



Research paper

Low cost, mobile radiation anomaly detection with deep adversarial auto encoders on the edge

Charles Sayre^{a,*}, William Bjorndahl^b, Eric C. Larson^a, Joseph Camp^b, Rodolfo Rodriguez-Davila^c, Manuel Quevedo-Lopez^c, Bruce Gnade^c^a Computer Science, Southern Methodist University, Dallas, 75205, TX, USA^b Electrical and Computer Engineering, Southern Methodist University, Dallas, 75205, TX, USA^c Department of Materials Science and Engineering, University of Texas at Dallas, Dallas, 75080, TX, USA

ARTICLE INFO

Keywords:

Anomaly detection
Adversarial networks
Radiation sensing

ABSTRACT

The detection and localization of radiation sources using low-cost, mobile detectors is a challenging application, necessitating new research into sensing devices and detection algorithms. While new sensing that employs small detectors for detecting γ -rays has emerged, the decreased sensitivity of the sensor makes it challenging to maintain reliability compared to larger detectors. Machine learning could be a viable method for enhancing sensitivity by classifying background radiation spectra from anomalous spectra, but this approach can struggle to identify novel radioactive sources or identify sources in dynamic background environments. To address these challenges, we propose the use of adversarial auto encoders (AAEs) for anomaly detection in radiation sensing systems. With the use of our AAE architecture, we eliminate the need for obtaining examples of radiation anomalies for training data and increase the resilience of the sensing when the background radiation is dynamic. We evaluate the system in various contexts using a custom designed detector, showing the AAE model generalizes to various locations and radiation sources. We also show a real-time field test with the detection system in both handheld and drone mounted testing.

1. Introduction

Numerous agencies and groups, such as the National Academy of Engineering and the National Science Foundation, have designated the prevention of nuclear terror as a “grand challenge” in engineering. At the heart of the problem is the detection of the unlawful movement of radioactive materials. However, this is a complex problem because radiation detection is expensive and easy for an adversary to recognize and avoid. In addition, its accuracy and efficiency depends significantly on multiple parameters, such as distance with respect to the source, energy, background, and shielding, among others.

Making radiation detectors smaller and cheaper while maintaining performance is a key focus in national security, where the United States Department of Homeland Security (DHS) has sought initiatives that can rapidly deploy specialized solutions for a given area. Presently radiation detectors used by the DHS consist of large stationary devices (Kouzes, 2010) at ports-of-entry designed to handle passenger vehicles or handheld devices (Milbrath et al., 2007) designed to be used manually by an operator. While these devices are useful for their given role, they do not have the flexibility to be used in any given task. Reducing the size of the detectors obviously improves the device’s flexibility

but comes at the cost of performance in radiation detection (El Hamli et al., 2022). This is where machine learning can come into play, and these smaller but less sensitive detectors can be utilized to create a system that is both flexible and accurate. In our work, we fabricate low-cost cesium iodide, thallium-doped (CsI(Tl))-based γ -ray detector systems which use machine learning of the sensed data to increase its sensitivity while maintaining low cost and portability.

Previous work using machine learning for radiation detection has focused on both supervised and unsupervised methods. Supervised learning for radiation detection has primarily focused on the use of support vector machine (Morgan et al., 2022) and multi-layer perceptrons (Kangas et al., 2008). While unsupervised learning has primarily focused on the use of convolutional auto encoders (Kangas et al., 2008). In this work, we augment previous research using an Adversarial Auto encoder (AAE), surpassing state-of-the-art performance in a range of tasks. An early version of this model was also investigated on weak radiation sources (Sayre et al., 2023). In this work, we design several scenarios for data collection that characterize the viability of the approach in a real mobile application and refine the model for superior performance. Our model successfully generalizes to unseen radiation

* Corresponding author.

E-mail address: sayrec@smu.edu (C. Sayre).

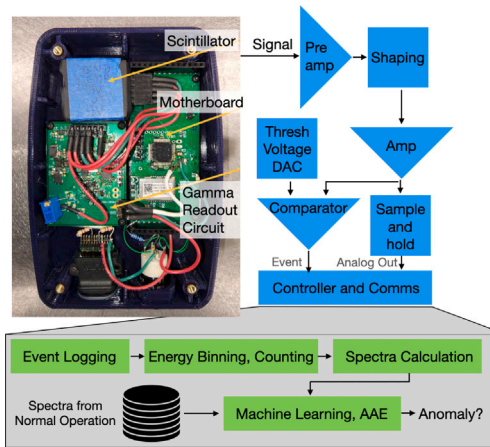


Fig. 1. The custom detector used in the project and the corresponding circuit diagram. The output energy levels are used to construct a spectrum of γ -ray energies that feed into our machine learning anomaly detection algorithm.

sources with remarkable sensitivity. Throughout our experiments, we aim to answer the following research questions:

1. When dealing with diverse background radiation levels, how does the performance of our AAE model in detecting anomalies and avoiding false alarms compare to other anomaly detection systems?
2. How does our AAE model fare against a competing autoencoder when tested with a simulated real-world dataset?
3. In a real-world field test conducted in a mobile environment, can our AAE model effectively identify strong, previously unseen sources of radiation?

Each research question is designed to investigate if our AAE model can generalize outside the data that it is trained with and if the model can detect various sources of radiation. While the complete task of anomaly detection would include identification of potentially hazardous materials, especially compared to more common isotopes used in applications like medicine, here we focus exclusively on detecting any radioactive material that raises radiation above the background. This serves as a screening that there is something further to investigate in the area for radioactive material. We design the memory footprint and complexity of our model such that it is can be run on a small microcomputer, and therefore can easily be connected to a detector for use in a handheld configuration or mounted to a drone. This kind of adaptability could prove very useful to radiation anomaly detection in a variety of situations—for example, when detectors are mounted to a swarm of drones seeking a radiation source.

1.1. Background in radiation sensing

Traditional radiation sensing is compounded by several issues including: (1) the intensity of the emitted radiation falls off at least as $\frac{1}{r^2}$, where r is the distance between the source and the detector; (2) the emission signature can be reduced further by shielding, requiring sensors to be highly sensitive; (3) background radiation varies both spatially and temporally, potentially leading to false positives if the background is not well understood; (4) detectors and associated electronics are generally expensive, and the cost increases significantly with improved performance and size.

To help mitigate and address these issues, our research team built a custom radiation detector that responds to γ -ray radiation, providing a low cost sensing mechanism for radiation, as shown in Fig. 1. This detector is based on CsI(Tl) and is connected to a silicon

photomultiplier with custom electronics to detect γ -rays and provide their energy. Moreover, the scintillator used is designed and fabricated by our research team. A diagram of the system and a picture of the manufactured prototype are shown in Fig. 1. As γ -rays deposit energy into the detector, the response is proportional to the energy of the γ -ray. By sensing over a period of time, our sensor can construct a histogram equivalent to the response function of the detector of all the γ -ray energies that strike the scintillator. Different radiation sources will exhibit different signatures of in γ -ray spectra, which can then be used to identify the radiation sources. Fig. 2 shows the response function of the detector under different individual control sources ($<1\mu\text{Ci}$) to assess its performance under a controlled environment. The vertical axis shows a normalized channel count (the number of sensed γ -rays at each energy) calculated as the mean channel count for each 5-second window present in the data. We show spectra for background and three radiation sources: ^{60}Co , ^{137}Cs , and ^{54}Mn . Each radiation source emits different γ -ray energies, and combined with the shape of the response function, helps us identify the source. It becomes obvious that in real scenarios with multiple sources present, the problems become difficult, and machine learning algorithms can allow us to identify anomalies under variable environments. Fig. 2 shows the response function of the detector when the sources are located at 0, 15 and 30 cm. It is clear the effect of $\frac{1}{r^2}$ reducing the collected signal from the source by the detector. The farther the distance the smaller the intensity until it can be compared to the background level. This motivates our team to use machine learning methods to increase the detection sensitivity of the sensor, while maintaining the use of low cost of CsI(Tl) scintillators.

It is also important to note that the baseline expected γ -radiation in an area differs. To illustrate this point, the background radiation levels were measured in the Dallas-Fort Worth Metro Area (Location A) as well as the Nine Mile Training Center at ETG ranch in Eldorado, Texas (Location B). The differences between the average spectra in Dallas and the average spectra at the ETG ranch is shown in Fig. 2 (right). This figure is zoomed on the lowest energies from the detector, as counts in higher energies are negligible. We can see that the lower energies have the largest counts for both locations, but the ETG location has more baseline γ -ray activity.

2. Related work

An overview of the related works and their major components are outlined in Table 1. Anomaly detection using auto encoders dates back to 1995 where Japkowicz et al. designed the HIPPO model to complete a classification task on the CH46 helicopter gearbox fault detection, molecular biology promoter recognition, and sonar target classification (Japkowicz et al., 1995). The HIPPO model was constructed using the auto encoder architecture created by Hinton et al. featuring input and output layers of the same size and a singular hidden layer that is smaller than the input and output layers (Hinton, 1990). Through this process they found that the HIPPO model had lower error rates than decision tree and neural classifier approaches.

As new auto encoder architectures have been developed those have also been used for anomaly detection. The Variational Auto Encoder (VAE) architecture (Kingma and Welling, 2022) has been used to create an anomaly detection scheme using reconstruction probability, instead of reconstruction error, as a more principled and objective anomaly detection metric (An and Cho, 2015). Here they found that the VAE and reconstruction probability outperformed other anomaly detection methods including traditional auto encoders. The AAE architecture has also been shown to be an effective way to perform anomaly detection tasks from image (Beggel et al., 2020) and video (Li and Chang, 2019) classification to detection on acoustic signals (Principi et al., 2017-05) and on electrocardiograms (Shan et al., 2022). This work as well as our previous work (Sayre et al., 2023) seek to apply the use of AAE anomaly detection to the field of radiation source detection.

Radiation detection falls into three main categories, radioisotope identification, anomaly detection, and source localization. Radioisotope

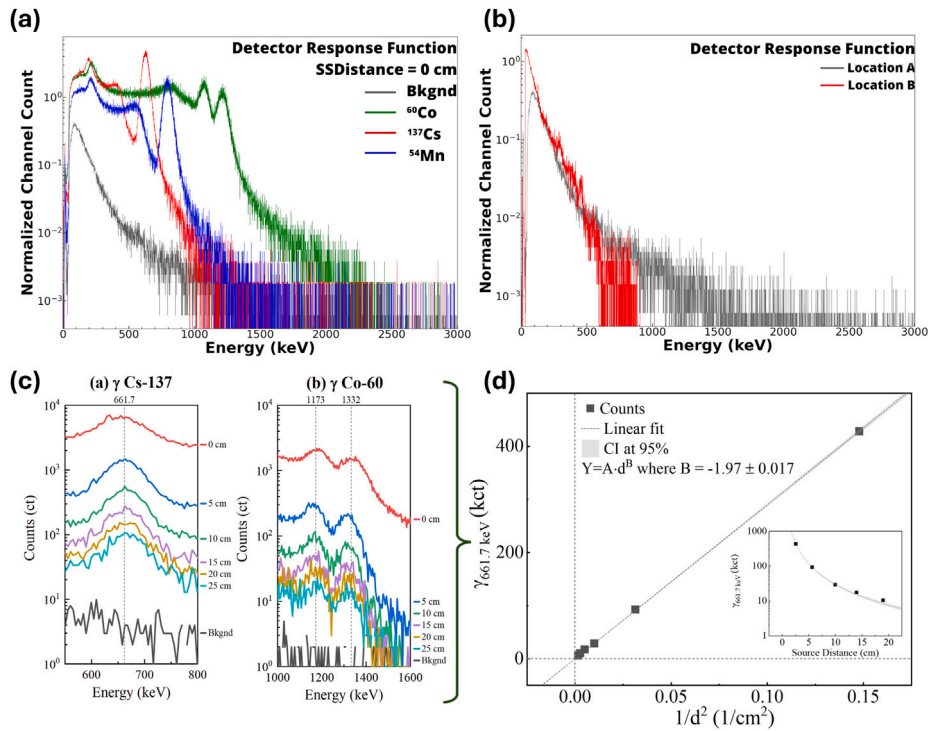


Fig. 2. (a) The mean 5-second representative response function of the γ -ray detector with sources ($<1\mu$ Ci) right on top of the sensor. (b) The 5-second average background spectra at two different locations. Location A is in the Dallas Fort Worth Area and Location B is in Southwest Texas. The two spectra are similar, but have slightly different base counts of γ -radiation. (c) and (d) Illustrate the detection dependence of the detector as a function of the distance for the same low activity sources. The detected γ counts decrease proportional to $\frac{1}{r^2}$.

identification is discerning the different kinds of radioactive sources present in a given source and from the background radiation. Anomaly detection is detecting the presence of radiation of any type from the background. Traditionally both of these tasks have been completed using statistical methodologies like peak matching, distance measures, and gross count rates. These tasks have seen advancements through the use of machine learning techniques recently (Kamuda et al., 2020) Mitchell et al. (2020) Ghawaly et al. (2022b) Bilton et al. (2021) Sayre et al. (2023).

While radioisotope identification is outside the scope of this project it is a natural extension of this work. Previous work has investigated the use of PCA over Mahalanobis Distance to detect differences between Naturally occurring radioactive materials (NORMs) and anomalous sources of ^{113}Ba and depleted uranium (Runkle et al., 2006). Many have expanded this to work on radioisotope identification using neural networks. Using a multi-layer perception (MLP), researchers from the French Alternative Energies and Atomic Energy Commission were able to predict the enrichment level of uranium (Vigneron et al., 1996). MLPs have also been used on isotope identification to ID individual isotopes present in a given mixture of elements (Kamuda et al., 2017). Other model architectures have also been used on similar problems from abductive machine learning (Abdel-Aal and Al-Haddad, 1997) to convolutional neural networks (CNNs) (Kamuda et al., 2020). Numerical methodologies are also advancing in this field with the use of dynamic quantum clustering on isotope identification (Weinstein et al., 2014).

For anomaly detection there is a use of a mix of statical models and machine learning based models. Statistical models used for anomaly detection include N-sigma, Mahalanobis Distance, and Spectral Comparison Ratios (SCR) which have shown to work better than simple gross count rates (Jarman et al., 2008). Distance based measures have also been used to combat high false positive rates from shielding by physical boxes or from NORMs in shielded nuclear source test-

ing (Omotaomu et al., 2009). Early work with neural networks on this task was a linear associative memory neural network to detect the isotopes present in mixtures which deviated from the peak analysis typically used for anomaly detection (Olmos et al., 1991). Machine learning algorithms that have been used for anomaly detection of radioactive sources recently include a Gaussian learning kernel (Alamianiotis et al., 2015), neuromorphic architecture aimed to reduce consumption (Mitchell et al., 2020; Ghawaly et al., 2022b), and auto encoders (Ghawaly et al., 2022a; Bilton et al., 2021; Sayre et al., 2023; Klasky et al., 2023). The use of auto encoders in previous works for γ -ray anomaly detection have primarily examined models with unrestricted latent space (Ghawaly et al., 2022a; Bilton et al., 2021). Auto encoders have also become a part of the testing suite for Los Alamos National Laboratory for evaluating the efficacy of new scintillators. Klasky et al. (2023) In our previous work we examined an AAE design to restrict the latent representation of background radiation data to a normal distribution, comparing it to simple supervised classifiers. Sayre et al. (2023) Anomaly detection has also been paired with isotope identification typically through the use of two networks, one for anomaly detection and a second for source ID (Sharma et al., 2012; Bilton et al., 2021). Additionally, anomaly detection on real world data has become of particular interest. From detection ^{137}Cs and weapons grade plutonium using an SCR-based approach in Seattle (Pfund et al., 2010) or using a static edge system in the NOVArray, a γ -ray detector array in Northern Virginia, with K-means clustering to organize spectra into background, rain, and anomalous (Bandstra et al., 2023). This allowed them to classify radiation events across various weather events in a location with radon fallout affecting the background radiation levels. As outlined in Table 1, there remains a gap in the contributions of related works in the application area of radiation anomaly detection, especially considering the usage of deep autoencoder approaches.

Table 1
Overview of approaches to radiation anomaly detection.

Authors	Year	Citation	Sim data	Real data	Deep learning	Autoencoder	Adversarial
Olmos et al.	1991	Olmos et al. (1991)		✓	✓		
Jarman et al.	2008	Jarman et al. (2008)		✓			
Omitaomu et al.	2009	Omitaomu et al. (2009)	✓				
Alamaniotis et al.	2015	Alamaniotis et al. (2015)		✓			
Bilton et al.	2021	Bilton et al. (2021)	✓		✓	✓	
Ghawlay et al.	2022	Ghawlay et al. (2022b)	✓		✓		
Ghawlay et al.	2022	Ghawlay et al. (2022a)	✓	✓	✓	✓	
Sayre et al.	2023	Sayre et al. (2023)		✓	✓	✓	✓
This Work	–	–	✓	✓	✓	✓	✓

Table 2
Sources available in each of the datasets used.

Source	⁶⁰ Co	⁵⁴ Mn	¹³⁷ Cs	Pu	HEU	^{99m} Tc	¹³¹ I
Small sources (1 μ Ci)	✓	✓	✓				
Large sources (>1mCi)			✓	✓	✓		
ORNL simulated data	✓			✓	✓	✓	✓

3. Approach and methods

3.1. Datasets

Over the course of this project three different datasets were used to test the AAE model as well as the various baseline classifiers. Our research team collected two of the datasets using our custom detector and a third dataset was made available by Oak Ridge National Laboratory (ORNL) ([Nicholson et al., 2020](#)). The differences in these datasets will be vital to showing the efficacy of the different models in being able to handle a changing environment.

In our data collection, the detectors are connected to a computer which records and aggregates these γ -ray events into spectra that can then be passed into the machine learning models for prediction. One of the datasets that utilized these detectors was created by placing the detector and small radioactive sources in a stationary position at various distances to create spectra with a wide range of intensities. Cobalt-60 (⁶⁰Co), Cesium-137 (¹³⁷Cs), and Manganese-54 (⁵⁴Mn) were the radioactive sources used in this dataset each with a strength of 1 μ Ci. This dataset has been used in our previous work that investigated an early version of our AAE model ([Sayre et al., 2023](#)). This laboratory collected dataset can be used to probe the sensitivity of the system with known sources, distances, and γ -ray energy.

Another dataset was also collected by our research team using larger radioactive sources and the system running in real-time on a small portable computer. Radioactive sources collected in this manner consisted of ¹³⁷Cs, uranium, and plutonium. These sources had strengths of 26 mCi for the cesium source and 3 mCi for each of the uranium and plutonium sources. This dataset was collected in collaboration with the DHS and the Rio Robotico event held in central Texas and allowed for us to test our AAE model on sources close to what could be experienced in real-world deployment.

Finally, we also used a synthetic data source created by Nicholson et al. from ORNL in Tennessee ([Nicholson et al., 2020](#)). This dataset is meant to mimic a 2in \times 4in \times 16in NaI(Tl) detector moving at a constant speed through a city block. The simulated data consists of background radiation and six radioactive sources placed in 15 different locations. This dataset contains many sources both in isotope and scale that we would otherwise not be able to work with. These sources are highly enriched uranium (HEU), weapons grade plutonium (WGPu), Technetium-99 metastable isotope (^{99m}Tc), Iodine-131 (¹³¹I), ⁶⁰Co. Additionally, it includes samples of HEU which have been masked with ^{99m}Tc. In this work, we treat the masked sample as a source, separate from both HEU and ^{99m}Tc, for the purposes of anomaly detection. The ORNL dataset is structured to be in list mode, timestamped γ -ray events with the associated energy level. For each of the

runs of the detector down the city block, the time of closest approach to the radioactive source, if present, was denoted in the file meta data. A breakdown of each source present in each of the datasets can be found in [Table 2](#)

To format the data to work in our system we truncated the energy values of the γ -ray events to be only whole numbers. The occurrences of each of these energy levels were aggregated into 5 s spectra as we had done with the data collected from the CsI(Tl) scintillators in the previous experiments. The array length, or number of channels, for these spectra was set at 4,096, however this may have been unnecessary or an active hindrance to the model as there were functionally only about 3,000 channels that could be filled due to the resolution of the simulated detector. For training using only the background training data provided in the dataset which resulted in 95,736 spectra, each corresponding to a 5-second collection time. Additionally for evaluation of the training process, we formatted the spectra of the radioactive sources in a similar way with the only difference being that we selected these to be only the spectra recorded inside of a 30 s window around the time of closest approach to the source.

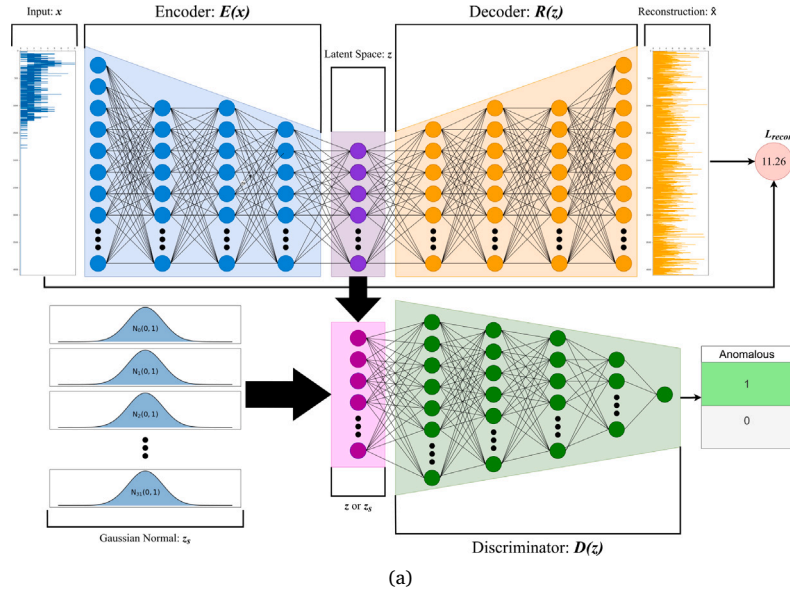
3.2. DHS standards of practice

Radiation source detection using energy analysis is guided by standards developed by the Department of Homeland Security ([Department of Homeland Security, 2019](#)), specifically through the Countering Weapons of Mass Destruction (CWMD) Office in collaboration with the National Institute of Standards and Technology (NIST). These guidelines are primarily designed for fixed-site deployments, such as portal monitors that scan vehicles at known distances ([Department of Homeland Security, 2019](#)). In contrast, our system is intended for mobile applications — such as drone-based sensing — where the distance between the source and detector is variable and often unknown. As a result, many of the existing standards do not directly apply to our use case.

Nonetheless, some performance benchmarks from these standards are still relevant. For example, false alarm rates are typically limited to one per 1,000 occupancies (with each occupancy defined as the time it takes a vehicle to pass through a detection platform), which roughly equates to one false positive per hour. Thus, we adopt this threshold in our own evaluation. While additional standards exist for source type identification and analysis procedures, their application in dynamic or aerial sensing environments remains unclear. Therefore, our focus is on addressing foundational research questions related to the use of machine learning in this new, mobile detection context.

3.3. Architecture and modeling

The AAE is a machine learning model that merges the ideas of generative adversarial networks and auto encoders to match the latent space vector some prior distribution ([Makhzani et al., 2016](#)). The model consists of three components, the encoder (E), decoder (R), and discriminator (D) as shown in [Fig. 3](#). The encoder-decoder pair works as they would in other auto encoder models to create a latent space embedding that input data can be transformed into via the encoder and then recovered by the decoder. That is, we denote an input signal to the encoder as x and the output of the decoder as \hat{x} such that $E(x)$ encodes



Encoder: $E(x)$	Decoder: $R(z)$	Discriminator: $D(z)$
Input Size = 4096	Input Size = 32 (from Latent Space)	Input Size = 32 (from Latent Space)
Linear, 2000 Neurons	Linear, 1000 Neurons	Linear, 1000 Neurons
LeakyReLU, $\alpha = 0.3$	LeakyReLU, $\alpha = 0.3$	ReLU
1-D Batch Normalization	Linear, 2000 Neurons	Linear, 500 Neurons
Linear, 2000 Neurons	LeakyReLU, $\alpha = 0.3$	Linear, 264 Neurons
LeakyReLU, $\alpha = 0.3$	Linear, 2000 Neurons	ReLU
1-D Batch Normalization	LeakyReLU, $\alpha = 0.3$	Dropout, Percentage = 0.25
Linear, 1000 Neurons	Linear, 4096 Neurons	Linear, 1 Neurons
LeakyReLU, $\alpha = 0.3$	ReLU	Sigmoid
1-D Batch Normalization		
Linear, 32 Neurons (Latent Space)		

(b)

Fig. 3. Overview of the neural network model architecture. Note: α controls the slope of the LeakyReLU function when inputs are negative.

the data into a latent space and $\hat{x} = R(E(x))$ gives the reconstruction of the signal from the encoder-decoder pair. In this way, we can define the first loss function of the encoder-decoder pair, which seeks to minimize the reconstruction error of the signal from samples drawn from the distribution of real signal data, X :

$$\mathcal{L}_{recon} = \frac{1}{|X|} \sum_{x \in X} (R(E(x)) - x)^2$$

Both E and R are updated to minimize this loss function. Without any other losses, this would be identical to the mean-squared error formulation of a traditional auto-encoder. However, without additional losses there is too much freedom for the models to construct a variable or discontinuous latent space. We desire the latent space to be distributed in a known way for the samples in X . Thus, a discriminator is added that can process the latent representation of real data, $E(x)$, or a similar sized vector, z_s where each element of the vector is drawn from a Gaussian normal distribution of zero mean and unit variance, $\mathcal{N}(0, 1)$.

The discriminator acts to predict the difference between the encoded embedding and the random normal vector z_s . This results in two

competing loss functions that work in an adversarial manner. Each is based on binary cross entropy as follows:

$$\mathcal{L}_{disc} = \sum_{z \sim \mathcal{N}(0,1)} \log(D(z_s)) + \sum_{x \in X} \log(1 - D(E(x)))$$

Intuitively, this loss function updates the discriminator model, D to detect the difference between sampled points, z_s and real encoded points $E(x)$. The encoder is then updated with a separate competing loss function:

$$\mathcal{L}_{enc} = \sum_{x \in X} \log(D(E(x)))$$

that optimizes the encoder to fool the discriminator by making the latent space of the encoded samples as similar as possible to a normal distribution. Only the encoder is optimized using this loss function to ensure that the discriminator maintains detection accuracy. Thus, after proper training, the real latent samples should occupy a distribution near the origin of the latent space, with standard deviation of unity in all directions. Anything outside of this area can be considered anomalous.

Table 3
Optimal hyper parameters found on the AAE network using optuna.

Hyper parameter	Optimal value
Epochs	75
α	0.98
Latent Space Dim.	32
E(x) Learning Rate	0.01
R(x) Learning Rate	0.01
D(x) Learning Rate	0.004
E(x) Optimizer	SGD
R(x) Optimizer	Adam
D(x) Optimizer	SGD

Finally, to combine the loss functions for the encoder, we use a simple weighting:

$$\mathcal{L}_{tot} = \alpha \cdot \mathcal{L}_{recon} + (1 - \alpha) \cdot \mathcal{L}_{enc}$$

where α is a hyper parameter tuned during training.

Through the training process the encoder–decoder pair works together to produce a latent space that is indistinguishable from the prior distribution to the discriminator. Here we utilize this to make an unsupervised classifier for radiation anomaly detection. The design of the AAE model allows for detection of unknown sources if they differ from its training data sufficiently. Moreover, the encoder of the model is trained to compactly compress training data near the origin of the latent space—thus even subtle input differences from unknown sources can result in large deviations in the latent space.

For each of the outlined experiments the AAE is trained exclusively on spectra of background radiation. Therefore, the model should learn that background radiation falls into a normal distribution in the latent space, $E(x)$, while any spectra with nearby radioactive sources should fall outside of that distribution and be marked as anomalous. Given this, the model is designed to predict for any novel radioactive source as they all should be noticeably different from the background environment in the latent space. Additionally, because the model is trained to enforce that the latent representation of the background radiation lays in a normal distribution it can be trained on variable background environments so that the range of background radiation is fully captured in the latent representation. Therefore, after the model converges, the output $D(E(x))$ can be used directly as a probability that the data point x is anomalous. Thus, the AAE can naturally be used for anomaly detection. We also explore the use of reconstruction error as another way to determine the severity of the anomaly, discussed later.

The architecture is shown in Fig. 3. The input into this model is a 5 s spectrum of radiation from the detector, which is comprised of a series of channels ranging from 0 to 4096. The outputs consist of a probability output from the discriminator and a reconstruction of the original input spectra from the decoder. The discriminator output is a single probability for the likelihood of a radioactive substance being detected. The decoder output is used to find the mean squared error with the original input vector. This value can then be used to create a measure for how far a radioactive source is to the detector.

3.4. Hyper parameter tuning

Hyper parameter tuning was performed on the model using the optuna framework (Akiba et al., 2019). The continuous variables searched were learning rates for each of the sub-networks, α , latent space size, and epochs. The optimization functions for each of the sub-networks were also searched using either the Adam optimizer or Stochastic Gradient Descent. Hyper parameters were searched across 100 models choosing the optimal values based on the highest receiver operating characteristic. Data used for this task came from the stationary small-source collection described above. The final results of the hyper parameter search can be found in Table 3.

4. Experimental results

4.1. Comparison to unsupervised learners

This experiment centers around testing how the AAE model and other unsupervised classifiers handle background measurements from two different background environments. The background radiation levels were measured in the Dallas-Fort Worth Metro Area (Location A) as well as the Nine Mile Training Center at ETG ranch in Eldorado, Texas (Location B). At the ETG ranch site the research team was able to utilize large radioactive sources for the first time, as such we collected spectra in the area of those sources but far enough away that the detector was not detecting any radiation. The differences between the average spectra in Dallas and the average spectra at the ETG ranch is shown in Fig. 2 (right). Effectively this means we have created two sets of training data for these tests. One that contains only background radiation collected from Location A and another that contains background from both Location A and Location B. This is done to emulate scenarios when a small amount of calibration data is available from a known background environment.

The unsupervised models chosen for this test were One Class SVM (OCSVM), Isolation Forest (IF), and Local Outlier Factor (LOF). These methods were chosen as they have a history of being used for radiation anomaly detection (Klasky et al., 2023). Additionally, we modified a version of the GANomaly model to work with radiation spectra. The original GANomaly was designed to be used on 2-D images (Akay et al., 2018)—as such, we adjusted the architecture to use 1-D convolutions across the spectra. In this way, we can compare our approach to both traditional unsupervised anomaly detectors and more contemporary detection methods. Along with the AAE, these models were trained and tested on two different datasets. The first dataset consists of half background collected only from the DFW area and half ^{137}Cs with the source being located at 0, 15, and 30 cm away from the detector with each distance comprising one-third of the radioactive source data. The second dataset augmented the first by substituting around one-quarter of the background radiation spectra with background spectra from the ETG ranch site. Doing this should result in background spectra that are harder for the unsupervised classifiers to neatly separate from the radioactive spectra, while the AAE should perform better as it is designed to compress all of the null spectra to fit into a latent space that is normally distributed. For each of the models the False Positive Rate (FPR) was held to 0.015 to allow us to make fair comparison between the models.

The experimental results in Fig. 4 show how the different models handle the two different datasets. On the testing dataset for the DFW background the AAE comes second out of all of the models with an accuracy of 90.41% with the IF model surpassing it in this case with an accuracy of 93.23%. The GANomaly model struggles to generalize to the testing data with an accuracy of 71.04%. Furthermore, when augmenting the dataset to include background spectra from Location B (the ETG ranch) the advantages of the AAE architecture come through, maintaining an accuracy nearly 90%. With accuracies at 80% and the AUC at 0.5 for both the OCSVM and LOF models were unable to draw a decision boundary between radioactive sources and background. The GANomaly model did see an improvement on the Location B training data with test accuracy of 81.28%. The IF model saw a slight decline in accuracy on the ETG ranch augmented dataset dropping to 87.80% accuracy. When looking at the AUC values for the IF and AAE models we found that the AAE has a superior AUC compared to all models at all locations. The AAE had AUC values of 0.914 and 0.899 for the testing data from locations A and B respectively while the IF model had 0.831 and 0.701. The GANomaly model had an increase in the AUC from Location A to Location B which went from 0.764 to 0.880. This further shows how the AAE model can better represent different background radiation profiles in the latent space.

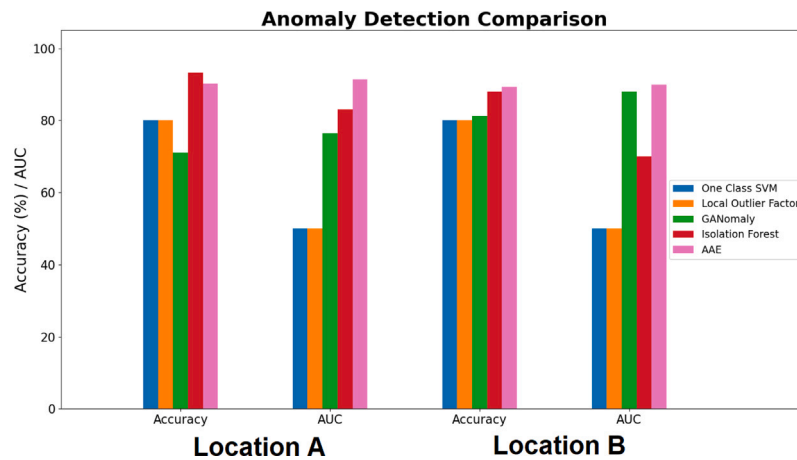


Fig. 4. Accuracy and AUC of our AAE model compared to other anomaly detection techniques. AUC is shown as a percentage of the maximum possible AUC ($AUC \times 100$). Results are shown for two locations. FPR held at 0.015 for each model. McNemar tests were performed on the predictions of each of the baseline models compared to the predictions from the AAE, finding that all performance improvements are significant ($p < 0.05$).

Because the AAE and IF models were strong performers, we select these for further analysis. Fig. 5 shows a qualitative comparison of the results of the AAE model and the IF model using principal components analysis (PCA) of the latent representation from the AAE model and colored by the outputs of each model. That is, Fig. 5 displays a series of 2D plots across three dimensions of the PCA space from the AAE encoder network. These plots are colored by the predictions, background or in the presence of radiation from the IF and AAE models. The shapes encode the distance from the detector to the radiation sample when it was collected. Three graphs are shown as a scatter matrix and a zoomed version of the clusters in the PCA1-PCA2 graph. The rightmost plot is in PCA1-PCA3 space and can be thought of as rotating the first plot (in the 3D PCA space) to look down on the data. While this visualization does not capture all of the differences present in the embeddings (as many differences are in dimensions beyond PCA3), reducing the dimensions to this three-dimensional space does help visualize some of the main differences and clustering.

By looking at the embedded latent space in this way, we can see how the AAE model works to capture the essential differences in the spectra. The first PCA dimension accounts for the majority of the differences in the spectra with all of the data collected at 0 cm clustering away from all other data. This is due to the magnitudes of the spectra being so much greater than when they are placed further away, following the inverse square law of radiation energy. The IF model predicts if the test data is a part of the inlier class (background radiation), and the AAE model is predicting if there is an outlier (anomalous γ -rays) while both perform similarly at classifying the strongest radioactive sources (triangles), when looking at the insets on the left side of Fig. 5 corresponding to the background radiation and the sources place 15 cm or 30 cm away from the detector, we can see that the radiation does tend to fall to the right on the PCA1 axis for the AAE model, but not for the IF model. Thus the AAE has a slightly better clustering of background and radiation sources that are further away. When rotating the graph to look at PCA1-PCA3 space we can see the background spectra tend to fall near the origin with the radioactive spectra tending to be further removed from that point. We also note that the IF model specifies inlier, while our algorithm specifies outliers, such that hue of points in each figure is reversed.

4.2. Comparison to auto encoding

With the goal of this model to be deployed in an operational environment we expanded testing beyond just the small, single microcurie sources that we had been using. We therefore analyzed the system using the previously described ORNL dataset which contains simulated

spectra meant to mimic a 2in \times 4in \times 16in NaI(Tl) detector moving at a constant speed through a city block. This data allows for us to test how the model generalizes to other types of detectors. It is also a popular dataset in the radiation detection community, allowing for additional comparison to previous works.

When compared to the data we collected from our custom detector, the channel counts are less distinct, especially when looking at background radiation. These high levels of background radiation can have detrimental effects on the ability for the AAE to train properly. As we showed in Fig. 5 the embedding space for the AAE model seems to rely on the magnitude of the spectra as a fundamental part of making a decision plane for anomalies. If the background it is trained on has a magnitude similar to that of the radioactive sources that it is trying to detect this will lead to many incorrect predictions. As we were limited to where we could test our model (especially in urban environments), this dataset help us understand how the detector would work in environments that might have higher levels of background radiation like those simulated here.

For testing, we followed the methodology outlined in Characterization of the Autoencoder Radiation Anomaly Detection (ARAD) model by Ghawaly et al. (2022a). Here accuracy of the model on testing background radiation is measured simply by the percentage of spectra correctly identified by the model. For accuracy on the radioactive source, it was measured as the percentage of files where the model correctly identified the radiation at any point along the detectors path within a 30 s window centered on the time of closest approach by the detector. This was done to account for the drop-off in radiation levels as the detector enters and leaves the area with the source. An ideal detection algorithm would display the strongest anomalies at the apex of the approach.

The initial results for accuracy are shown in Table 4 (top row) which, at first glance, appear quite positive. However, these results have an issue of very high False positive rate, which would render the system useless in most applications because the false identifications would overwhelm the operator's time. Without tuning the probability thresholds for the model we have a false positive rate of 197.3 per hour which correlates to a background accuracy of 72.6%, which is unacceptable. In working with the UT-Dallas detectors we found that our false positive rate was closer to 10 per hour. To find a better operational point, we raised the probability threshold of the discriminator incrementally until finding one that seems reasonable for our application with limited false positive rates. Table 4 outlines the results of this threshold tuning.

As expected, as we increased the discriminator's probability threshold the false positive rate decreased but the accuracies also decreased

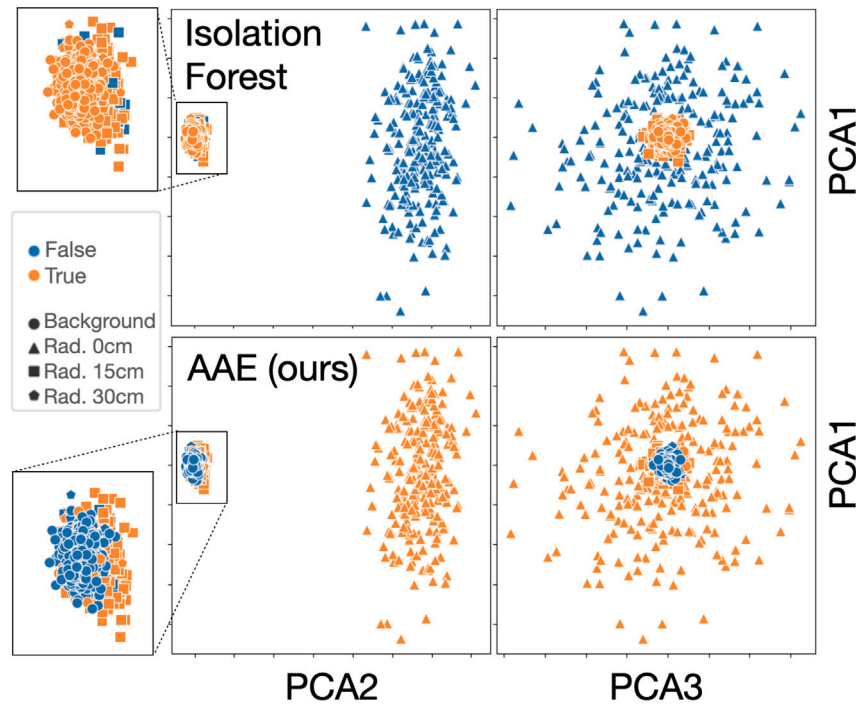


Fig. 5. Scatter matrix of PCA dimensions calculated from the latent space of the AAE model. They are colored by their prediction of an anomaly. The colors between the two are reversed because the IF model predicts if it is an inlier (Background) and the AAE predicts if it is an outlier (Radioactive). Note that the zoomed cluster for the AAE has a more obvious clustering compared to the isolation forest.

Table 4

True positive rate and false alarms per hour for our AAE model provided for various types of radiation sources in the ORNL dataset from (Ghawaly et al., 2022a). The final two rows show the comparative performance of our model and the ARAD model which have each been properly tuned to have a False Positives per hour rate below 0.5, with the best results shown in bold. Larger is better for all measures except the false alarms column, where lower is better.

Threshold	HEU	WGPu	^{99m} Tc	¹³¹ I	⁶⁰ Co	HEU+ ^{99m} Tc	FP per hour
T=0.5 (base)	0.925	0.888	0.963	0.931	0.9	0.963	197.3
T=0.6	0.65	0.6	0.806	0.669	0.669	0.813	61.7
T=0.7	0.331	0.388	0.631	0.438	0.356	0.675	13.5
T=0.8	0.194	0.281	0.438	0.244	0.113	0.538	1.66
T=0.841 (Tuned)	0.157	0.231	0.400	0.194	0.081	0.513	0.483
Tuned Comparison	HEU	WGPu	^{99m} Tc	¹³¹ I	⁶⁰ Co	HEU+ ^{99m} Tc	FP per hour
ARAD ((Ghawaly et al., 2022a), Tuned)	0.234	0.180	0.268	0.189	0.130	0.362	0.5
AAE (Ours, Tuned)	0.157	0.231	0.400	0.194	0.081	0.513	0.483

accordingly. In order to get to a false positive rate similar to what we saw in the other experiments we would have to raise the threshold to 0.7 on the discriminator, increasing it to 0.85 sees a decrease the false positive rate to 0.392 per hour which is in line with what was done with the ARAD model in Ghawaly et al. (2022a) where they set their rate to be 0.5 false positives per hour. This false positives per hour threshold is set by Department of Homeland security (Department of Homeland Security, 2016). The bottom two rows of Table 4 shows this comparison in full detail.

Our AAE model outperformed the ARAD model on the WGPu, ^{99m}Tc, ¹³¹I, and HEU masked by ^{99m}Tc, while it lags behind on HEU and ⁶⁰Co. However, the AAE does so at a slightly lower false positive rate per hour. That is, only one false positive is expected every two hours. This shows that both the ARAD and our AAE model have merit to be used as radiation detection algorithms, with the AAE showing slightly better resilience to more varied sources. This dataset is designed to be challenging, with many weak simulated sources that are far away from the simulated mobile detection vehicle. Even so, our AAE methodology did attain results that are comparable to (and in many cases surpass) other state-of-the-art radiation anomaly detection models.

4.3. Deployed model evaluation

We now shift our evaluation to the efficacy of our AAE model when used in a real-time system. To run the model in real-time we placed the model inside of an API created using the gRPC API framework. We use a lightweight hardware system, where the model and API are run locally on a UP Xtreme Series board with an Intel Celeron 4305UE at 2 GHz and 4 GB of system memory. Events are read directly from the detector and passed through the API to the model which returns a probability of the detected radioactivity being anomalous. It was desired for the system to provide an estimate of the strength of the anomaly to help create a way to evaluate how far the detector is from the radioactive source. The primary measurement we chose for this is the mean squared error (MSE), however this alone would not be enough. To help linearize the outputs we introduce log scaling the MSE, as follows:

$$MSE = \frac{1}{n} \sum_{i=0}^n (x_i - \hat{x}_i)^2, \quad D_{error} = \ln(10 \times MSE)$$

where \hat{x} denotes the output of the decoder model and x denotes the input to the encoder model. We have called this log-scaled MSE the D_{error} . This measure was created from the observation that our model typically produces MSE values below 0.1 for background observations,

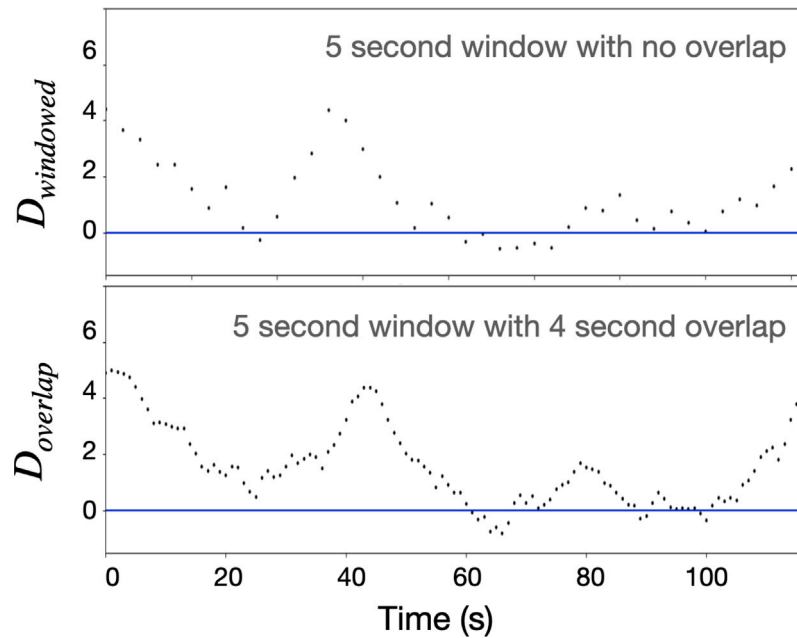


Fig. 6. Deployed output running with and without overlap in the windows. All windows are processed using a five second long window and overlap in bottom figure is set to 4 s. The value D is shown as a source is moved closer to the system (larger magnitudes) and further away from the system (lower magnitudes). To distinguish the D_{error} calculated from the windows with and without overlap they are called $D_{overlap}$ and $D_{windowed}$ respectively. The radiation source used in experiments is ^{137}Cs .

while radioactive spectra produce values are much greater. This means that using the MSE in this way will tend to produce D_{error} values that are less than 0 for background spectra and above 0 for radioactive spectra.

Additionally, to run this in real-time we needed to slightly modify the input into the model. In previous experiments we used exclusive 5-second windows, however since the model is meant to be used to update the direction of travel of a drone this needs to be smaller. To account for this we made a slight alteration to uses 5-second windows that contain 4 s of overlap with the previous window submitted to the model.

The effect of using these overlapping windows is highlighted in Fig. 6. Here a $1\mu\text{Ci } ^{137}\text{Cs}$ is moved along a similar path towards and away from the detector with the model running on spectra with no overlap and spectra with 4-second overlap. As the source moves away from the detector there are relatively lower D_{error} values and D_{error} increases as the source gets closer. With no overlap the model is responsive to the source moving, but requires the source to move slowly to prevent large changes in the D_{error} . Using an overlap, we can more effectively tell how the source is moving relative to the detector.

4.4. Field testing with large radiation sources

As a part of the project we were able to work with large sources at the Rio Robotic event at the ETG ranch, Nine Mile facility. Over the course of two events in Fall 2023 and Spring 2024, our research team completed several experiments with the model and detector working in real-time. Initially, we worked with the system as a handheld device recording background and radioactive sources in an environment significantly different from those that we had worked with in the DFW area.

An approximate map of the ETG ranch radiation set up is provided in Fig. 7(left). Three radioactive sources were provided for us to work on, those being 3mCi uranium, 3mCi plutonium, and 26mCi ^{137}Cs . Background radiation was collected from yellow area demarked on the satellite image and was approximately 300 ft from the nearest source and 900 ft from the farthest. Background radiation was collected during

the day underneath a shaded tent, and we were allowed to walk as close as possible to the source for their relative radiation levels. For these tests the model was running on a laptop with the same detector designed by UT-Dallas, highlighted in Section 1.1 and used in the earlier laboratory data collection. This provided an opportunity to not only test how the model performed in real-time, but also to work with the types of sources that the system might encounter in an operational environment.

While at the ETG site we exclusively used the AAE model that was trained on data collected from the DFW area. This first plot in Fig. 7 shows how the AAE model performed while recording background at the ETG site. In the post field test period, we were able to train a new model that used a background radiation dataset that consisted of 25% background from the test site and 75% background from DFW. This new version of the model, which is called the ETG model was then retested on the background from the ETG site, which is shown in the second row of Fig. 7(right). This greatly reduced false positives that the model was creating for that location and goes to show the adaptability of the model in the training period. For each of these models the threshold for detecting radiation was set to 0.5 probability. While this detection can be improved by properly adjusting this threshold, the primary purpose of these field tests is confirming the efficacy of the D_{error} and $D_{overlap}$ scores. Fig. 7 illustrates that the AAE model is sensitive to training data. We observed that the model, when trained only on baseline data from DFW, produced false positives when used at the ETG site. To alleviate this drawback, a small amount of calibration data was required for tuning the model. This underlines a potential limitation of the AAE model (and unsupervised models in general)—the diversity of inlier examples must be sufficient to capture baseline variations. In our specific application this background data is easy to acquire and fine tuning can be carried out on modest hardware.

With uranium and plutonium, we were allowed to be within about 1 foot of the source while with the ^{137}Cs the closest we were able to get was about eight feet. The results of walking around these sources are shown in Fig. 8 using the ETG model created after completing the field test. The peaks of each of these plots correspond to the point of closest approach to the source. For uranium and plutonium, we




Fig. 7. The evaluation area and location of radiation sources (left). The background recordings for a model trained in location A (top right) and for a model augmented with background labels from the same area, Locations B (bottom right).

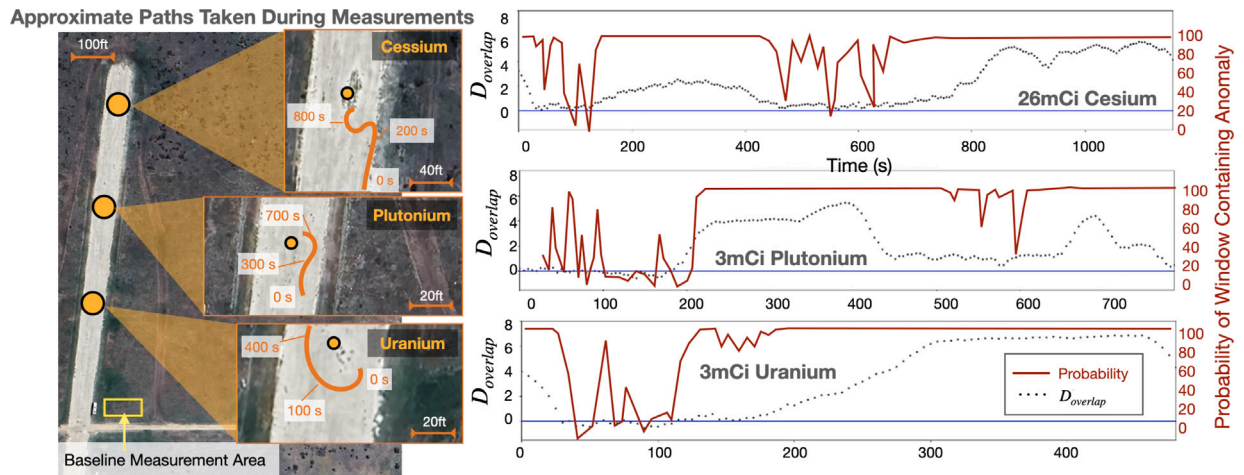


Fig. 8. Left: The approximate paths used for the data collection of the figures on the right, marked with times. Right: The visualization of the approach of the system to each radiation source. Dots illustrate the quantity D_{overlap} derived from reconstruction error with overlapping windows and solid lines convey the probability of the window containing an anomalous point. As can be seen from the figure, the probability is quite sensitive to the approach to the source, while the D magnitude changes in response to proximity.

detected positive radiation values up to about 8 ft away from the source while with ^{137}Cs we detected radiation as far away as 60 ft from the source. Additionally, this shows why we need to create a metric for the model to use which we can evaluate the distance between the detector and the source. When using just the probability there are much more stark drops in the probability as the distance changes which are not as useful as the much smoother changes seen in the D_{error} .

The only difference between the results shown in Fig. 8 and the initial results that we had from the AAE model trained on DFW data is that the DFW model was shifted higher on both the D_{error} and the probability axes. This can be thought of as higher levels of background being added to the dataset as being subtracted from the predictions of the reanalyzed model. This leads credence to our working hypothesis that the input data can be modified in any number of ways to exclude location specific variance that might be in the background radiation. And that doing these sorts of modifications will not have much of an impact on the performance of the model for detecting the likelihood of and the distance to radioactive sources.

Ultimately, the field tests at ETG were extremely useful in testing how the model performs in environments similar to those it might see in operation—albeit with reduced communication frequency. The model proved to be resilient to the background environment being changed from when it was trained in the DFW area.

5. Conclusion

In this work, it was shown how small form factor γ -ray detectors can be combined with adversarial auto encoders to produce radiation sensors with key advantages. In particular, we found that this methodology allows for background radiation to be modeled efficiently and trained without specific examples of radioactive isotopes. That is, the system is trained only with background radiation. In our experiments this background was collected for a few hours, but was consistent for experiments occurring days after the initial calibration. We found several advantages of this methodology, including: (1) the system is highly sensitive, showing better sensitivity than other methods applied to the same data; (2) the type of γ -ray emission signature does not confound the sensor, (3) low intensity sources are still detectable from a range of distances and larger emission sources are detectable from long distances up to 80 m; (4) the spatial and temporal background variation is effectively captured by the adversarial model leading to better sensitivity and fewer false positives; the AAE model generalizes beyond just the CsI(Tl) created for this project by UT-Dallas and is potentially applicable to other detectors as show by the ORNL dataset which is based on NaI(Tl) detector. and (6) the materials and electronic equipment are orders of magnitude less expensive than current solutions. Each of these advantages is addressed and evaluated in the paper. We also note a few limitations of our approach. Firstly, our analysis indicates that some calibration data may be needed when background radiation changes

significantly. Secondly, the AAE model is a deep learning approach and therefore requires a significant amount of background data for the initial training. Once trained, the fine tuning can be performed with a relatively modest dataset. We compare to other similar auto-encoding methods, showing that adversarial encoding has a few distinct advantages in modeling background variability. We hypothesize that this is the main reason we are able to achieve greater sensitivity and fewer false positives. Our work demonstrates an advancement in the detection of unlawful movement of radioactive materials, allowing for mobile swarms of detectors to canvas an area and locate unauthorized sources.

CRedit authorship contribution statement

Charles Sayre: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Formal analysis, Data curation, Conceptualization. **William Bjorndahl:** Writing – review & editing, Software, Methodology, Formal analysis, Data curation, Conceptualization. **Eric C. Larson:** Writing – review & editing, Supervision, Conceptualization. **Joseph Camp:** Writing – review & editing, Supervision, Conceptualization. **Rodolfo Rodriguez-Davila:** Writing – review & editing, Visualization, Formal analysis, Conceptualization. **Manuel Quevedo-Lopez:** Writing – review & editing, Supervision. **Bruce Gnade:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Charles Sayre reports financial support was provided by U.S. Department of Homeland Security Countering Weapons of Mass Destruction Office. William Bjorndahl reports financial support was provided by US Department of Homeland Security Countering Weapons of Mass Destruction Office. Eric C. Larson reports financial support was provided by US Department of Homeland Security Countering Weapons of Mass Destruction Office. Joseph Camp reports financial support was provided by U.S. Department of Homeland Security Countering Weapons of Mass Destruction Office. Rodolfo Rodriguez-Davila reports financial support was provided by U.S. Department of Homeland Security Countering Weapons of Mass Destruction Office. Manuel Quevedo-Lopez reports financial support was provided by U.S. Department of Homeland Security Countering Weapons of Mass Destruction Office. Bruce Gnade reports financial support was provided by U.S. Department of Homeland Security Countering Weapons of Mass Destruction Office. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This work was funded by the United States Department of Homeland Security, Division on Countering Weapons of Mass Destruction (CWMD), grant 18DNARI000029-05-00.

Data availability

The authors do not have permission to share data.

References

- Abdel-Aal, R.E., Al-Haddad, M.N., 1997. Determination of radioisotopes in gamma-ray spectroscopy using abductive machine learning. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detect. Assoc. Equip.* 391 (2), 275–288.
- Akay, Samet, Atapour-Abarghouei, Amir, Breckon, Toby P., 2018. GANomaly: Semi-supervised anomaly detection via adversarial training.
- Akiba, Takuya, Sano, Shotaro, Yanase, Toshihiko, Ohta, Takeru, Koyama, Masanori, 2019. Optuna: A next-generation hyperparameter optimization framework.
- Alamaniotis, Miltiadis, Choi, Chan K., Tsoukalas, Lefteri H., 2015. Anomaly detection in radiation signals using kernel machine intelligence. In: 2015 6th International Conference on Information, Intelligence, Systems and Applications. IISA, pp. 1–6.
- An, Jinwon, Cho, Sungzoon, 2015. Variational autoencoder based anomaly detection using reconstruction probability. In: 2015-2 Special Lecture on IE.
- Bandstra, M.S., Abgrall, N., Cooper, R.J., Hellfeld, D., Joshi, T.H.Y., Negut, V., Quiter, B.J., Salathe, M., Sankaran, R., Kim, Y., Shahkarami, S., 2023. Background and anomaly learning methods for static Gamma-ray detectors. *IEEE Trans. Nucl. Sci.* 70 (10), 2352–2363.
- Beggel, Laura, Pfeiffer, Michael, Bischl, Bernd, 2020. Robust anomaly detection in images using adversarial autoencoders. In: *Machine Learning and Knowledge Discovery in Databases*. Springer International Publishing, pp. 206–222.
- Bilton, Kyle J., Joshi, Tenzing H.Y., Bandstra, Mark S., Curtis, Joseph C., Hellfeld, Daniel, Vetter, Kai, 2021. Neural network approaches for mobile spectroscopic Gamma-ray source detection. *J. Nucl. Eng.* 2 (2), 190–206.
- Department of Homeland Security, 2016. American national standard for evaluation and performance of radiation detection portal monitors for use in homeland security. pp. 1–70.
- Department of Homeland Security, 2019. Technical capability standard for radiation portal monitor systems with energy analysis capability – 2019. pp. 1–33.
- El Hamli, A., Bazza, A., Moussa, A., Hamal, M., Zerfaoui, M., Ouchrif, M., 2022. Data and simulation studies on the influence of scintillation crystal dimensions on spectrometric parameters. *Appl. Radiat. Isot.* 181, 110053.
- Ghawaly, James M., Nicholson, Andrew D., Archer, Daniel E., Willis, Michael J., Garishvili, Irakli, Longmire, Brandon, Rowe, Andrew J., Stewart, Ian R., Cook, Matthew T., 2022a. Characterization of the autoencoder radiation anomaly detection (ARAD) model. *Eng. Appl. Artif. Intell.* 111, 104761.
- Ghawaly, James, Young, Aaron, Archer, Dan, Prins, Nick, Witherspoon, Brett, Schuman, Catherine, 2022b. A neuromorphic algorithm for radiation anomaly detection. In: *Proceedings of the International Conference on Neuromorphic Systems 2022*. Association for Computing Machinery, pp. 1–6.
- Hinton, Geoffrey E., 1990. Connectionist learning procedures. *artificial intelligence*, 40 1-3: 185–234, 1989. reprinted in J. Carbonell, editor. *Mach. Learn.: Parad. Methods*, MIT Press.
- Japkowicz, Nathalie, Myers, Catherine, Gluck, Mark, 1995. A novelty detection approach to classification. In: *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, vol. 1, pp. 518–523.
- Jarman, Kenneth D., Runkle, Robert C., Anderson, Kevin K., Pfund, David M., 2008. A comparison of simple algorithms for gamma-ray spectrometers in radioactive source search applications. *Appl. Radiat. Isot.* 66 (3), 362–371.
- Kamuda, M., Stinnett, J., Sullivan, C.J., 2017. Automated isotope identification algorithm using artificial neural networks. *IEEE Trans. Nucl. Sci.* 64 (7), 1858–1864.
- Kamuda, Mark, Zhao, Jifu, Huff, Kathryn, 2020. A comparison of machine learning methods for automated gamma-ray spectroscopy. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detect. Assoc. Equip.* 954, 161385.
- Kangas, Lars J., Keller, Paul E., Siciliano, Edward R., Kouzes, Richard T., Ely, James H., 2008. The use of artificial neural networks in PVT-based radiation portal monitors. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detect. Assoc. Equip.* 587 (2), 398–412.
- Kingma, Diederik P., Welling, Max, 2022. Auto-encoding variational Bayes.
- Klasky, Marc Louis, McCann, Michael Thompson, Haack, Jeffrey Robert, Matheny, Jason Ray, 2023. Comparison of supervised and unsupervised machine learning algorithms for threat detection and scintillator performance for radiation portal monitoring.
- Kouzes, Richard T., 2010. Neutron and gamma ray detection for border security applications. In: 2010 1st International Nuclear & Renewable Energy Conference. INREC, pp. 1–3.
- Li, Nanjun, Chang, Faliang, 2019. Video anomaly detection and localization via multivariate gaussian fully convolution adversarial autoencoder. *Neurocomputing* 369, 92–105.
- Makhzani, Alireza, Shlens, Jonathon, Jaitly, Navdeep, Goodfellow, Ian, Frey, Brendan, 2016. Adversarial autoencoders.
- Milbrath, B.D., Choate, B.J., Fast, J.E., Hensley, W.K., Kouzes, R.T., Schweppe, J.E., 2007. Comparison of LaBr₃:Ce and NaI(Tl) scintillators for radio-isotope identification devices. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detect. Assoc. Equip.* 572 (2), 774–784.
- Mitchell, J. Parker, Schuman, Catherine D., Potok, Thomas E., 2020. A small, low cost event-driven architecture for spiking neural networks on FPGAs. In: *International Conference on Neuromorphic Systems 2020*. ACM, pp. 1–4.

- Morgan, Dane, Pilania, Ghanshyam, Couet, Adrien, Uberuaga, Blas P., Sun, Cheng, Li, Ju, 2022. Machine learning in nuclear materials research. *Curr. Opin. Solid State Mater. Sci.* 26 (2), 100975.
- Nicholson, Andrew D., Peplow, Douglas E., Ghawaly, James M., Willis, Michael J., Archer, Daniel E., 2020. Generation of synthetic data for a radiation detection algorithm competition. *IEEE Trans. Nucl. Sci.* 67 (8), 1968–1975.
- Olmos, P., Diaz, J.C., Perez, J.M., Gomez, P., Rodellar, V., Aguayo, P., Bru, A., Garcia-Belmonte, G., de Pablos, J.L., 1991. A new approach to automatic radiation spectrum analysis. *IEEE Trans. Nucl. Sci.* 38 (4), 971–975.
- Omitaomu, Olufemi A., Ganguly, Auroop R., Patton, Bruce W., Protopopescu, Vladimir A., 2009. Anomaly detection in radiation sensor data with application to transportation security. *IEEE Trans. Intell. Transp. Syst.* 10 (2), 324–334.
- Pfund, David Michael, Jarman, Kenneth D., Milbrath, Brian D., Kiff, Scott D., Sidor, Daniel E., 2010. Low count anomaly detection at large standoff distances. *IEEE Trans. Nucl. Sci.* 57 (1), 309–316.
- Principi, Emanuele, Vesperini, Fabio, Squartini, Stefano, Piazza, Francesco, 2017-05. Acoustic novelty detection with adversarial autoencoders. In: 2017 International Joint Conference on Neural Networks. IJCNN, pp. 3324–3330.
- Runkle, R.C., Tardiff, M.F., Anderson, K.K., Carlson, D.K., Smith, L.E., 2006. Analysis of spectroscopic radiation portal monitor data using principal components analysis. *IEEE Trans. Nucl. Sci.* 53 (3), 1418–1423.
- Sayre, Charles, Larson, Eric C., DiLiegro, Gabs, Camp, Joseph, Gnade, Bruce, 2023. Radiation anomaly detection using an adversarial autoencoder. In: 2023 57th Asilomar Conference on Signals, Systems, and Computers. pp. 1010–1014.
- Shan, Lianfeng, Li, Yu, Jiang, Hua, Zhou, Peng, Niu, Jing, Liu, Ran, Wei, Yuanyuan, Peng, Jiao, Yu, Huizhen, Sha, Xianzheng, Chang, Shijie, 2022. Abnormal ECG detection based on an adversarial autoencoder. *Front. Physiol.* 13.
- Sharma, Shiven, Bellinger, Colin, Japkowicz, Nathalie, Berg, Rodney, Ungar, Kurt, 2012. Anomaly detection in gamma ray spectra: A machine learning perspective. In: 2012 IEEE Symposium on Computational Intelligence for Security and Defence Applications. pp. 1–8.
- Vigneron, V., Morel, J., Lépy, M.C., Martinez, J.M., 1996. Statistical modelling of neural networks in γ -spectrometry. *Nucl. Instrum. Methods Phys. Res. Sect. A: Accel. Spectrometers Detect. Assoc. Equip.* 369 (2), 642–647.
- Weinstein, Marvin, Heifetz, Alexander, Klann, Raymond, 2014. Detection of nuclear sources in search survey using dynamic quantum clustering of gamma-ray spectral data. *Eur. Phys. J. Plus* 129 (11), 239.