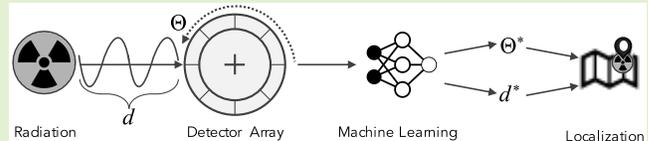


Improving Radioactive Material Localization by Leveraging Cyber-Security Model Optimizations

Ryan Sheatsley¹, Graduate Student Member, IEEE, Matthew Durbin, Student Member, IEEE, Azaree Lintereur, Member, IEEE, and Patrick McDaniel², Fellow, IEEE

Abstract—One of the principal uses of physical-space sensors in public safety applications is the detection of unsafe conditions (e.g., release of poisonous gases, weapons in airports, tainted food). However, current detection methods in these applications are often costly, slow to use, and can be inaccurate in complex, changing, or new environments. In this paper, we explore how machine learning methods used successfully in cyber domains, such as malware detection, can be leveraged to substantially enhance physical space detection. We focus on one important exemplar application—the detection and localization of radioactive materials. We show that the ML-based approaches can significantly exceed traditional table-based approaches in predicting angular direction. Moreover, the developed models can be expanded to include approximations of the distance to radioactive material (a critical dimension that reference tables used in practice do not capture). With four and eight detector arrays, we collect counts of gamma-rays as features for a suite of machine learning models to localize radioactive material. We explore seven unique scenarios via simulation frameworks frequently used for radiation detection and with physical experiments using radioactive material in laboratory environments. We observe that our approach can outperform the standard table-based method, reducing the angular error by 37% and reliably predicting distance within 2.4%. In this way, we show that advances in cyber-detection provide substantial opportunities for enhancing detection in public safety applications and beyond.

Index Terms—Machine Learning, security, gamma-ray detectors.



I. INTRODUCTION

THE integration of computation and sensing has revolutionized the management of physical spaces [12]. For example, new capabilities enable smart buildings that reduce energy use and lessen carbon footprints, smart homes which ease our personal lives, and smart infrastructures which support semi-autonomously secured spaces. Collectively, these Cyber-Physical Systems (CPS) are driving massive innovation, and in particular, advancing public safety in many domains. Security—the protection of physical spaces from adversaries who wish to manipulate or harm the space or those who reside

in it—is one of the important areas being advanced. Specifically, detection of adversarial entities, actions, or dangerous states is the principal use of physical-space sensors.

One of the most well known detection problems in physical security is the localization of radioactive materials. The use of nuclear technology has grown since the discovery of radiation [1] and it is now found in many applications, including power, medicine, and space. As the use of nuclear technology increases, so does the potential for misuse of radioactive materials. The canonical domain for discussing the detection of rogue radioactive materials is the container shipping industry (specifically, cargo inspection). Shipping containers are critical infrastructure, yet represent an ideal mode of transportation for adversaries due to their low cost [20] and the presence of large amounts of metal (and other shielding materials) which attenuates radioactive signals, substantially reducing the efficacy of radiation detectors to extract relevant signals from the surrounding noise [10]. As yet another example, urban search applications during large public gatherings (e.g., the Super Bowl, or Times Square during New Year's Eve) face similar challenges; the surrounding buildings can cause severe signal attenuation and impede search effectiveness.

Over the past few decades, a conventional technique for source localization (known as *Directional Gamma-ray Detection*) has relied on the use of pre-populated datasets (i.e., *reference tables*) calibrated at specific distances in

Manuscript received November 19, 2020; revised January 12, 2021; accepted January 13, 2021. Date of publication February 1, 2021; date of current version March 17, 2021. This work was supported in part by the U.S. Army Combat Capabilities Development Command Army Research Laboratory and was accomplished under Cooperative Agreement Number W911NF-13-2-0045 (ARL Cyber Security CRA), in part by the National Science Foundation under Grant CNS-1805310, in part by the U.S. Army Research Laboratory, and in part by the U.S. Army Research Office under Grant W911NF-19-1-0374. The associate editor coordinating the review of this article and approving it for publication was Dr. Ravibabu Mulaveesala. (Corresponding author: Ryan Sheatsley.)

Ryan Sheatsley and Patrick McDaniel are with the Department of Computer Science and Engineering, The Pennsylvania State University, State College, PA 16802 USA (e-mail: sheatsley@psu.edu).

Matthew Durbin and Azaree Lintereur are with the Department of Nuclear Engineering, The Pennsylvania State University, State College, PA 16802 USA.

Digital Object Identifier 10.1109/JSEN.2021.3055778

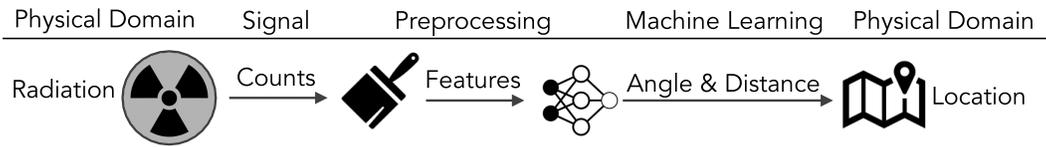


Fig. 1. Localizing Radioactive Materials - We apply data curation techniques, shown to be successful in cyber-detection applications, to counts collected from radioactive signals to predict both the angle to and distance from a radioactive source.

laboratory environments [11], [28]. In a process thematically similar to collecting malware signatures from known samples, the table is built using location templates. When a location needs to be screened, gamma ray detectors, placed in a fixed geometry, are used to acquire counts. The distribution of counts across the detectors are compared against these reference tables to predict if a source is present, and the corresponding detector-to-source angle. Even with these methods, a secondary phase is often necessary, where responders (paired with portable detectors) manually search for the radioactive material on foot [15], [26], [36]. A central limitation to this process is comparing the readings from the detectors to the reference tables: in non-trivial cases, attenuation, scattering, shielding, and other naturally occurring phenomena can significantly deviate the characteristics of the acquired gamma-ray signals. These factors limit the utility of reference tables, and are further exacerbated as the true distance to the radioactive source diverges from the calibration distance of the reference table. Localization of radioactive materials has a striking similarity to one of the most fundamental problems in cyber-security: detection. Similar to malware, spam, or network intrusion detection, an adversary hides a malicious artifact, and the challenge falls upon the defender to use information from the environment (often overburdened by noise) to detect and locate that artifact. For the past four decades, the security community has developed techniques that reveal information-rich artifacts and methods to amplify desirable signals from the surrounding noise. Our insight is that the application of techniques from cyber-security to radioactive material localization has the potential to produce significant advancements in physical security.

From a computer security perspective, the physical phenomena described earlier (e.g., attenuation, scattering, and shielding) inject noise into the readings. In fact, the presence of this noise causes naive application of machine learning to yield results worse than even the relatively inaccurate table-based approaches. We posit that applying and adapting feature scaling techniques used in cyber detection domains will mitigate the impact of noise and amplify the signal of interest. Specifically, we apply unit norm scaling (regularly used in spam detection) & robust feature standardization (often used in domains that have features with broad scales, i.e., network intrusion detection) to be able to discern the signal of interest from the noisy environment. With these techniques, we significantly improve the ability to localize radioactive materials. Further, we explore the abilities of machine learning models to estimate the distance to radioactive material, thus quantifying both angle and distance.

One of the new opportunities afforded by the novel application of cyber-security techniques to this physical domain is

that we are no longer bound to predicting directionality exclusively; we explore the abilities of machine learning models to estimate the distance to radioactive materials and show that they are effective in a suite of different environments. The application and adaptation of these techniques present a new capability that has not been achieved for radiation detection applications. This paper represents a significant step forward in Directional Gamma-ray Detection with the development of a novel framework to predict distance as well as direction (and thus location) with a stationary detection system.

In this paper, we present techniques which have the potential to advance the current capabilities for locating radioactive materials. We apply and adapt data curation techniques used successfully in cyber detection domains, and tune machine learning models to localize radioactive sources. We assess the approach using the Monte-Carlo-based radiation transport framework (Monte Carlo N-Particle Transport Code [9]) and physical experiments using radioactive sources in laboratory settings. We design the experiments to include obstructions that affect radioactive signals, which can serve as a proof-of-concept for cargo inspection and urban search scenarios. An overview of our approach is shown in Figure 1.

We evaluate this approach with six different models on seven datasets, of which five are simulated and two are experimental (collected in a laboratory environment). With the simulated experiments we find that the cyber-inspired approach reduces the angular error by 37% (4.9° , 95% CI ± 0.07 from the reference table to 3.1° , 95% CI ± 0.04 with our approach) and we can predict distance within 2.4%, 95% CI ± 0.54 of the source's location (up to 15 m). For the laboratory experiments, we reduce the angular error by 26% (8.5° , 95% CI ± 0.22 from the reference table to 6.3° , 95% CI ± 0.17 with unit norm scaling) and predict distance within 13.0%, 95% CI ± 3.74 of the radioactive materials (up to 3 m). Our contributions are:

- We present techniques adapted from cyber-security detection to exploit the use of gamma-ray signals for accurate radioactive material localization.
- We demonstrate that our approach surpasses the traditional table-based approach, incurring an average angular error of 3.1° , 95% CI ± 0.04 vs. 4.9° , 95% CI ± 0.07 , respectively.
- We extend the standard definition of localization to include distance. Here, the posited approach can predict distance within 2.4%, 95% CI ± 0.54 of a simulated radioactive source when the source strength is known.
- We perform experiments with real, radioactive sources to validate our findings in complex laboratory environments. Our approaches surpass the table-based method,

incurring an average angular error of 6.3° , $95\% \text{ CI} \pm 0.17$ vs. 8.5° , $95\% \text{ CI} \pm 0.22$, respectively. Moreover, we can predict distance within 13.0% , $95\% \text{ CI} \pm 3.74$ of real radioactive materials.

- We provide seven new datasets (including simulated and real data) that we make public for future research in this important domain, curated for use with machine learning.

II. PROBLEM DEFINITION

A. Threat Model

Our problem is essentially a game of hide and seek: an adversary places a radioactive source and the objective, as the defender, is to confirm its existence and determine its location. We assume a complex environment, that is, it contains obstructions (e.g., buildings) that interfere with (and thus, obfuscate) the signal produced by the radioactive source. Further, we assume a stationary environment: the adversary is non-adaptive and the radioactive source and the radiation detector are stationary (however, the distance and angle between the source and the radiation detector can change across different experimental scenarios), as are the surrounding obstructions in the environment. We also assume that the adversary has no ability to intervene with the operation of the radiation detector and that it is operating optimally (and therefore trust the produced readings to be as accurate as physical phenomena permit).

To some extent, we also assume no a priori knowledge of the radioactive material being used by the adversary (which we detail in Section IV). While the experiments are done exclusively with one material (due to availability, safety, and applicability), the method with which we detect and localize a radioactive source is agnostic to any gamma ray emitting isotope used by an adversary. The intuition behind this is straightforward: while different gamma ray emitting isotopes exhibit unique radioactive signatures, they all emit quanta within certain energy regions. Therefore, for this approach, detecting a particular isotope is simply a function of which portion of the energy region is scanned. Thus, a takeaway of this work is that responders can be “blind”, in some sense, to the specific isotope used by an adversary.

B. Detector Setup

The detection of a radioactive source relies upon the use of a detector array. The radiation produced by the source interacts with the detectors to generate a signal. This signal is then captured and used to produce a histogram corresponding to the energy deposited via different interactions. The detection system is comprised of four or eight detectors, each of which collect individual energy histograms. Here we sub-sample counts over a range of energies which are the inputs (i.e., features) to the machine learning models.

As mentioned in Section I, there are a suite of environmental factors that can negatively affect the readings of the detector (aiding the adversary). Figure 2 highlights some of the main sources of these environmental factors, and how they impact the signal. Notably, there are three central phenomena produced by obstructions: attenuation, scatter, and shielding. We describe in more detail how these phenomena affect the

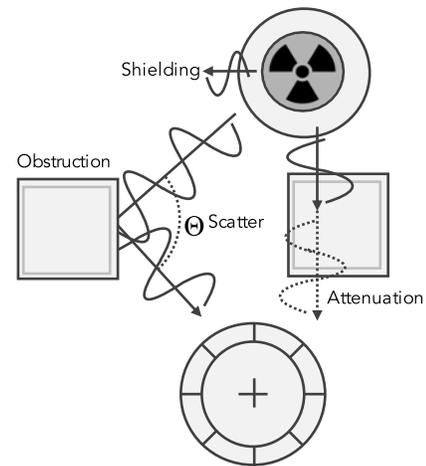


Fig. 2. Detection in a complex environment - The principal objective is to detect and localize radioactive sources in complex environments that induce undesirable phenomena, e.g., shielding, attenuation, scatter, and background noise.

readings in Section III, but for the purposes of the problem definition, these broadly just reduce the signal or amplify noise.

C. Machine Learning for Security

Machine learning has been successful in computer security detection applications including network intrusion detection, malware, and zero-day vulnerabilities [17], [29], [32]. However, machine learning has not been applied to radioactive source search scenarios, which shares many parallels with domains within computer security. A central factor for deploying machine learning in security-sensitive domains is feature scaling [23], [30], [31]. In network intrusion detection, there are many different kinds of features, which can contain outliers that can negatively affect standardization [30]. We find that detecting radioactive materials faces a similar burden, in that the distributions of gamma-ray counts can contain strong outliers (i.e., sources of noise). By applying robust feature standardization techniques that account for these outliers, we improve the accuracy of many learning algorithms.

As a second optimization, we take inspiration from techniques used traditionally for spam detection: unit norm scaling. Specifically, we observe that much like analyzing the relative frequency of words in emails, learning algorithms are likely to be more accurate in localizing radioactive materials with relative detector counts rather than gross signals.

III. RADIOACTIVITY IN THE PHYSICAL WORLD

Radiation, which is energy in transit, can be either electrically charged (e.g., electrons, protons, and alpha particles) or uncharged (gamma rays, x-rays, and neutrons). Uncharged radiation poses unique detection challenges, but is not easily shielded [16]. In this work, gamma-rays are the principal phenomena of interest as they are not readily shielded by thin metals (i.e., shipping containers), unlike charged particles [33]. Also, gamma-rays produce unique energy signatures [14] which can be used to classify the radioactive source, analogous to signatures produced by malware in intrusion detection systems. Generally, gamma-rays that pass through a detector

interact in one of three ways: the photoelectric effect, Compton scattering, or pair production. These interactions produce readings that eventually become the features of the approach.

Poisson Statistics: Absent of physical phenomena (and any detector deficiencies), one of the most fundamental challenges in interpreting detector readings is that they are burdened by Poisson statistics. The underlying uncertainty complicates accurate interpretation of the readings; the stochastic nature of radioactive decay means that **the exact same experiment repeated twice in a row will yield different results**. This fact gives a fundamental insight into what makes localizing radioactive sources a challenging problem.

In the simplest scenario (no obstructions, line-of-sight to the radioactive source, and ideal detector characteristics), the phenomena above, coupled with Poisson statistics, can have a notable affect on detector results, which table-based analysis approaches have difficulty rectifying. Much like network adversaries who obfuscate their signature to frustrate detection systems by mixing benign requests in the midst of malicious ones, noise produced by scattering and attenuation (as well as the fundamental uncertainty) can have a non-trivial negative effect on this approach with respect to localizing radioactive sources.

A. Existing Approaches

The two phase search procedure, which relies on detection and localization, is a demanding process, both in time and labor, given that responders must triangulate the radioactive source manually. Currently deployed techniques aim to combine both phases: by analyzing minute differences between counts received across detectors in an array of fixed geometry, the angle to the radioactive source can be determined (within some error). This problem is known as *directional gamma-ray detection*. Most conventional techniques use pre-populated datasets (i.e., *reference tables*) of known source locations, calibrated at a specific distance in lab environments [28] (shown in Figure 4). However, these approaches suffer in non-trivial cases where attenuation, scattering, shielding, and other naturally occurring phenomena affect the detected signals. More, these methods still require responders to manually search in the suspected direction of the radioactive source. Thus, they are susceptible to human-error, inaccuracies of the reference tables, and are bound by the number of responders that can be equipped with portable detectors to triangulate the source.

For this work, we measure the net counts for each individual detector and compute their differences to predict both angle and distance (which gives us localization). This combines two previously disjoint stages, enabling responders to quickly identify the location of the radioactive material. We hypothesize there is latent information that characterizes the environment which enables the location of the radioactive material to be determined. Consider that if one detector receives more counts than another, it is likely the source is in the direction of the detector with the highest counts. However, if an obstruction is directly in front of this detector, then the neighboring detectors may receive more counts. The challenge here is to capture these subtle situations—tools used in detection in

computer security environments are effective at pulling out this embedded information, and we exploit this observation in our analysis.

IV. APPROACH

Localizing radioactive materials in noisy environments shares many of the same challenges observed in the cyber-security detection space. Here, we briefly detail the radioactive source used, some relevant characteristics of the detector, the framework used in the simulated experiments, describe the feature scaling adaptations, and present the machine learning algorithms used.

A. Material Detected

Cobalt-60 (^{60}Co) is the radioactive isotope used in the simulations and laboratory experiments. Cobalt-60 is a relevant isotope to study, as it can be found in many domains, including medicine, industry, food, and nuclear power [13], [18], [34]. The widespread use of ^{60}Co means that, in practice, it is an isotope responders often wish to locate.. Most importantly, while we use ^{60}Co in the experiments, we emphasize that these techniques are not specific to this isotope; many radioactive isotopes have characteristic peaks similar to ^{60}Co , simply at different energies [16].

B. The Detectors

In gamma-ray spectroscopy, there are two main detector types: scintillators and semiconductors. Though most semiconductors offer better resolution and improved intrinsic efficiency (i.e., a high probability of interaction with gamma-rays), they are expensive, and some requiring cooling to liquid nitrogen temperatures ($-196\text{ }^\circ\text{C}$) [27]. Such requirements were impractical for this work.

Thus, we use thallium-doped sodium iodide (NaI(Tl)) scintillation detectors, popular in many field applications [19]. There are a handful of properties that make NaI(Tl) detectors useful for experimentation, namely: room-temperature operation, high efficiency, and large photofraction (i.e., the fraction of incident photons fully absorbed). The popularity and accessibility of the detector makes it an attractive choice for evaluating the applicability of this approach.

C. Monte Carlo N-Particle Transport Code

MONTE CARLO N-PARTICLE TRANSPORT CODE (MCNP) is a Monte Carlo method simulator for radiation transport¹. It uses Monte Carlo methods to simulate interactions (i.e., absorption, and scattering) as radiation propagates through a medium. Monte Carlo methods are considered to be the de facto standard for applications in radiation analysis due to their ability to accurately model radiation transport and interactions [9]. Due to its ability to simulate nearly any environment, MCNP is used in many fields, including medicine, detector design, reactor design, radiography, material penetration tests, radiation dosimetry, among others [9]. Figure 3 showcases the capacity of MCNP to model real radioactive phenomena—the simulated detector responses map nearly 1:1 onto the laboratory detector readings. Differences between simulated and laboratory readings are largely attributed to the effects of *gain-shift* [4], [25], stemming from temperature and

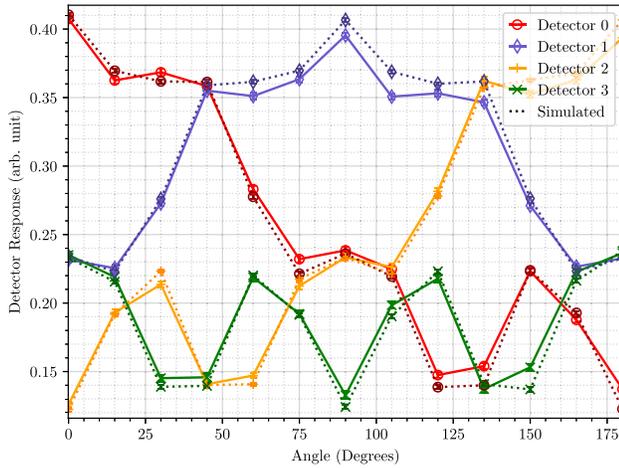


Fig. 3. Detector responses as a function of angle for laboratory experiments and simulations - MCNP is capable of modeling radioactive phenomena with striking accuracy and precision. Error bars on response values are included, but comparable in size to the plot markers. MCNP is the de facto tool used for radiation analysis, particularly large-scale experiments that are difficult to safely execute in live laboratory settings.

other environmental influences. Gain shift has the effect of “shifting” the spectra, which can lead to counts artificially being added or subtracted to the reading. As each detector experiences differing amounts of gain shift, these small effects can propagate to notable differences in the normalized input features. To partially mitigate this, energy calibrations are regularly performed.

D. Feature Selection

With a popular isotope and effective detector, we return to the central goal: localization of radioactive sources in complex environments. To achieve this goal, we first ask a basic question: *what can we measure?* We defer to historically successful techniques to answer this question.

As described in Section III, prior work has combined detection and direction (i.e., angle) into a single phase by analyzing slight variances in the responses of each detector in a fixed geometry array [6]. Here, we use gamma-ray counts as features, and extend the analysis to localizing (i.e., predict both angle and distance) radioactive sources. In this way, we seek to combine all phases of detection and localization into one step, averting cost for any special equipment and saving time by avoiding a manual foot search.

E. Feature Optimization & Cyber-Security

The abundance of research in localizing radioactive sources has many answers to the question asked above. However, there is yet another question that we must ask: *how do we filter noise without dampening the signal of interest?* To answer this question, our intuition leads us to our central hypothesis: techniques used in the cyber-security detection space can be useful in localizing radioactive materials. We describe the techniques and relevant adaptations below².

¹“Radiation Transport” software simulates the propagation of radiation through space and its interactions in media.

²It is worth noting that we tried other data manipulation techniques that had marginal (or sometimes even negative) effect on the accuracy of the models, namely: 0-1 rescaling, normalization (i.e., mean centered at 0), and standardization (i.e., mean centered at 0 and a standard deviation of 1).

1) **Robust Feature Standardization:** For many learning algorithms, standardizing feature scales is an important prerequisite. A common technique is to subtract the mean and scale to unit variance. However, standard techniques can be negatively affected by features that have highly skewed distributions, like those seen in network intrusion detection [30]. Instead, subtracting the median and scaling features according to some quantile range has been shown to produce better results for features with outliers. We standardize the features via:

$$\hat{x} = \frac{x - \tilde{x}}{\max Q_3 - \max Q_1} \quad (1)$$

where x is the original feature, \hat{x} is a standardized feature, \tilde{x} is the median value, and $\max Q_i$ is the maximum value for the i^{th} quantile. By scaling based on the maximum value in some quantile, we mitigate the negative influence outliers may have on the accuracy of some learning techniques.

2) **Unit Norm Scaling:** l_p normalization is an arguably uncommon feature scaling technique (l_2 regularization on model parameters is a common use of l_p norms in machine learning). If we represent a dataset as an $M \times N$ matrix of M samples and N features, then most feature scaling techniques operate across *all samples* (i.e., $M \times 1$)—that is, one particular feature for all samples is scaled in some manner. However, unit norm differs in that it operates across *one sample* (i.e., $1 \times N$). We can formulate unit norm scaling as:

$$\hat{x} = \frac{x}{\|x\|_{l_p}} \quad (2)$$

where \hat{x} is a l_p norm scaled input. Our intuition for using unit norm scaling follows successful applications of natural language processing towards spam detection: a natural objective for spam detection is to determine the *term frequency* of words (or n-grams) in an email. For example, it is difficult to draw any conclusions if the bigram “free money” appears a handful of times in email. However, stronger conclusions can be drawn if “free money” was the *most common* bigram. We apply this same reasoning to predicting the angle to a radioactive source: whether or not a detector receives 10 or 100 counts is hardly useful; what is more important is that a detector received *the most* counts (which is then the most likely direction to the radioactive source). This insight leads to substantial increases in angle prediction accuracy.

F. Reference Tables & Machine Learning

Here, we describe the reference tables and learning algorithms used in this work. The techniques presented here were specifically chosen, as they offer unique advantages over one another, such as interpretability, scalability, and accuracy. Further details are presented in Appendix .

1) **Reference Tables:** As described in Section III, reference tables are commonly used for directional gamma-ray detection. To build a reference table: detectors are setup in a fixed geometry, a known source is selected and placed at a fixed distance from the detector array, and the relative counts for each detector are recorded at varying source angles. These reference tables are ostensibly a closed-form approximation of the phenomena as they encode the response of a detector as a function of angle to the source, as shown in Figure 4.

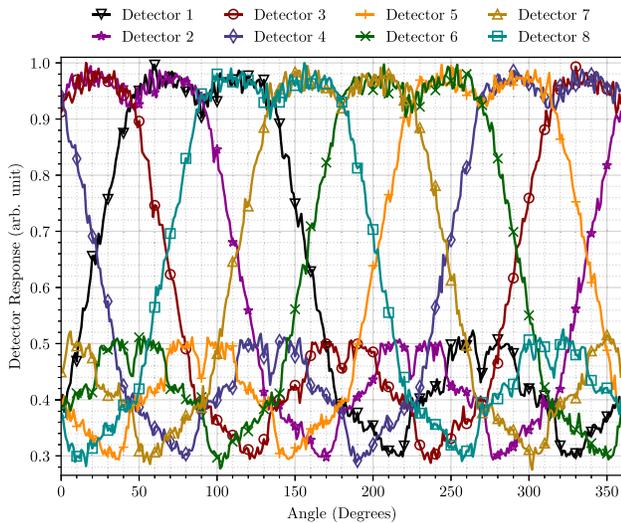


Fig. 4. MCNP simulated detector responses as a function of source angle - reference tables are built from these responses and are then used to approximate unknown sources. Calculated error bars are comparable in size to line thickness, and are thus excluded.

Once the reference table is calibrated, the responses of the detectors to an unknown source are compared to the reference table. Often with a least-squares regression analysis, where the angle to the unknown source is predicted by:

$$\theta_y = \arg \min_{\theta} (\Gamma(\theta) - x)^2 \quad (3)$$

where θ_y is the angle predicted by the reference table, $\Gamma(\theta)$ are the calibrated detector responses for some angle θ , and x is a vector of detector counts for a particular observation. Equation 3 leads us to two observations: 1) differences between the calibration environment and real world environment (such as the presence of obstructions) will lead to discrepancies in the relative detector responses, potentially leading responders in the wrong direction, and 2) often more severely, slight variations in distance can produce profoundly different distributions of counts than any observation for which the reference table was calibrated.

2) *Logistic Regression*: Logistic regression (LR), akin to linear regression, computes a weighted sum of input features with an additional bias term, and applies the logistic function to this sum [35]. While more sophisticated techniques have emerged, each have limitations. We include logistic regression models in this work to investigate if simpler models suffice to perform localization tasks accurately and quickly.

3) *Support-Vector Machines*: Prior to the inception of deep learning, support-vector machines (SVMs) dominated machine learning benchmarks across many domains [3]. SVMs are attractive as they can form non-linear decision boundaries, which may be necessary given the noisiness of this domain. k-Nearest

4) *k-Nearest Neighbors*: Neighbors (kNN) is a non-parametric approach [2]. As the observed count distributions in this domain can change rapidly in a variety of unique environments, kNN is particularly useful as it does not make any assumptions about the underlying data.

5) *Decision Trees*: Decision Trees (DT) are flexible machine learning algorithms that are commonly used today [24].

They require minimal data preparation, have low performance overheads, and offer intuitive explanations of the formed decision boundaries. Decision trees are appealing in this domain not only as another non-parametric technique, but also because of their white-box design: the learned decisions are easy to interpret, which can be useful in understanding subtle changes in the decision process as a function of the environment.

6) *Deep Neural Networks*: Deep neural networks represent a state-of-the-art class of learning techniques that have demonstrated success in the most challenging machine learning benchmarks [7]. Definitions vary, but generally speaking, deep neural networks often refer to any class of artificial neural networks with multiple layers between the input and output layers. For this work, the “deep neural network” is a fully-connected feedforward network, with some number of hidden layers, trained with back-propagation, using the rectifier (colloquially, “ReLU”) as the activation function.

V. EVALUATION

In this section, we evaluate the approach on seven datasets: five were simulated and two were measured³. The experiments were performed on a Dell Precision T7600 with Intel Xeon E5-2630 and NVIDIA Geforce TITAN X. We used Scikit-learn [22] for data curation and for instantiating the machine learning models. We ask:

- Is the naive machine learning solution sufficient for localizing radioactive sources?
- Can our optimizations approximate or exceed the performance of existing table-based angle predictions?
- When source strength is known, can distance be predicted and, if so, how is the accuracy affected as a function of distance from the radioactive source?

Summary: The simulation and physical laboratory experiments demonstrate that the developed techniques outperform the reference table for predicting angle by 37% (down from 4.9° to 3.1°) and estimate distance within 2.4% of the distance to a radioactive source.

A. Experimental Scale & Parameters

The simulated and laboratory datasets contain a radioactive source that is located 1–15 m and 1–3 m away from the detector array center, respectively. Simulated acquisitions were approximated as 14 s counts of a 1 and 10 μ Ci ^{60}Co source, while laboratory acquisitions were 5 minute counts of an approximately 1 μ Ci ^{60}Co source. With scaling (discussed below), these are common parameters in this field, and serve as a proof-of-concept for this analysis method.

The experiments contain three factors that are scaled in a manner that makes the datasets especially challenging: attenuation, measurement time, and distance. These datasets serve as benchmarks for evaluating approaches on real radioactive material subject to all forms of physical phenomena and environmental factors.

³The radioactive ^{60}Co source used in the laboratory experiments was a low activity source. At approximately 1 μ Ci, it is safe to handle with the appropriate safety procedures, which were defined and strictly followed for all the measurements.

1) *Attenuation*: Radioactive sources used by medical or industrial entities can be on the order of ~ 100 Ci (around eight orders of magnitude stronger than the real source used in this work) [13], [21]. These sources have the potential to be used by the adversary, and thus have activities similar to those which these methods could be used to localize. Separate simulations were conducted with the four detector array and ^{60}Co sources of various activities to gauge the scaling of obstruction thickness and attenuation. Results showed that the attenuation effect of a single cinder block (10 cm of solid concrete) on a $1 \mu\text{Ci}$ source is approximately the same as the attenuation effect of 150 cm of solid concrete on a 1 Ci source. Thus, the experiments model a challenging scenario where a weak radioactive source is in a thick concrete building.

2) *Sample Time*: We also modeled the measurement time in the simulated datasets to be comparable to the times used in realistic search scenarios [16]. Recall, radioactive sources decay at a particular *rate*—that is, a 1 Ci source sampled for 5 minutes will produce (theoretically) identical results to a $\frac{1}{2}$ Ci source sampled for 10 minutes (assuming a half-life much greater than the measurement time). Since these sampling times are comparable to realistic scenarios, but with a source potentially 8 orders of magnitude weaker than those in real scenarios, predicting angle and distance is challenging from a scaling perspective.

3) *Distance*: Most sources radiate in an isotropic manner, and thus fall subject to many laws in signal processing, notably the inverse-square law. This means that the intensity perceived by a detector decreases squarely with distance. Additional simulations were conducted to gauge the scaling effects on distance, to put the laboratory measurements in perspective to real world expectations. Results isolating the effects of geometry alone showed that the counts we receive with the $1 \mu\text{Ci}$ ^{60}Co source in five minutes at 3 m are comparable to the counts we would get from a 1 Ci ^{60}Co source in a single second at 100 m. So while the sources used in the laboratory experiments are challenging to detect at 0.5 to 3 m away (with short data acquisition times), they more than scale to parameters useful for actual source search applications.

B. Datasets

The first task was to create datasets of simulated and real-world measured gamma-ray counts in various settings to enable evaluation of the detection algorithm. We generated seven, which we will refer to as *S-Dataset 1*, *S-Dataset 2* $10^{6/7}$, *S-Dataset 3* $10^{6/7}$ for the simulated experiments (from MCNP), *L-Dataset 1*, and *L-Dataset 2* for the laboratory experiments. The 10^6 and 10^7 variants of S-Dataset 2 and 3 describes the number of simulated gamma-rays per trial. The different numbers of simulated gamma-rays correspond to different source strengths or measurements times, and represent differing levels of statistics. All scenarios use the same radioactive source (^{60}Co) at varying distances, angles, and obstruction locations, as described below and summarized in Table I.

1) *Simulated Datasets*: The datasets were generated with an 8-detector array and a source at varying distances between 1 and 15 m. S-Datasets 1 & 2 contain 72,000 samples where

TABLE I
DATASET STATISTICS FOR THE EXPERIMENTS

Dataset	Obstruction Size (m)	Obstruction Location	Angle ($^\circ$)	Distance (m)	Num. Trials
S-Dataset 1			0–360	1–15	72,000
S-Dataset 2	$1 \times 2 \times 5$	Fixed	0–360	1–15	72,000
S-Dataset 3	$0.5 \times 2 \times 5$	Moving	0–90	1–15	27,000
L-Dataset 1			0–90	0.5–3	125
L-Dataset 2	$0.5 \times 2 \times 5$	Fixed	0–90	0.5–3	125

the source is uniformly rotated a full 360° around the detector array (at roughly 1° increments), while S-Dataset 3 contains 27,000 samples and the radioactive source is only rotated 90° . For each trial, either 10^6 or 10^7 gamma-rays were simulated, corresponding to a 14 s count of a $1 \mu\text{Ci}$ or a $10 \mu\text{Ci}$ ^{60}Co source. S-Datasets 2 & 3 contained a solid concrete obstruction that mimics the effects of a concrete building. In S-Dataset 2, the obstruction was stationary; in S-Dataset 3, the obstruction was randomly placed between ten locations. Table I provides experiment details.

While one million gamma-rays may sound significant, recall the isotropic nature of radiation and the variety of physical phenomena described in Section III; in reality, less than 0.1% of these gamma-rays will cause some interaction (either positively or negatively) with the detector array 3 m from the source.

2) *Laboratory Datasets*: The two laboratory datasets were acquired with a 4-detector array setup (shown in Figure 6) and a radioactive source at varying distances between 0.5 and 3 m. Both datasets contain 125 samples where a $3 \mu\text{Ci}$ ^{60}Co radioactive source is rotated 90° at (roughly) 15° increments around the detectors. L-Dataset 1 has no obstructions and L-Dataset 2 contains concrete obstructions at fixed locations. A summary of the data is presented in Table I, and a photograph of the detector setup and accompanying block diagram are shown in Figures 5 & 6.

While 125 samples per dataset may seem small, it is both significant in this context and sufficient. A single sample often requires approximately 5 minutes to collect (i.e., nearly 11 hours of data collection for one dataset). Moreover, throughout this entire process, we regularly performed energy calibrations on the detector, and periodically acquired separate background radiation spectra to make the readings as accurate as possible⁴. Finally, some learning models (e.g., SVMs) are performant on small datasets. Thus, these relatively small laboratory datasets represent a challenge in localizing radioactive sources when data may be severely limited.

C. Experiment Overview

This section details experiments exploring how the proposed models predict angle and distance as compared to the reference tables. In the following figures, the accuracy is the number of samples where the exact angle (or distance⁵) was predicted

⁴Detectors in reality experience what is known as “gain shift”—the energy spectrum for radiation slowly changes overtime from a variety of environmental factors. Additionally, the background radiation can vary with time and location. For radioactive source search scenarios, a single calibration is often sufficient, however, data collection for scientific use, such as this work, requires recalibrating the detector for gain shift and background radiation regularly to obtain accurate measurements.

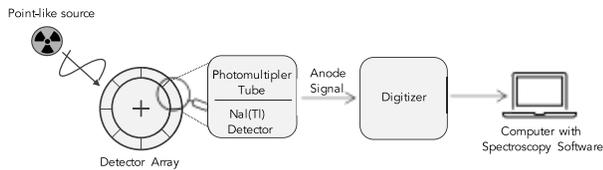


Fig. 5. Detection in Radioactive Environments - The gamma-rays from the ^{60}Co source interact with the detector array, consisting of NaI detectors which emit light upon gamma-ray interactions, photomultiplier tubes (PMTs) to convert the light into an analog pulse, and a digitizer to convert the analog pulse into a digital signal. The digital signal is then processed by spectroscopy software to convert the signals into the “counts”, which are used as input to the machine learning models.



Fig. 6. Setup of L-Dataset 2 - This is a photograph of the detector array in the laboratory. The ^{60}Co source is attached to the cardboard tube. The concrete blocks are the obstructions and cause the behavior of the gamma-rays produced by the ^{60}Co source to be representative of an urban environment.

correctly over the total number of samples. We highlight these results as, in real scenarios, even $\pm 5^\circ$ angular tolerance may be unacceptable (particularly if the radioactive source is estimated to be far away).

1) Naive Machine Learning: The naive approach to this problem is to directly use the raw counts collected by the detector as inputs to learning models. As shown in [Figure 7 \(a\)](#), the results are average: the reference table had average angular error of 4.9° , $95\% \text{ CI} \pm 0.07$ while the best model (k-Nearest Neighbor) had an average angular error of 4.2° , $95\% \text{ CI} \pm 0.11$. The logistic regression model and deep neural network were unable to predict neither angle nor distance correctly. For the distance at which the reference table was calibrated (around 200 cm; where the peaks are), the reference table outperformed all of the models. As the distance increased, the reference table quickly became inaccurate, unlike Decision Trees and k-Nearest Neighbors. However, neither of these algorithms eclipsed the accuracy of the reference table at any particular distance compared to the maximal accuracy of

⁵The distances are binned (and approximated) via the Freedman–Diaconis Estimator, which is an outlier-resilient, optimal binning algorithm [8]. Importantly, the estimator suggests bin widths so that the difference between the empirical and theoretical probability distributions are minimal. Reported accuracy is the number of samples where the bin was predicted over the total number of samples. For the simulations and laboratory experiments, 42 and 8 bins were created, i.e., 35 and 3.75 cm per bin, respectively.

the reference table (i.e., the distance it was calibrated for). For the laboratory experiments, the best models performed *worse*: the reference table had an average angular error of 8.5° , $95\% \text{ CI} \pm 0.22$, while the best model had an average angular error of 12.9° , $95\% \text{ CI} \pm 0.37$. Thus, for predicting angle, the simple application of machine learning yields results worse than reference tables at their calibrated distance, marginally better at other distances for the simulated experiments, and explicitly worse for the measured datasets.

For predicting distance ([Figure 7 \(b\)](#)), the naive approach produced impressive results with some of the algorithms for scenarios in which the source strength is assumed to be known. The best model (k-Nearest Neighbors) could predict distance within 2.5%, $95\% \text{ CI} \pm 0.55$ of the distance to a source. However, for the simulated data, most of the learning techniques were not able to predict distance at all, while decision trees and k-Nearest Neighbor produced a sigmoid-like curve for the ^{60}Co 10^6 simulation. Perhaps not surprisingly, these models can estimate (nearly perfectly) the distance to sources that are exceedingly close, yet struggle for sources that are relatively far away.

2) Applying Cyber-Security Detection Techniques: As detailed in Section IV, the radioactive environment is inherently burdened by noise, similar to intrusion detection domains, and thus we suspected that a cyber detection approach would readily apply in this domain. [Figure 8 \(a\)](#) demonstrates the results. Immediately, we can see significant improvements: many of the models now exceed the reference table accuracy, even at the distance at which the reference table was calibrated. After applying unit norm scaling, the average angular error for the best model (k-Nearest Neighbors) is 3.1° , $95\% \text{ CI} \pm 0.04$ (down from 4.2°), an improvement from the reference table by 37% (down from 4.9°). For the laboratory datasets, we also see improvements: we reduce the angular error to 6.3° , $95\% \text{ CI} \pm 0.17$ (down from 12.9°), an improvement from the reference table by 26%.

For predicting distance, we applied robust feature standardization. Like network intrusion detection, this domain is inherently noisy and contains outliers that may negatively influence standard feature scaling techniques. [Figure 9](#) shows a sample distribution of detector counts with the quartile ranges we scale from in the experiments and [Figure 8 \(b\)](#) shows the results. A small improvement is made to the overall accuracy after applying robust feature standardization (from 2.5%, $95\% \text{ CI} \pm 0.55$ to 2.4%, $95\% \text{ CI} \pm 0.54$ for the best models) and substantial gains for the other models (e.g., logistic regression, support vector machines, and deep neural networks) as shown in [Figure 8](#). For the laboratory experiments, robust feature standardization also yielded small improvements (from 13.5%, $95\% \text{ CI} \pm 3.66$ to 13.0%, $95\% \text{ CI} \pm 3.74$).

We highlight key takeaways of this work:

- Our approach can far surpass the capabilities of reference tables *even for the distance at which the table was calibrated*. This demonstrates that: 1) calibrations do not lend themselves well to the complex nature of problems in real environments, and 2) model-based approaches, paired with cyber-security detection techniques, are effective tools for localizing radioactive sources.

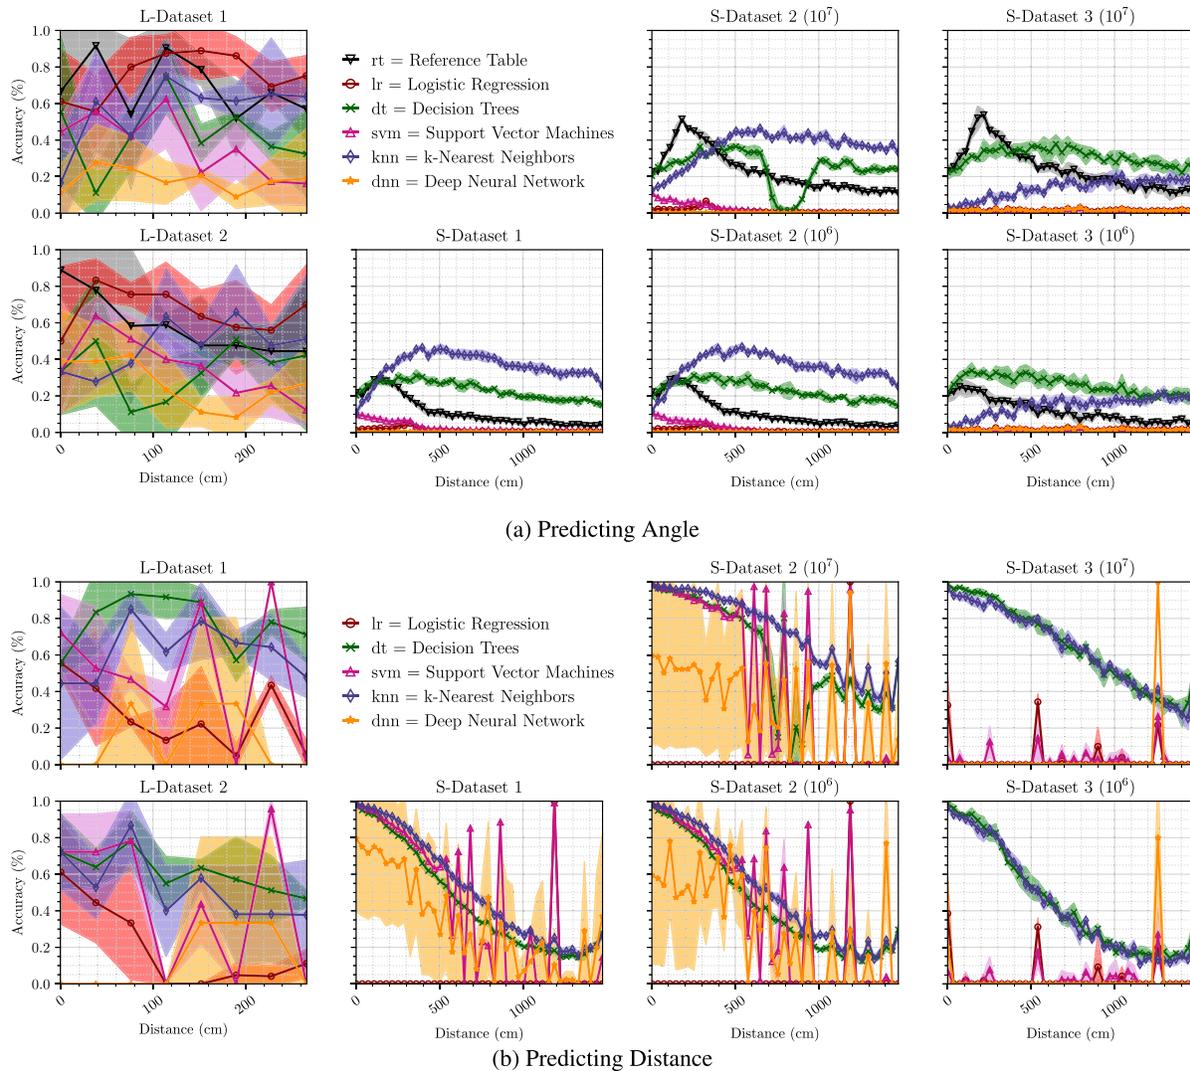


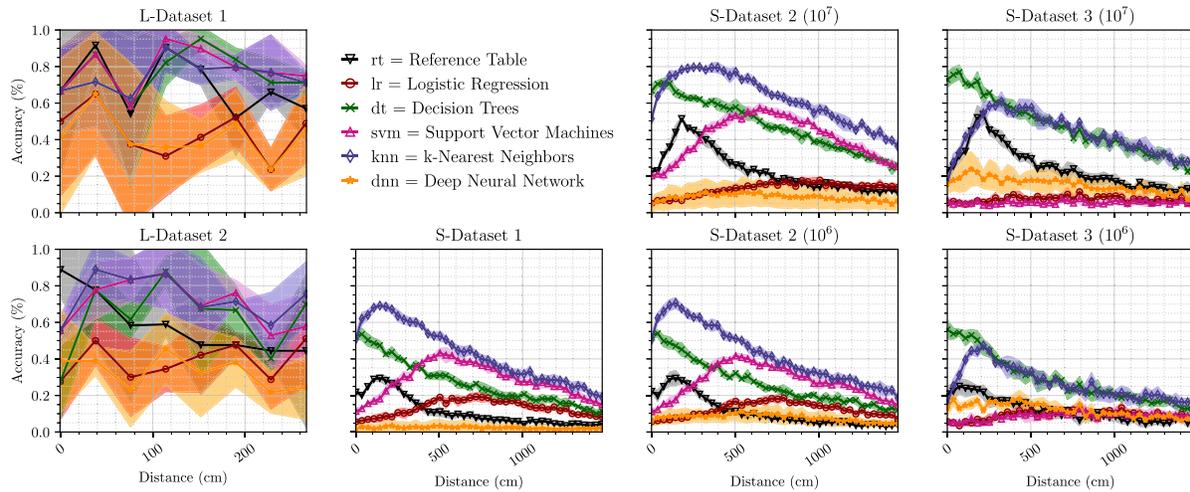
Fig. 7. Localizing radioactive sources with naive machine learning - Applying machine learning to raw detector counts produced average results for predicting angle. Estimating distance was acceptable for some learning techniques.

- Our approach accurately estimates distance to a radioactive source. Prior to this work, techniques either required mobile detectors (either first responders on foot or vehicles) to triangulate radioactive sources manually; now, we are void of these limitations. We can localize a radioactive source simultaneously at the time it is detected. While initial trials benefited from an a priori knowledge of the source strength, similar standardization and additional cyber-physical security techniques are being explored to apply distance predictions to sources of unknown strength.
- Perhaps not surprisingly, obstructions have a tangible impact on modeling radioactive behavior: most of the models observed an $\sim 10\%$ decrease in accuracy in the most challenging datasets where the locations of obstructions varied. Radioactive source search scenarios in reality will likely observe similar challenges given that no environment is identical.
- An order of magnitude increase in particle counts (10^6 to 10^7) was especially helpful to increase model accuracy at the longest distances (i.e., greater than 8.4 m). In other

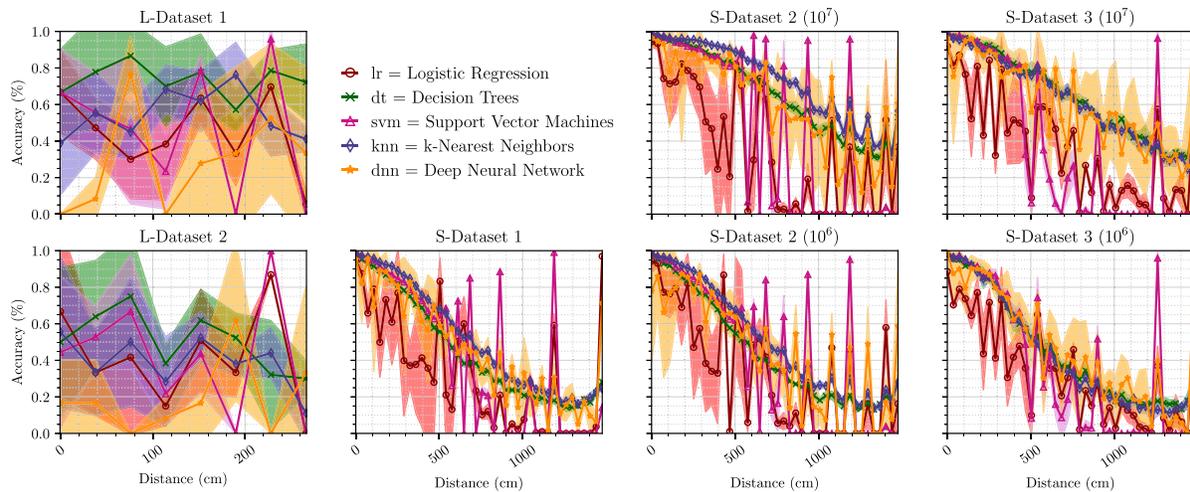
words, high activity sources (or longer acquisition times) can be localized accurately over long distances.

VI. OBSERVATIONS

3) *Unit Norm Scaling*: One of the most significant improvements we observed for angle prediction was the application of unit norm scaling. There are multiple reasons why this technique was so effective: much like detecting spam in emails, the absolute frequency of words is hardly useful; instead, it is often more interesting to see how frequent some words are used *relative to one another*. The intuition is straightforward: if the bulk of an email contains words that are commonly associated with spam, then the email is most likely spam as well; that is to say, long emails that contain 80% “spam words,” for example, are fundamentally no different (in terms of spam or not) than short emails with a similar relative amount of spam words. We follow this same reasoning for localizing radioactive materials: when receiving a sum total of 1000 counts or 100 counts, if a particular detector receives the majority then, in both cases, the radioactive source is most likely in front of this particular detector. This has the added



(a) Predicting Angle



(b) Predicting Distance

Fig. 8. Localizing radioactive sources with cyber-security detection techniques - applying unit norm scaling and robust outlier standardization lead to significant improvements for most of the models.

benefit of augmenting the training set—learning approaches no longer have to disentangle that 1000 counts or 100 is relatively meaningless for predicting angle, as those two situations are being treated identically. These observations give insight into why this feature scaling technique was so effective.

4) Robust Outlier Standardization: We found that scaling the features in a manner robust to outliers was effective for predicting distance. While using the raw counts was acceptable for sources that were close to the detectors, we noticed that the accuracy of the models decreased quickly as the distance linearly increased (i.e., the Inverse-square law in practice). We observed that robustly scaling features helped maintain the accuracy of the models over longer distances. Figure 9 led us to our insight: we aim to emphasize the signal from within the two dotted lines as the bulk of the counts indicated that the source was directly behind the detector in this example. However, due, in part, to the physical phenomena described in Section III, the detector observed a small increase in counts directly in front of it (i.e., at 0°). Thus, we hypothesized that mitigating the influence of these outliers would aid in

predicting distance. The results demonstrate that this insight was indeed helpful.

5) Estimating Distance: While these approaches are relatively accurate at predicting distance, estimating distance is difficult, especially with a stationary system. Today, there is a focus on using mobile systems to localize radioactive sources (e.g., Mobile Urban Radiation Search (MURS) [5]). By taking multiple samples at different locations, mobile systems can exploit basic triangulation algorithms to estimate the distance to a source (perhaps with more accuracy than our initial approaches). A natural limitation of this approach is the cost, requirement of a mobile environment, and dependency on multiple samples for triangulation. However, the fact that the success of our approaches are agnostic to these requirements yields a unique opportunity; much like how predictions with the stationary system studied benefited from the application of cyber techniques, the same is likely true for the techniques used in systems such as MURS. Broadly speaking, the overlap between radiation detection and detection in computing environments suggests that a large class of approaches are likely

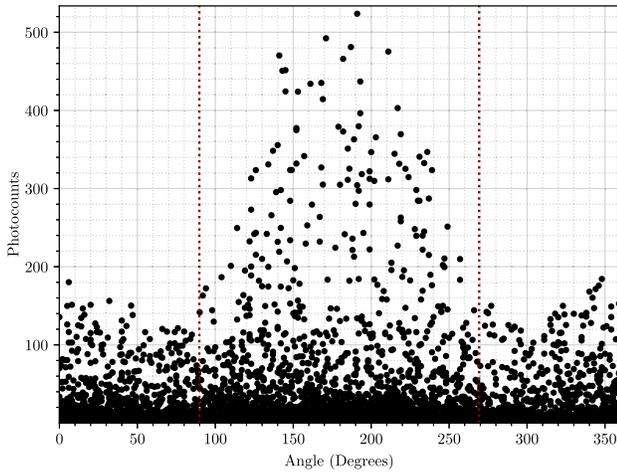


Fig. 9. Demonstrating Outliers - This figure demonstrates a photocount distribution for one of the detectors. The dotted lines mark the quantile ranges of the data that we scale to. This gives intuition into why robust feature standardization is helpful: in this figure, the radioactive source is directly behind the detector (as shown by the peak at 180°), yet there is a noticeable increase in photocount readings directly in front of the detector (i.e., 0° and 359°)—these increased counts are manifestations of noise in part from the physical phenomena.

to benefit from cyber-inspired approaches, as observed with the system used in this work.

VII. CONCLUSION

In this paper, we investigated new analysis approaches for radioactive source localization. We explored how techniques from the cyber-security detection domains can surpass traditional table-based approaches and extend the standard definition of localization to include distance. We observed through both simulated and physical laboratory experiments that our techniques surpassed the angular accuracy of table-based approaches, reducing the angular error by 37% and reliably predicting distance within 2.4%.

Moreover, we observed how naive applications of machine learning either produced inaccurate predictions or failed to eclipse the accuracy of table-based approaches at distances for which the table was calibrated. Yet, curated applications, such as unit-norm scaling and robust standardization, can produce results which surpass table-based approaches across all evaluated distances. This work demonstrates that the complex signals produced during source localization efforts can benefit from machine learning approaches with curated application of signal amplification techniques. Future efforts will focus on more robust distance prediction techniques and greater collection of physical measurements in controlled laboratory settings. Through our application of cyber-security principles, we introduced a state-of-the-art approach in localizing radioactive sources in complex physical scenarios.

APPENDIX

We provide a brief summary of the evaluated learning techniques and show our hyperparameters used in the evaluation.

Logistic Regression: Logistic regression (LR), akin to linear regression, computes a weighted sum of input features with

an additional bias term, and applies the logistic function to this sum [35]. We can define a logistic regression model as follows:

$$\ell = (1 + e^{(-\theta^\top \cdot x)})^{-1} \quad (4)$$

where θ represents the weights of the model and x represents our input vector. While more sophisticated techniques have emerged, i.e., deep learning, each have their own limitations. We include logistic regression models in this work to investigate if simpler models suffice to perform localization tasks accurately and quickly

Support-vector Machines: Prior to the inception of deep learning, support-vector machines (SVMs) dominated machine learning benchmarks across many domains [3]. For binary classification, SVMs seek to find two hyper-planes that satisfy:

$$\theta^\top \cdot x + b \leq \{-1, 1\} \quad (5)$$

where θ represents the weights of the model, x represents our input vector, and $\{-1, 1\}$ encode the two classes. The algorithm then seeks to maximize the difference between the two hyperplanes, while satisfying the above constraints. SVMs are attractive as they can form non-linear decision boundaries, which may be necessary given the noisiness of this domain.

k-Nearest Neighbors: The k-nearest neighbors algorithm (kNN) is a non-parametric approach, used ubiquitously in academia and industry. As the observed count distributions in our domain can change rapidly in a variety of unique environments, kNN is particularly useful as it does not make any assumptions about the underlying data. We can define the kNN algorithm as follows (using the Euclidean metric for distance):

$$\hat{y} = \arg \min_{y:(x,y) \sim \mathcal{D}} \sqrt{(\hat{x} - x)^2} \quad (6)$$

where \hat{y} is predicted class, $(x, y) \sim \mathcal{D}$ represents the input-class pairs for the observed distribution, and \hat{x} represents a new observation at the time of inference. As a non-parametric model, we expected kNN to exhibit adequate accuracy, amortized across a variety of unique scenarios.

Decision Trees: Decision Trees (DT) are flexible machine learning algorithms that are commonly used today [24]. They require minimal data preparation, have low performance overheads, and offer intuitive explanations of the formed decision boundaries. Decision trees often follow a binary if-then-else structure whose rules are built by minimizing:

$$T(i, t) = \frac{|x| : x_i \leq t}{|\mathcal{D}_T|} \cdot G(i, \leq t) + \frac{|x| : x_i > t}{|\mathcal{D}_T|} \cdot G(i, > t) \quad (7)$$

where i a feature, t is the threshold for i , x are input vectors from distribution \mathcal{D}_T partitioned at node T , and $G(i, t)$ (i.e., the Gini impurity score) is:

$$G(i, \leq t) = 1 - \sum_y \left(\frac{|x| : x_i \leq t, (x, y) \sim \mathcal{D}_T}{|x| : x_i \leq t} \right)^2$$

TABLE II
MODEL HYPERPARAMETERS

Logistic Regression	Solver: L-BFGS Epochs: 100 1.0 L_2 Regularization
Decision Trees	Criterion: Gini Max Depth: 10
Support Vector Machines	Kernel: RBF Gamma: $1/(\# \text{ features} * \text{ variance})$ 1.0 L_2 Regularization
k-Nearest Neighbors	Neighbors: 5 Algorithm: Ball Tree (Leaf size of 30) Distance: Minkowski
Deep Neural Network	Layers: 3 with 15 neurons (fully-connected) Activation: Relu Epochs: 300 Optimizer: Adam 10^{-4} L_2 Regularization

where $(x, y) \sim \mathcal{D}_T$ represents the input-class pairs for the observed distribution. Decision trees are appealing in this domain not only as another non-parametric technique, but also because of their white-box design: the learned decisions are easy to interpret, which can be useful in understanding subtle changes in the decision process as a function of the environment.

Deep Neural Networks: Deep neural networks represent a state-of-the-art class of learning techniques that have demonstrated success in the most challenging machine learning benchmarks [7]. Definitions vary, but generally speaking, deep neural networks often refer to any class of artificial neural networks with multiple layers between the input and output layers. We can formalize a deep neural network (with ReLU as the activation function) as:

$$P(x) = S(P^\ell(\max(0, \theta_\ell^\top \cdot x + b))) \quad (8)$$

where x is an input vector, P^ℓ is the ℓ^{th} iterate of P (i.e., function composition) where ℓ is the number of layers in the network, θ_ℓ are the weights for the ℓ^{th} layer, b is a vector of biases, and S is defined as the softmax layer. Given the computational complexity required of artificial neural networks with many hidden layers, we interested if deep neural networks were expressive enough to be effective in our most complex scenarios.

ACKNOWLEDGMENT

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Combat Capabilities Development Command Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation and the U.S. Army Research Office.

REFERENCES

- [1] *Nuclear Development: Nuclear Energy Today*, OECD Nuclear Energy Agency, Paris, France, p. 114.
- [2] N. S. Altman, "An introduction to kernel and nearest-neighbor nonparametric regression," *Amer. Statistician*, vol. 46, no. 3, pp. 175–185, Aug. 1992.
- [3] E. B. Boser, M. I. Guyon, and N. V. Vapnik, "A training algorithm for optimal margin classifiers," in *Proc. 5th Annu. Workshop Comput. Learn. Theory (COLT)*, New York, NY, USA. Pittsburgh, PA, USA: Association for Computing Machinery, 1992, pp. 144–152.
- [4] R. Casanovas, J. J. Morant, and M. Salvadó, "Temperature peak-shift correction methods for NaI(Tl) and LaBr₃(Ce) gamma-ray spectrum stabilisation," *Radiat. Meas.*, vol. 47, no. 8, pp. 588–595, Aug. 2012.
- [5] J. C. Curtis *et al.*, "Simulation and validation of the mobile urban radiation search (MURS) gamma-ray detector response," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 954, Feb. 2020, Art. no. 161128.
- [6] M. Durbin, R. Sheatsley, C. Balbier, T. Grieve, P. McDaniel, and A. Lintereur, "Development of machine learning algorithms for directional gamma ray detection," in *Proc. Inst. Nuclear Mater. Manage.*, Jun. 2019, pp. 1–10.
- [7] L. Terrence Fine, S. L. Lauritzen, M. Jordan, J. Lawless, and V. Nair, *Feedforward Neural Network Methodology*, 1st ed. Berlin, Germany: Springer-Verlag, 1999.
- [8] D. Freedman and P. Diaconis, "On the histogram as a density estimator: L2 theory," *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, vol. 57, no. 4, pp. 453–476, Dec. 1981.
- [9] T. J. Goorley *et al.*, "Initial MCNP6 release overview—MCNP6 version 1.0," Los Alamos Nat. Lab., Los Alamos, NM, USA, Tech. Rep. LA-UR-13-22934, 1086758, Jun. 2013.
- [10] M. D. Grypp, C. M. Marianno, J. W. Poston, and G. C. Hearn, "Design of a spreader bar crane-mounted gamma-ray radiation detection system," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 743, pp. 1–4, Apr. 2014.
- [11] D. Hanna, L. Sagnières, P. J. Boyle, and A. M. L. MacLeod, "A directional gamma-ray detector based on scintillator plates," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 797, pp. 13–18, Oct. 2015.
- [12] A. Humayed, J. Lin, F. Li, and B. Luo, "Cyber-physical systems security—A survey," *IEEE Internet Things J.*, vol. 4, no. 6, pp. 1802–1831, Dec. 2017.
- [13] A. Chmielewski, "Gamma irradiators for radiation processing," *Ind. Appl. Chem. Sect. Int. At. Energy Agency, IAEA, Vienna, Austria*, Tech. Rep. 83-909690-6-8, 2006.
- [14] H. E. Johns and J. R. Cunningham, *The Physics of Radiology*, 4th ed. Springfield, IL, USA: Charles C. Thomas, 1983.
- [15] R. T. Klann, J. Shergur, and G. Mattesich, "Current state of commercial radiation detection equipment for homeland security applications," *Nucl. Technol.*, vol. 168, no. 1, pp. 79–88, Oct. 2009.

- [16] F. G. Knoll, *Radiation Detection and Measurement*, 3rd ed. New York, NY, USA: Wiley, 2000.
- [17] M. Kubat, R. C. Holte, and S. Matwin, "Machine learning for the detection of oil spills in satellite radar images," *Mach. Learn.*, vol. 30, no. 2, pp. 195–215, 1998.
- [18] G. R. Malkoske and J. Slack, "Cobalt-60 production in CANDU power reactors," *Can. Nucl. Soc. Bull.*, vol. 23, no. 3, pp. 16–21, 2002.
- [19] S. Mukhopadhyay, R. Maurer, R. Wolff, E. Smith, P. Guss, and S. Mitchell, "Networked gamma radiation detection system for tactical deployment," *Proc. SPIE*, vol. 9593, Aug. 2015, Art. no. 959316.
- [20] B. Kempinski, "Scanning and imaging shipping containers overseas: Costs and alternatives," Congr. United States, Congressional Budget Office, Washington, DC, USA, Tech. Rep. 51478, Jun. 2016.
- [21] P. Ortiz, M. Oresegun, and J. Wheatley, "Lessons from major radiation accidents," *Safety*, vol. 21, no. 1, p. 230, 2002.
- [22] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [23] D. Protić and M. Stanković, "Anomaly-based intrusion detection: Feature selection and normalization influence to the machine learning models accuracy," *Eur. J. Eng. Formal Sci.*, vol. 2, no. 3, p. 101, Dec. 2018.
- [24] J. R. Quinlan, "Induction of decision trees," *Mach. Learn.*, vol. 1, no. 1, pp. 81–106, Mar. 1986.
- [25] P. L. Reeder and D. C. Stromswold, "Performance of large NaI (TI) gamma-ray detectors over temperature-50 degrees C to+ 60 degrees C," Pacific Northwest Nat. Library, Richland, WA, USA, Tech. Rep. 14735, 2004.
- [26] L. A. Remick, L. J. Crapo, and R. C. Woodruff, "US national response assets for radiological incidents," *Health Phys.*, vol. 89, no. 5, pp. 471–484, Nov. 2005.
- [27] G. B. Saha, *Physics and Radiobiology of Nuclear Medicine*. New York, NY, USA: Springer, 2001.
- [28] C. Schrage, N. Schemm, S. Balkir, M. W. Hoffman, and M. Bauer, "A low-power directional gamma-ray sensor system for long-term radiation monitoring," *IEEE Sensors J.*, vol. 13, no. 7, pp. 2610–2618, Jul. 2013.
- [29] R. Sommer and V. Paxson, "Outside the closed world: On using machine learning for network intrusion detection," in *Proc. IEEE Symp. Secur. Privacy*, Oakland, CA, USA, 2010, pp. 305–316.
- [30] J. Song, H. Takakura, Y. Okabe, and Y. Kwon, "A robust feature normalization scheme and an optimized clustering method for anomaly-based intrusion detection system," in *Advances in Databases: Concepts, Systems and Applications*, vol. 4443, R. Kotagiri, P. R. Krishna, M. Mohania, and E. Nantajeewarawat, Eds. Berlin, Germany: Springer, 2007, pp. 140–151.
- [31] M. Tahir, M. Li, N. Ayoub, and M. Aamir, "Efficacy improvement of anomaly detection by using intelligence sharing scheme," *Appl. Sci.*, vol. 9, no. 3, p. 364, Jan. 2019.
- [32] C.-F. Tsai, Y.-F. Hsu, C.-Y. Lin, and W.-Y. Lin, "Intrusion detection by machine learning: A review," *Expert Syst. Appl.*, vol. 36, no. 10, pp. 11994–12000, Dec. 2009.
- [33] *NRC: Radiation Basics*, U.S. NRC, Rockville, MD, USA, 2020.
- [34] S. Walker, "EPA facts about cobalt-60," U.S. EPA, Washington, DC, USA, Tech. Rep 176332, 2014.
- [35] E. R. Wright, "Logistic regression," in *Reading and Understanding Multivariate Statistics*. Washington, DC, USA: American Psychological Association, 1995, pp. 217–244.
- [36] K. P. Ziocok, "The lost source, varying backgrounds and why bigger may not be better," in *Proc. AIP Conf.*, Washington, DC, USA, vol. 632, 2002, pp. 60–70.