tests of statistical hypotheses," *Philos. Trans. Roy. Soc. London A, Containing Papers Math. Phys. Character*, vol. 231, nos. 694–706, pp. 289–337, Feb. 1933. [Online]. Available: <https://royalsocietypublishing.org/doi/10.1098/rsta.1933.0009>
- [20] J. M. Ghawaly *et al.*, "Data for training and testing radiation detection algorithms in an urban environment," *Sci. Data*, vol. 7, no. 1, p. 328, Dec. 2020. [Online]. Available: <http://www.nature.com/articles/s41597-020-00672-2>
- [21] A. Nicholson *et al.* (2020). *Data for Training and Testing Radiation Detection Algorithms in an Urban Environment*. [Online]. Available: <https://www.osti.gov/servlets/purl/1597414/>
- [22] A. D. Nicholson, D. E. Peplow, J. M. Ghawaly, M. J. Willis, and D. E. Archer, "Generation of synthetic data for a radiation detection algorithm competition," *IEEE Trans. Nucl. Sci.*, vol. 67, no. 8, pp. 1968–1975, Aug. 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9121756/>
- [23] (2019). *Detecting Radiological Threats in Urban Areas*. [Online]. Available: <https://www.topcoder.com/challenges/30085346>
- [24] T. Joshi, B. Quiter, J. Curtis, M. Bandstra, and J. Ghawaly, "Radiological anomaly detection and identification (RADAI) v1.0," Jul. 2021. [Online]. Available: <https://www.osti.gov/servlets/purl/1821386> and <https://www.osti.gov/biblio/1821386>, doi: [10.11578/dc.20210921.3](https://doi.org/10.11578/dc.20210921.3).
- [25] J. Curtis, M. Bandstra, T. Joshi, and K. Bilton, "Berkeley anomaly detection (BAD) v1.0," Lawrence Berkeley Nat. Lab. (LBNL), Berkeley, CA, USA, Tech. Rep. BAD v1.0, Dec. 2018. [Online]. Available: <https://www.osti.gov/biblio/1492928-berkeley-anomaly-detection-bad-v1>
- [26] P. Tandon, P. Huggins, R. Maclachlan, A. Dubrawski, K. Nelson, and S. Labov, "Detection of radioactive source using Bayesian aggregation of data from mobile spectrometers," ResearchGate, 2015. [Online]. Available: https://www.researchgate.net/publication/275584746_Detection_of_Radioactive_Source_Using_Bayesian_Aggregation_of_Data_from_Mobile_Spectrometers, doi: [10.13140/RG.2.1.2050.7049](https://doi.org/10.13140/RG.2.1.2050.7049).
- [27] P. Tandon, "Bayesian aggregation of evidence for detection and characterization of patterns in multiple noisy observations," *AI Matters*, vol. 2, no. 4, pp. 25–26, Jun. 2016.
- [28] 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: <http://ieeexplore.ieee.org/document/4291778/>
- [29] 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," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 784, pp. 339–345, Jun. 2015. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0168900215000716>
- [30] 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://linkinghub.elsevier.com/retrieve/pii/S0969804315304024>
- [31] D. M. Pfund, "Radiation anomaly detection and classification with Bayes model selection," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 904, pp. 188–194, Oct. 2018. [Online]. Available: <https://linkinghub.elsevier.com/retrieve/pii/S0168900218308829>
- [32] K. J. Bilton, "Data-driven approaches to spectral gamma-ray source detection and identification," Ph.D. dissertation, Dept. Nucl. Eng., UC Berkeley, Berkeley, CA, USA, 2020. [Online]. Available: <https://escholarship.org/uc/item/59k407r4>
- [33] *NVIDIA Jetson Nano Developer Kit | NVIDIA Developer*. Accessed: Aug. 5, 2021. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>
- [34] *Particle*. Accessed: Aug. 5, 2021. [Online]. Available: <https://docs.particle.io/photon/>