

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Data-driven Approaches to Spectral Gamma-ray Source Detection and Identification

Permalink

<https://escholarship.org/uc/item/59k407r4>

Author

Bilton, Kyle James

Publication Date

2020

Peer reviewed|Thesis/dissertation

Data-driven Approaches to Spectral Gamma-ray Source Detection and Identification

by

Kyle James Bilton

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Nuclear Engineering

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Kai Vetter, Chair

Doctor Tenzing Joshi

Professor Fernando Pérez

Professor Karl van Bibber

Fall 2020

Data-driven Approaches to Spectral Gamma-ray Source Detection and Identification

Copyright 2020
by
Kyle James Bilton

Abstract

Data-driven Approaches to Spectral Gamma-ray Source Detection and Identification

by

Kyle James Bilton

Doctor of Philosophy in Nuclear Engineering

University of California, Berkeley

Professor Kai Vetter, Chair

The ability to detect and identify gamma-ray sources by means of analyzing gamma-ray spectra is essential for nuclear safety and security, and accurately performing these tasks in environments with varying background radiation remains a challenge. One common approach to enhancing detection capabilities is directing research and development at novel detection materials and systems. Alternatively, detection sensitivity can be enhanced by making use of more sophisticated data processing methods on existing detection systems. Leveraging advances in data analysis methods, this dissertation introduces and characterizes novel data-driven approaches to spectral detection and identification. An emphasis is made on methods that can potentially be augmented with complementary non-radiological data (e.g., video streams), with the objective of enhancing performance by constraining models using information about the local environment.

Two general data analysis methods are examined for both detection and identification: non-negative matrix factorization (NMF) and neural networks. When applied to gamma-ray spectra, NMF yields accurate and interpretable models of background and sources using relatively few parameters. Neural networks are considered for their flexibility in design, the significant amount of active research in the area, and the ease with which models can be augmented with additional data sources. For both the NMF and neural network models, detection and identification methods are introduced, the performance of each is evaluated relative to benchmarks from the literature, and an assessment on tradeoffs, specifically as they relate to practical considerations, is discussed. The methods introduced in this work provide improvements over the examined benchmarks, and each method can be applied to existing systems. Additionally, discussion is provided on the potential to extend each method further using complementary non-radiological data.

Contents

Contents	i
List of Figures	iv
List of Tables	xi
List of Abbreviations	xii
1 Introduction	1
1.1 Nuclear Safety and Security	1
1.2 Radiation Signatures for Nuclear Safety and Security	3
1.3 Detection Systems for Nuclear Safety and Security	3
1.4 Dissertation Outline	4
2 Elements of Spectral Gamma-ray Source Detection and Identification	6
2.1 Gamma-ray Physics	6
2.1.1 Gamma-ray Interactions with Matter	7
2.1.2 Attenuation	9
2.2 Gamma-ray Detection and Measurement	9
2.2.1 Scintillation Detection	10
2.2.2 Spectral Binning	11
2.3 Sources of Gamma Radiation	11
2.3.1 Gamma-ray Background	11
2.3.2 Gamma-ray Sources of Interest	13
2.4 Model for Spectral Detection and Identification	14
2.4.1 Anomaly Detection	14
2.4.2 Identification	15
2.5 Maximum Likelihood Estimation	15
2.5.1 Deviance	17
2.6 Decision Theory	17
3 Spectroscopic Data Sources for Algorithm Development and Evaluation	20
3.1 Source Injection	20

3.1.1	Background Data for Source Injection	21
3.1.2	Source Injection using Background-Subtracted Data	22
3.1.3	Source Injection using Simulated Data	23
3.2	Source Injection using RadMAP Data	29
3.2.1	Background Data Preparation	30
3.2.2	Source Injection Data Preparation	32
3.3	Source Injection using Urban Radiation Search Competition Data	33
3.3.1	URSC Background Data	34
3.3.2	Source Simulations	35
3.4	Performance Metrics	36
4	Non-negative Matrix Factorization of Gamma-ray Spectra	40
4.1	Dimensionality Reduction	41
4.1.1	Principal Component Analysis	41
4.1.2	Poisson Principal Component Analysis	42
4.2	Non-Negative Matrix Factorization	42
4.3	NMF for Background Modeling	45
4.4	Anomaly Detection using NMF	46
4.4.1	Benchmark: Gross Counts $K\sigma$	47
4.4.2	AD_i in the Limit of Known Background	47
4.4.3	Performance Evaluation	48
4.4.4	Estimating Uncertainties due to Empirical Thresholds	50
4.5	Source Identification using NMF	51
4.5.1	Likelihood Ratio Tests	52
4.5.2	Analytical Modeling of Likelihood Ratio Test Statistics	53
4.5.3	Benchmark: Region of Interest	56
4.5.4	ADI_i in the Limit of Known Background	57
4.5.5	Performance Evaluation	57
4.6	NMF for Source Spectrum Modeling	59
4.6.1	Model Comparisons	61
4.6.2	Performance Evaluation: Bare Source	63
4.6.3	Performance Evaluation: 50 mm Steel Shielding	63
4.7	Concluding Remarks	63
5	Neural Networks for Detection and Identification	67
5.1	Overview of Artificial Neural Networks	68
5.1.1	Fully-connected Layers	68
5.1.2	Convolutional Layers	69
5.1.3	Recurrent Layers	70
5.1.4	Network Training	71
5.2	Related Work	71
5.2.1	Early Work: 1990s - 2000s	71

5.2.2	Recent Work: Late 2010s - 2020	72
5.2.3	Starting Points for Improvement and Innovation	73
5.3	Anomaly Detection using Autoencoders	74
5.4	Source Identification	75
5.5	Model Training and Evaluation	77
5.5.1	Performance Evaluation and Data	78
5.5.2	Model Optimization	79
5.5.3	Benchmarking	80
5.6	Results	80
5.6.1	Anomaly Detection	80
5.6.2	Identification	83
5.7	Conclusion	83
6	Concluding Remarks	87
6.1	Summary	87
6.2	Future Work	88
A	Sigmoid Maximum Likelihood Estimation Derivation	90
B	Probability of Detection Sampling	92
B.1	Fixed Activity Sampling	92
B.2	Adaptive Sampling	92
B.3	Adaptive Bayesian MDA Estimation	93
C	Spectral Anomaly Detection Derivation	94
C.1	Principal Component Analysis	94
C.2	Spectral Anomaly Detection	95
	Bibliography	97

List of Figures

2.1	Example gamma-ray background spectrum from an array of NaI(Tl) scintillator detectors on the RadMAP mobile detection system, introduced in Section 3.2. Overlaid is the mean spectrum from 10,000 randomly-sampled spectra, and the most prominent background lines from Table 2.1 are annotated. The spectrum was formed using a 1-s integration time and 128 square root bins between 50 and 3000 keV. Additional details on how the spectrum was formed are given in Section 3.2.	13
3.1	Diagram showing the processing steps to produce training and testing spectra for use in source injection. Background events, spectral bin edges \mathcal{E} , and the integration time Δt are used as inputs to generate background spectra, which are then split into data for model training and evaluation.	22
3.2	Detector and source setup for angular response simulations. A 2" \times 4" \times 16" NaI(Tl) bar is centered at the origin, while monoenergetic photons from a distant source, drawn as parallel lines, make an angle θ with the x -axis. Note that by symmetry, angular points only need to be sampled in the range $\theta \in [0, 90^\circ]$	24
3.3	Detector response $\eta(E, \theta)$ for four different energies and two incident angles. The response is slightly lower at $\theta = 75^\circ$ due to having a lower effective area from this incident angle.	25
3.4	Diagram summarizing the procedure used to create source angular templates $\psi_s(\theta)$ for arbitrary radionuclide types using simulated list mode data. List mode data from Monte Carlo simulation, along with other parameters, are fed into the processing chain, ultimately resulting in source templates.	27
3.5	Source templates $\psi(\theta = 0^\circ)$ and $\psi(\theta = 85^\circ)$ for ^{131}I . There is a noticeable difference in the magnitude of the two templates due to the lower effective area at higher angles.	28
3.6	Diagram showing the kinematics of a source pass. The detector system, represented by the small rectangle, moves in the direction of the arrow and passes the source, represented by the circle. During the pass, the source spectrum is computed at each time step t_i based on the angle θ_i and source-detector distance $\ \mathbf{r}_i\ $	29

- 3.7 Figure from [16]. The RadMAP mobile detection system and its various radiological and contextual sensors. In this work, data from the 10×10 NaI(Tl) array and the GPS sensor are used. 30
- 3.8 (Left pane) A "waterfall" plot of a sequence of spectra past a source during data collection in Berkeley. The source, seen as a spot in the low-energy region of the spectrum, appears shortly after the 20 s mark. (Right pane) The mean spectrum over the waterfall plot compared to the mean spectrum over the entire background dataset. The measured spectrum is consistent with a ^{99m}Tc source, which is used in SPECT imaging. Note that this sequence was measured near Alta Bates Hospital in Berkeley, consistent with the hypothesis that these anomalous gamma-ray events were emitted from a patient. 31
- 3.9 Figure from [46]. The left panel of the figure represents a top view of a single street in the simulated URSC environment. The rectangles in this figure represent buildings, each of which has a unique geometry and background emission signature, resulting in a highly variable sequence of measurements as the system moves down the street. The right three panels of the figure show the effect of placing a source at three different locations on the same street, where the color indicates the source intensity as a function of position. These three panels illustrate the variation in spectral measurements due to environmental scattering. 33
- 3.10 (Left pane) Gross count rate as for the first 30 seconds of three randomly-sampled background-only runs from the URSC training set. Transient peaks are seen, resulting from the detector system passing a building or environmental feature with elevated emission rate. In addition to variability in gross count rate, there is temporal variation in each spectral bin, reflecting variability of the type of radionuclide emission encountered as a function of time. (Right pane) Probability density function of gross counts across all background spectra in the URSC training dataset compared to the probability mass function of a Poisson distribution having a mean equal to the mean from background samples. Due to the variability in background environments, the variance of gross counts from samples greatly exceeds that of a Poisson distribution with the same mean. 35
- 3.11 (Left pane) Normalized two-dimensional histogram of background spectra from the URSC dataset, illustrating the variability in spectra encountered in the URSC dataset. A high degree of variability is seen around the mean spectrum, drawn as a solid black line. (Right pane) Correlation matrix of background dataset, illustrating the relationship between energy bins that is not captured in the figure in the left pane. 36

- 3.12 Sigmoid functions used for modeling probability of detection as a function of source activity. At each activity, 100 samples are drawn from the true distribution with parameters $\mu^* = 155 \mu\text{Ci}$, $\sigma^* = 15 \mu\text{Ci}$ and fit to the model in Equation (3.16), resulting in estimated parameters $\hat{\mu}^* = 154.6 \mu\text{Ci}$, $\hat{\sigma}^* = 16.5 \mu\text{Ci}$. In characterizing detection algorithms for particular source scenarios, the underlying parameters μ^* and σ^* are unknown and are estimated by fitting Equation (3.16) to measured points. Using Equation (3.17), this example would result in an MDA of $203.2 \mu\text{Ci}$ using the estimated parameters $\hat{\mu}^*$ and $\hat{\sigma}^*$ 38
- 4.1 Three Poisson NMF basis vectors \mathbf{V} formed from background spectra collected by the RadMAP mobile detection system (upper left) and histograms of weights \mathbf{A} corresponding to each component (lower left). The first three principal components formed from background spectra (upper right) and histograms of weights corresponding to each component (lower right). For both NMF and PCA, models were formed using 86,400 randomly-sampled spectra, each with one-second integration time, and for NMF, training was done over 30,000 fit iterations. Note that there is no significance in the numbering of NMF components, unlike in PCA, however, NMF components were numbered to match the order of corresponding PCs. 46
- 4.2 Probability of detection of ^{133}Ba (top) and ^{137}Cs (bottom) sources at 20-m standoff using NMF-based detection algorithms, a PCA-based detection algorithm, and both NMF and PCA-based algorithms in the case of a known background, each model using $k = 2$ components at $1/8 \text{ hr}^{-1}$ FAR and 1-s integration time. Solid lines show the measured performance, and dashed lines show performance in the KBL. Note that only anomaly detection is performed here, meaning sources are not identified by the algorithms. The 95% probability of detection is indicated by the dashed horizontal line. Error bars indicate the 68% Jeffreys interval. Not shown is a curve for gross counts, which has a significantly higher MDA. 49
- 4.3 Comparison of a fit to a spectrum containing ^{133}Ba using only the background model \mathbf{V} (top) and using the background and source model \mathbf{V}' . The upper pane shows the weighted background components, as well as their sum. The lower pane shows the sum of the background components, the weighted source template, and the sum of the background and source components. By including the source template, the model is able to fit the 356 keV ^{133}Ba peak. An $85\text{-}\mu\text{Ci } ^{133}\text{Ba}$ source at 20-m standoff distance is used. 52

- 4.4 Distributions of likelihood ratio test statistic $D_s(\mathbf{x})$ used in testing for the presence of ^{133}Ba (top) and ^{137}Cs (bottom). The distributions are formed by computing the difference of negative log-likelihoods between source and background models \mathbf{V}_s and background only models \mathbf{V} over a background dataset containing 55 hours of NaI spectra collected by RadMAP. The distribution is used to empirically select a threshold based on a target FAR. The vertical lines show thresholds for particular FARs, given in the legend. 54
- 4.5 Histograms of likelihood ratio test statistics for ^{60}Co (left column) and ^{137}Cs (right column) computed using background data compared to the analytical model in Equation (4.15) for $m = 1$ (top row) and $m = 2$ (bottom row). The reduced χ^2 goodness-of-fit measure is reported for each fit, showing that the models fit reasonably well, suggesting Equation (4.15) can be used for modeling the probability distribution of test statistics for hypothesis testing. In each case, the inverse CDF of the probability distribution can be computed and used to yield the threshold for a given probability of false alarm. 55
- 4.6 Thresholds used in likelihood ratio tests, computed using Equation (4.16). The left column shows thresholds for $m = 1$, where sources are approximated with a single vector, and the right column shows thresholds for $m = 2$, used for the two-component NMF source representation. N is used to represent the effective number of independent hypotheses being tested. In the analysis in this work, $N_s = 1$, meaning a single source is tested at a time, so N represents the number of shielding configurations considered. This plot shows that for $N_c \geq 5$, a two-component NMF source model results in a lower threshold, assuming both models are able to describe variation in the source sufficiently well. 56
- 4.7 Probability of detection and identification for ^{133}Ba (top) and ^{137}Cs (bottom) using the ADI_P , ADI_{L2} , NMF-based algorithms in the case of a known background, each using, $k = 2$ components, and an ROI algorithm. Solid lines show the measured performance, and dashed lines show performance in the KBL. Spectra were formed using 1-s integration time for 99 NaI detectors on RadMAP, and source injection was performed using 20-m standoff distance and 6.7 ± 1.3 m/s vehicle speed. For both sources, a $1/8 \text{ hr}^{-1}$ FAR was used. Both the ADI_P and ADI_{L2} methods are able to achieve a 95% probability of detection, indicated by the dashed horizontal line, at a much lower activity than the ROI algorithm. Error bars indicate the 68% Jeffreys interval. 58

- 4.8 Normalized detector responses, referred to as spectral templates, for both ^{60}Co (left) and ^{137}Cs (right) under five source shielding configurations: no shielding, 50 mm of concrete, 50 mm of lead, 50 mm of steel, and 100 mm of steel. The templates were formed by creating spectra from simulated 1 mCi sources at a 10-m standoff distance, then normalizing the spectra to 1. For both sources, variability in spectral shape is seen to depend on the amount and type of shielding surrounding the source. An ideal gamma-ray source identification algorithm is able to accurately identify a gamma-ray source in the presence of various shielding types. 60
- 4.9 NMF components \mathbf{V}_s generated for ^{60}Co (left pane) and ^{137}Cs (right pane). For both sources, NMF generates physically-interpretable components that can be related to the full-energy peaks and the downscattering continuum associated with shielding. Intuitively, a source under a given shielding configuration can be thought of as a linear combination of the full-energy peak(s) of the source, and a downscattering contributions due to the particular shielding configuration present. 61
- 4.10 Probability of detection curves for unshielded ^{60}Co (top) and ^{137}Cs (bottom). For each activity, 500 background spectra were randomly sampled, and down-sampled source spectra for a source at a 10-m standoff were injected into the background. While each of the three models results in similar performance, there is an advantage to using the bare source model, as it is able to accurately model the source using fewer parameters than the other two approaches. . . . 64
- 4.11 Probability of detection curves for ^{60}Co (top) and ^{137}Cs (bottom) shielded by 50 mm of steel. Both the learned and library models provide improvements over the bare template model due to their ability to model variation in the spectral shape due to the shielding. 65
- 5.1 A dense connection between layer $\mathbf{h}_i \in \mathbb{R}^4$ and layer $\mathbf{h}_{i+1} \in \mathbb{R}^2$, which can be understood from left to right. The lines between element h_{ij} in the left (i.e., first) layer and element $h_{i+1,k}$ in the right (i.e., second) layer are represented as a column i in the matrix $\mathbf{W} \in \mathbb{R}^{2 \times 4}$. Not shown is an additional bias term $\mathbf{b}_i \in \mathbb{R}^2$ and nonlinear activation function $\sigma(\cdot)$ used to create the state $\mathbf{h}_{i+1} = \sigma(\mathbf{W}_i \mathbf{h}_i + \mathbf{b}_i)$ 69

- 5.2 (Left panel) Diagram showing the dimensionality of features $\mathbf{h}^{(i)}$ at each layer i for an example dense autoencoder architecture with five hidden layers. A 128-bin spectrum is input into the autoencoder, and dense layers are computed by performing nonlinear transformations on each preceding layer. The inverse of each operation is then performed to decode the latent features, resulting in a smoothed spectrum. (Right panel) An input background spectrum \mathbf{x} and corresponding autoencoder reconstruction $\hat{\mathbf{x}}$ are shown. When trained on background, the autoencoder learns spectral features such as background peaks and the associated downscattering continuum. Both the input and output spectra shown here contain 128 bins that scale with the square root of energy. Note that any apparent deviations between the input and output spectra (e.g., at the 1460 keV peak) are due to low-statistics, as the bins of the measured spectrum \mathbf{x} are discrete random samples of the mean Poisson rate $\hat{\mathbf{x}}$ 76
- 5.3 Example architecture of a convolutional identification network, in a similar fashion to refs. [69, 81]. A 2-dimensional feature map resulting from convolutional operations is flattened into a single feature vector of length 1024, and this is reduced down to the output size of 18 (17 sources, 1 background channel). A max-pooling operation is applied to the features resulting from the convolutional operation, reducing feature size from 128 to 64. In the case of an RNN, the dense layer with size (1, 128) at time t is fed back to combine with the previous layer at time $t + 1$. Not shown here is a softmax function that the output is fed into. 78
- 5.4 Histogram of mean MDA for autoencoders evaluated on the validation set. Each model, 40 models in total, was trained using a random value of the L2 regularization coefficient λ , and random configuration of number of neurons in the dense layers of the network. This figure shows that despite being trained with different parameters, initial weights, and mini-batches, most were able to yield similar performance. The model corresponding to the lowest mean MDA from this figure is examined further on the test set. 81
- 5.5 Comparison of MDA for the three detection methods across all 17 sources at a $1/8 \text{ hr}^{-1}$ FAR. Sources are sorted in ascending order MDA for the baseline NMF method. Each model was evaluated by injecting each source type across activities into each run of the background test set and computing the MDA. The background used, the test set, was separate from the training and validation background, and thus gives a sense of how well each model generalizes to unseen background data. Note that the discrepancy between the PCA-based method and the other two is likely due to the detection metric used for the PCA-based approach, which comes from the literature. The error bars shown were computed from Equation (3.18). Note that there is an overlap between DAE and PCA for ^{111}In and between NMF and PCA for ^{75}Se 82

- 5.6 Histogram of mean MDA for RNN and feedforward (FF) ID networks evaluated on the validation set. Each model, 40 in total, was trained using a random value of the L2 regularization coefficient λ , number of kernels in the convolutional layer, and number of neurons in the first dense layer. This distribution shows a general trend of improvement when using recurrent layers. The models corresponding to the lowest mean MDA for both feedforward and recurrent networks are examined further on the test set. 84
- 5.7 Comparison of the three methods evaluated on the test set: NMF-based identification, a feedforward network (FF), and an RNN-based identification method. Sources are sorted in ascending order MDA for the baseline NMF method. A total FAR of approximately $1/8 \text{ hr}^{-1}$ across all sources was achieved by setting a threshold for each source individually set based on an effective FAR of $1/(8 \times 17) \text{ hr}^{-1}$. The RNN is seen to generally provide an improvement over its feedforward counterpart, though there are a few notable examples, such as ^{133}Ba . Note that there is an overlap of points between NMF and FF for ^{60}Co , ^{123}I , ^{131}I , and ^{192}Ir 85

List of Tables

2.1	Prominent gamma-ray energies of primordial radionuclides [24]	12
3.1	Parameters affecting gamma-ray source detection	21
3.2	Simulated gamma-ray sources and their most prominent gamma-ray energies .	37
4.1	MDA_{95} (μCi + statistical error) for anomaly detection	50
4.2	Comparison of empirical anomaly detection thresholds	51
4.3	Lower and upper estimates of MDA_{95} (μCi + statistical error) for anomaly detection	51
4.4	Energy windows (keV) used in ROI-based source identification	57
4.5	MDA_{95} (μCi + statistical error) for simultaneous detection and identification . .	59
4.6	Comparison of MDA (μCi) for NMF-based identification using different source approaches	63

List of Abbreviations

AIC Akaike Information Criterion

ANN Artificial Neural Network

CDF Cumulative Distribution Function

CNN Convolutional Neural Network

FAR False Alarm Rate

KBL Known Background Limit

MDA Minimum Detectable Activity

MLE Maximum Likelihood Estimation

NMF Non-negative Matrix Factorization

PCA Principal Component Analysis

ROI Region of Interest

SAD Spectral Anomaly Detection

SVD Singular Value Decomposition

URSC Urban Radiation Search Competition

Acknowledgments

The research described in this dissertation was primarily facilitated by members of the Applied Nuclear Physics (ANP) program at Lawrence Berkeley National Laboratory, and I would like to thank all of those involved. First, to my advisor, Professor Kai Vetter. Kai's support and encouragement has enabled me to explore a range of projects and research topics over the years, ultimately leading to this dissertation. To my mentor Tenzing Joshi, who I could always come to for sound research advice and insightful feedback. To Mark Bandstra, for being a major source of inspiration regarding creative problem solving. To everyone from ANP over the years, including, but not limited to, Ren Cooper, Joey Curtis, Andy Haefner, Victor Negut, Ryan Pavlovsky, Marco Salathe, Erika Suzuki, and Brian Quiter. To ANP member and former classmate Daniel Hellfeld, with whom I've had countless interesting conversations with over bowls of clam chowder at The Stuffed Inn (on that note, I also thank the entire staff of The Stuffed Inn). To Grey Batie for being a good friend and colleague, and always asking interesting questions during my talks. Also, thanks for going with me to see Waka Flocka.

At UC Berkeley, I would like to thank Professors Fernando Pérez and Karl van Bibber, for taking time to serve on my qualifying exam and review this dissertation, and for the engaging courses I have taken with them. Also, to my colleagues at the Department of Nuclear Engineering and the Berkeley Applied Research on the Imaging of Neutrons and Gamma-rays (BeARING) group.

Outside of LBNL and UC Berkeley, there are many others I could not have done this without. Colin Mickle and Hannah Rohret, thank you both for your close friendship throughout my time in Berkeley. And lastly, thank you to my family, and everyone else who has been here along the way.

Chapter 1

Introduction

Technologies involving nuclear reactions and radioactive decay continue to serve humanity in numerous ways, such as generating electricity, enabling the diagnosis and treatment of disease, and enhancing crop yields [1]. Radiological materials and radiation-producing machines do, however, pose a risk to human health, especially when precautions are not taken or if materials are misused. These potential hazards present the need for technologies to detect and localize sources of radiation. This dissertation introduces novel methods which serve to aide in the detection of radioactive sources by means of gamma-ray spectral analysis, specifically by building mathematical models which attempt to discern anomalous sources from benign gamma-ray background radiation. The remainder of this chapter introduces the foundation for this work, including a detailed description of the source search problem, general remarks about current approaches, and the introduction of new concepts that may enhance detection capabilities of radioactive sources outside of regulatory control.

1.1 Nuclear Safety and Security

The concepts and methods introduced in this work are discussed in the context of *nuclear safety* and *nuclear security*. The International Atomic Energy Agency (IAEA) Safety Glossary [2] defines nuclear safety as

"The achievement of proper operating conditions, prevention of accidents and mitigation of accident consequences, resulting in protection of workers, the public and the environment from undue radiation risks."

This glossary also defines nuclear security as

"The prevention and detection of, and response to, criminal or intentional unauthorized acts involving nuclear material, other radioactive material, associated facilities or associated activities."

While nuclear safety and security differ on their focus, one common facet between the two is the need for the ability to detect radioactive sources. Regarding nuclear safety, detection is essential to activities such as radiation contamination remediation following a radiological incident. For example, following the 2011 accident at the Fukushima Daiichi Nuclear Power Plant [3], spectroscopic radiation detectors were used to characterize dose rates and geospatial distributions of particular radionuclides released in the accident. A wide variety of sensors and analyses were used to measure these quantities in the air and ground [4], and to even detect radiological anomalies on sea floor [5]. In each of these cases, the measurements can be used to inform actions, for example, by providing information on areas to avoid.

Though nuclear safety includes the safe use of radiation and limiting risks to public health, nuclear security is focused on reducing the possibility of misuse of radiological materials. There are a number of ways to approach the goals of nuclear security. In addition to safeguarding radiological materials at their point of intended use (e.g., nuclear power plants), measures are taken to detect the transit or presence of such materials outside of these places. For instance, radiation detection systems are sometimes deployed at high-risk events, such as sport matches or political events. Prior to UEFA EURO 2012 football championship in Kiev, the IAEA transferred ownership of a mobile radiation detection system to Ukraine [6] in order to "strengthen the country's nuclear security efforts." At the time of this writing, the National Nuclear Security Administration (NNSA) is conducting aerial radiation detection surveys in Washington, D.C. to establish background levels prior to the Presidential Inauguration scheduled for January 20, 2021 [7]. In both of these examples, mobile systems enable the detection of anomalous sources, providing a line of defense to attendees and important public figures. Such efforts by the NNSA are considered to be part of the broader Global Nuclear Detection Architecture [8], an effort by the U.S. government to detect illicit radioactive material in three domains: the U.S. interior, U.S. borders, and exterior regions (e.g., points of origin in foreign territories).

Though some may question the necessity of such approaches, especially considering the costs and effort involved, it is an unfortunate reality that nuclear material has been accessible outside of regulatory control in the recent past. The IAEA Incident and Trafficking Database is a record of interdictions of Special Nuclear Material (SNM), which is nuclear material that can, in sufficient quantities, be used in a weapon. While the entire database is restricted to all but a few at IAEA and certain government agencies, a number of incidents are publicly known to some degree, for example, one involving the confiscation of nearly 3 kg of highly-enriched uranium in St. Petersburg, Russia in 1994 [9]. More recently, a 3000-curie ^{60}Co source, considered by the IAEA to be in the most dangerous class of sources due to its strength, was stolen from a truck in Mexico City [10]. While the theft is not believed to have been intended as part of a terrorist act, a source of that magnitude presents a potentially lethal hazard to those near it. As unlikely as such threats may seem, the potential impact of the distribution of SNM and other radiological material is severe enough to consider approaches for mitigating these risks.

1.2 Radiation Signatures for Nuclear Safety and Security

With nuclear safety and security in mind, the possible signatures that can be used to detect radiological sources are examined. Broadly speaking, there are two classes of technologies used for detecting radioactive sources: active and passive detection systems. In active methods [11], such as radiography, an external radiation source (e.g., an x-ray generator) is used to interact with potential radiation sources, generating an output signal that indicates the presence or absence of a source. Passive systems instead detect and measure ionizing radiation emitted directly from radiological materials. Due to the cost, complexity, and potential hazards presented by active interrogation methods, passive systems are often more appropriate in many practical situations and are exclusively considered in this work, though the methods presented here could be augmented for use in some active methods. Though the process of nuclear decay often results in a variety of particles, including photons, electrons, neutrons, and heavy charged particles, due to attenuation effects, photons and neutrons are the only particles which can be detected at the distances required for safety and security purposes (i.e., up to tens of meters). Neutrinos and antineutrinos are currently being considered as a candidate for nuclear proliferation detection [12] but because of the low probability of interaction, neutrino detection requires massive detection volumes (e.g., kilotons of detection material) and reactor-scale sources, making them infeasible for most applications. Due to neutron emission in spontaneous fission seen in SNM, the detection of neutrons is particularly useful in the detection of SNM. However, because of the rich information they provide for performing source identification, gamma-ray spectra are exclusively examined in this work.

The aim of gamma-ray detection for nuclear safety and security is to detect and identify as weak sources as possible, and to do so using sensors that are relatively simple to operate and are also relatively inexpensive in order to increase the accessibility to systems. A common approach to enhancing sensitivity to weak gamma-ray sources is to improve the spectrometers themselves. In particular, more advanced materials and detector electronics are often developed, promising enhanced detector properties (e.g., energy resolution and efficiency) relative to standard fielded instruments. While many promising candidates for novel detection materials have been introduced in recent years, thallium-activated sodium iodide (NaI(Tl)), remains the standard in most applications. As a result, it is worth considering how existing systems can be improved by means of sophisticated data processing. This dissertation introduces concepts and methods for enhancing detection performance by means of improved spectroscopic gamma-ray data processing.

1.3 Detection Systems for Nuclear Safety and Security

To achieve the objectives of nuclear safety and security by means of gamma-ray spectroscopy, complete systems for detecting and analyzing gamma radiation measurements are used. A broad distinction is made between stationary and mobile gamma-ray de-

tection systems, both of which have their own use cases and challenges. Stationary systems are generally placed at locations such as international border crossings and shipping ports [13, 14]. Mobile detection systems can be managed to fit a wide range of scales, ranging from handheld [15], to van-based [16, 17], up to helicopters [18], each having its own set of tradeoffs. The methods introduced in this work apply to both stationary and mobile detection systems, however, mobile detection is emphasized in this work, as it is more general.

Detecting radioactive sources for safety and security purposes means detecting these sources in the presence of background radiation, discussed further in Section 2.3. While stationary systems face a relatively constant background, neglecting temporal variations, mobile systems encounter a wide range of radiation background environments. Detection methods must then be able to adapt to new environments, particularly being able to estimate the background in a region that has not been previously surveyed. Failing to adapt to new environments may result in ineffective detection algorithms, as the system may produce false positives on benign background, reducing the effectiveness of an algorithm.

Existing gamma-ray detection and identification methods only consider spectral features within gamma-ray spectra, meaning no relationship is made between measurements and the local environment. There is, however, information that could potentially be useful in interpreting measurements in an automated algorithm. Ref. [19] shows the background count rate of a mobile detection system dramatically drop as the vehicle crosses a steel bridge over a channel of water, presumably due to a decrease in terrestrial gamma-ray background sources in the area. While this phenomenon can be understood in the context of the local environment, current spectroscopic algorithms are incapable of accounting for this variation automatically, and to adjust operating parameters accordingly. Computer vision methods, for example, semantic segmentation [20], allow for automatic identification of visual features that could be fed into spectroscopic algorithms. For instance, visual features can be used to adjust predictions of the gamma-ray background, and detection thresholds can be adjusted accordingly. This work introduces novel spectral detection and identification methods that can potentially be coupled to complementary contextual sensors with the intention of improving performance. In developing these detection methods, procedures for assessing algorithm performance are established, which may serve as useful guidelines for quantifying enhancements from algorithms enhanced with data from complementary sensors.

1.4 Dissertation Outline

To achieve the objectives described above, the remainder of this dissertation is outlined as follows:

- Chapter 2 reviews relevant physics, radiation detection principles, and mathemati-

cal frameworks used in the introduction and analysis of detection and identification algorithms.

- Chapter 3 describes the methods and metrics used to evaluate the performance of detection and identification algorithms, as well as a detailed description of the data used later in this work.
- Chapter 4 introduces the use of non-negative matrix factorization for gamma-ray spectral modeling, detection, and identification, as well as a comparison of the method to existing algorithms.
- Chapter 5 examines and evaluates the performance of novel methods for neural-network based detection and identification.
- Chapter 6 concludes the dissertation, along with a review of future research directions enabled by the methods introduced here.

Chapter 2

Elements of Spectral Gamma-ray Source Detection and Identification

Gamma-ray spectra are measurements that provide a snapshot of the local gamma-ray environment. This chapter sets the stage for a detailed discussion on the design and analysis of advanced spectral analysis methods, including a review of the relevant physics, detection instrumentation, and providing a concise description of the gamma-ray data life cycle, from data collection to preparation for further analysis. Additionally, the mathematical frameworks in which source detection and identification algorithms are analyzed are reviewed. Using these fundamental concepts, this chapter outlines general principles for spectral detection and identification approaches, which are useful in discussing the advanced methods described in chapters 4 and 5.

2.1 Gamma-ray Physics

Most radiation sources of interest for nuclear safety and security undergo a nuclear decay (beta, alpha, or spontaneous fission) that, some fraction of the time, leaves the daughter nucleus in an excited state. The relaxation of this excited state, to a lower-energy state, results in emission of a gamma-ray. Neglecting the small uncertainties in energy levels of nuclear states, the resulting gamma rays are treated as discrete particles of a characteristic energy. See ref. [21] for a more detailed treatment of gamma-ray emission.

A gamma-ray source \mathcal{S} produces gamma rays from a set of one or more discrete energies. For example, ^{60}Co , a commonly-used laboratory source, undergoes a β^- decay to ^{60}Ni , which primarily de-excites by emitting a pair of gamma rays at 1173 keV and 1332 keV with emission probabilities 99.85 % and 99.9826 %, respectively [22]. The identity of a source \mathcal{S} is then inferred by the detection of the characteristic gamma rays associated with \mathcal{S} .

An ideal gamma-ray detector would be able to measure the exact energy of each gamma-ray, resulting in a spectrum of discrete energies. In practice, however, gamma-

ray spectra depart from the ideal case. For one, gamma rays can undergo interactions between emission and detection. Additionally, detection systems are fundamentally noisy instruments, providing additional variability in the measured energy. This section reviews interactions between gamma radiation and matter in order to understand their effect on measured spectra, and ultimately detection and identification algorithms.

2.1.1 Gamma-ray Interactions with Matter

Between emission from the nucleus and interacting with a detector, a gamma ray may undergo one or more interactions with intervening material¹. Most of these interactions, including with the detection material itself, affects the measured gamma-ray energy. This section reviews interactions between gamma rays and matter, specifically, the interactions most relevant to gamma-ray spectroscopy – the photoelectric effect, Compton scattering, and pair production. When occurring within detector material, each of these interactions liberates electrons which produce signals that act as the basis for generating spectra, described further in Section 2.2.

Photoelectric Effect

In photoelectric interactions, the energy of an incident gamma ray is transferred directly to an atomic electron. To be liberated from the nuclear electrostatic potential, the electron must gain a nonzero amount of energy ϕ from the incident gamma-ray. For a gamma ray with energy E_γ , the energy E_e of the resulting electron is

$$E_e = E_\gamma - \phi. \quad (2.1)$$

The photoelectric cross section as a function of energy, $\sigma_{\text{PE}}(E)$, for a given absorber material with atomic number Z scales as

$$\sigma_{\text{PE}}(E) \propto \frac{Z^n}{E^m}, \quad (2.2)$$

where n and m range from 3 to 5 [22]. The probability of photoelectric absorption decreases drastically with incident gamma ray energy, and the photoelectric effect is often dominant over other effects in the low-energy regime (i.e., less than around 250 keV).

In gamma-ray spectroscopy, the photoelectric effect offers a best case scenario for energy deposition, as the electron energy is simply a constant offset ϕ from the true gamma-ray energy. For a source emitting gamma rays with energy E , in the ideal case, the calibrated energy spectrum would be a peak centered at E with a width dependent on the detector energy resolution. However, due to Compton scattering, the measured distribution forms a continuum of energies in addition to peaks.

¹Strictly speaking, a new gamma ray is produced in Compton interactions, not simply the original gamma ray with reduced energy.

Compton Scattering

In Compton scattering, an incident gamma ray interacts with an atomic electron, however, the gamma ray does not fully impart its energy to the electron. This interaction results in a gamma ray with reduced energy E' and an electron with kinetic energy E_e , given by

$$E_e = E_\gamma \left(1 - \frac{1}{1 + \frac{E_\gamma}{m_e c^2} (1 - \cos \theta)} \right) \quad (2.3)$$

where θ is the angle of the scattered photon relative to the incident photon, and $m_e c^2 \approx 511$ keV is the rest energy of an electron. The scattering angle θ resulting from the interaction is a random variable sampled from a distribution described by Klein-Nishina formula [23]. The cross section as a function of energy for Compton scattering goes as [22]

$$\sigma_{\text{CS}}(E) \propto \frac{Z}{E}. \quad (2.4)$$

This cross section decreases with energy at a lower rate than that of the photoelectric effect, meaning that Compton scattering becomes the more dominant interaction as the incident gamma ray energy increases. Note that at low photon energies, such a collision corresponds to Thomson scattering, in which only the direction of the incident photon changes.

With the inclusion of Compton scattering, the distribution of detected gamma-ray energies now departs from the peak resulting from the photoelectric effect, resulting in a continuum of energies. In particular, gamma rays emitted from a source can Compton scatter with the detector material, or intervening materials between the source and detector, adding variability to the measured spectrum.

Pair Production

Pair production occurs when a high-energy photon, specifically with energy higher than twice the electron rest energy (1022 keV), enters the Coulomb field of a high- Z material. In this scenario, the gamma ray imparts momentum on the nucleus, allowing the electron-positron pair to separate, as opposed to annihilating, by momentum conservation. This effect results in a number of observable features in gamma-ray spectra. Suppose an incident gamma-ray interacts with a detector via pair production, and that the resulting positron annihilates with an electron within the detector material, resulting in two 511 keV photons at 180° apart. If both 511 keV photons are detected, the entire incident energy E will be reconstructed. If one or both of those photons exit undetected, this results in spectral features known as single and double escape peaks, in which peaks appear at $E - m_e c^2$ and $E - 2m_e c^2$, respectively.

The cross section for pair production goes as

$$\sigma_{\text{PP}}(E) \propto Z^2 \ln(E - 2m_e c^2). \quad (2.5)$$

Relative to the previous two interactions discussed, pair production dominates at higher energies, typically on the MeV scale. This effect is not as significant as the others with respect to mobile detection, however. In particular, short measurements (e.g., 1 s) are generally used, which is often insufficient for observing phenomena source as single and double escape peaks. However, this effect, in part, explains the presence of 511 keV photons often seen in background spectra, explained further in Section 2.3.

2.1.2 Attenuation

It is useful to describe these gamma-ray interactions on a macroscopic scale, and to consider their combined effect. One concept used to describe these effects together is *attenuation*, specifically, the reduction in amplitude of a gamma-ray signal by an attenuating material. Suppose that the source flux at a reference point on one side of an 1-dimensional attenuating material with thickness Δx is I_0 . The flux on the other side of the material will be exponentially attenuated as

$$I(\Delta x) = I_0 e^{-\mu \Delta x}, \quad (2.6)$$

where μ is the attenuation coefficient for the material. The coefficient μ is the sum of attenuation coefficients for each type of interaction between the incident radiation and the material, each of which can be estimated using their cross sections [23]:

$$\mu = \mu_{PE} + \mu_{CS} + \mu_{PP}. \quad (2.7)$$

Attenuation is particularly important for understanding the effects of source shielding. Specifically, Equation (2.6) suggests that the gamma ray flux can be significantly reduced for certain combinations of μ and Δx , meaning that sources can be *shielded* using a moderately thin but dense material (e.g., a layer of steel or lead). Note that the attenuation coefficient μ is a function of photon energy, and that photons are attenuated at different rates depending on energy. This means that in addition to reducing the amplitude of a detected signal, shielding affects the shape of a source spectrum. Though attenuation fundamentally limits what can be detected, knowing how shielding affects measurements can potentially be used to improve identification performance, described later in Section 4.6.

2.2 Gamma-ray Detection and Measurement

There are two main types of detectors used in practical applications: scintillation detectors and semiconductor detectors. Due to having a lower cost per volume and fewer operational constraints than semiconductor detectors, scintillation detectors are more widely used in fielded applications, and in this work, scintillation detectors are exclusively considered. Note, however, that the methods developed in this work can readily be applied to data from semiconductor detectors. A brief description of gamma-ray detection in

scintillation detectors, from gamma ray interactions to data preprocessing, is given in this section. Refs. [22, 23] provide thorough treatments of the conceptual and practical elements of gamma-ray detection and measurement.

2.2.1 Scintillation Detection

A scintillation crystal is a material that, when struck with charged particles, emits photons, generally in the visible range. These secondary photons are converted to electrons, and the resulting current is used to infer the energy of the incident gamma ray. Suppose an incident photon interacts with a scintillation crystal via one of the three interactions previously discussed, producing an electron in the detector volume. The resulting electron has sufficient energy to ionize other atoms in the crystal, producing additional charges as it traverses the detector. In terms of energy states, electrons which are originally in the valence band are promoted to the conduction band. Instead of de-exciting directly back to the valence band, scintillation materials generally rely on intermediate *activation states* between the valence and conduction bands. For instance, Tl is included in forming the crystal NaI, yielding the material NaI(Tl), which is widely used in practical applications. During de-excitation, electrons can temporarily occupy the intermediate state, allowing for the emission of a lower-energy photon once it decays down to the valence band. This process produces many visible photons for an incident gamma ray, for example, approximately 38,000 photons per MeV for NaI(Tl) [23].

The scintillation photons must be converted to an electrical signal to produce useful measurements, for example, by means of a photocathode and photomultiplier tube (PMT). The photocathode, commonly made of alkali metals, converts the secondary photons to electrons via the photoelectric effect, and the PMT applies a high voltage to a chain of elements called *dynodes*, each of which provides a multiplication factor to the signal current. Resulting from the detector-photocathode-PMT chain is an analog electronic signal which must be processed further before it produces useful measurements. Specifically, an electronic circuit is used to convert a charge signal produced by the PMT to a voltage. It is assumed that the magnitude of the raw voltage signals are proportional to the gamma-ray energy – higher energy incident radiation results in more secondary photons, which in turn generates a large current through the PMT. A calibration using sources of known energies is then performed, providing a mapping between voltage signals and gamma-ray energies.

One of the key characteristics of gamma-ray spectrometers is the energy resolution. Energy resolution is influenced by several factors in the detection system, including statistical fluctuations in the number of scintillation photons, inefficiencies (e.g., in conversion at the photocathode), and electronic noise. Ultimately, all these effects are manifested as *spectral blurring*, which is seen as a spread of energies with variance $\sigma^2(E)$ for a full-energy peak with energy E . More commonly, the energy resolution is described using the full-width at half max (FWHM), more specifically, the percent resolution at a given

energy E :

$$\frac{\text{FWHM}(E)}{E} = \frac{2\sqrt{2\ln 2}}{E}\sigma(E). \quad (2.8)$$

2.2.2 Spectral Binning

A calibrated spectroscopic detector produces a stream of energies E_i occurring at a time t_i , referred to as list-mode data. It is common to aggregate list-mode data to form a gamma-ray *spectrum* out of a collection of measurements $\{(E_0, t_0), (E_1, t_1), \dots, (E_n, t_n)\}$. Creating a gamma-ray spectrum can be thought of as discretizing gamma-ray energies – the detection of a gamma ray with energy $E_i < E < E_{i+1}$ is reduced to a single count in the energy range between E_i and E_{i+1} . When forming spectra out of many events, recognizable features begin to appear – full-energy peaks (i.e., peaks corresponding to the photoelectric effect) begin to appear for nearby sources, along with Compton continua, as well as background radiation, described in Section 2.3.

One common approach to performing the energy discretization described above is *linear binning*, in which the gamma-ray energy range $[E_{\min}, E_{\max}]$ is evenly divided into d energy bins. In general, detector energy resolution increases with the square root of energy, meaning for linear binning, peaks span a number of bins proportional to the square root of the peak energy. *Square root binning*, which is used throughout this work, linearly divides the square root of gamma-ray energy range, resulting in bins with widths that increase with energy. Square root binning is useful for spectral analysis, as it puts spectral features on a similar scale.

Detected gamma-ray events are aggregated over some time interval, referred to as the *integration time*. Increasing the integration time has the effect of reducing statistical uncertainties in count rate per energy bin as more counts are accumulated. However, in mobile detection, there is generally relative motion between the detector system and sources of interest, meaning a longer integration time would reduce the signal-to-background ratio. As a result, a relatively short integration time (e.g., 1 s) is used and spectra are frequently processed, resulting in a higher signal-to-background ratio.

2.3 Sources of Gamma Radiation

A given gamma-ray spectrum consists of photons from naturally-occurring background radiation and potentially additional sources in the area. This section introduces key contributing factors to background gamma radiation, as well as an overview of sources of interest and sources commonly seen by fielded instruments.

2.3.1 Gamma-ray Background

Background gamma radiation, referred to hereafter simply as *background*, results primarily from terrestrial radionuclides and radiation of a cosmic origin [24], and is virtually

Table 2.1: Prominent gamma-ray energies of primordial radionuclides [24]

Parent Radionuclide	Child Radionuclide	Gamma-ray Energy
^{40}K	N/A	1460 keV
^{235}U	^{231}Th	185.7 keV
^{238}U	^{214}Pb	351.9 keV
	^{214}Bi	609 keV
	^{214}Bi	1120.3 keV
	^{214}Bi	1764.5 keV
	^{214}Bi	2204.2 keV
^{232}Th	^{212}Pb	238.6 keV
	^{208}Tl	2614.5 keV
	^{228}Ac	911.2 keV

always present in gamma-ray spectra, though it can be reduced with specialized instrumentation setups (e.g., a detector housed in a lead structure). Since gamma-ray spectra invariably contain background, the ability to detect and identify gamma-ray sources then relies on separating naturally-occurring radiation from other sources.

The primordial radionuclides observed are ^{40}K , ^{235}U , ^{238}U , and ^{232}Th , as well as the daughter nuclides of the uranium and thorium isotopes [22, 24]. The most commonly-observed gamma-ray energies from these are listed in Table 2.1. Cosmic background radiation, however, primarily consists of a 511 keV peak from positron annihilation following pair production and a power law continuum [23, 24]. The abundance of each of these radionuclides, as well as the local environment for gamma-rays to lose energy via Compton scattering, varies with location, meaning observed gamma-ray background vary substantially.

Gamma-ray background can also vary with time, for example, following changes in weather patterns. ^{222}Rn , which is produced in the ^{238}U decay chain, and its progeny in the atmosphere are "scavenged" by clouds and re-deposited in the ground during rain, resulting in higher emissions from these radionuclides at the surface [25]. The effect of this is that even a stationary detection system, which does not measure any spatial variability in the gamma-ray background, will occasionally experience temporal variability due to radon decay.

A background spectrum measured by a NaI(Tl) system is shown in Figure 2.1, annotated with the gamma-ray lines in Table 2.1. Overlaid on this spectrum is the mean spectrum resulting from a collection of 10,000 randomly-sampled spectra, which shows background peak features more clearly. Note that there are more counts at the low-energy end of the spectrum due to Compton downscattering as previously described.

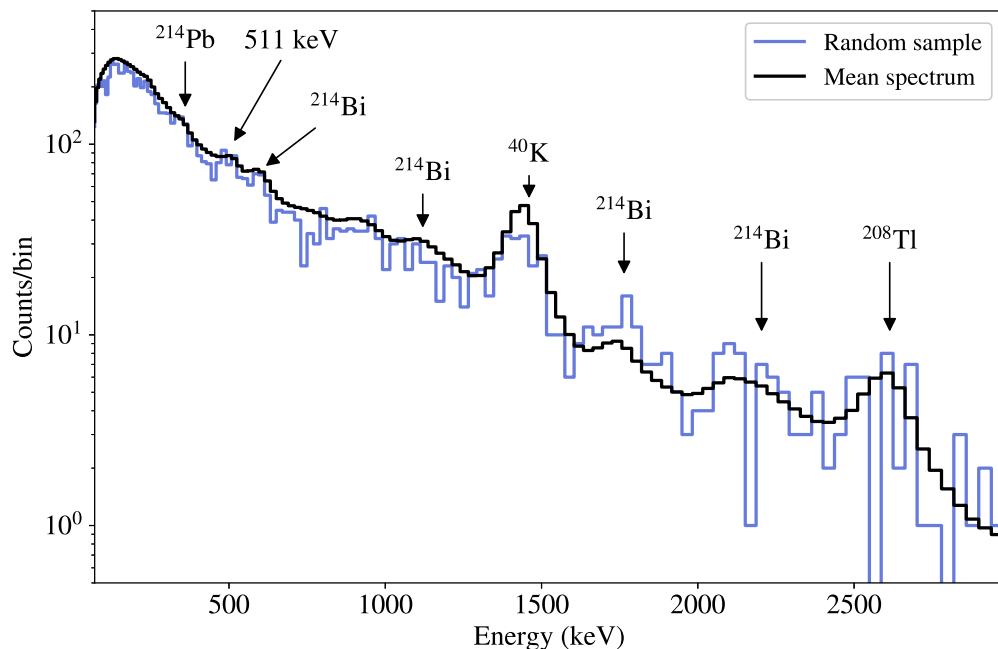


Figure 2.1: Example gamma-ray background spectrum from an array of NaI(Tl) scintillator detectors on the RadMAP mobile detection system, introduced in Section 3.2. Overlaid is the mean spectrum from 10,000 randomly-sampled spectra, and the most prominent background lines from Table 2.1 are annotated. The spectrum was formed using a 1-s integration time and 128 square root bins between 50 and 3000 keV. Additional details on how the spectrum was formed are given in Section 3.2.

2.3.2 Gamma-ray Sources of Interest

The remainder of this section discusses sources of interest for nuclear safety and security, and sources commonly seen in the field.

Special Nuclear Material

Special Nuclear Material (SNM), a term defined by the Nuclear Regulatory Commission in Title I of the Atomic Energy Act of 1954 [24], refers to material consisting of uranium enriched in ^{233}U and ^{235}U , and any isotope of plutonium. SNM is of particular interest to nuclear security, as these materials are capable of producing nuclear explosions when sufficient quantities are assembled. Though SNM is of significant interest to practical detection systems, it is neither measured nor simulated here, and is not discussed further in this work.

Medical and Industrial Gamma-ray Sources

One class of sources that is commonly seen by fielded detectors are those used in nuclear medicine. Several types of sources are used for medical imaging [26], namely Positron Emission Tomography (PET) and Single Photon Emission Computed Tomography (SPECT), as well as other sources for therapeutic purposes [27]. In each of these cases, patients that have recently undergone imaging or treatment give off gamma rays from sources injected for the procedures. Additionally, gamma-ray sources show up in diverse applications such as equipment sterilization, oil well logging, and leak detection in plumbing [1]. A table of medical and industrial sources, along with their gamma-ray energies, is later provided in Chapter 3 in Table 3.2.

2.4 Model for Spectral Detection and Identification

This section is meant to act as a bridge between the discussion on physical processes and the mathematical models used for anomaly detection and source identification. To begin, note that a gamma-ray spectrum with d energy bins can be viewed as a vector $\mathbf{x} \in \mathbb{R}_+^d$. From this perspective, a source detection algorithm, of which the output indicates the presence or absence of an anomalous source, can be thought of as a function providing a mapping from \mathbb{R}_+^d to $\{0, 1\}$. Similarly, an identification algorithm, which identifies the sources present in a spectrum using a library of N known sources, transforms an input spectrum from \mathbb{R}_+^d to $\{0, 1\}^N$. The remainder of this section outlines properties for both detection and identification, and relates these properties to existing approaches.

2.4.1 Anomaly Detection

Generally speaking, anomaly detection aims to determine whether or not measured data is consistent with expected behavior. Common use cases of anomaly detection include bank fraud detection [28], network security [29], and manufacturing [30]. In gamma-ray spectroscopy, anomaly detection is used establish whether or not a spectrum is consistent with background without attempting to identify radiological sources.

One common approach to spectral anomaly detection is to generate an estimate $\hat{\mathbf{x}}$ of the gamma-ray background contained in a measurement \mathbf{x} , and computing a metric to quantify the difference $D(\mathbf{x}, \hat{\mathbf{x}})$ between the two. Anomaly detection, and identification as we will see next, can then be thought of as a two-part processes, referred to here as *feature extraction* and *classification*.

Feature extraction refers to the process of drawing salient information contained in a spectrum needed to perform detection or identification. In the approach outlined above, the process of generating the background estimate $\hat{\mathbf{x}}$ can be considered feature extraction. Methods for generating the background estimate $\hat{\mathbf{x}}$ is a central issue examined in this dissertation, and is explored further in Section 4.2 and Section 5.3.

Classification refers to the process of taking extracted features and making predictions. In this context, features extracted from gamma-ray spectra are used to predict the presence or absence of sources. As mentioned above, the primary element of this classification procedure is defining a detection metric. Examples of the metric $D(\mathbf{x}, \hat{\mathbf{x}})$ are based on the L2 norm [31, 32] and Mahalanobis distance [13, 33]. With a metric defined, a threshold for when to consider a spectrum anomalous or not must then be set, which is the topic of Section 2.6.

2.4.2 Identification

In this work, identification generally refers to *simultaneous detection and identification*, which, for N different sources of interest, can be thought of as N simultaneous detection problems. As a result, identification operates in a similar fashion to the detection approach outlined above. The primary difference is that instead of searching for any anomalies in a spectrum, identification attempts to find spectral features corresponding to particular radionuclides. In this approach, an identification algorithm must then know spectral features about the sources which it aims to identify. This approach to feature extraction is often done by defining spectral regions of interest [34], providing the algorithm with spectral source *templates*, or by learning representations of the source spectra [35]. Similar to anomaly detection, metrics are defined to indicate the presence or absence of a particular source.

2.5 Maximum Likelihood Estimation

To distinguish weak sources from background, detection and identification algorithms must be robust to the variability seen in spectral measurements. Additionally, the decay of a radionuclide is an inherently random process, meaning gamma-ray spectra \mathbf{x} are subject to statistical fluctuations. Accurate models for detection and identification must then be able to account for variability in spectra. In this work, variability is modeled using parameterized probabilistic models, and using the statistical principle of Maximum Likelihood Estimation (MLE) is used to estimate model parameters. This section provides a brief overview of MLE, and derives an expression used for modeling Poisson processes, which is later used for accurately modeling gamma-ray spectra.

Suppose a process resulting in a random variable \mathbf{x} can be described by a probability distribution $p(\mathbf{x}|\boldsymbol{\theta})$, parameterized by $\boldsymbol{\theta} \in \Theta$ for some parameter space Θ . The probability $p(\mathbf{x}|\boldsymbol{\theta})$, a function of the parameters $\boldsymbol{\theta}$, is referred to as a likelihood. The true parameters $\boldsymbol{\theta}^*$ of the data-generating distribution are not known, however, an estimate of the parameters, denoted by $\hat{\boldsymbol{\theta}}$, is found by a set of measurements \mathcal{D} of random variable \mathbf{x} . MLE [36] is the process of finding estimate parameters $\hat{\boldsymbol{\theta}}$ that maximize the likelihood of the data

\mathcal{D} . Concretely, $\hat{\theta}$ is defined as

$$\hat{\theta} = \arg \max_{\theta \in \Theta} p(\mathcal{D}|\theta). \quad (2.9)$$

In practice, the log-likelihood, which increases monotonically with the likelihood, is used as it often simplifies calculations and prevents numerical underflow. The likelihood (and equivalently, log-likelihood) can be thought of as a measure of how well a dataset fits to a model. It is common to define *error* functions between data and models, and thus it is typical to instead minimize the negative log-likelihood, which is a convention used for the remainder of this work. Using this convention, MLE is equivalently described as

$$\hat{\theta} = \arg \min_{\theta \in \Theta} \{ -\ln p(\mathcal{D}|\theta) \}. \quad (2.10)$$

The number of gamma rays $x \in \mathbb{Z}_+$ measured over some interval of time Δt is a random variable from a Poisson process with rate $\lambda \in \mathbb{R}_+$:

$$x \sim \text{Poisson}(\lambda). \quad (2.11)$$

The likelihood of x under this model is then

$$p(x|\lambda) = \frac{e^{-\lambda} \lambda^x}{x!}. \quad (2.12)$$

A gamma-ray spectrum \mathbf{x} can then be thought of as a d -dimensional random vector from a d -dimensional Poisson distribution with rate $\boldsymbol{\lambda} \in \mathbb{R}_+^d$, with likelihood

$$p(\mathbf{x}|\boldsymbol{\lambda}) = \prod_{j=1}^d p(x_j|\lambda_j), \quad (2.13)$$

where the product rule of probability is used, as the number of counts in each bin is independent. Now suppose that the Poisson rate in a given bin across spectra is not assumed to be equal, meaning each spectrum \mathbf{x}_i is associated with a rate $\boldsymbol{\lambda}_i$ forming a matrix $\mathbf{\Lambda} \in \mathbb{R}_+^{n \times d}$. An example of when this behavior would be seen is a moving detector, in which background rates vary with location. The likelihood of a collection of n gamma-ray spectra in a matrix $\mathbf{X} \in \mathbb{R}_+^{n \times d}$ is then

$$p(\mathbf{X}|\mathbf{\Lambda}) = \prod_{i=1}^n \prod_{j=1}^d p(x_{ij}|\lambda_{ij}). \quad (2.14)$$

The associated negative log-likelihood is

$$-\ln p(\mathbf{X}|\mathbf{\Lambda}) = -\sum_{i=1}^n \sum_{j=1}^d \ln p(x_{ij}|\lambda_{ij}) = \sum_{i=1}^n \sum_{j=1}^d \lambda_{ij} - x_{ij} \ln \lambda_{ij} + \ln x_{ij}!. \quad (2.15)$$

Section 4.2 and Section 5.3 model gamma-ray spectra by introducing methods to estimate the Poisson rate Λ in Equation (2.14). Using these models of spectra, additional operations can be performed to detect and identify sources.

One aspect of MLE is that it conveniently allows for the incorporation of model constraints, known as *prior probabilities*, often referred to simply as *priors*. In particular, priors, formed on the basis of Bayes' Theorem [37], are probability distributions on model parameters. These prior probabilities then influence the parameter estimates, giving a modified version of MLE known as *maximum a posteriori*. Note that the use of priors generally has the same effect as the concept of regularization used in optimization. In this work, regularization is used to enforce model behavior such as sparsity, primarily seen in Chapter 5.

2.5.1 Deviance

Likelihood values can be used to compute a goodness-of-fit statistic, for example, the *deviance* statistic [38]. The deviance between data $\mathbf{y} \in \mathbb{R}^d$ and corresponding model $\boldsymbol{\theta} \in \mathbb{R}^d$ is computed as

$$D(\mathbf{y}, \boldsymbol{\theta}) = 2(\ln p(\mathbf{y}|\mathbf{y}) - \ln p(\mathbf{y}|\boldsymbol{\theta})) = 2 \sum_{i=0}^d \ln p(y_i|y_i) - \ln p(y_i|\theta_i). \quad (2.16)$$

In essence, the deviance gives a measure of the difference between an ideal model $p(\mathbf{y}|\mathbf{y})$ where each parameter is known exactly, and the given model $\ln p(\mathbf{y}|\boldsymbol{\theta})$. In this work, the Poisson deviance between a measured spectrum \mathbf{x} and corresponding modeled spectrum $\hat{\mathbf{x}}$ is computed using

$$D(\mathbf{x}, \hat{\mathbf{x}}) = 2 \sum_{i=1}^d \hat{x}_i - x_i + x_i \ln \frac{x_i}{\hat{x}_i}. \quad (2.17)$$

Because of its foundation in MLE, deviance is a natural choice for anomaly detection using likelihood-based models, and is examined further in Chapter 5.

2.6 Decision Theory

A spectrum \mathbf{x} is equal to the sum of both the source and background spectra, (i.e., $\mathbf{x} = \mathbf{x}_s + \mathbf{x}_b$), and an ideal algorithm would be able to exactly attribute the counts in the spectrum belonging to both source and background. One fundamental limitation preventing the realization of ideal algorithms such as this is statistical variability and limited statistics from sources. Suppose a weak source is nearby, and it contributes very few counts to the spectrum, less than the statistical variations in background. Though a source is present, the relative contributions of source and background cannot be decomposed over the integration times relevant to this work. Fundamentally, there are limits on what can be

detected, and this section discusses the tradeoffs involved with setting detection thresholds, knowing that ideal detection algorithms cannot be realized in practice.

Suppose a detection method f_d performs operations on \mathbf{x} yielding a scalar metric $D(\mathbf{x}, \hat{\mathbf{x}})$ that measures how anomalous a spectrum is. A threshold T can be set on the value of $D(\mathbf{x}, \hat{\mathbf{x}})$, and any spectra resulting in a metric greater than T would indicate the presence of a source. The choice of T is then related to the desired accuracy of the detection method; if T is too high, sources can go undetected, and if T is too low, the method will attribute benign background to anomalous sources. This behavior is quantitatively described using the following measures:

- **True positive rate:** proportion of tests that correctly detect a source that is present
- **True negative rate:** proportion of tests that correctly identify background spectra as not containing a source
- **False positive rate:** proportion of tests that incorrectly detect a source being present
- **False negative rate:** proportion of tests that incorrectly miss a source that is present

An ideal system would have high true positive and negative rates, and low false positive and negative rates. Practical systems are not ideal, and a value of T is chosen that provides a favorable tradeoff between these measures. In practice, one of these measures is held constant, and the threshold T is solved for using this constraint. In particular, the False Alarm Rate (FAR), which is proportional to the false positive rate, is often held at a fixed value. This choice is rooted in the fact that there is a certain FAR that operators can tolerate – too many false alarms, and the output may be disregarded. As a result, low limits are usually placed on FAR, for example, one alarm in eight hours, which for 1 second spectra translates to 3.5×10^{-5} , a rate much lower than considered in many applications of statical analysis and machine learning.

Due to variability in the expected range of background spectra, there is a wide range of values of D , characterized by the probability distribution $p(D)$. Setting T based on the desired FAR requires knowing $p(D)$ for background spectra. If the functional form of $p(D)$ is known, T can be solved for analytically. By definition, any background spectrum \mathbf{x} with metric $D(\mathbf{x}, \hat{\mathbf{x}})$ higher than T is a false positive, meaning the area under $p(D)$ above T is equal to the probability of false alarm:

$$p_{\text{FA}} = \int_T^{\infty} p(D)dD. \quad (2.18)$$

The probability of false alarm p_{FA} is equal to the false alarm rate (e.g., $1/8 \text{ hr}^{-1}$) multiplied by the measurement interval (e.g., 1 s). Since $p(D)$ integrates to unity, Equation (2.18) can be rewritten as

$$\int_0^T p(D)dD = 1 - p_{\text{FA}}. \quad (2.19)$$

The threshold can then be expressed as

$$T = F^{-1}(1 - p_{\text{FA}}), \quad (2.20)$$

where F^{-1} is the inverse Cumulative Distribution Function (CDF) of $p(D)$.

If the functional form of $p(D)$ is not known, the distribution can be approximated empirically using measured data. In the empirical case, the resulting distribution can be fit to a model, and the threshold can be computed from this model, or a threshold can be chosen based on samples using an expected number of false alarms.

Chapter 3

Spectroscopic Data Sources for Algorithm Development and Evaluation

A quantitative evaluation of a detection or identification algorithm is essential in understanding the utility of the approach. As discussed in Chapter 2, detection algorithms must operate in a wide range of gamma-ray background environments, and be able to detect and identify sources that span the energy range roughly of 40 to 3000 keV. Algorithm performance is assessed using data consisting of known sources in a variety of background environments in a process referred to as *source injection*. This chapter introduces the various elements of source injection, and describes two urban radiation data sources that are used later in this work for quantitative algorithm analysis.

3.1 Source Injection

The performance of a detection or identification algorithm is determined by several factors, which are summarized in Table 3.1. Broadly speaking, the ability to detect radioactive sources is determined by variables related to the source (e.g., activity), detector parameters, coupled source-detector variables (e.g., distance between source and detector), the environment that the detector system is operating in, and operational parameters. Of these, detector and operational parameters are generally fixed, meaning for a given scenario, the most important variables are related to the source and the local environment. Despite having many parameters influencing the performance of an algorithm, performance can be assessed by varying the signal-to-background ratio. As a result, all but one of these parameters can be held fixed, and performance can be quantified as a function of this single variable.

In this work, algorithm performance is evaluated by analyzing many sequences of spectra, referred to as *runs*. To form spectra \mathbf{X} for a given run, source spectra \mathbf{X}_s are *injected* into background spectra \mathbf{X}_b (i.e., $\mathbf{X} = \mathbf{X}_b + \mathbf{X}_s$). The number of background runs, which are either measured or simulated, is generally fixed due to the amount of effort

Table 3.1: Parameters affecting gamma-ray source detection

Source parameters	Source type
	Source activity
	Source shielding
Detector parameters	Detector type
	Detector size
	Integration time
Source-detector variables	Source-detector distance
	Relative motion
	Obstructions between source and detector
Environmental variables	Background energy distribution
	Background radiation flux
	Gamma-ray scatterers
Operational parameters	False alarm rate
	Library of sources of interest

required to produce the data. The source spectra \mathbf{X}_s , however, can easily be re-generated for a new set of parameters, providing an efficient way to estimate the performance of algorithms. One limitation of this approach is that the source data are decoupled from the background environment, meaning effects such as environmental scattering, which are observed in practice, are not captured by this procedure.

By injecting variable sources, detection methods can be assessed on a wide range of parameters without requiring to take extensive measurements. The remainder of this section details how both \mathbf{X}_b and \mathbf{X}_s are generated, either from measured or simulated data.

3.1.1 Background Data for Source Injection

The background spectra $\mathbf{X}_b \in \mathbb{R}_+^{n \times d}$, having n measurements with d bins, can either consist of sequential measurements or be time-independent, depending on injection scenario. For time-independent background, \mathbf{X}_b is a random sample of n background spectra, and the source injected corresponds to the peak signal-to-background measurement in the run. This approach is useful in cases where the algorithm has no time-dependence, meaning each input spectrum is treated independently. In the sequential approach, the rows of \mathbf{X}_b correspond to a time series of measurements, which is essential for algorithms that have dependence on recent measurements, such as a moving average of the background

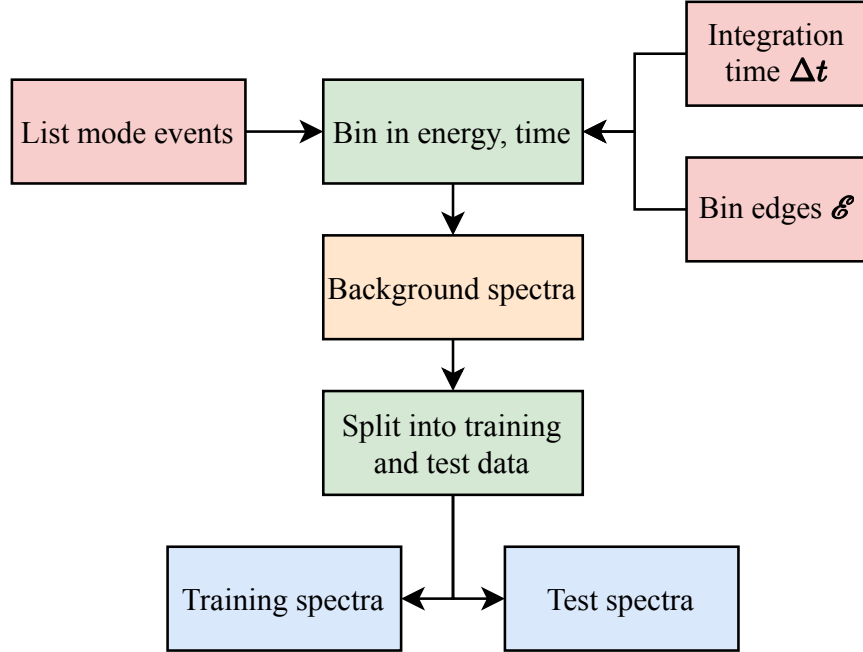


Figure 3.1: Diagram showing the processing steps to produce training and testing spectra for use in source injection. Background events, spectral bin edges \mathcal{E} , and the integration time Δt are used as inputs to generate background spectra, which are then split into data for model training and evaluation.

estimate.

Suppose \mathcal{X} is a dataset produced by a detection system. Based on the discussion from Section 2.2, spectra are formed using list mode gamma-ray data from \mathcal{X} by specifying an integration time Δt and energy binning structure \mathcal{E} . The resulting background data can then be separated into training and testing sets. The training set is used to learn model parameters and set values such as thresholds, and the test set is used to perform source injection on. This split is performed to assess the performance of algorithms on previously-unseen data. Figure 3.1 gives a graphic summary of the procedure for producing background data for both training and evaluation.

3.1.2 Source Injection using Background-Subtracted Data

When a time series of background spectra are used in source injection, the source spectra X_s are produced in such a way to simulate the effect of the detector moving past a source. Using a mobile detection system, a source can be set in a particular configuration (i.e., standoff distance, shielding, etc.), and the system can be driven past this source, collecting spectral data along the way. Of course, the system measures the combination of background and source, and the background needs to be removed to produce an estimate

of the source contribution. To get a reliable estimate of the source contribution to measured spectra, and in order to later downsample to a wide range of activities, a source with a high-activity A (e.g., several mCi) is used.

Suppose measurements of background were collected around the same time as the source data collection, meaning the background in both sets are likely to be consistent. The spectra containing sources can be background-subtracted, yielding noisy estimates of the source contributions. The resulting spectra can be averaged over many runs, resulting in relatively smooth source spectra Λ_s .

Using Λ_s , binomial downsampling is used to achieve a particular source activity A' . The binomial probability of success g is computed as the ratio between A' and A , and random samples \mathbf{X}_s can be generated as

$$\mathbf{X}_s \sim \text{B}(\Lambda_s, g), \quad (3.1)$$

where B refers to the binomial distribution.

One major drawback to this approach is that it is subject to many sources of uncertainty, such as the vehicle speed and time alignment of spectra between runs. Furthermore, this approach requires many repeated measurements even for a single source, requiring a significant amount of effort. Though simple in design, this approach does not easily scale to many types of sources and source configurations such as shielding. To assess algorithm performance on a wide range of source types and configurations, simulated data is more commonly used.

3.1.3 Source Injection using Simulated Data

Monte Carlo simulations are capable of generating accurate source spectra that can be used in source injection. In this context, source simulation refers to the process of simulating spectra from arbitrary gamma-ray sources using known physical values, particularly, gamma-ray energies and branching ratios. A detector response function $\eta(E, \Omega)$ at gamma-ray energy E and source-detector angle Ω is computed from simulated events, and $\eta(E, \Omega)$ can be used to generate source spectra. This section describes the procedure to generate $\eta(E, \Omega)$, and in turn, generate source spectra \mathbf{X}_s .

To make the discussion here concrete, a 2" \times 4" \times 16" NaI(Tl) crystal covered in a 1 mm thick aluminum case is used to demonstrate the concepts in this section. Figure 3.2 illustrates the setup of the simulation, with the NaI(Tl) detector in the center, and photons from a distant monoenergetic source originating at an angle $\Omega = \theta$ with respect to the x -axis. The popular framework Geant4 [39] is used to carry out simulations.

Detector Response Simulations

The detector response $\eta(E, \theta)$ can be thought of as the effective area of a detector as a function of incident photon energy E from an angle θ . In this work, $\eta(E, \theta)$ is computed using a Monte Carlo simulation of N particles with energy E incident to the detector at




Figure 3.2: Detector and source setup for angular response simulations. A 2" \times 4" \times 16" NaI(Tl) bar is centered at the origin, while monoenergetic photons from a distant source, drawn as parallel lines, make an angle θ with the x -axis. Note that by symmetry, angular points only need to be sampled in the range $\theta \in [0, 90^\circ]$.

an angle θ . For each simulated photon, the simulation returns the deposited energy E' in the detector. In these simulations, $N = 10^8$ photons were simulated for each pair of energy E and angle θ . The initial position of the simulated photon is randomly drawn from a simulation volume which encompasses the entire detector, which in this case was a cylinder with a radius of 30 cm.

The detector response simulations do not take into account detector energy resolution, and detected events must be manually blurred. To blur spectra, a detected event with energy E' is sampled from a Gaussian distribution with mean E' and variance $\sigma(E')^2$, resulting in a blurred energy E'' :

$$E'' \sim \mathcal{N}(E', \sigma(E')^2). \quad (3.2)$$

For this particular model of the NaI(Tl) bar, the standard deviation as a function of energy, $\sigma(E)$, takes the functional form [40]

$$\sigma(E) = \frac{\text{FWHM}(E)}{2\sqrt{2\ln 2}} = \frac{10^{-2} * w_0 E_0}{2\sqrt{2\ln 2}} \left(\frac{E}{E_0} \right)^{w_1}, \quad (3.3)$$

where $w_0 = 7.5$, $w_1 = 0.7$, and $E_0 = 661$ keV. Following blurring, a low-energy threshold of 30 keV was applied, discarding all events with energies below this.

A spectrum is then generated for all events at the given incident energy and angle. When normalized, this spectrum can be thought of as a probability distribution $p(E'|E)$, which is read as *the probability of measuring a deposition with energy E' given an incident photon with energy E* . The quantity $p(E'|E)$ does not take into account the probability of




Figure 3.3: Detector response $\eta(E, \theta)$ for four different energies and two incident angles. The response is slightly lower at $\theta = 75^\circ$ due to having a lower effective area from this incident angle.

an incident photon interacting with the detector, which is described by the *effective area* A_{eff} . The effective area is computed as

$$A_{\text{eff}} = \frac{N_{\text{int.}}}{N} A, \quad (3.4)$$

where $N_{\text{int.}}$ is the total number of simulated source photons that interacted with the detector and A is the area of the simulation volume in which photons were generated. The quantity A_{eff} can be thought of as the product between the surface area of the detector, as seen from angle θ , and the probability of an incident photon with energy E interacting with the detector. Using these values, the detector angular response is computed as:

$$\eta(E, \theta) = A_{\text{eff}} p(E'|E) \quad (3.5)$$

Figure 3.3 shows the detector response computed for five different energies at two different incident angles. At each photon energy, an associated full-energy peak and Compton continuum is observed.

Detector Response Interpolation

Figure 3.2 shows that there is symmetry about the z -axis for this particular simulation setup, meaning only a quarter of the angular space needs to be simulated. Within this parameter space, sampling points must be chosen, as simulations are performed at discrete angles θ_j and incident gamma rays are simulated at discrete energies E_i . However, gamma-ray sources result in a wide variety of energies, meaning the detector response

at arbitrary energies is needed. One approach to meet this need is to simulate a set of angles Θ and a set of energies \mathbf{E} , and interpolate between simulated values to obtain the response at an arbitrary pair of energy and angle. In this context, spectral interpolation can be thought of as an approach to generating a function $f(E; \theta)$ which returns the detector response $\eta(E, \theta)$. Using the known incident energies \mathbf{E} and the corresponding detector responses $\eta(E, \theta)$ for a fixed θ , the function $f(E; \theta)$ is approximated using linear interpolation.

Generating Source Templates

A given source s (e.g., ^{60}Co) can be described by pairs (E_i, B_i) of gamma-ray energies E_i and associating branching ratios B_i . A source template $\psi_s(\theta)$ is the sum of the detector response $\eta(E_i, \theta)$ for each gamma-ray line at energy E_i , weighted by the branching ratio B_i :

$$\psi_s(\theta) = \sum_i B_i \eta(E_i, \theta). \quad (3.6)$$

The values E_i and B_i for a particular source are queried from the National Nuclear Data Center (NNDC) [41, 42] via the Python package `becquerel` [43]. When multiplied with an appropriate scaling factor, described in the following section, these spectra can be used to generate spectral samples. Figure 3.5 shows templates for ^{131}I at $\theta = 0^\circ$ and $\theta = 85^\circ$. Though both templates show roughly the same shape, primarily full-energy peaks at the source gamma-ray energies, the magnitude of the template is lower at $\theta = 85^\circ$ due to a lower effective area. The entire procedure from generating source templates from simulation data is summarized in Figure 3.4. Note that this approach does not include scattering within the source material or in the nearby environment, and a more thorough treatment would include these factors.

Mean Rate Calculations

The mean Poisson rate $\lambda_s(\theta)$ (i.e., the full spectral response at angle θ) for a given source is a sum over the mean rate $\lambda_i(\theta)$ from each gamma-ray line at energy E_i :

$$\lambda_s(\theta) = \sum_i \lambda_i. \quad (3.7)$$

The rate λ_i is the contribution to the rate from the gamma-ray line at E_i alone, and can be expressed in terms of the source activity \mathcal{A}_s , source-detector distance $\|\mathbf{r}\|$, effective area $\eta(E_i, \theta)$, and branching ratio B_i . For a spectral integration time Δt , the mean rate takes the form

$$\lambda_i = \frac{\mathcal{A}_s \eta(E_i, \theta) B_i \Delta t}{4\pi \|\mathbf{r}\|^2}. \quad (3.8)$$

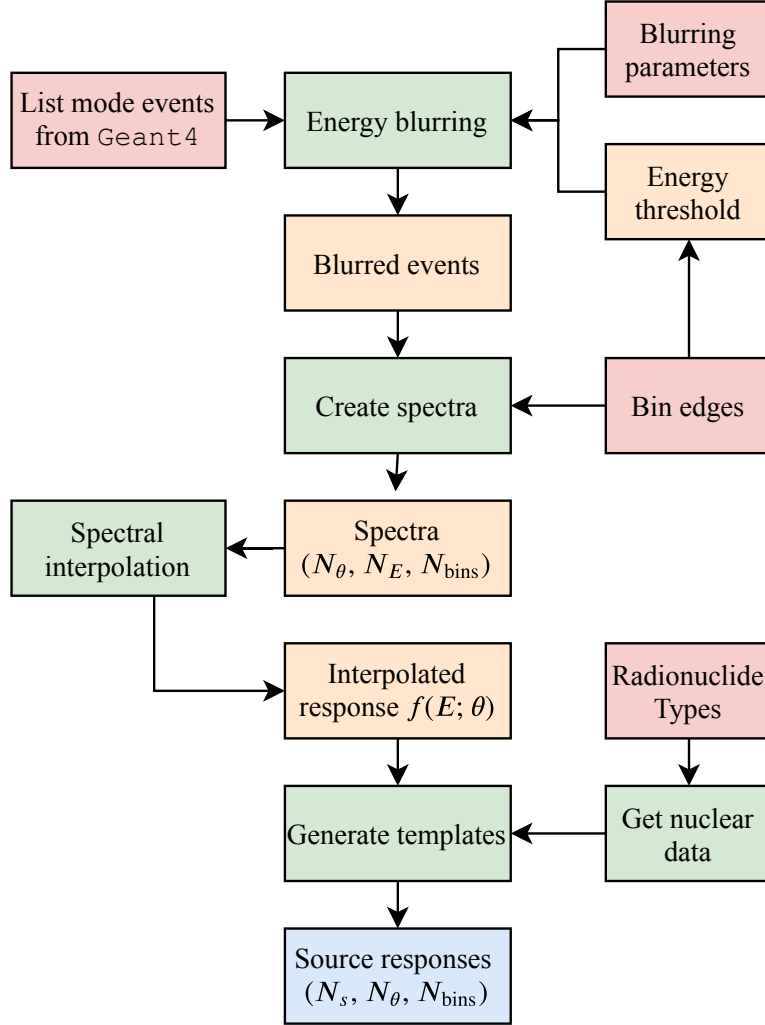


Figure 3.4: Diagram summarizing the procedure used to create source angular templates $\psi_s(\theta)$ for arbitrary radionuclide types using simulated list mode data. List mode data from Monte Carlo simulation, along with other parameters, are fed into the processing chain, ultimately resulting in source templates.

Note that air attenuation between the source and detector, which would take the form $e^{-\mu_i \|\mathbf{r}\|}$ for an attenuation coefficient μ_i , is neglected here. The total mean rate is then

$$\lambda_s(\theta) = \sum_i \frac{\mathcal{A}_s \eta(E_i, \theta) B_i \Delta t}{4\pi \|\mathbf{r}\|^2} \quad (3.9)$$

$$= \frac{\mathcal{A}_s \Delta t}{4\pi \|\mathbf{r}\|^2} \sum_i \eta(E_i, \theta) B_i \quad (3.10)$$

$$= \mathcal{S} \psi_s(\theta). \quad (3.11)$$

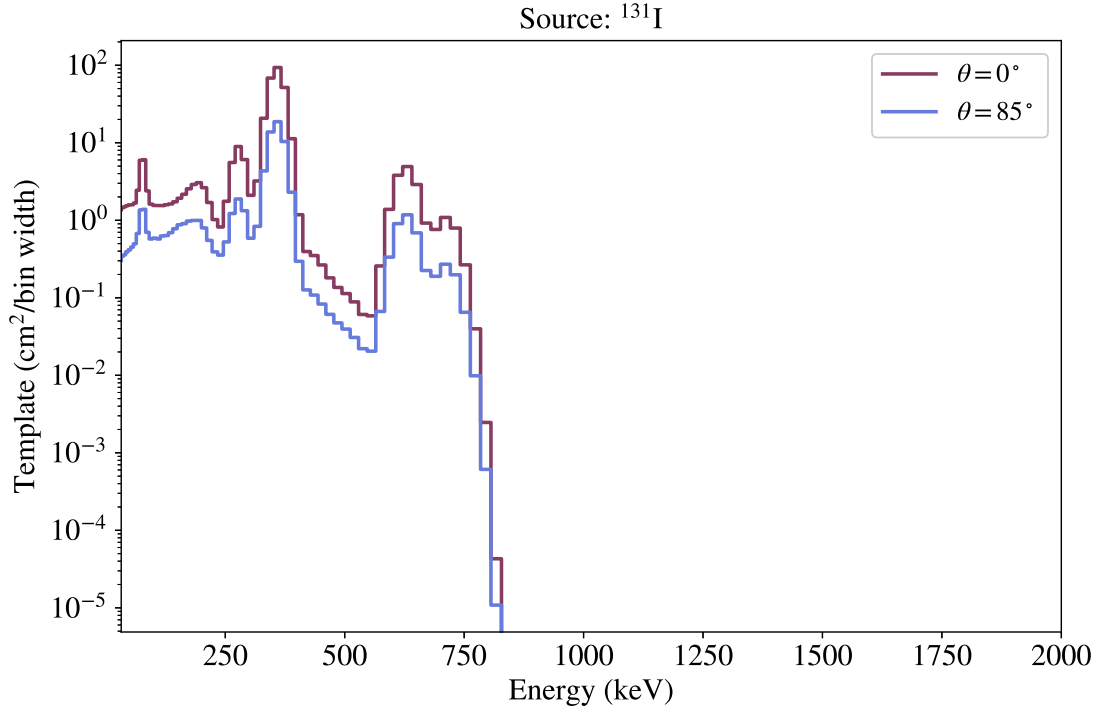


Figure 3.5: Source templates $\psi(\theta = 0^\circ)$ and $\psi(\theta = 85^\circ)$ for ^{131}I . There is a noticeable difference in the magnitude of the two templates due to the lower effective area at higher angles.

The mean rate $\lambda_s(\theta)$ is then the product of a scaling factor \mathcal{S} and the source template $\psi_s(\theta)$, and it is used to generate a spectrum $s \sim \text{Poisson}(\lambda_s(\theta))$. This concludes the procedure for generating spectra for arbitrary radionuclides. However, for injection scenarios using temporal modeling, the kinematics of the detection system past the source must be included, which is examined next.

Kinematic Modeling

Suppose the detector system moves in a straight path at a fixed speed v , has a distance of closest approach d with the source, and aggregates gamma rays at an integration time of Δt . Figure 3.6 provides a diagram illustrating the scenario. For a given run, a time of closest approach, denoted by t^* , is randomly sampled in the range $[\delta, t_{\text{tot.}} - \delta]$, where $t_{\text{tot.}}$ is the total length of the run, and δ is a number chosen to prevent too few source spectra from being present in the run. Given t^* , v , and d , the source-detector distance $\|\mathbf{r}\|$ is computed for each time step in the run as

$$\|\mathbf{r}(t_i)\|^2 = v^2(t_i - t^*)^2 + d^2, \quad (3.12)$$

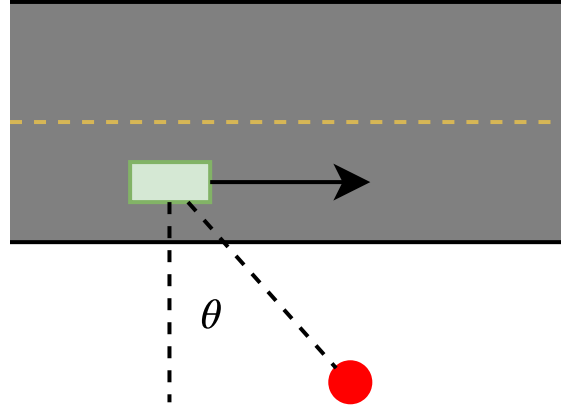


Figure 3.6: Diagram showing the kinematics of a source pass. The detector system, represented by the small rectangle, moves in the direction of the arrow and passes the source, represented by the circle. During the pass, the source spectrum is computed at each time step t_i based on the angle θ_i and source-detector distance $\|\mathbf{r}_i\|$.

where the discrete time steps t_i are defined by the spectral integration time Δt . This distance is computed at each time step, and used in computing the mean rate given by Equation (3.9). Additionally, the source-detector angle θ_i is computed at each time step as

$$\theta_i = \arctan \left(\frac{|x(t_i) - x(t^*)|}{d} \right). \quad (3.13)$$

This concludes the discussion on procedures for generating source injection data. Specific details about datasets produced using these procedures are discussed next.

3.2 Source Injection using RadMAP Data

The Radiological Multi-sensor Analysis Platform (RadMAP) [16], shown in Figure 3.7 was an experimental mobile detection system comprised of a suite of radiation detectors and non-radiological contextual sensors, including Global Positioning System (GPS) and cameras, among others. RadMAP was operated by Lawrence Berkeley National Laboratory between 2011 and 2017, during which the system was used in several data collection campaigns around the San Francisco Bay Area. The availability of highly variable gamma-ray background data in urban environments makes RadMAP an excellent candidate for algorithm development and evaluation. Of particular interest to this work is a 10×10 array of $10 \text{ cm} \times 10 \text{ cm} \times 5 \text{ cm}$ NaI(Tl) scintillation detectors. While such a large detection array is not common in operational scenarios, it can be assumed that the relative performance across algorithms will likely remain constant between NaI(Tl) detectors of varying sizes.




Figure 3.7: Figure from [16]. The RadMAP mobile detection system and its various radiological and contextual sensors. In this work, data from the 10×10 NaI(Tl) array and the GPS sensor are used.

3.2.1 Background Data Preparation

This section describes the preprocessing steps performed in preparing background data for model training and evaluation. This work makes use of RadMAP data, consisting of calibrated NaI(Tl) data and GPS coordinates, collected between December 2011 and December 2013. Detected events from the 10×10 NaI array were aggregated using a 1-s integration time, and binned in energy using 128 "square root" bins (see Section 2.2) 50 keV and 3000 keV. Due to low-energy noise, the first three bins were removed, resulting in 125 bins between 67 keV and 3000 keV.

Source Encounter Removal

Due to the uncontrolled nature of spectral measurements in urban environments, the dataset contains several source encounters. In order to be able to build accurate models of the gamma-ray background, these source encounters must be removed. This presents a challenge – how can anomalies be confidently removed without knowing the distribution of background? This section briefly provides heuristics to remove source encounters from the dataset.

The first, and simplest, method used for removing source encounters is using the total number of counts per spectrum. As a first pass, a histogram of the total number of counts in each spectrum is generated, obvious outliers were manually inspected, and




Figure 3.8: (Left pane) A "waterfall" plot of a sequence of spectra past a source during data collection in Berkeley. The source, seen as a spot in the low-energy region of the spectrum, appears shortly after the 20 s mark. (Right pane) The mean spectrum over the waterfall plot compared to the mean spectrum over the entire background dataset. The measured spectrum is consistent with a ^{99m}Tc source, which is used in SPECT imaging. Note that this sequence was measured near Alta Bates Hospital in Berkeley, consistent with the hypothesis that these anomalous gamma-ray events were emitted from a patient.

measurements close to the anomaly in time were removed. To remove nearby measurements, which may contain the same source but are not as pronounced, measurements are divided into segments that are contiguous in time. The system was often times not measuring continuously, and as a result, there are breaks in data collection, which were used to define these contiguous segments. When an anomalous spectrum is found, the entire contiguous segment is then removed, which is admittedly conservative. Figure 3.8 shows an example segment containing a source encounter. The segment of data, measured near Alta Bates Hospital in Berkeley, appears to contain a ^{99m}Tc source, used in SPECT imaging.

Following the removal of gross count anomalies, source encounters are found and removed using a spectral-based method, namely, an approach based on Non-negative Matrix Factorization (NMF) described in Chapter 4. The procedure is as follows:

1. Draw a random sample of spectra from the dataset and fit an NMF model with this.
2. Analyze all spectra in the entire set for spectral anomalies. Specifically, as discussed later in Chapter 4, an anomaly detection metric is defined for NMF-based spectral models, and this metric is used here.
3. Inspect anomalies and remove entire segments containing the anomalies.

One issue with this approach is that sources are possibly being included in the training set, thus reducing the ability to detect anomalies of that source type. As a result, this procedure is repeated several times, each time taking a new sample into the training set. After performing this several times, most obvious anomalies are removed, though many benign spectra are removed as well. An automated approach for simultaneously identifying spectral anomalies while learning background models is an important area of research, but is not examined further in this work.

3.2.2 Source Injection Data Preparation

Two sources of data are used to perform source injection with RadMAP: background-subtracted ^{137}Cs spectra, and simulated ^{133}Ba spectra. For both radionuclide types, list mode source data was provided by ref. [44].

Background-subtracted Data

Following the procedure described in Section 3.1.2, background-subtracted spectra were generated from a set of 10 drive-bys past a 4 mCi ^{137}Cs source. In this measurement scenario, RadMAP moved on a straight path at a speed of approximately 15 mi/h (6.7 m/s), while the source was positioned 20 m off the road on the starboard side of the vehicle. The 10 runs were aligned in time using the background-subtracted count rate. That is, the time at which the background-subtracted count rate was highest was assumed to be the time at which the vehicle was closest. Using the time of closest approach, three spectra before and after the maximum were kept, yielding a sequence of seven spectra for each drive-by. The mean spectrum for each of the seven spectra is computed over all 10 drive-bys, yielding a matrix $\Lambda_s \in \mathbb{R}_+^{7 \times 125}$, representing the mean rate over a run.

RadMAP Simulations

Simulations of ^{133}Ba were performed using a toolkit known as SoftWare for Optimization of Radiation Detectors (SWORD) [45]. As with the ^{137}Cs background-subtracted measurements, the simulations used seven sampling points, three on each side of the position at the distance of closest approach. The simulated spectra correspond to 1 s measurements of a ^{133}Ba source with 101.4 mCi activity. The resulting spectra at each of the seven positions are linearly scaled down to the maximum source used in measured data, 4 mCi. As with ^{137}Cs , this procedure results in a matrix of mean spectra $\Lambda_s \in \mathbb{R}_+^{7 \times 125}$.

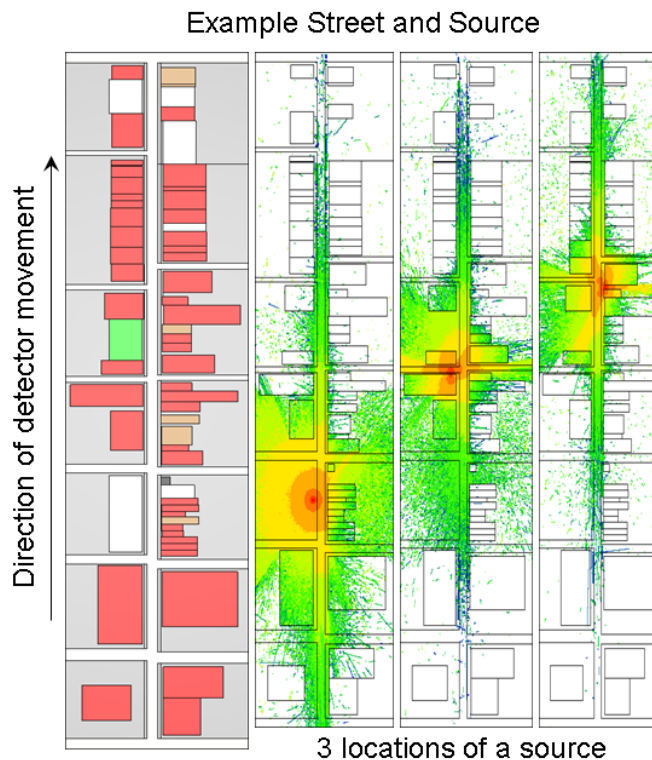


Figure 3.9: Figure from [46]. The left panel of the figure represents a top view of a single street in the simulated URSC environment. The rectangles in this figure represent buildings, each of which has a unique geometry and background emission signature, resulting in a highly variable sequence of measurements as the system moves down the street. The right three panels of the figure show the effect of placing a source at three different locations on the same street, where the color indicates the source intensity as a function of position. These three panels illustrate the variation in spectral measurements due to environmental scattering.

3.3 Source Injection using Urban Radiation Search Competition Data

A public data competition for detecting radiation sources in urban environments, sponsored by the United States Department of Energy, was held between March 8th, 2019 and April 24th, 2019. The objective of this competition, referred to here as the Urban Radiation Search Competition (URSC), was to detect, identify, and localize gamma-ray sources in a series of runs of spectral gamma-ray data. Each run consisted of simulated list-mode data produced by a 2" \times 4" \times 16" NaI(Tl) gamma-ray detector moving through an urban environment in a straight path at a fixed speed. Each run also provided a unique envi-

ronment, similar to how streets in an urban environment vary in both building materials and shapes, as well as the presence of open spaces. Variation in building materials results in variation in background radiation emissions, while variation in geometries yields variable scattering environments for gamma rays. The result is highly variable background data, and when a source is present, a complex scattering environment that varies with source position. Figure 3.9 shows an example street consisting of different buildings drawn as rectangles, along with a source placed at three different locations, illustrating variability resulting from source positioning.

The training dataset provided for the competition consisted of list mode data and labels which indicated the source type and time of closest approach, if applicable. A separate testing dataset was provided for perform predictions on, though it is not used in this work. A given run either contains no source, or one of the following possibilities: ^{60}Co , ^{131}I , $^{99\text{m}}\text{Tc}$, WGPu, highly-enriched uranium (HEU), or a combination of HEU and $^{99\text{m}}\text{Tc}$, each of which can be shielded or unshielded. While the dataset is useful for comparing algorithms in challenging variable environments, there is a dearth of source information (e.g., activity, shielding type) makes it difficult to use in performing a quantitative evaluation of algorithms. However, source can be injected into background data provided with the URSC dataset, making it an excellent resource for analyzing algorithms. In this work, background data from the URSC dataset is used as the basis for source injection, and simulations are performed using the procedure described in Section 3.1.3.

3.3.1 URSC Background Data

The URSC training data consists of 9700 runs, 4900 of which only contain events resulting from background radiation. In this work, only the 4900 background runs are utilized for source injection. The background runs provide a total of 170 h of list mode data, with the mean run being 125 s in length. Of this data, a random sample of 90% of the runs (4410 runs) are used for training data, and the remaining 10% (490 runs) are used for evaluation. Note that in this work, all spectra are formed using energy bins with edges between 30 keV and 3000 keV, using 128 bins with "square root" bins.

Figure 3.10 illustrates the variability in the count rate in the dataset by showing the count rate of the first 60 seconds of three randomly-sampled runs, and compares this to a histogram of gross counts across the entire URSC dataset. To show the variability in spectral information, Figure 3.11 shows a two-dimensional histogram of background spectra. The the contour of spectra alone does not convey information about the relationship between different spectral bins, so the right pane of Figure 3.11 shows the correlation matrix computed over the entire background dataset. Both of these confirm that, despite being simulated, the URSC dataset contains high variability, and makes it useful for developing and assessing detection and identification algorithms.




Figure 3.10: (Left pane) Gross count rate as for the first 30 seconds of three randomly-sampled background-only runs from the URSC training set. Transient peaks are seen, resulting from the detector system passing a building or environmental feature with elevated emission rate. In addition to variability in gross count rate, there is temporal variation in each spectral bin, reflecting variability of the type of radionuclide emission encountered as a function of time. (Right pane) Probability density function of gross counts across all background spectra in the URSC training dataset compared to the probability mass function of a Poisson distribution having a mean equal to the mean from background samples. Due to the variability in background environments, the variance of gross counts from samples greatly exceeds that of a Poisson distribution with the same mean.

3.3.2 Source Simulations

The procedure and detector model described in Section 3.1.3 is used here for simulating source spectra to be used in source injection. Specifically, the detector response for the NaI(Tl) bar is simulated and used to generate source templates for arbitrary gamma-ray sources. The kinematics for a specific setup are then used to model the motion of the detector past the source. In particular, this work uses a standoff distance of 10 m and a vehicle speed of 5 m/s.

Table 3.2 lists the sources simulated, along with their most prominent gamma-ray energies, specifically, gamma-ray energies with intensities over 1%. In Chapter 5, these source templates are used for both training and assessing detection and identification models.

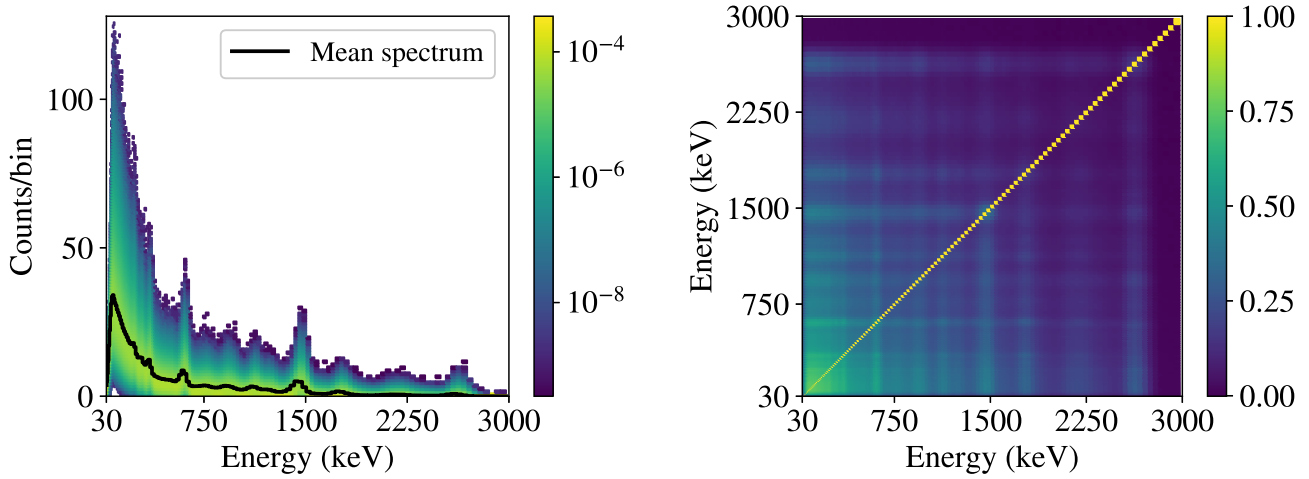


Figure 3.11: (Left pane) Normalized two-dimensional histogram of background spectra from the URSC dataset, illustrating the variability in spectra encountered in the URSC dataset. A high degree of variability is seen around the mean spectrum, drawn as a solid black line. (Right pane) Correlation matrix of background dataset, illustrating the relationship between energy bins that is not captured in the figure in the left pane.

3.4 Performance Metrics

Algorithms that are evaluated using source injection must be compared on a common set of metrics in order to understand relative performance. This section introduces metrics that are used in the remainder of this work for evaluating performance.

Detection algorithms result in a binary output (i.e., a source is detected or not), and as a function of the signal-to-background ratio, or in this case, activity, there exists a transition between these two binary states. The quantitative assessment of algorithm performance then includes characterizing this transition. Averaging the outputs of a detection algorithm over several trials as a function of activity results in the *probability of detection* $\hat{p}_D(A)$. At a given false alarm rate, the probability of detecting a particular source at activity A is the proportion of times the source was detected k out of n trials:

$$\hat{p}_D(A) = \frac{k}{n}. \quad (3.14)$$

In this context, a trial can either refer to the number of spectra analyzed or the number of source passes analyzed, depending on how the analysis is defined. Note that this procedure ideally defines a binomial process with true probability of success (i.e., detection) p^* , and that $\hat{p}_D(A)$ is the maximum likelihood estimate of p^* .

Table 3.2: Simulated gamma-ray sources and their most prominent gamma-ray energies

Source	Energies (keV)	Source	Energies (keV)
¹⁹⁸ Au	411.8	¹¹¹ In	171.3, 245.4
¹³³ Ba	81.0, 276.4, 302.9, 356.0, 383.8	¹⁹² Ir	296.0, 308.5, 316.5, 468.0
⁸² Br	554.3, 619.1, 698.4, 776.5, 827.8, 1044.0, 1317.5, 1474.9	⁵⁴ Mn	838.8
⁵⁷ Co	122.0, 136.5	¹²⁴ Sb	602.7, 722.8, 1691.0
⁶⁰ Co	1173.2, 1332.5	⁴⁶ Sc	889.3, 1120.5
¹³⁷ Cs	661.7	⁷⁵ Se	121.1, 136.0, 264.6, 279.5, 400.6
¹⁵² Eu	121.8, 344.3, 778.9, 964.1, 1085.8, 1112.0, 1408.0	¹¹³ Sn	255.1, 391.7
¹²³ I	159.0, 529.0	²⁰¹ Tl	68.9, 70.8, 80.2, 167.4
¹³¹ I	284.3, 364.5, 637.0, 722.9		

At low activities, $\hat{p}_D(A)$ tends to 0, and as source activity is increased, $\hat{p}_D(A)$ approaches 1. This behavior can be modeled using a sigmoid function of the form

$$q(A; \mu, \sigma) = \frac{1}{1 + e^{-(A-\mu)/\sigma}}, \quad (3.15)$$

where μ and σ are parameters that determine the shape. Equation (3.15) is fit to empirical values of $\hat{p}_D(A)$ using a binomial likelihood function of the form

$$p(\mathbf{k}|n, q(A; \mu, \sigma)) = \prod_{i=1}^N \binom{n}{k_i} q_i^{k_i} (1 - q_i)^{n-k_i}, \quad (3.16)$$

where \mathbf{k} refers to a collection of N detections made at corresponding activities A_i , and q_i is an abbreviation for $q(A_i; \mu, \sigma)$ for activity A_i . In particular, MLE is used to estimate values of μ and σ which result in the sigmoid function that best models the measured data points. Appendix A provides a full derivation of the objective function used in estimating μ and σ .

Figure 3.12 shows an example of a sigmoid function representing a "true" detection process with parameters $\mu^* = 155 \mu\text{Ci}$, $\sigma^* = 15 \mu\text{Ci}$. This distribution is sampled at activities between $10 \mu\text{Ci}$ and $270 \mu\text{Ci}$ in steps of $10 \mu\text{Ci}$. At each sampled activity A_i , k_i detections are made, and the entire collection of points are fit using the model in Equation (3.16), and the resulting model is also shown in the figure.

The probability of detection curve for a given algorithm under a particular detection scenario can be summarized using a single statistic, the Minimum Detectable Activity

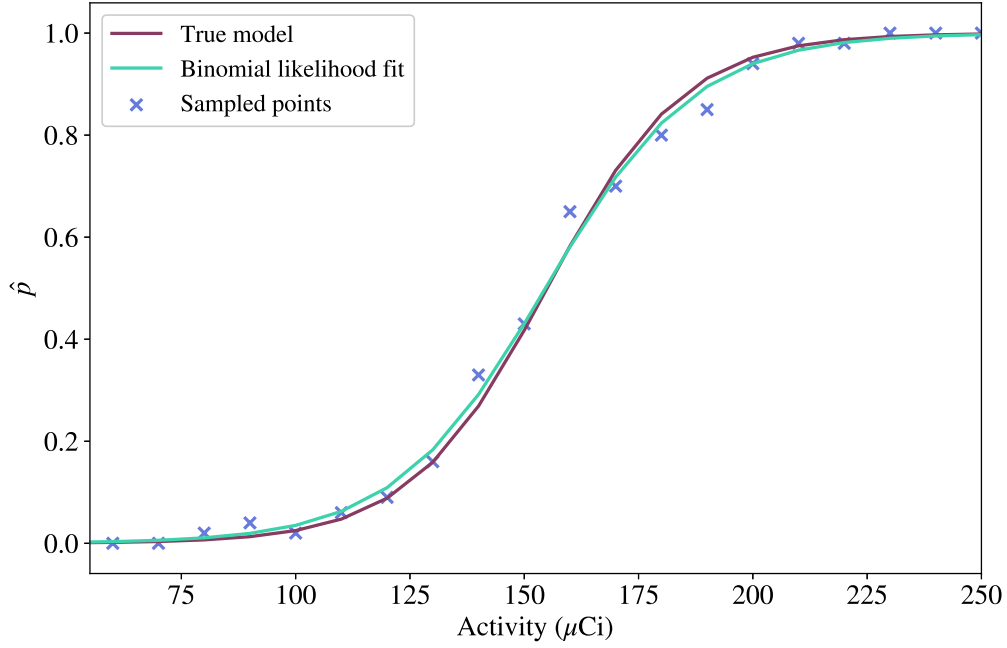


Figure 3.12: Sigmoid functions used for modeling probability of detection as a function of source activity. At each activity, 100 samples are drawn from the true distribution with parameters $\mu^* = 155 \mu\text{Ci}$, $\sigma^* = 15 \mu\text{Ci}$ and fit to the model in Equation (3.16), resulting in estimated parameters $\hat{\mu}^* = 154.6 \mu\text{Ci}$, $\sigma^* = 16.5 \mu\text{Ci}$. In characterizing detection algorithms for particular source scenarios, the underlying parameters μ^* and σ^* are unknown and are estimated by fitting Equation (3.16) to measured points. Using Equation (3.17), this example would result in an MDA of $203.2 \mu\text{Ci}$ using the estimated parameters $\hat{\mu}^*$ and σ^* .

(MDA), which is the activity at which a source can be detected with a given probability (e.g., 95%). With all other parameters are held fixed, including false alarm rate, an algorithm with a lower MDA than other candidate algorithms is preferred, as it is able to detect weaker sources. From the analytical model given by Equation (3.15), the MDA at probability of detection q (e.g., 95%) can be found by solving for A :

$$\text{MDA}(q) = \mu - \sigma \ln(q^{-1} - 1) \quad (3.17)$$

and the standard error on the MDA is estimated as

$$\sigma_{\text{MDA}}^2 = \sigma_{\mu}^2 + \sigma_{\sigma}^2 (\ln(q^{-1} - 1))^2, \quad (3.18)$$

where σ_{μ} and σ_{σ} are the uncertainties of parameters μ and σ , respectively. The parameter uncertainties are estimated from the diagonal elements of the inverse Hessian matrix

computed during optimization. Using Equation (3.17), the example shown in Figure 3.12 would result in an MDA of $203.2 \mu\text{Ci}$ using the estimated parameters $\hat{\mu}^*$ and σ^* .

When computing the probability of detection, a set of sampling activities must be chosen. In the simplest form, sample points can be linearly spaced, covering some wide range of activities that is likely to contain the transition point from the source being undetected to detected. This approach is computationally expensive when the transition is unknown, for example, in a new data injection configuration for an algorithm. Adaptive approaches, which are able to pinpoint the transition region and sample activities around the transition can also be used for sampling. See Appendix B for a comparison of adaptive sampling techniques that are introduced in this dissertation.

Chapter 4

Non-negative Matrix Factorization of Gamma-ray Spectra

Gamma-ray spectra contain rich information about the local radiological background, nearby gamma-ray sources, and even the surrounding environment from which gamma rays scatter. Ideally, spectral models used for detection and identification capture this information in an interpretable manner, possibly allowing for introspection into these physical processes. This chapter introduces the use of Non-negative Matrix Factorization (NMF) for approximating gamma-ray spectra. NMF models of gamma-ray spectra often have a clear physical interpretation, extracting spectral features that can be understood in terms of the local gamma-ray environment. Furthermore, this approach generally yields accurate spectral models using relatively few parameters, and these models can be used as the basis for anomaly detection and source identification.

This chapter begins with a brief discussion on dimensionality reduction, a concept which underlies NMF, then follows with a detailed description of NMF is given, along with procedures for detecting and identifying gamma-ray sources using these models. NMF-based detection and identification algorithms are then evaluated, and generally show improvements over the benchmark methods. The chapter closes with a discussion on how NMF-based models can potentially enable more sophisticated detection capabilities which leverage additional, non-radiological data streams.

Much of the content in this chapter originally appeared in the following publications:

- K. J. Bilton, T. H. Joshi, M. S. Bandstra, J. C. Curtis, B. J. Quiter, R. J. Cooper, and K. Vetter, "Non-negative Matrix Factorization of Gamma-Ray Spectra for Background Modeling, Detection, and Source Identification." In: *IEEE Transactions on Nuclear Science* vol. 66, no. 5 (May 2019), pp. 827-837. doi: 10.1109/TNS.2019.2907267 [47].
- K. J. Bilton, M. S. Bandstra, T. H. Joshi, J. C. Curtis, R. J. Cooper, and K. Vetter, "Modeling Shielded Gamma-ray Source Spectra using Non-negative Matrix Factorization." *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, Manchester, United Kingdom, 2019, pp. 1-7, doi: 10.1109/NSS/MIC42101.2019.9059733 [48].

4.1 Dimensionality Reduction

Dimensionality reduction is a procedure for encoding a vector $\mathbf{x} \in \mathbb{R}^d$ as a vector $\mathbf{a} \in \mathbb{R}^k$, where k is typically much less than d , yielding a compact representation of the original data. For linear dimensionality reduction techniques, a data matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ consisting of n d -dimensional vectors can be approximated as the inner product of two low-rank matrices:

$$\hat{\mathbf{X}} = \mathbf{A}\mathbf{V} \quad (4.1)$$

where $\mathbf{A} \in \mathbb{R}^{n \times k}$ are weights encoding the data samples \mathbf{X} in the low-dimensional space, and $\mathbf{V} \in \mathbb{R}^{k \times d}$ is a set of k basis vectors spanning a subspace of \mathbb{R}^d . This allows measurements to be approximated as a linear combination of the rows of \mathbf{V} , which are components that capture structure in the data. One use of dimensionality reduction is data visualization [49], in which high-dimensional data \mathbf{X} are projected to two or three dimensions using $k = 2$ or $k = 3$, respectively. The coordinates \mathbf{A} in the low-dimensional subspace can then be visualized, possibly allowing one to see patterns in the data captured by the components in \mathbf{V} that may not have been transparent at the full dimensionality. Before examining NMF more closely, two dimensionality reduction methods that have been previously applied to gamma-ray spectral analysis are discussed.

4.1.1 Principal Component Analysis

The underlying idea of Principal Component Analysis (PCA) [50] is to find an orthogonal basis \mathbf{V} for representing data, where the vectors in \mathbf{V} are ordered by the amount of variance in the data accounted for by each vector. Initially, d basis vectors are produced, and the top k basis vectors are kept, reducing the dimensionality.

In standard PCA, basis vectors are found by computing the Singular Value Decomposition (SVD) on the covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ computed from samples of data \mathbf{X} . Due to a higher number of counts, and thus variance, in the low-energy region of gamma-ray spectra, the standard PCA approach will result in basis vectors that correspond to low-energy bins, disregarding valuable information at higher energies. To address this behavior, it is more appropriate to perform the SVD on the correlation matrix \mathbf{C} , which is found by dividing each feature (i.e., energy bin) by the variance of the feature. The PCA formalism used in gamma-ray spectroscopy is presented in detail in Appendix C.

PCA can also be understood from a maximum likelihood perspective. Specifically, PCA can be treated as a latent variable model with negative log-likelihood [51]

$$-\ln p(\mathbf{X}|\boldsymbol{\mu}, \mathbf{W}, \sigma^2) = \frac{nd}{2} \ln(2\pi) + \frac{n}{2} \ln \det(\mathbf{S}) + \frac{1}{2} \sum_{i=1}^n (\mathbf{x}_i - \boldsymbol{\mu})^\top \mathbf{S}^{-1} (\mathbf{x}_i - \boldsymbol{\mu}) \quad (4.2)$$

where $\mathbf{S} = \sigma^{-2}\mathbb{I} - \sigma^{-2}\mathbf{W}\mathbf{M}^{-1}\mathbf{W}^\top$, $\mathbf{M} = \mathbf{W}^\top\mathbf{W} + \sigma^2\mathbb{I}$, and $\boldsymbol{\mu}$, \mathbf{W} , and σ^2 are model parameters. Note that Equation (4.2) is consistent with a Gaussian likelihood model, suggesting that PCA implicitly assumes a Gaussian statistical model. As previously discussed in

Section 2.5, gamma-ray spectroscopy, particularly at low count rates, is appropriately described using Poisson statistics.

Figure 4.1 shows PCA weights and components computed from background data. The basis vectors \mathbf{V} resulting from PCA are orthogonal, and thus contain both positive and negative elements. Although the first several components may resemble spectra from known background contributions such as ^{40}K and uranium and thorium series radioisotopes, the orthogonality of the components rules out true physical solutions, which must be non-negative.

Previous work has used PCA as the basis for anomaly detection and identification [13, 52]. The PCA-based Spectral Anomaly Detection (SAD) alarm metric [32, 53, 31] uses the first k PCs $\hat{\mathbf{V}}$ found from PCA to measure deviations between a spectrum and its low-dimensional PCA approximation

$$\text{SAD}(\mathbf{x}) = \frac{\|\mathbf{x} - \hat{\mathbf{x}}\|_2}{\sqrt{\|\mathbf{x}\|_1}}, \quad (4.3)$$

where $\hat{\mathbf{x}}$ is the PCA reconstruction of the input spectrum \mathbf{x} . The normalization in the denominator is included to account for larger residuals at higher count rates. See Section C for additional details about the SAD method.

4.1.2 Poisson Principal Component Analysis

More recently, Poisson PCA (PPCA) has been used to accurately model the underlying Poisson statistics of photon detection [53]. PPCA is a specific case of Exponential family Principal Component Analysis [54], and it has shown an improvement over standard PCA in source detection at low signal-to-noise, presumably due to more accurately modeling the statistics at low counts [55]. Although PPCA assumes the correct statistical model and generates non-negative spectral components, the components still lack an intuitive, physical interpretation because they are defined in logarithmic space and are therefore multiplicative, not additive. While PCA and PPCA often result in models containing features that one may recognize as resulting from physical processes (e.g., a component containing a 1460 keV peak from ^{40}K), these methods do not fully capture the additive, non-negative nature of gamma-ray spectra.

4.2 Non-Negative Matrix Factorization

Non-negative Matrix Factorization (NMF) [56, 57] is a linear dimensionality reduction technique that approximates a matrix \mathbf{X} in terms of the matrices \mathbf{A} and \mathbf{V} , where all three matrices are constrained to have non-negative entries. Prior to performing NMF, the rank k must be specified, which is in contrast to PCA, where one can vary k without recomputing the decomposition. The factorization in Equation (4.1) can be performed using MLE,

and given the discrete nature of gamma-ray spectra, the Poisson log-likelihood in Equation (2.15), which relates the measured data x_{ij} in \mathbf{X} and the low-rank approximation \hat{x}_{ij} in $\hat{\mathbf{X}}$, is used in the optimization. From an MLE perspective, each element \hat{x}_{ij} can be thought of as the mean rate per energy bin per spectrum.

Maximizing the log-likelihood in Equation (2.15) is equivalent to minimizing the Kullback-Leibler (KL) divergence between \mathbf{X} and $\hat{\mathbf{X}}$. The KL divergence is a quantity from information theory [58] that measures how well a given "true" probability distribution is approximated by an estimate of the distribution. Here, the "true" distribution is given by the measured data \mathbf{X} , while the estimate is $\hat{\mathbf{X}}$. Though performing NMF is generally a non-convex problem with no exact solution, meaning solutions tend to be local minima, the error between the data \mathbf{X} and the estimate $\hat{\mathbf{X}} = \mathbf{A}\mathbf{V}$, as measured by the KL divergence, has been shown to be monotonically non-increasing under repeated application of the following iterative update rules [59]:

$$a_{il} \leftarrow a_{il} \frac{\sum_j v_{lj} x_{ij} / \hat{x}_{ij}}{\sum_j v_{lj}} \quad (4.4)$$

$$v_{lj} \leftarrow v_{lj} \frac{\sum_i a_{il} x_{ij} / \hat{x}_{ij}}{\sum_i a_{il}} \quad (4.5)$$

$$v_{lj} \leftarrow \frac{v_{lj}}{\sum_m v_{lm}} \quad (4.6)$$

$$a_{il} \leftarrow a_{il} \sum_m v_{lm} \quad (4.7)$$

Equation (4.4) is used to update each element a_{il} of \mathbf{A} , then Equation (4.5) is used to update each element v_{lj} of \mathbf{V} . The rows of \mathbf{V} are then normalized using Equation (4.6), and the weights are updated again using Equation (4.7) based on the normalization of \mathbf{V} . The procedure of using the update rules in Equation (4.4) – Equation (4.7) to perform NMF is hereafter referred to as Poisson NMF (PNMF).

As a comparison to PNMf, NMF performed using a least squares objective function, referred to as L2NMF, is also examined. The objective function for L2NMF takes the form:

$$\mathcal{L}(\mathbf{X}, \hat{\mathbf{X}}) = \|\mathbf{X} - \hat{\mathbf{X}}\|_F^2 \quad (4.8)$$

where F denotes the Frobenius norm. From an MLE perspective, the objective function in Equation (4.8) implicitly assumes a Gaussian distribution in each bin with mean \hat{x}_{ij} and

having equal variance. The iterative update rules for L2NMF [59] are given by:

$$\begin{aligned}
 a_{il} &\leftarrow a_{il} \frac{(\mathbf{X}\mathbf{V}^\top)_{il}}{(\mathbf{A}\mathbf{V}\mathbf{V}^\top)_{il}} \\
 v_{lj} &\leftarrow v_{lj} \frac{(\mathbf{A}^\top \mathbf{X})_{lj}}{(\mathbf{A}^\top \mathbf{A}\mathbf{V})_{lj}} \\
 v_{lj} &\leftarrow \frac{v_{lj}}{\sum_m v_{lm}} \\
 a_{il} &\leftarrow a_{il} \sum_m v_{lm}
 \end{aligned} \tag{4.9}$$

As with the PNMF update rules, the L2NMF update rules are used to sequentially update weights \mathbf{A} , components \mathbf{V} , then a normalization of bases is performed.

Given a data matrix \mathbf{X} , a desired rank k , and initial values for \mathbf{A} and \mathbf{V} (e.g., using random initialization), Equation (4.4) – Equation (4.7) and Equation (4.9) yield approximate NMF solutions using generalized KL divergence and least squares objective functions, respectively. After initializing \mathbf{A} and \mathbf{V} , both matrices are updated in an alternating fashion by applying the multiplicative rules on one while the other is held constant, which is referred to as training. Similar to other iterative methods, the multiplicative updates on \mathbf{A} and \mathbf{V} are typically repeated for a fixed number of steps, or until the value of the reconstruction error, as measured by Equation (2.15) or Equation (4.8), converges within some specified tolerance.

As a result of the iterative nature, NMF is more computationally intensive than techniques such as PCA. Additionally, the non-convexity of NMF means that solutions may yield different combinations of weights and basis vectors for different initializations. For the applications described in this work, however, these drawbacks have not been prohibitive. For example, while the training procedure generally requires significantly more time to complete than for PCA, training is typically performed offline, meaning it does not present an operational burden. Given a trained model \mathbf{V} , the procedure for encoding a single spectrum with NMF weights is carried out in near real-time, meaning anomaly detection and identification can be performed as spectra are being collected.

In performing dimensionality reduction, the rank k of the low-dimensional subspace must be specified. In general, increasing k will result in enhanced fits. For complete decomposition methods such as PCA and NMF, as $k \rightarrow d$, the reconstruction error will vanish, resulting in the loss of ability to detect anomalies. As a result, the aim is to choose a value of k that allows models to describe background variations in data sufficiently well, but not at the cost of losing the ability to detect anomalies within spectra.

The Akaike Information Criterion (AIC) [60] is used for determining the number of components used, and is defined as:

$$\text{AIC} = 2\kappa - 2 \ln p(\mathbf{X}|\hat{\mathbf{X}}) \tag{4.10}$$

where $\kappa = kn + kd - k$ is the total number of free parameters in the model $\hat{\mathbf{X}} = \mathbf{AV}$, with k degrees of freedom being removed due to the normalization constraint on \mathbf{V} . The value of k that results in the lowest AIC on the training data is used, as it yields the best fit with the fewest parameters.

4.3 NMF for Background Modeling

To perform background modeling, the process of generating the basis \mathbf{V} that captures background variation, a matrix $\mathbf{X} \in \mathbb{R}_+^{n \times d}$ of n d -dimensional spectra is constructed. For the factorization $\hat{\mathbf{X}} = \mathbf{AV}$, the rows of \mathbf{V} are normalized, such that the weights \mathbf{A} can be interpreted as the number of gamma-ray counts in the spectrum from the corresponding NMF component.

The left side of Figure 4.1 shows three components and the associated weights resulting from performing 30,000 PNMF fitting iterations on 86,400 one-second RadMAP NaI(Tl) spectra randomly-sampled from various locations around the San Francisco Bay Area. Note that the ordering of resulting NMF components is arbitrary, unlike PCA where lower component numbers capture more variance in the data. However, the NMF components in Figure 4.1 were numbered to match the ordering of the corresponding PCA components shown on the right side of Figure 4.1. To address the non-uniqueness of solutions described in Section 4.2, note that in this section, each NMF decomposition is initialized using non-negative double singular value decomposition (NNDSVD) [61], which allows for deterministic NMF approximations. NMF solutions generated using NNDSVD-initialized matrices do not necessarily result in optimal performance, with respect to detecting and identifying sources of interest, but these solutions have found to work sufficiently well for the analyses performed in this work.

The NMF components in Figure 4.1 capture features that have physical meaning. Component 1 appears to be the mean spectral shape. Component 2 has a higher magnitude than the other two at energies below 125 keV, which may be indicative of distant sources or multiple scatters. Component 3 is a combination of the ^{40}K 1460 keV line and the ^{208}Tl 2615 keV line, among others, and generally shows more pronounced peaks than Component 1. As a result of the physical interpretation, when encoding a spectrum using NMF, the component weights yield insight to the composition of the gamma-ray background. This interpretation is particularly of interest for the integration of gamma-ray detectors with other environmental sensors, where the latter could potentially provide information about the composition of the local gamma-ray background.

As a comparison with components formed using NMF, the upper right panel of Figure 4.1 shows the first three principal components (PCs) found by performing PCA using the correlation matrix generated from \mathbf{X} , using the procedure described in Section 4.1.1. The comparison of components \mathbf{V} and distributions of weights \mathbf{A} from both NMF and PCA shows that there are similarities in the structure that is captured within components. The distributions of weights for component 1 for NMF and PCA, which contain the mean

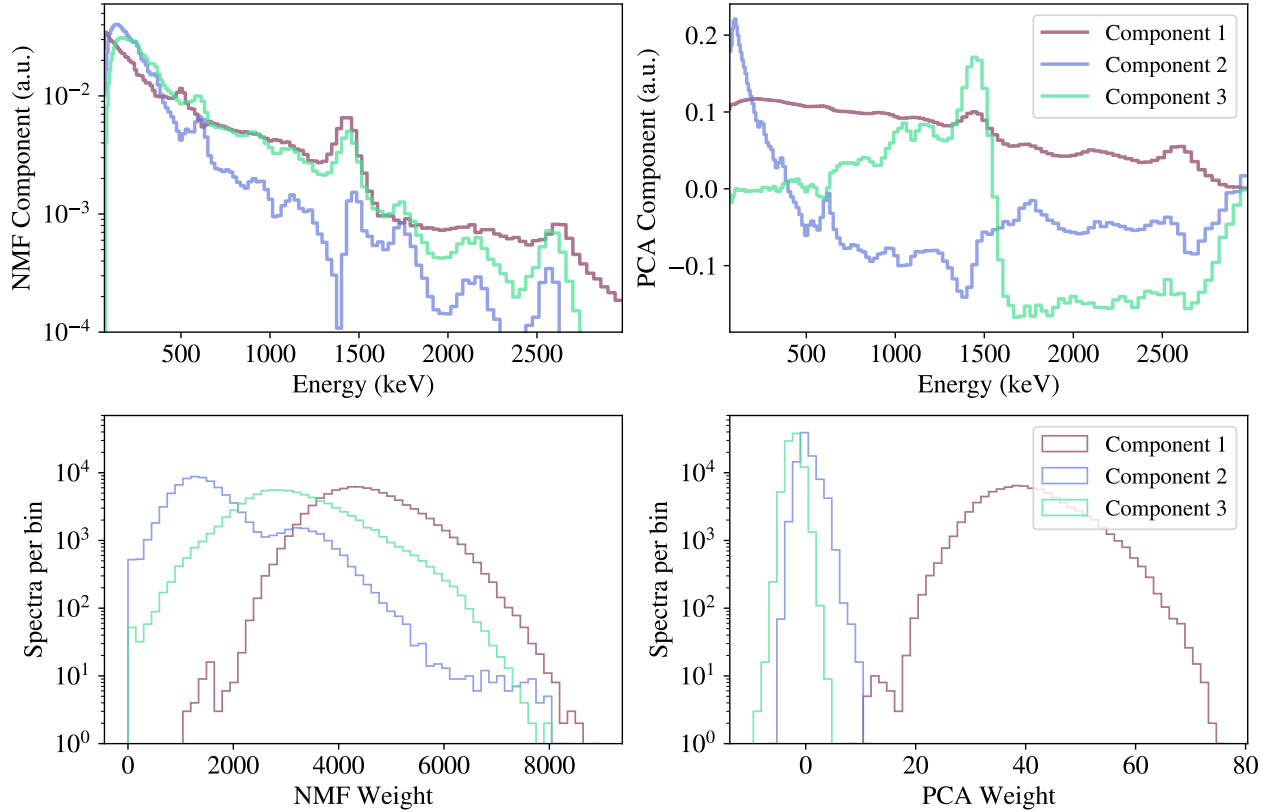


Figure 4.1: Three Poisson NMF basis vectors \mathbf{V} formed from background spectra collected by the RadMAP mobile detection system (upper left) and histograms of weights \mathbf{A} corresponding to each component (lower left). The first three principal components formed from background spectra (upper right) and histograms of weights corresponding to each component (lower right). For both NMF and PCA, models were formed using 86,400 randomly-sampled spectra, each with one-second integration time, and for NMF, training was done over 30,000 fit iterations. Note that there is no significance in the numbering of NMF components, unlike in PCA, however, NMF components were numbered to match the order of corresponding PCs.

spectral shape, are similar both in shape and also in that both capture the most variance. Differences in spectral shapes are also apparent, the most obvious being that components 2 and 3 from PCA capture both positive and negative spectral features.

4.4 Anomaly Detection using NMF

The NMF formulation described above can be used for anomaly detection by finding the reconstructed low-rank approximation $\hat{\mathbf{x}} = \mathbf{a}^T \mathbf{V}$ of a spectrum \mathbf{x} , and treating $\hat{\mathbf{x}}$ as the

mean rate per energy bin per spectrum. First, the p-value p_i is calculated in each bin i . The intuition behind using a p-value is to estimate how consistent the measurement x_i is with the background model \hat{x}_i . To account for measurements that are either much higher or lower than the mean rate, consider the two-tailed p-value

$$p_i = 2 \min(F(x_i; \hat{x}_i), 1 - F(x_i; \hat{x}_i)) \quad (4.11)$$

where $F(x_i; \hat{x}_i)$ is the Cumulative Distribution Function (CDF) of the Poisson distribution with mean rate \hat{x}_i . P-values are aggregated to form the NMF anomaly detection score:

$$\text{AD}_{\{P,L2\}}(\mathbf{x}) = -\frac{1}{d} \sum_{i=1}^d \ln p_i. \quad (4.12)$$

Using the procedure described in Section 2.6, a threshold for a given model at a set false alarm rate can be produced by evaluating the metric over a set of background data. Two existing methods, Gross Counts $K\sigma$, referred to here simply as $K\sigma$, and the PCA-based anomaly detection SAD, are compared to AD_P and AD_{L2} .

4.4.1 Benchmark: Gross Counts $K\sigma$

$K\sigma$ measures how many standard deviations σ the gross counts in a spectrum are away from the mean number of background counts μ as:

$$K = \frac{\|\mathbf{x}\|_1 - \mu}{\sigma}, \quad (4.13)$$

where the 1 subscript in the denominator denotes the L1 norm of \mathbf{x} , or gross counts. Here, the mean μ and standard deviation σ of the total number of counts within a spectrum are determined from the training set.

4.4.2 AD_l in the Limit of Known Background

To understand the limitations of the NMF-based anomaly detection and identification methods, the performance of each algorithm in the ideal case, where the mean background rate for each spectrum is known, is examined. That is, for each set of injected spectra \mathbf{X}_{test} consisting of background and source contributions (i.e., $\mathbf{X}_{\text{test}} = \mathbf{X}_{\text{bkg}} + \mathbf{X}_{\text{src}}$), it is assumed that the mean background rate $\hat{\mathbf{X}}_{\text{bkg}}$ resulting from performing NMF on each spectrum is known. Note that the background used here \mathbf{X}_{bkg} are the 2 h of background data referred to previously in Section 3.2, which are separate measurements from the training data used in generating models.

In this best case scenario, algorithm performance is primarily limited by statistical fluctuations in the background term \mathbf{X}_{bkg} , referred to as the Known Background Limit (KBL). One reason for exploring the KBL is that this procedure lends insight to how much

performance can be improved with an enhanced understanding of the local radiological background.

To determine KBL performance, the following procedure is used:

1. Project the measured background contribution \mathbf{X}_{bkg} onto the subspace generated by \mathbf{V} to yield the NMF weights \mathbf{A}_{bkg} .
2. Estimate the mean background rate as $\hat{\mathbf{X}}_{\text{bkg}} = \mathbf{A}_{\text{bkg}} \mathbf{V}$.
3. Use $\hat{\mathbf{X}}_{\text{bkg}}$ as the reconstructed spectra to which the measured spectra \mathbf{X} are compared.

4.4.3 Performance Evaluation

To determine the number of components to use for NMF and PCA-based anomaly detection models, Equation (4.10) was used using $n = 86,400$ 1-s spectra from the RadMAP system described in Section 3.2, and both PNMf and PCA yielded a value of $k = 2$ components. The AIC for L2NMF was seen to decrease monotonically with k , and as a result $k = 2$ was also used to be consistent with the other models, though additional values of k are also considered for completeness. For readers familiar with PCA, the proportion of variance explained as a function of k was also examined for PCA and suggested a value of $k = 2$, in agreement with the AIC. Each PNMf, L2NMF, and PCA model described in the following sections using the RadMAP dataset will then use $k = 2$ components trained over a single subset of 24 hours of background data sampled from the 55-hour training set. Each model is then evaluated over the entire 55-hour training set in empirically determining thresholds.

Values for NMF-based and PCA-based anomaly detection scores Equation (4.12) and Equation (4.3) were calculated over the background training set \mathbf{X} , and a threshold on scores were empirically determined using a $1/8 \text{ hr}^{-1}$ ($3.5 \times 10^{-5} \text{ s}^{-1}$) FAR. The same 24 hour sample of spectra used in training the NMF models was used in training PCA models, and the entire 55-hour set of data was used for threshold setting.

Figure 4.2 shows a comparison of probability of detection curves for NMF and PCA-based anomaly detection algorithms, each using $k = 2$. Table 4.1 shows the corresponding MDA_{95} resulting from a fit of P_D to a sigmoid for the $k = 2$ models, as well as models using $k = 4$. Note that models with $k > 4$ were also examined, but performance did not improve significantly, and results are omitted. From Figure 4.2 and Table 4.1, AD_P and AD_{L2} are shown to maintain an MDA near $200 \mu\text{Ci}$ for both sources, while AD_{PC} has an increased MDA for ^{137}Cs , which is likely resulting from the difference in anomaly detection metrics between the NMF and PCA-based models. Note, however, that for both sources and for both $k = 2$ and $k = 4$, AD_P outperforms AD_{L2} , which is perhaps attributed to the use of a more accurate statistical model in AD_P . Furthermore, in detecting ^{137}Cs , AD_P has a similar MDA to AD_{PC} KBL. That is, by using a more accurate statistical model and computing a different anomaly detection score, AD_P is able to detect




Figure 4.2: Probability of detection of ^{133}Ba (top) and ^{137}Cs (bottom) sources at 20-m standoff using NMF-based detection algorithms, a PCA-based detection algorithm, and both NMF and PCA-based algorithms in the case of a known background, each model using $k = 2$ components at $1/8 \text{ hr}^{-1}$ FAR and 1-s integration time. Solid lines show the measured performance, and dashed lines show performance in the KBL. Note that only anomaly detection is performed here, meaning sources are not identified by the algorithms. The 95% probability of detection is indicated by the dashed horizontal line. Error bars indicate the 68% Jeffreys interval. Not shown is a curve for gross counts, which has a significantly higher MDA.

weaker sources than the PCA-based method under a near-perfect understanding of the gamma-ray background. However, AD_{PC} offers a slight advantage over AD_{L2} for ^{133}Ba , suggesting that these performance enhancements may primarily exist for sources with gamma-rays at higher energies.

Not shown in Figure 4.2 are curves for $K\sigma$. As Table 4.1 shows, the MDA of $K\sigma$ is much higher than any of the NMF and PCA-based methods. The poor performance of $K\sigma$ reflects the background variability seen by RadMAP, as high thresholds must be set to maintain the FAR of $1/8 \text{ hr}^{-1}$.

Table 4.1: MDA_{95} (μCi + statistical error) for anomaly detection

Source Method	^{133}Ba	^{137}Cs
AD_P KBL, $k = 2$	119.9 ± 0.3	132.7 ± 0.4
AD_{L2} KBL, $k = 2$	132.2 ± 0.3	150.3 ± 0.6
AD_{PC} KBL, $k = 2$	131.2 ± 0.2	212.0 ± 0.5
AD_P , $k = 2$	188.5 ± 0.7	215.9 ± 0.5
AD_P , $k = 4$	181.3 ± 0.9	209.4 ± 0.7
AD_{L2} , $k = 2$	195.9 ± 0.4	240.3 ± 0.6
AD_{L2} , $k = 4$	182.4 ± 0.5	231.1 ± 0.8
AD_{PC} , $k = 2$	192.9 ± 0.4	336.5 ± 0.7
AD_{PC} , $k = 4$	189.1 ± 0.5	325.8 ± 0.9
$K\sigma$	1971.0 ± 13.6	2302.2 ± 16.1

When comparing the performance of the dimensionality reduction-based anomaly detection methods with their respective KBLs, it is apparent that there is room for improvement. By including additional information about the local gamma-ray environment, it is possible that these methods approach the case of the known background. However, the NMF-based methods are better suited for coupling with environmental sensors, due to the physical interpretation of NMF.

4.4.4 Estimating Uncertainties due to Empirical Thresholds

In using empirical thresholds, results are ultimately dependent on the data used to estimate thresholds. To understand the effect that the choice of dataset used for determining thresholds has on performance, an approach which estimates thresholds based on random subsets of data is examined here. In particular, the data is split into many contiguous segments of measurements, and a number of these segments are picked at random, without replacement, and used for computing thresholds. Such random sampling was performed 100 times, with each iteration having a combined measurement time of approximately 24h. Thresholds for AD_P , AD_{L2} , and AD_{PC} were computed for each of the 100 trials, resulting in a distribution of thresholds for each anomaly detection method. The estimated mean and standard deviation of each metric, along with thresholds generated using the entire 55-hour dataset, are given for $k = 2$ models in Table 4.2. This table shows that thresholds estimated using the entire dataset fall within one standard deviation of the mean of threshold generated using random subsets. The mean and standard deviation of these thresholds are then used to generate a lower and upper estimates for

Table 4.2: Comparison of empirical anomaly detection thresholds

Threshold	AD_P	AD_{L2}	AD_{PC}
Full dataset	1.337	1.420	1.611
Random subsets	1.341 ± 0.007	1.415 ± 0.012	1.613 ± 0.017

Table 4.3: Lower and upper estimates of MDA_{95} (μCi + statistical error) for anomaly detection

Method \ Source	^{133}Ba	^{137}Cs
	$AD_P, k = 2$	($187.5 \pm 0.6,$ 191.9 ± 0.6)
$AD_{L2}, k = 2$	($194.5 \pm 0.4,$ 199.9 ± 0.4)	($241.1 \pm 0.7,$ 251.9 ± 0.8)
$AD_{PC}, k = 2$	($190.4 \pm 0.4,$ 196.3 ± 0.5)	($331.4 \pm 0.7,$ 343.4 ± 0.7)

thresholds, yielding a lower and upper estimate on MDA. Lower and upper values for MDA using thresholds equal to a standard deviation below and above the mean, respectively, are reported in Table 4.3. The range of MDA values presented in Table 4.3 suggest that the procedure for computing AD_P may result in empirical thresholds that are less sensitive to data than AD_{L2} and AD_{PC} .

4.5 Source Identification using NMF

With the additive, parts-based interpretation of NMF in mind, the background model \mathbf{V} can be augmented to account for contributions from gamma-ray sources of interest by including a source template $\mathbf{t}_s \in \mathbb{R}_+^d$ corresponding to source s (e.g., ^{133}Ba) with the matrix \mathbf{V} to form a combined background and source model $\mathbf{V}' = \begin{bmatrix} \mathbf{V} \\ \mathbf{t}_s \end{bmatrix}$. By fitting a spectrum to \mathbf{V}' , the spectrum is approximated as a linear combination of the background components and the source. The source template, representing the detector system's response to a particular source, can be generated from measurements, simulation, or by performing NMF on data containing sources.

When examining new spectra \mathbf{X}' , the model \mathbf{V}' is applied and held fixed, resulting in weights \mathbf{A}' . Figure 4.3 shows the result of performing an NMF fit to a spectrum using both the background-only model \mathbf{V} and the combined source and background model \mathbf{V}' . The measurement is a one-second NaI spectrum of background data from RadMAP that

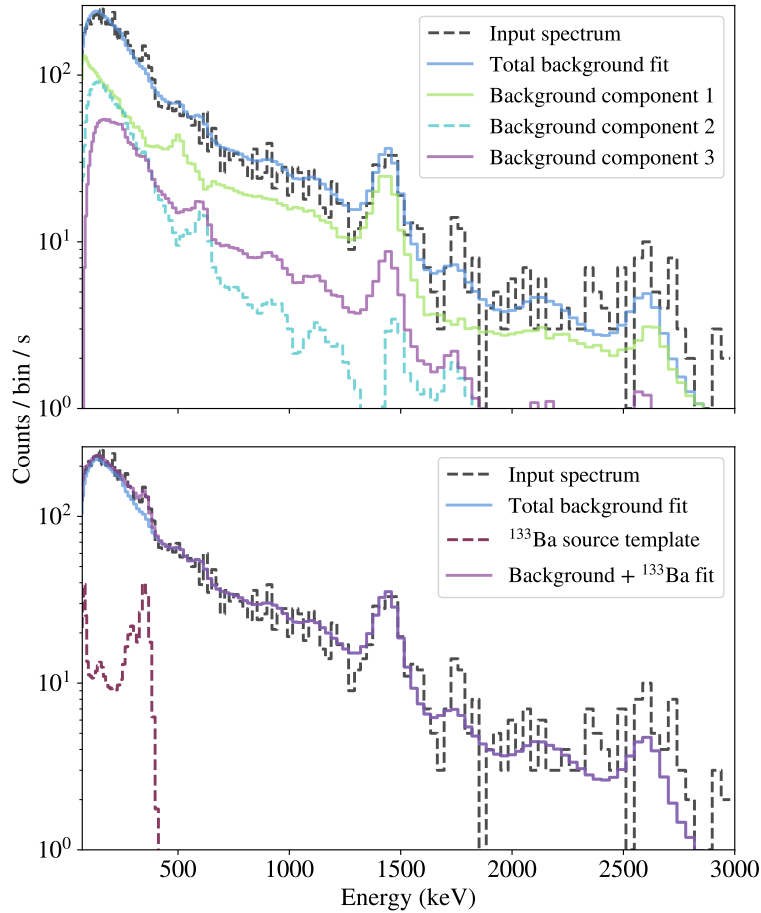


Figure 4.3: Comparison of a fit to a spectrum containing ^{133}Ba using only the background model V (top) and using the background and source model V'. The upper pane shows the weighted background components, as well as their sum. The lower pane shows the sum of the background components, the weighted source template, and the sum of the background and source components. By including the source template, the model is able to fit the 356 keV ^{133}Ba peak. An 85- μCi ^{133}Ba source at 20-m standoff distance is used.

includes source contribution from an injected 85- μCi ^{133}Ba source at 20-m standoff distance. Together, the background components generated from NMF and the ^{133}Ba source template perform a better approximation of the spectrum than the background model alone.

4.5.1 Likelihood Ratio Tests

To detect and identify sources using the augmented model V', a test statistic for determining the presence or absence of a source within a spectrum is examined. In particular,

testing the null hypothesis that a spectrum is consistent with background, relative to the alternate model containing source s , is done by a likelihood ratio test. A given spectrum \mathbf{x} is fit to both the background model \mathbf{V} and augmented model \mathbf{V}' , yielding the negative log-likelihoods $-\ln P(\mathbf{x}|\mathbf{V})$ and $-\ln P(\mathbf{x}|\mathbf{V}')$, respectively, which are then used to form a log-likelihood ratio test statistic

$$D_s(\mathbf{x}) = -2\{\ln P(\mathbf{x}|\mathbf{V}) - \ln P(\mathbf{x}|\mathbf{V}')\}. \quad (4.14)$$

The test statistic $D_s(\mathbf{x})$ quantifies the amount by which the fit is improved by including the additional source template \mathbf{t}_s ; if the fit is not significantly enhanced, $D_s(\mathbf{x}) \approx 0$, and $D_s(\mathbf{x})$ increases as including the template improves the fit. Using the procedure described in Section 2.6, a threshold T_s for source s at a given FAR is estimated by computing the test statistic over background data. Figure 4.4 shows histograms of $D_s(\mathbf{x})$ for ^{137}Cs and ^{133}Ba from 55 hours of NaI spectra collected by RadMAP. The distribution peaks at $D_s(\mathbf{x}) = 0$ for both sources, and decreases with increasing $D_s(\mathbf{x})$. The procedure of training models using NMF, fitting spectra using background models and source templates, and performing likelihood ratio tests to perform identification is hereafter referred to as ADI_P and ADI_{L2} .

4.5.2 Analytical Modeling of Likelihood Ratio Test Statistics

Using such an empirical approach, one needs more than $N = p_{FA}^{-1} = R_{FA} * \Delta t$ data points per template to estimate a threshold for a single source, where R_{FA} is the FAR (e.g., $1/8 \text{ hr}^{-1}$) and Δt is the integration time (e.g., 1 s). In some cases, there may be an insufficient amount of data available to estimate thresholds at a particular FAR, meaning such an approach is not possible. While the empirical approach for threshold estimation provides one method to account for actual variability seen by a detector system, analytical estimates of test statistic thresholds are possible, and eliminate the need for a large amount of data for threshold estimation. Specifically, Wilks' theorem [62] states that likelihood ratio test statistics, when the null hypothesis is true, are asymptotically (i.e., as $d \rightarrow \infty$) distributed as χ_m^2 , where m indicates the difference in the number of degrees of freedom between the alternate and null models. In the context of this work, m is the number of parameters used in modeling the source contributions, meaning $m = 1$ for the case of a single source template. Wilks' theorem assumes certain conditions about the parameters involved in the MLE being performed, namely, that parameters are Gaussian-distributed and that estimated parameters are at the interior of the parameter space, neither of which are met in the NMF-based hypothesis testing described above. When parameters are estimated on the boundary on the parameter space (i.e., close to zero), as with source weights in NMF, the distribution of likelihood ratio test statistics, when the null hypothesis is true, is modeled as a mixture of χ^2 distributions with different degrees of freedom [63]. For m source components added to an NMF-derived background model, the distribution of test

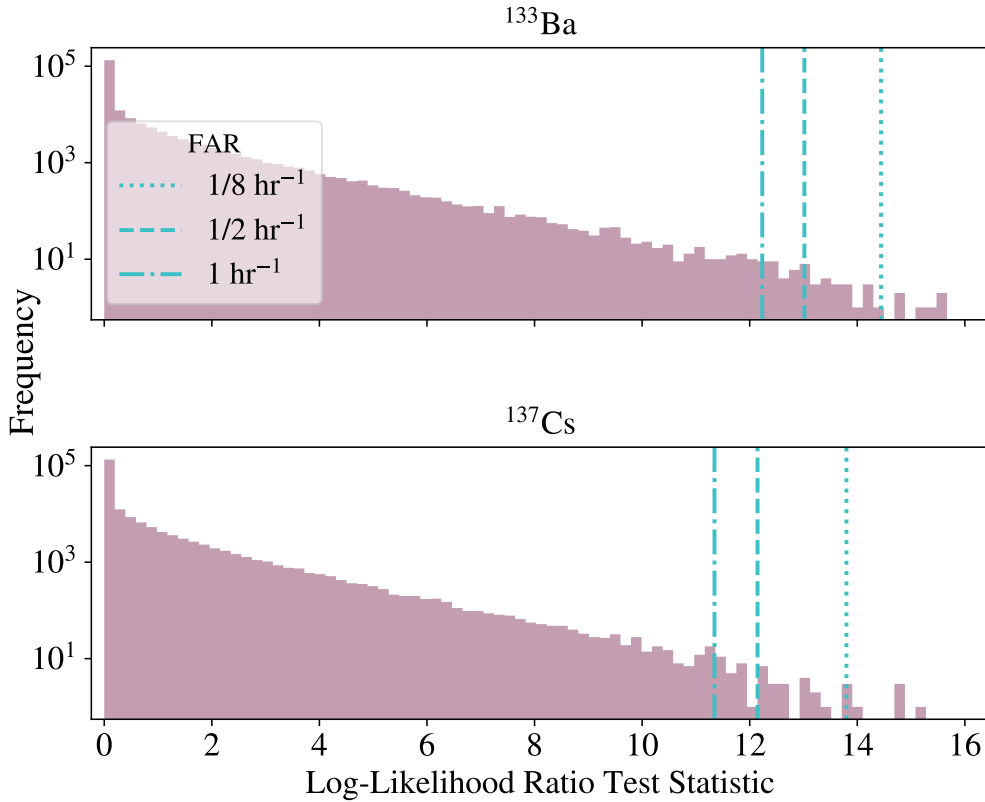


Figure 4.4: Distributions of likelihood ratio test statistic $D_s(\mathbf{x})$ used in testing for the presence of ^{133}Ba (top) and ^{137}Cs (bottom). The distributions are formed by computing the difference of negative log-likelihoods between source and background models \mathbf{V}_s and background only models \mathbf{V} over a background dataset containing 55 hours of NaI spectra collected by RadMAP. The distribution is used to empirically select a threshold based on a target FAR. The vertical lines show thresholds for particular FARs, given in the legend.

statistics can be modeled as

$$P(D; m) = \frac{1}{2^m} \sum_{i=0}^m \binom{m}{i} \chi_i^2. \quad (4.15)$$

Figure 4.5 shows a comparison between likelihood ratio test statistics computed using background data for both ^{60}Co and ^{137}Cs , as well as the distribution described by Equation (4.15) for $m = 1$ and $m = 2$. The maximum at $D = 0$ is accounted for by the term $2^{-m} \chi_0^2$, which results from the additional parameters being estimated (i.e., the NMF weights associated with the source components) being at the edge of the parameter space (i.e., often times close to being 0 when the null model is true). Additionally, the increased probability of larger test statistics for $m = 2$ can be understood as enhanced fitting due

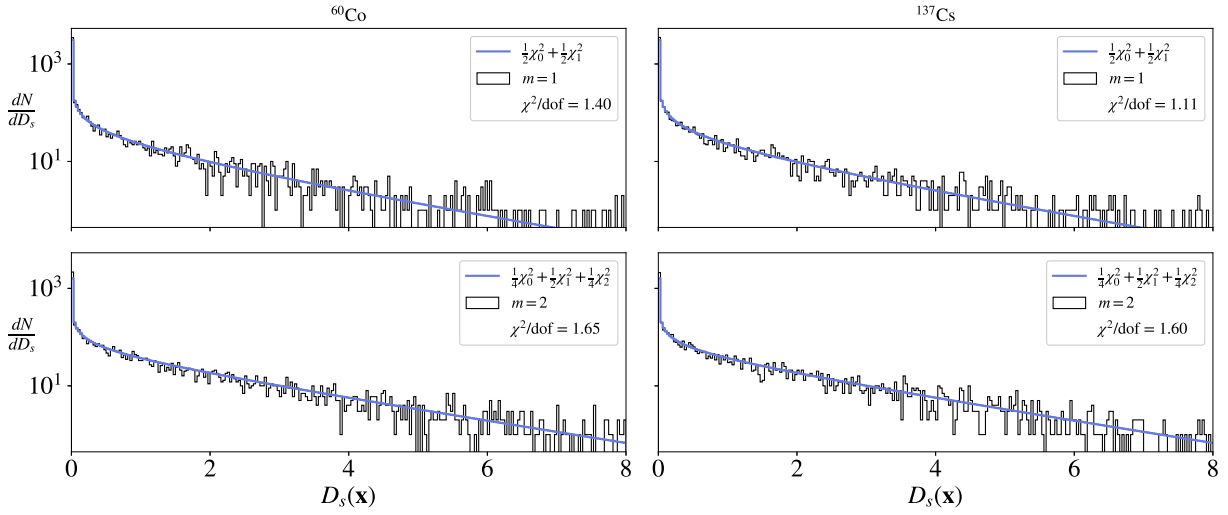


Figure 4.5: Histograms of likelihood ratio test statistics for ^{60}Co (left column) and ^{137}Cs (right column) computed using background data compared to the analytical model in Equation (4.15) for $m = 1$ (top row) and $m = 2$ (bottom row). The reduced χ^2 goodness-of-fit measure is reported for each fit, showing that the models fit reasonably well, suggesting Equation (4.15) can be used for modeling the probability distribution of test statistics for hypothesis testing. In each case, the inverse CDF of the probability distribution can be computed and used to yield the threshold for a given probability of false alarm.

to the additional degree of freedom. Note that the $m = 2$ case, in which the source is represented with two parameters, is described in Section 4.6.

Using the distribution given by Equation (4.15), thresholds associated with a particular p_{FA} or CFAR can be computed. To maintain a given CFAR when simultaneously performing multiple hypothesis tests, one must apply a correction factor to maintain a constant FAR. In this work, a Bonferroni correction, which scales the significance level inversely by the number of hypotheses tested, is used for simplicity in analyzing the behavior of thresholds. Note, however, that the Bonferroni correction results in relatively conservative threshold estimates, and that less stringent alternatives are available (e.g., false discovery rate methods). The threshold T associated with an m -parameter model and N different sources is computed from the cumulative density function of Equation (4.15):

$$\frac{p_{FA}}{N} = \int_T^\infty P(D; m) dD. \quad (4.16)$$

Equation (4.16) is solved numerically to yield T . Values of T for various m and N are shown in Figure 4.6.

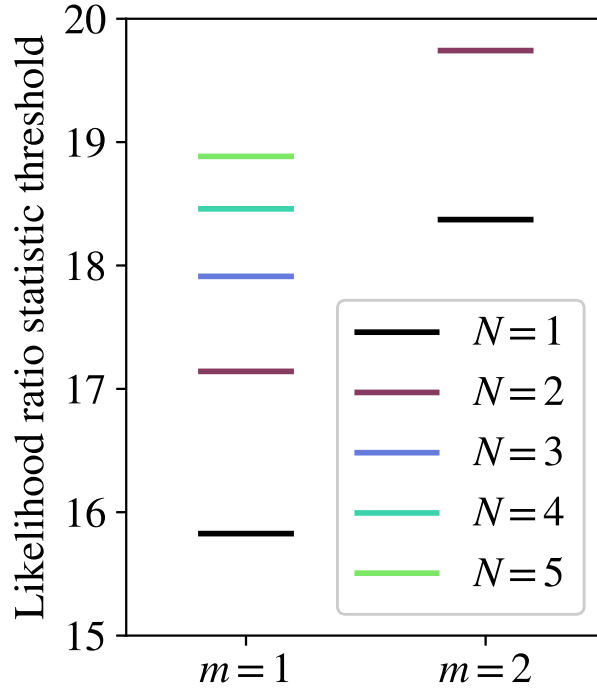


Figure 4.6: Thresholds used in likelihood ratio tests, computed using Equation (4.16). The left column shows thresholds for $m = 1$, where sources are approximated with a single vector, and the right column shows thresholds for $m = 2$, used for the two-component NMF source representation. N is used to represent the effective number of independent hypotheses being tested. In the analysis in this work, $N_s = 1$, meaning a single source is tested at a time, so N represents the number of shielding configurations considered. This plot shows that for $N_c \geq 5$, a two-component NMF source model results in a lower threshold, assuming both models are able to describe variation in the source sufficiently well.

4.5.3 Benchmark: Region of Interest

To assess the relative performance of the ADI_P and ADI_{L2} algorithms, a Region of Interest (ROI) algorithm was evaluated on the same data. The ROI algorithm, as described in [34], estimates the number of source counts within a spectral region in which gamma rays from a particular source are expected. To estimate source counts in the ROI \mathcal{R} , the algorithm first uses a linear relationship between the number of background counts in the ROI and the number of counts in neighboring regions \mathcal{B}_1 and \mathcal{B}_2 . Regions \mathcal{B}_1 and \mathcal{B}_2 are chosen to be at higher energies than \mathcal{R} so that they do not contain counts from downscattered source gamma rays, while being sufficiently close to \mathcal{R} to predict background counts in \mathcal{R} . The window edges defining the ROI and the background regions are shown in Table 4.4.

Table 4.4: Energy windows (keV) used in ROI-based source identification

Source	\mathcal{R}	\mathcal{B}_1	\mathcal{B}_2
^{133}Ba	340.0 - 397.2	397.2 - 443.0	443.0 - 474.9
^{137}Cs	632.3 - 670.4	670.4 - 749.8	749.8 - 812.2

The number of background counts in \mathcal{R} is estimated as $\hat{r} = \mathbf{b}^\top \mathbf{w}$, where \mathbf{b} is a vector containing counts within the background windows \mathcal{B}_1 and \mathcal{B}_2 , and \mathbf{w} are weights found via Poisson regression [38] on the same background training set used for training NMF models. When evaluating spectra for the presence of the source, the difference between the measured counts r and estimated background counts \hat{r} , or the residual, forms the ROI metric, defined as

$$\text{ROI}(\mathbf{x}) = \frac{r - \hat{r}}{\sqrt{\hat{r}}}. \quad (4.17)$$

This metric is then compared to a decision threshold to determine the presence or absence of a source. As with the other methods described in this analysis, the decision thresholds are computed empirically for a given FAR by calculating the residual over background data. Spectra in which the ROI has an excess of counts are then considered anomalous.

4.5.4 ADI_l in the Limit of Known Background

As with anomaly detection described in Section 4.4, the performance of the identification algorithms is also evaluated in the case of a known background. For NMF-based identification, the background contribution $\hat{\mathbf{X}}_{\text{bkg}}$ is held fixed, and the source contribution that best fits the source and background model to the spectra \mathbf{X} is found.

4.5.5 Performance Evaluation

Using a FAR of $1/8 \text{ hr}^{-1}$, thresholds for both sources were empirically determined, yielding the values $T_{\text{Cs}} = 14.1$ and $T_{\text{Ba}} = 14.4$. Additionally, an L2NMF model was formed and thresholds were calculated using the same training set as the PNMf model described above, though thresholds and figures specific to L2NMF are omitted for the sake of brevity. Note that in each case, only a single source is being searched for. In the case of multiple sources being separately tested for, false alarm rates should be adjusted accordingly (e.g., using the Bonferroni correction), ultimately increasing thresholds and reducing sensitivity.

Figure 4.7 shows the probability of detection for ^{133}Ba and ^{137}Cs using both NMF-based algorithms and an ROI algorithm, and Table 4.5 shows the corresponding MDA_{95} for each algorithm resulting from a fit of P_D to a sigmoid. Figure 4.7 shows that there is a significant improvement in detection capabilities by using NMF-based algorithms, and that there are additional improvements by using PNMf models over L2NMF.




Figure 4.7: Probability of detection and identification for ^{133}Ba (top) and ^{137}Cs (bottom) using the ADI_P , ADI_{L2} , NMF-based algorithms in the case of a known background, each using, $k = 2$ components, and an ROI algorithm. Solid lines show the measured performance, and dashed lines show performance in the KBL. Spectra were formed using 1-s integration time for 99 NaI detectors on RadMAP, and source injection was performed using 20-m standoff distance and 6.7 ± 1.3 m/s vehicle speed. For both sources, a $1/8 \text{ hr}^{-1}$ FAR was used. Both the ADI_P and ADI_{L2} methods are able to achieve a 95% probability of detection, indicated by the dashed horizontal line, at a much lower activity than the ROI algorithm. Error bars indicate the 68% Jeffreys interval.

Similar to results shown in Section 4.4.3, there is a difference between measured performance of the NMF-based detection and identification methods and the performance in the case of a known background. By informing these methods with data about the local environment, more accurate background models could potentially be created, enhancing detection and identification performance.

Table 4.5: MDA_{95} (μCi + statistical error) for simultaneous detection and identification

Source Method	^{133}Ba	^{137}Cs
ADI_P KBL, $k = 2$	38.9 ± 0.1	46.3 ± 0.2
ADI_{L2} KBL, $k = 2$	40.8 ± 0.1	29.1 ± 0.1
ADI_P , $k = 2$	79.6 ± 0.3	82.4 ± 0.4
ADI_{L2} , $k = 2$	92.7 ± 0.3	88.4 ± 0.3
ROI	168.5 ± 0.6	145.2 ± 0.8

4.6 NMF for Source Spectrum Modeling

In addition to background variation, there is variability in spectral shape due to shielding and environmental scattering of source gamma rays. Interactions of source gamma rays with materials between the source and detector (e.g., steel shielding surrounding a source) reduces the number of gamma rays incident with the detector for a given source-detector configuration, but will also result in a different spectral shape than in the case of an unshielded source. Figure 4.8 illustrates the variation in spectral shape for the sources ^{60}Co and ^{137}Cs under four different types of shielding which are commonly used in transporting industrial sources, using six $2'' \times 4'' \times 16''$ NaI(Tl) scintillation detectors. Gamma-ray source identification algorithms that rely on spectral shape, especially those using template matching, often suffer in detection sensitivity due to this modulation of spectral shape. In this section, a method for accurate modeling of variability in spectral shape under various source shielding configurations is discussed, and identification performance using these models is quantified.

Approaches for mathematically modeling the effect of source attenuation on gamma-ray spectra have been previously introduced. For example, the multiple isotope material basis set (MIMBS) method [64] approximates a given spectrum as a linear combination of spectra resulting from attenuated sources. MIMBS folds attenuation in by analytically computing the spectrum for a source under a particular activity and shielding density, and in fitting a spectrum, source activities and the absorber densities for each material are simultaneously solved for. Instead of an analytical approach such as the MIMBS method, this chapter uses a data-driven approach for modeling the effects of source attenuation on gamma-ray spectra. One reason for considering a data-driven approach over analytical calculations is the possibility for reducing the number of parameters involved with fitting a given spectrum.

In this section, the concept of source spectrum modeling using NMF, first introduced in ref. [65], is studied further and discussed in the context of the identification method introduced in Section 4.5. Instead of assuming a library of shielding types, as with meth-




Figure 4.8: Normalized detector responses, referred to as spectral templates, for both ^{60}Co (left) and ^{137}Cs (right) under five source shielding configurations: no shielding, 50 mm of concrete, 50 mm of lead, 50 mm of steel, and 100 mm of steel. The templates were formed by creating spectra from simulated 1 mCi sources at a 10-m standoff distance, then normalizing the spectra to 1. For both sources, variability in spectral shape is seen to depend on the amount and type of shielding surrounding the source. An ideal gamma-ray source identification algorithm is able to accurately identify a gamma-ray source in the presence of various shielding types.

ods such as MIMBS, the NMF approach learns common modes of variation from examples of shielding types. To demonstrate the concepts in this section, measured background gamma-ray spectra and simulated source spectra from a mobile detection system equipped with six $2'' \times 4'' \times 16''$ NaI(Tl) scintillation detectors are used.

The general procedure for learning a representation for a source s using NMF is to form a matrix \mathbf{X}_s from spectral samples of s in different shielding configurations, as well as different source positions relative to the detector, and to perform NMF on \mathbf{X}_s . The components \mathbf{V}_s generated from NMF then can be understood as d -dimensional vectors that combine linearly to approximate source spectra. Using \mathbf{V}_s and the NMF-derived background model \mathbf{V}_b , new spectra are encoded with the matrix $\mathbf{V}' = \begin{bmatrix} \mathbf{V}_b \\ \mathbf{V}_s \end{bmatrix}$, yielding weights \mathbf{A} that can then be used for identification.

For each radionuclide and shielding type considered here, the source and shielding combination was simulated at 19 different positions in a 180° arc around the detector system in 10° increments, each at a 10-m standoff distance and 1.3 m off the ground, and each being measured with a 1 s integration time [17]. Due to symmetry in the detector array, the detector response was mirrored about the axis of symmetry, resulting in 36 source positions covering 360° around the detector system. Spectra from all six detectors were summed, resulting in a single spectrum for each of the 36 sample points around the detector system. Spectra for a given source under all shielding configurations were stacked to form matrix \mathbf{X}_s .




Figure 4.9: NMF components V_s generated for ^{60}Co (left pane) and ^{137}Cs (right pane). For both sources, NMF generates physically-interpretable components that can be related to the full-energy peaks and the downscattering continuum associated with shielding. Intuitively, a source under a given shielding configuration can be thought of as a linear combination of the full-energy peak(s) of the source, and a downscattering contributions due to the particular shielding configuration present.

NMF was performed on X_s , using a random initialization for the matrices A and V , with the update rules in refs. Equation (4.4) - Equation (4.7) until the mean objective function Equation (2.15) remained constant within a tolerance of 10^{-7} . As discussed in Section 4.2, the AIC is again computed to choose an optimal number of components, yielding $k = 2$ for both sources used in this section.

Figure 4.9 shows a comparison of the two source components for ^{60}Co and ^{137}Cs generated using the shielding types shown in Figure 4.8. As with for background models, NMF yields interpretable modes of variation for both source types used in Figure 4.9. For ^{60}Co , Component 1 primarily captures the 1173 keV and 1332 keV peaks and the high-energy edge of the Compton continuum, while a large portion of the low-energy continuum, as well as a downscattering feature associated with the 1332 keV peak seen between the two peaks, is captured by Component 2. Similarly, for ^{137}Cs , Component 1 contains the 662 keV full-energy peak, and downscattering is captured by Component 2. These two independent components can then be weighted in different amounts to closely match the shielding configurations shown in Figure 4.8.

4.6.1 Model Comparisons

Three methods are compared using the injection scheme described above. The first approach, introduced in Section 4.5, uses a single bare source template to represent source contributions to spectra, and is hereafter referred to as the *bare* approach. The second method, referred to as the *library* approach, makes use of a library of shielded sources

and fits a given spectrum to each shielding type in the library. Lastly, the *learned* NMF component approach is used for representing sources. Each of the three approaches results in a different number of fitting parameters, and as a result, each has a different threshold according to Equation (4.16).

In the discussion below, the number of sources being simultaneously searched for is denoted by N_s , and the number of shielding configurations used in the library approach is N_c . When each source and shielding configuration is treated as an independent hypothesis test, $N = N_s N_c$. In each of these three cases, the threshold then depends on three parameters: the number of sources being tested for N_s , the number of shielding configurations N_c , and the number of parameters used in representing the source m . In this analysis, only a single source is being searched for at a time, meaning $N_s = 1$, and that N in Equation (4.16) only depends on N_c . Using the two remaining parameters, N_c and m , thresholds can be computed analytically using Equation (4.16) and compared, shown in Figure 4.6. Assuming a source can be accurately described by a given model for particular values of N_c and m , the ordering of thresholds gives an ordering of MDA for different values of N_c and m . Each of the three models are discussed further here.

Bare Template

In this case, only a single source shielding configuration (i.e., bare) is considered, meaning $N = N_c = 1$. Also, since only a single parameter is used in modeling the source contribution to spectra, $m = 1$. As seen in Figure 4.6, this combination of N and m results in the lowest possible threshold, meaning that when a bare source is present, this method will perform better than the other two methods.

Library of Templates

Here, two or more templates form a library of templates, corresponding to different shielding configurations. For example, Figure 4.8 shows $N_c = 5$ shielding configurations for both ^{60}Co and ^{137}Cs . With respect to the number of parameters used in a specific likelihood ratio test, $m = 1$ parameters are used here as well. The difference between the library of templates approach and bare template approach is then the number of hypotheses being tested for $N > 1$.

NMF-Derived Templates

In this case, the number of parameters associated with the likelihood ratio test is $m = k$, where k is the number of NMF components used in representing the source, and that $N = 1$. As previously noted, $k = 2$ components are used, meaning $m = 2$ in Equation (4.15). Figure 4.6 shows that if sufficient variability can be captured with the learned model, then the learned model is preferred over a five-template library with equivalent modeling capability, as the learned model results in a lower threshold.

Table 4.6: Comparison of MDA (μCi) for NMF-based identification using different source approaches

Method \ Source	Bare template	Library of templates	Learned components
^{60}Co , bare	24.5 ± 0.3	26.0 ± 0.3	25.7 ± 0.3
^{60}Co , 50 mm steel	154.4 ± 1.5	146.2 ± 1.9	145.0 ± 1.9
^{137}Cs , bare	49.2 ± 0.6	51.6 ± 0.7	50.8 ± 0.7
^{137}Cs , 50 mm steel	652.4 ± 9.4	615.4 ± 6.4	613.4 ± 6.4

For each of the three classes of source spectrum models, the same background model V_b was used. The matrix V_b was generated by performing NMF on a collection of 6019 1 s background spectra, using the same procedure described in Section 4.6. Due to low variability in spectral shape across samples, a single (i.e., $k = 1$) NMF component resulted in the lowest AIC for the background model, and is used in the remainder of the analysis.

4.6.2 Performance Evaluation: Bare Source

While the methods discussed in this work are ultimately meant to identify shielded sources, they must also be able to identify bare sources, which is examined here. Figure 4.10 shows the probability of detection curves for the three models in the case of a bare source, and the MDA for each is given in Table 4.6. Each model is able to accurately model the bare source, however, the bare source model resulted in the lowest MDA due to the threshold penalty for additional parameters used in the other two models, as expected from the threshold analysis in Section 4.6.1.

4.6.3 Performance Evaluation: 50 mm Steel Shielding

The probability of detection was again computed for the two sources using the three models, but here the source was shielded by 50 mm of steel, shown in Figure 4.11. The MDA for both sources under each of the three methods is given in Table 4.6. Both the learned and library approaches provide improvements over the bare model, however, the learned and library approaches fall within error of one another.

4.7 Concluding Remarks

Detecting and identifying gamma-ray sources relies on the ability of an algorithm to distinguish source and background. Here, the use of NMF was proposed, under the assumption of Poisson statistics, to model both gamma-ray background and source spectra.

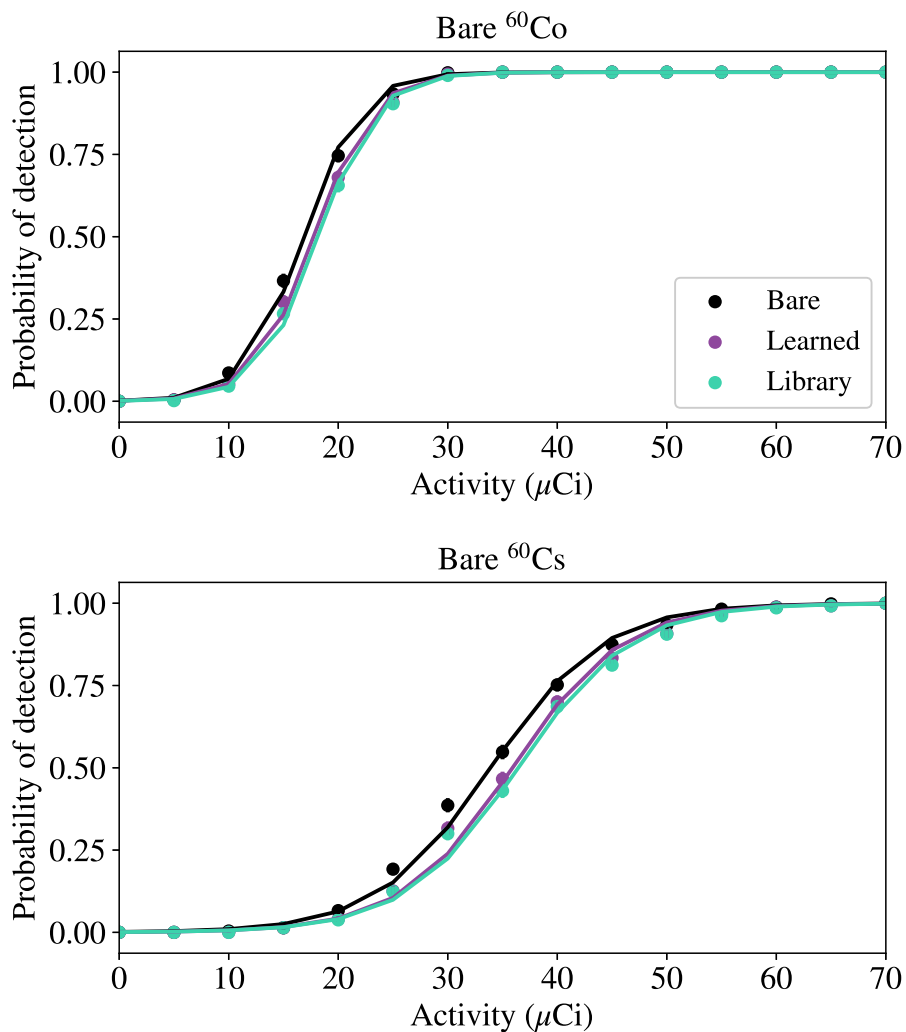


Figure 4.10: Probability of detection curves for unshielded ^{60}Co (top) and ^{137}Cs (bottom). For each activity, 500 background spectra were randomly sampled, and downsampled source spectra for a source at a 10-m standoff were injected into the background. While each of the three models results in similar performance, there is an advantage to using the bare source model, as it is able to accurately model the source using fewer parameters than the other two approaches.

These models can be used to find radiological anomalies, and by using spectral templates, source identification can be performed. In addition to performing identification for single sources, multiple sources can be simultaneously tested for by appending additional source templates to the matrix V_s . Under the given parameters, the anomaly detec-

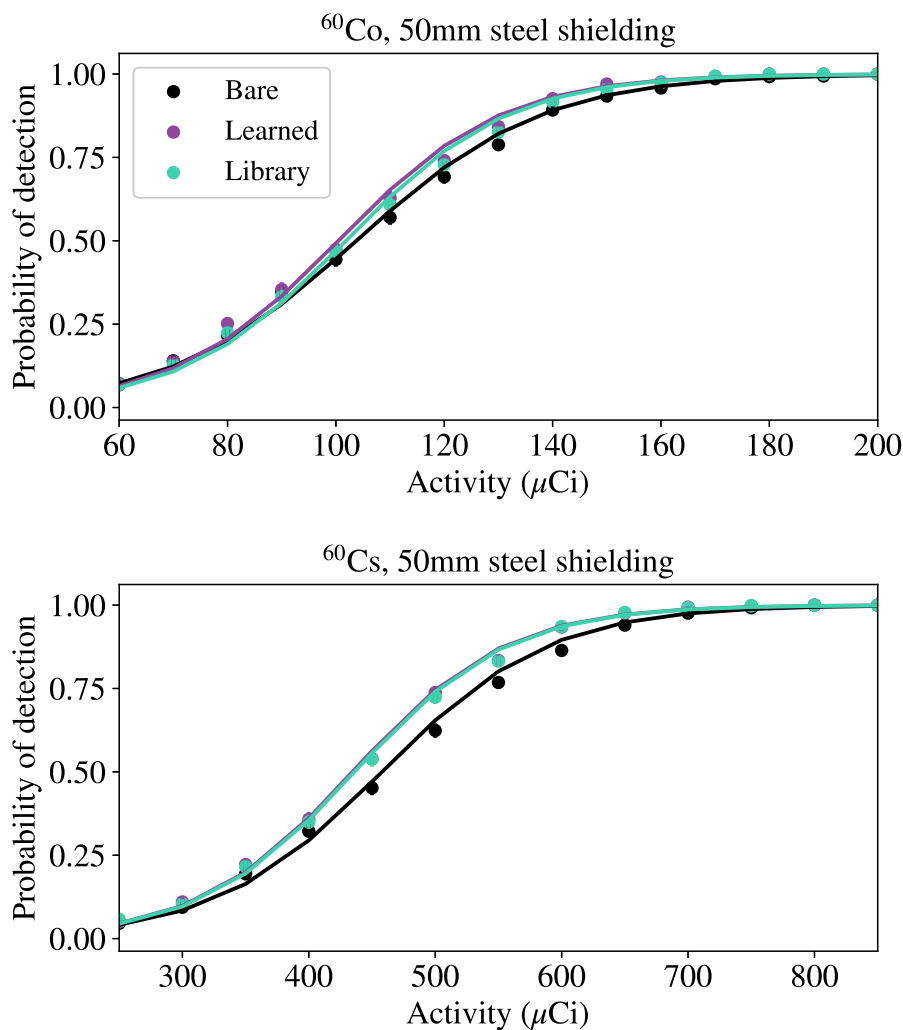


Figure 4.11: Probability of detection curves for ^{60}Co (top) and ^{137}Cs (bottom) shielded by 50 mm of steel. Both the learned and library models provide improvements over the bare template model due to their ability to model variation in the spectral shape due to the shielding.

tion and identification algorithms outperform standard algorithms. Note, however, that the decomposition-based methods here were only optimized with respect to the number of degrees of freedom, and that the performance of each method, including the benchmark algorithms, may be enhanced by optimizing integration time, performing spectral smoothing, normalizing input variables, and optimizing energy windows in the case of ROI.

This method has also been extended to model source spectra, which is particularly useful for modeling the effects of shielding on spectral shape. The NMF models resulting from training on source spectra under different shielding configurations generated components that can be understood as separately containing full-energy peaks and down-scattering continua. Section 4.6 showed that a marginal improvement can be achieved when using NMF-derived source components over using bare templates when detecting shielded sources. Furthermore, the interpretable components resulting from these models may aid in determining shielding type.

While the methods in this chapter were demonstrated using relatively large and unique detector systems, these methods are general enough to be deployed on a wide range of systems. Furthermore, these methods can be extended to detect and identify gamma-ray sources other than the two studied in this analysis, albeit with additional care. One consideration is the similarity in spectral shape between the source of interest and the NMF components used in estimating background – sources with a high degree of similarity to background components may be incorrectly attributed to background, ultimately reducing the ability to detect the particular source. If a source is seen to exhibit such behavior, one may devise a means of reducing the similarity, for example leveraging the non-uniqueness and selecting background components that are the most distinct from the sources of interest (e.g., with respect to cosine similarity or KL divergence). Similarly, when searching for multiple sources, one must be aware of the similarity between sources, as source templates that contain overlapping gamma-ray lines may result in misidentification.

Beyond anomaly detection and identification, Poisson NMF can be considered a general framework for approaching gamma-ray spectroscopic analyses, as it lends itself to a useful physical interpretation due to its additive, non-negative nature. For example, the geospatial distribution of NMF weights may yield the environmental composition of particular radioisotopes, which could potentially find use in applications such as contamination mapping.

In Figure 4.2 and Figure 4.7, it is shown that there is a significant difference between the measured performance and the performance in the limit of known background. Additional information about the environment, for example, in the form of Bayesian priors and regularization, may potentially enhance detection and identification performance. Prior probabilities of background weights \mathbf{A} , determined from previous measurements, could be introduced to provide constraints to background models by means of maximum a posteriori, as opposed to the MLE used here, which may allow for a more accurate background estimate. Additionally, performance can be enhanced further by aggregating test statistics over several spectra, or by sequentially estimating weights (e.g., using a Kalman filter) to create time-dependent models.

Chapter 5

Neural Networks for Detection and Identification

The NMF spectral models introduced in the previous chapter are often physically intuitive, provide accurate approximations, and are relatively simple to implement and deploy on fielded systems. The simplicity of these models, however, can limit their extensibility, specifically in performing data fusion with additional non-radiological sensors. In attempting to leverage correlations between spectral features and image features, previous attempts using linear models have relied on ad-hoc approaches [67]. It is then worth considering additional approaches to perform detection and identification, specifically methods which are capable of being augmented with additional data sources. This chapter explores approaches to spectral modeling using neural networks, a general class of non-linear models which can be used for spectral detection and identification. In addition to being compelling candidate methods for spectral analysis, there is a clear path for exploring the use of data fusion with non-radiological sensor data.

The remainder of this chapter is outlined as follows. Section 5.1 gives an overview of neural network elements, specifically as they pertain to spectral detection and identification. Section 5.2 provides a review of related research, as well as directions of research following this past work. Spectral anomaly detection using neural networks is introduced in Section 5.3, and Section 5.4 discusses networks for source identification. A description of the data and metrics used in analyzing the performance of the methods described in this chapter is provided in Section 5.5. Section 5.6 quantifies the performance of various models, including benchmark methods. Lastly, Section 5.7 concludes with a discussion on considerations when choosing between methods to use in practice, along with additional directions of research to consider.

The content of this chapter originates from the following publication: K. J. Bilton, T. H. Joshi, M. S. Bandstra, J. C. Curtis, D. Hellfeld, and K. Vetter, "Neural Network Approaches for Mobile Spectroscopic Gamma-ray Source Detection." Submitted to *IEEE Transactions on Nuclear Science* (under review), 2020 [66].

5.1 Overview of Artificial Neural Networks

Artificial neural networks (ANNs), or simply neural networks, are a general data-driven method for performing function approximation, and are capable of producing functions for performing detection and identification. Neural networks are used to form a function f which operates on an input \mathbf{x} (e.g., an image or gamma-ray spectrum), producing an output $\hat{\mathbf{y}} = f(\mathbf{x})$. The function f consists of a series of relatively simple operations, which are parameterized by a set of learned model parameters \mathcal{P} . To perform identification, ANNs are used to determine a function which maps a given gamma-ray spectrum to the types of radionuclides, or lack thereof, that are observed in the spectrum. This section reviews the fundamental concepts underlying neural networks, specifically as they relate to detection and identification.

Neural networks can generally be considered a composition of l different functions, that is,

$$f(\mathbf{x}) = f_l(f_{l-1}(\dots(f_2(f_1(\mathbf{x}))))), \quad (5.1)$$

where the output of the composed function up to f_i is referred to as the i^{th} layer, denoted by $\mathbf{h}^{(i)}$. Using these layers, complete networks can be created to perform a variety of tasks. Often times, networks will have an architecture that is several layers "deep", giving rise to the name *deep learning*.

For a defined functional form of f , the parameters \mathcal{P} are estimated from examples of pairs of inputs \mathbf{x} and corresponding target values of output \mathbf{y} , such that $\hat{\mathbf{y}} = f(\mathbf{x}) \approx \mathbf{y}$. The functional form of f is one of the key elements to using neural networks, as this determines the *model capacity*. With too little model capacity, the learning procedure will fail to yield f that fits the training data well, whereas too much model capacity will cause the model to *overfit* training data, meaning it learns parameters that effectively allow the network to remember training examples.

Equation (5.1) specifically defines a *feedforward* neural network, which can be represented as a directed acyclic graph. This is in contrast to a *recurrent* neural network (RNN), which contains cycles and is used for modeling sequential data. In addition to examining feedforward networks for detection and identification, this chapter introduces the use of RNNs for performing spectral identification using temporal sequences of spectra. The remainder of this section provides common layer types which are useful for accomplishing both detection and identification.

5.1.1 Fully-connected Layers

A *fully-connected* or *dense* layer connects each element, or *neuron*, in a given layer to each element in the following layer. Suppose the i^{th} layer of a network consists of k elements, denoted by $\mathbf{h}^{(i)} \in \mathbb{R}^k$, and suppose the following layer has m elements, denoted by $\mathbf{h}^{(i+1)} \in \mathbb{R}^m$. A fully-connected operation provides a mapping $\mathbb{R}^k \rightarrow \mathbb{R}^m$ using a matrix $\mathbf{W}_i \in \mathbb{R}^{m \times k}$. A bias term \mathbf{b}_i can be added, resulting in the linear transformation




Figure 5.1: A dense connection between layer $\mathbf{h}_i \in \mathbb{R}^4$ and layer $\mathbf{h}_{i+1} \in \mathbb{R}^2$, which can be understood from left to right. The lines between element $h_{i,j}$ in the left (i.e., first) layer and element $h_{i+1,k}$ in the right (i.e., second) layer are represented as a column i in the matrix $\mathbf{W} \in \mathbb{R}^{2 \times 4}$. Not shown is an additional bias term $\mathbf{b}_i \in \mathbb{R}^2$ and nonlinear activation function $\sigma(\cdot)$ used to create the state $\mathbf{h}_{i+1} = \sigma(\mathbf{W}_i \mathbf{h}_i + \mathbf{b}_i)$.

$\mathbf{W}_i \mathbf{h}^{(i)} + \mathbf{b}_i$. Since the composition of linear operations will result in a linear model, nonlinearities are added in the form of an *activation function* $\sigma(\cdot)$ to increase modeling capacity. One common activation function, used in much of this work, is the Rectified Linear Unit (ReLU) function, which is defined as

$$\text{ReLU}(z) = \max(0, z). \quad (5.2)$$

In summary, a fully-connected operation from a layer i with state $\mathbf{h}^{(i)}$ to layer $i+1$ is given by

$$\mathbf{h}^{(i+1)} = \sigma(\mathbf{W}\mathbf{h}^{(i)} + \mathbf{b}), \quad (5.3)$$

where the matrix \mathbf{W}_i and vector \mathbf{b}_i are learned from data. Figure 5.1 shows an example dense connection between two layers, illustrating connections between the neurons in each layer.

5.1.2 Convolutional Layers

For data in which salient features are localized (e.g., images), a good approximation of f can often be made using far fewer connections between layers than in dense connections.

This behavior can be accomplished by means of convolutional layers, forming a Convolutional Neural Network (CNN). Convolutional layers relate local features in data by means of applying the convolution operation to input data for a given convolutional kernel (i.e., a function over a limited domain). CNNs have led to significant advances in the field of computer vision and are at the core of many state-of-the-art approaches to tasks such as classification, detection, and semantic segmentation [68]. As originally noted in ref. [69], gamma-ray spectra also contain local features (e.g., peaks and continua), as do images, and CNNs are an appropriate choice for this application.

In convolutional layers, one or more convolutional kernels, or filters, are applied to a given layer, yielding a new set of features. The form of the convolution operation varies slightly with the number of dimensions used, and in this work, only 1-dimensional convolutions are considered, treating a gamma-ray spectrum as features along a single axis. Suppose that the i^{th} layer contains n feature maps (i.e., n sets of elements resulting from n different convolutional kernels from the previous layer), each of length d , represented with a 2-dimensional tensor $\mathbf{h}^{(i)} \in \mathbb{R}^{n \times d}$. Also suppose that the convolutional operation to be applied to $\mathbf{h}^{(i)}$ consists of k different convolutional kernels, each with size l (typically a small number), represented by a 3-dimensional tensor $\mathbf{K} \in \mathbb{R}^{k \times n \times l}$. The m^{th} element of the j^{th} feature map of layer $\mathbf{h}^{(i+1)}$ resulting from convolving $\mathbf{h}^{(i)}$ with kernels \mathbf{K} is given by

$$\mathbf{h}_{j,m}^{(i+1)} = \sum_{x=0}^n \sum_{y=0}^l \mathbf{h}_{x,m+y}^{(i)} \mathbf{K}_{j,x,y}. \quad (5.4)$$

In addition to the number of kernels used and the size of the kernels, there are additional hyperparameters used in configuring the convolution operation, including padding, stride, and dilation, which affect the resulting feature map. See ref. [70] for a detailed discussion on each of these. In this work, a stride of 1 is used, and a size padding is used such that the resulting feature maps have the same dimension as the input feature maps.

As with dense layers, a nonlinear function is generally applied following a convolution operation. In addition, pooling operations are typically performed, which reduce the spatial dimension of the features by summarizing a group of features with a single value. For example, a feature map of size 128 passed through a *max pooling* function of size 2 will result in a feature map of size 64, where sequential groups of 2 features are replaced by the maximum of the two.

5.1.3 Recurrent Layers

RNNs are simply networks that contain one or more *recurrent* layers, which are layers that feed information from a hidden layer from one time step to another. There are a variety of recurrent layers, such as Elman layers [71], long short-term memory modules [72], and gated recurrent units [73]. Elman layers are among the simplest, as they are an extension of a fully-connected layers. In addition to transforming the state $\mathbf{h}^{(i)}$ to $\mathbf{h}^{(i+1)}$, as done in dense feedforward layers, Elman layers also feed the output at time t , $\mathbf{h}_t^{(i+1)}$, as an input

at time $t + 1$:

$$\mathbf{h}_{t+1}^{(i+1)} = \sigma(\mathbf{W}\mathbf{h}_{t+1}^{(i)} + \mathbf{U}\mathbf{h}_t^{(i+1)} + \mathbf{b}), \quad (5.5)$$

where \mathbf{W} , \mathbf{U} , and \mathbf{b} are learned parameters, and σ is a nonlinear activation function.

5.1.4 Network Training

A given network is a combination of layers which are parameterized by \mathcal{P} . In this chapter, *training* refers to the process of using a dataset \mathcal{X} , split into training and validation subsets, to update model parameters \mathcal{P} , generally by some variation of stochastic gradient descent, such that the loss evaluated on \mathcal{X} decreases with number of training iterations, or *epochs*. The training set \mathcal{X} is subdivided into a set used for updating model parameters, also referred to as training data, and a validation set used for assessing model performance during optimization. During each epoch, the model parameters are updated based on each *mini-batch* (i.e., random sample of a fixed size) of training data, and following these parameter updates, the error is computed on the validation data, giving a sense for how accurately the model is performing on data that was not used to update model parameters. The validation set is used to assess the generalization capabilities of the model, and in particular, it is used to indicate when the optimization procedure should cease. Initially, the loss from both the training and validation sets will decrease, however, there will often be a point at which the training loss continues to decrease, while the validation loss increases – a sign of overfitting. *Early stopping* is the method of stopping the training process once the validation loss begins to increase for some number of iterations, referred to as the *patience*.

5.2 Related Work

The use of neural networks to gamma-ray spectroscopy was introduced as early as the 1990s. Since then, many advancements have been made, both in the general field of machine learning, and the application to gamma-ray spectroscopy more specifically. This section reviews previous work in applying neural networks to this domain, and synthesizes previous research to provide directions for further improvements.

5.2.1 Early Work: 1990s - 2000s

Olmos et al. (1991) [74], is the first known application of neural networks to gamma-ray spectroscopy, specifically to perform spectral identification. The authors developed a neural network to predict the amount of a spectrum $\mathbf{x} \in \mathbb{R}^d$ attributed to each source in a library of N sources. That is, the network predicted a vector $\mathbf{c} = f(\mathbf{x}) \in \mathbb{R}^N$. The network used *linear associative memory*, giving the network the form

$$f(\mathbf{x}) = \mathbf{W}\mathbf{x} \quad (5.6)$$

for a set of weights $\mathbf{W} \in \mathbb{R}^{N \times d}$. \mathbf{W} was computed as the pseudoinverse of the matrix $\mathbf{X} \in \mathbb{R}^{d \times N}$ (i.e., $\mathbf{W} = \mathbf{X}^+$). This network was separately trained on Ge(Li) and NaI(Tl) spectra, and evaluated by predicting the amount of source present for known mixtures of radionuclides. This work was extended by the same authors in ref. [75] by accounting for gain drift by including derivative terms for the reference spectra in the matrix \mathbf{X} .

Despite the encouraging results, the work by Olmos et al. did not gain traction, and neural networks applied to gamma-ray spectroscopy did not appear again until the work by Pilato et al. (1999) [76]. Similar to Olmos, Pilato studied a method for computing relative contributions of different sources to a given spectrum, however, the method introduced in that work was primarily aimed at quantifying gamma-ray source activities in laboratory settings, meaning source identities were known *a priori*. Instead of training the network on entire spectra, the method used multiple regions of interest (ROIs) for the sources in question, and uses these ROIs as inputs to the network after performing PCA-based preprocessing. Similarly, Yoshida et al. (2002) [77] used ROIs to reduce the input size of networks, and then performed source activity estimation. The work by Yoshida relied on a spectral peak finding routine instead of analyzing the shape of the entire spectrum as with Olmos et al. A similar treatment is provided in more recent work by Medhat (2012) [78]. Lastly, Chen and Wei (2009) [79] developed an approach which applied the Karhunen-Loève transform to spectra, and used these inputs to a linear associative memory network which performs identification.

5.2.2 Recent Work: Late 2010s - 2020

It was not until the late 2010s that neural networks for gamma-ray spectroscopy would be revisited in the literature, beginning with Kamuda et al. (2017, 2019) [35, 80]. In these two papers, the authors approached the source identification problem using a modern applied neural network treatment. As with Olmos, the authors developed an approach for predicting the relative contribution of sources to a given input NaI(Tl) spectrum. The authors used relatively simple fully-connected architecture with two hidden layers which passed outputs to the softmax function, defined as

$$\text{softmax}(z_i) = \frac{\exp(z_i)}{\sum_j \exp(z_j)}, \quad (5.7)$$

which outputs values between 0 and 1. The resulting vector contains the relative spectral contribution from each source and background. The approach introduced by Kamuda showed promise, and was aligned with contemporary deep learning practices, including the use of dropout and early stopping, performing a random hyperparameter search, and more.

Up to this point, all spectral methods using neural networks were based on fully-connected layers. Kamuda et al. (2020) [69] introduced the use of convolutional layers for feature extraction. The authors recognized the possible benefits using convolutional

layers for recognizing regular patterns such as peaks and continua in gamma-ray spectra, similar to how CNNs are successful in the computer vision domain. Despite using a different structure, this approach was not a significant departure from the previous fully-connected approach in refs. [35, 80], both in methodology and performance. Daniel et al. (2020) [81] continued along the line of using CNNs for identification. The convolutional approach used by Daniel was largely similar to the one used by Kamuda, but introduced a few minor innovations, such as examining and interpreting spectra filtered with the kernels found following training. Such introspection is particularly interesting, as it may aide in understanding how a network is behaving and arriving at its outputs.

Recent work by Moore et al. (2019, 2020) [82, 83] presented a significant departure from the other recent work. In ref. [82], the authors made use of time series of spectra from the URSC dataset discussed in Section 3.3, treating them as monochromatic 2D images. Most notably, the authors introduced the use of successful architectures used in image classification, including CNNs based on AlexNet [84] and VGG, ResNet [85], and Inception modules [86], for performing identification on time series of spectra. In ref. [83], Moore et al. then applied these trained models to measured data from another system, demonstrating the concept of *transfer learning*. The introduction of transfer learning represents a significant contribution, as it demonstrates that models trained in one context (e.g., a simulated environment) can potentially be deployed in another environment (e.g., with measured data), which will be crucial in the large-scale deployment of such methods.

5.2.3 Starting Points for Improvement and Innovation

Nearly all spectral detection methods using neural networks face shortcomings that need to be addressed to be used in practical scenarios. Even the work by Kamuda et al., arguably the most appropriate neural network approach for urban detection up to its time, largely disregarded the effect of variable background, though it is briefly touched on in ref. [69]. To address the need for models that are robust across a variety of background environments, this chapter examines networks trained and evaluated over sets of variable background data.

Detection for nuclear safety and security generally involves the detection of weak sources, unlike most of the high-statistics spectra investigated by Kamuda and others so far. Such source encounters are generally brief, meaning robust neural network approaches must be able to perform detection and identification using measurements from short integration times. Additionally, practical systems must operate at low false alarm rates. Previous work reported results using high false alarm rates (e.g., 5 % in ref. [35]), which are unacceptable in practice – for spectra measured at 1 Hz, such a rate would produce an alarm every 20 s. To understand the operational relevance of these approaches, this chapter only considers false alarm rates in the operational regime (i.e., approximately 10^{-5}).

Another aspect that has been largely overlooked in other approaches thus far is the use of temporal information. As mentioned in Section 2.3, gamma-ray background varies with position, meaning that models that learn this temporal behavior may be able to generate better estimates of the mean background rate, allowing for better source-background separation. Section 5.4 examines the use of RNNs to learn temporal behavior of sequences of spectra.

Lastly, all approaches that have been introduced so far have been limited to identification, or in cases related to source search such as the work by Kamuda, simultaneous detection and identification. So far, there are no known applications of neural network techniques to pure spectral anomaly detection. The following section introduces neural networks for performing anomaly detection on spectra, and identification is examined later in the chapter.

5.3 Anomaly Detection using Autoencoders

As described in Section 2.4, spectral anomaly detection can generally be performed by generating an estimate $\hat{\mathbf{x}}$ of the background in an input spectrum \mathbf{x} , and computing an error measure $D(\mathbf{x}, \hat{\mathbf{x}})$ between the two. $D(\mathbf{x}, \hat{\mathbf{x}})$ is chosen to measure differences between the two inputs, and a threshold T is set, either empirically or analytically using statistical principles, to alarm on spectra that exceed this threshold. The NMF- and PCA-based methods discussed so far use linear models to generate the estimated background spectrum $\hat{\mathbf{x}}$. *Autoencoders* [70, 87] are a type of neural network suitable for performing the background estimation required for anomaly detection. Autoencoders are used to produce an output that is approximately equal to the input, meaning a function f is learned such that

$$\hat{\mathbf{x}} = f(\mathbf{x}) \approx \mathbf{x} \in \mathbb{R}_+^d. \quad (5.8)$$

A trivial function f would map the input using an identity matrix, however, as with NMF, we aim to learn a function f that de-noises the input spectrum. *Undercomplete* autoencoders, which perform dimensionality reduction to learn salient features about the input data, are examined in this chapter. Note that it is not expected for an undercomplete autoencoder reproduce the input exactly, but instead that it returns a denoised copy of the input. The general architecture for an undercomplete autoencoder is to reduce the dimensionality using an *encoder*, then increase to the input dimensionality using a *decoder*. For anomaly detection, the parameters of f are learned from background spectra, and ideally, spectra containing anomalous sources are reconstructed less accurately, resulting in higher detection metrics.

Counts in the input spectra are again assumed to be Poisson-distributed, leading to the use of a Poisson negative log-likelihood loss function of the form

$$-\ln p(\mathbf{X}|\hat{\mathbf{X}}) = \sum_{i=1}^n \sum_{j=1}^d \hat{X}_{ij} - X_{ij} \ln \hat{X}_{ij} + \ln X_{ij}!, \quad (5.9)$$

for a mini-batch $\mathbf{X} \in \mathbb{R}^{n \times d}$ of n spectra and corresponding autoencoder output $\hat{\mathbf{X}}$. For a spectrum \mathbf{x} and corresponding autoencoder reconstruction $\hat{\mathbf{x}}$, the Poisson deviance, defined in Equation (2.17), is used as a detection metric. As in previous chapters, a detection threshold is produced empirically from known background data. Specifically, $D(\mathbf{x}, \hat{\mathbf{x}})$ is computed on background spectra, giving a distribution of test statistics, and for a given FAR, a threshold is empirically set using this distribution.

Encoders and decoders can be built from arbitrary combinations of dense and convolutional layers. Both dense autoencoders (DAE) and convolutional autoencoders were seen to perform comparably during experimentation, but only DAEs are examined here due to their relative simplicity. In this work, only symmetric DAEs (i.e., decoders with architectures that mirror the encoders) are considered here, as it reduces the hyperparameter search space. A DAE with five hidden layers (seven layers total when including the input and output), each using a rectified linear unit (ReLU) activation function, is used to demonstrate spectral anomaly detection. This particular configuration was seen to perform sufficiently well for this initial assessment and the hyperparameter space explored. The number of neurons in each dense layer was found via a random optimization, described further below.

Figure 5.2 shows a diagram of an example dense autoencoder, along with a sample input spectrum \mathbf{x} and its corresponding reconstructed spectrum $\hat{\mathbf{x}}$. By training the autoencoder using the Poisson loss in Equation (5.9), \mathbf{x} can be seen as a sample from a Poisson distribution with mean rate $\hat{\mathbf{x}}$. As a result, $\hat{\mathbf{x}}$ contains smoothed spectral features corresponding to background peaks (e.g., ^{40}K at 1460 keV) and the associated downscattering continuum.

5.4 Source Identification

In performing identification, an input spectrum \mathbf{x} is mapped to an output vector $\hat{\mathbf{y}}$ indicating the presence or absence of a source. In many common applications of neural networks for classification, the mapping from \mathbf{x} to $\hat{\mathbf{y}}$ involves encoding one or more classes of instances present in \mathbf{x} as $\hat{\mathbf{y}}$ (e.g., an image \mathbf{x} containing a dog, encoded in $\hat{\mathbf{y}}$). In gamma-ray spectroscopy, however, background contributions will always be present in a given spectrum, and sources will appear in different proportions, or sources could potentially be shielded by attenuating material. As a result, the standard method of directly predicting a vector $\hat{\mathbf{y}} \in \{0, 1\}^N$, with a 1 at element i indicating the presence of class i , is generally not performed. Instead, the output $\hat{\mathbf{y}}$ is treated as the proportion of each source and background to the spectrum, such that $\sum_i \hat{y}_i = 1$, meaning the network is performing regression. To achieve this behavior, the output of the network $f(\mathbf{x})$ is passed to the softmax function given in Equation (5.7).

A common approach to such classification problems (e.g., in the form of AlexNet [84]) is to use a series of convolutional layers followed by dense layers. Specifically, one or more convolutional layers are used to produce convolutional feature maps, and these




Figure 5.2: (Left panel) Diagram showing the dimensionality of features $\mathbf{h}^{(i)}$ at each layer i for an example dense autoencoder architecture with five hidden layers. A 128-bin spectrum is input into the autoencoder, and dense layers are computed by performing non-linear transformations on each preceding layer. The inverse of each operation is then performed to decode the latent features, resulting in a smoothed spectrum. (Right panel) An input background spectrum \mathbf{x} and corresponding autoencoder reconstruction $\hat{\mathbf{x}}$ are shown. When trained on background, the autoencoder learns spectral features such as background peaks and the associated downscattering continuum. Both the input and output spectra shown here contain 128 bins that scale with the square root of energy. Note that any apparent deviations between the input and output spectra (e.g., at the 1460 keV peak) are due to low-statistics, as the bins of the measured spectrum \mathbf{x} are discrete random samples of the mean Poisson rate $\hat{\mathbf{x}}$.

feature maps are flattened into a 1-dimensional vector, as done in refs. [69, 81], which is then transformed using dense layers. This work makes use of a network with one convolutional layer followed by a max pooling layer and two dense layers, shown in Figure 5.3. Additional layers did not enhance performance for the experiments performed and hyperparameter search space used in this work. Identification networks are trained using mini-batches of spectra \mathbf{X} containing known proportions of source and background \mathbf{Y} . The cross-entropy loss function is used in optimizing network parameters, having the form

$$\mathcal{L}(\mathbf{Y}, \hat{\mathbf{Y}}) = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^{N+1} Y_{ij} \ln \hat{Y}_{ij}, \quad (5.10)$$

where Y_{ij} and \hat{Y}_{ij} are the elements of \mathbf{Y} and $\hat{\mathbf{Y}}$, respectively. Minimizing cross-entropy loss is equivalent to minimizing the Kullback-Leibler (KL) divergence, which is a measure between two probability distributions \mathbf{y} and $\hat{\mathbf{y}}$. The KL divergence, and thus cross-

entropy, is appropriate in this case since the true value of fractional source and background contributions \mathbf{y} and the corresponding estimate $\hat{\mathbf{y}}$ can be treated as probability distributions (i.e., $y_i \in [0, 1]$ and $\sum y_i = 1$). Cross-entropy is used here over the KL divergence, however, simply because it is more common for network-based applications.

Feedforward networks, which include the ANN-based identification methods from previous studies, treat sequential measurements as independent – no information from one measurement is passed to the following. In mobile detection, however, there is generally a relationship between sequential measurements – both background and source contributions to spectra generally do not vary abruptly. As a result, recent measurements can potentially be used to inform the current measurement being processed. For example, if source s_i was present in a spectrum at time t , it is more likely that the spectrum at $t + 1$ also contains source s_i than another source s_j . RNNs are capable of modeling this behavior. Due to the simple time dependence of the source models examined in this work, discussed in Section 3.3 and Section 5.5, this chapter only examines Elman layers for relating sequential spectra. In this work, the first dense layer following the flattening of convolutional features is replaced with an Elman layer, allowing for information from previous spectra $\mathbf{h}_{t+1}^{(i)}$ to be used in performing inference on new spectra (i.e., computing $\mathbf{h}_{t+1}^{(i+1)}$).

In training feedforward networks, mini-batches of spectra are used, where each spectrum contains a random radionuclide with a source activity randomly sampled uniformly from a predefined range. To provide additional variability in training data, a random source-detector angle θ (see Section 3.1) is used for each spectrum, but a fixed standoff distance of 10 m is used since the activity sampling provides variability in SNR. Additionally, pure background is included for the model to appropriately learn background features in the absence of source. RNNs, however, need to learn the temporal dynamics of sources, meaning that mini-batches cannot simply contain random samples of spectra with different physical parameters – the data must include series of spectral measurements of the detector moving past the source. Instead of training the network on random mini-batches of spectra in the form of a matrix $\mathbf{X}_i \in \mathbb{R}^{n \times d}$, 3-dimensional tensors $\mathbf{X}_i \in \mathbb{R}^{r \times n \times d}$ are used, where r refers to a number of runs of data which model the kinematics of a detector moving past a source. With this approach, the network simultaneously learns the mapping between input spectra and relative source contributions and also the temporal behavior of the detector past sources.

5.5 Model Training and Evaluation

This section describes how these networks for detection and identification are trained, evaluated, and benchmarked against the methods from Chapter 4.




Figure 5.3: Example architecture of a convolutional identification network, in a similar fashion to refs. [69, 81]. A 2-dimensional feature map resulting from convolutional operations is flattened into a single feature vector of length 1024, and this is reduced down to the output size of 18 (17 sources, 1 background channel). A max-pooling operation is applied to the features resulting from the convolutional operation, reducing feature size from 128 to 64. In the case of an RNN, the dense layer with size (1, 128) at time t is fed back to combine with the previous layer at time $t + 1$. Not shown here is a softmax function that the output is fed into.

5.5.1 Performance Evaluation and Data

As in Chapter 4, the performance of each method is evaluated using the Minimum Detectable Activity (MDA), computed by estimating the probability of detection as a function of source activity for a given detection scenario. In particular, the URSC dataset from Section 3.3 is used in training and evaluating algorithms. As before, a detection confidence $p_0 = 0.95$ is used for defining the MDA. Seventeen source types were generated using the procedure from Section 3.3: ^{198}Au , ^{133}Ba , ^{82}Br , ^{57}Co , ^{60}Co , ^{137}Cs , ^{152}Eu , ^{123}I , ^{131}I , ^{111}In , ^{192}Ir , ^{54}Mn , ^{124}Sb , ^{46}Sc , ^{75}Se , ^{113}Sn , and ^{201}Tl . In modeling the kinematics of the detector past the source, a vehicle speed of $v = 5$ m/s, along a straight line, and

a standoff distance $d = 10$ m are used. Note that the detector speed used in generating the background data was not provided as part of the original data competition, meaning the speed v used in modeling the source kinematics is not the same. While not ideal, this discrepancy is not believed to affect the conclusions drawn from these analyses, as the speed used here (5 m/s) is in the range of values used to produce the background data, which are between 1 m/s and 13.1 m/s, according to information provided as part of the competition.

5.5.2 Model Optimization

Data Preprocessing and Batch Normalization

Models often converge faster when performing input data preprocessing and feature rescaling within the network [88]. During the experimentation for this work, a linear rescaling based on the mean and standard deviation of training data, referred to as *standardization*, was found to perform well for both detection and identification networks. Standardization transforms an input spectrum \mathbf{x} to \mathbf{x}' as

$$\mathbf{x}' = \frac{\mathbf{x} - \boldsymbol{\mu}}{\boldsymbol{\sigma} + \epsilon}, \quad (5.11)$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}$ are the mean spectrum and standard deviation, respectively, and ϵ is a small positive constant to avoid division by 0. Furthermore, features in the network's hidden layers can be standardized, referred to as *batch normalization* [89] which additionally has a regularizing effect. In this work, DAEs and feedforward identification networks use batch normalization.

Optimizer and Regularization

This work uses the Adam optimizer [89] with an initial learning rate of 10^{-3} for performing parameter optimization. The learning rate is reduced by a factor of 10 when the validation loss does not decrease for 5 trials, reducing the maximum step size, as the model is presumably near a local minimum. To reduce overfitting, an L2 penalty is used, controlled by a coefficient λ . Additionally, *dropout* [90], in which neurons are randomly set to 0 with some probability p ($p = 0.5$ here), is used in the identification networks following the convolutional layer and after the first dense layer. Note that the initial learning rate and dropout probability were deemed to work sufficiently well during manual experimentation and are not optimized further.

Hyperparameter Optimization

The optimization procedure is performed for a given model architecture, dataset, and *hyperparameter* configuration. Hyperparameters refer to parameters that are not learned (i.e., they are configured prior to training), and generally determine a network's modeling

capacity. In the discussion so far, some hyperparameters include the number of neurons in a layer, number of convolutional kernels in a layer, and the L2 regularization coefficient λ . Because of the impact that hyperparameters have on model performance, care must be taken to choose optimal values. Common methods for performing hyperparameter optimization include grid search, random search, and Bayesian optimization. A joint optimization of all hyperparameters is beyond the scope of this work, and instead, this work does a partial optimization: a subset of hyperparameters for a model are fixed, determined from manual experimentation, and a random search [91] is performed with remaining hyperparameters.

To perform the random optimization, the following procedure is used. First, the variable hyperparameters are randomly sampled from a predefined space, and the model is trained and validated using the procedure previously described. Once a model has finished training, source injection is performed on background data from the validation set, giving an initial MDA for each source. Many models are trained using this procedure, and the model resulting in the lowest mean MDA across all sources is used to evaluate the final test set, as the model has shown the greatest generalization capabilities on unseen data. This optimal model is then evaluated on the test set and compared to benchmark algorithms. The specific hyperparameters tuned for each model are described in the following section with the results.

5.5.3 Benchmarking

The autoencoder-based detection method is benchmarked using the NMF- and PCA-based detection algorithms from Chapter 4. Additionally, the neural network-based identification methods are benchmarked using the NMF-based identification method from Chapter 4. As before, for both the detection and identification methods, thresholds are set based on distributions of metrics produced from evaluating on background data.

5.6 Results

5.6.1 Anomaly Detection

Values of λ were randomly sampled between 1 and 100 using a uniform distribution. The number of neurons in each hidden layer (three values in total for a symmetric DAE with five hidden layers), were randomly sampled under the constraint that each value decreased to center of the network. The number of neurons in the outermost values was sampled uniformly between 3 and 24. An example random configuration is (11, 7, 4), which means the number of features for all hidden layers in the network is (11, 7, 4, 7, 11). A total of 40 DAEs were trained using mini-batches of 512 randomly-sampled spectra. For each model, a threshold on the deviance test statistic was set using a FAR of $1/8 \text{ hr}^{-1}$.




Figure 5.4: Histogram of mean MDA for autoencoders evaluated on the validation set. Each model, 40 models in total, was trained using a random value of the L2 regularization coefficient λ , and random configuration of number of neurons in the dense layers of the network. This figure shows that despite being trained with different parameters, initial weights, and mini-batches, most were able to yield similar performance. The model corresponding to the lowest mean MDA from this figure is examined further on the test set.

Figure 5.4 shows a histogram of the mean MDA across all sources injected into the validation background set for the 40 different models. From this optimization, a model corresponding to $\lambda = 55.53$ and number of neurons per layer (7, 3, 2, 3, 7) is used for evaluation on the test set. Both the PCA- and NMF-based models were evaluated on the validation dataset, each indicating that a $k = 1$ component model was sufficient. Note that using a single component to represent either linear models reflects a lack of true variability in the background data for the number of spectral bins used, as many real-world datasets have been seen to yield higher numbers of components.

Figure 5.5 shows a comparison of the optimized DAE and linear models evaluated on each source injected into the background test set. These results show that the autoencoder-based detection method generally outperformed both the NMF- and the PCA-based methods. Specifically, on average, the autoencoder provided a 12% and 23% improvement over NMF- and PCA-based detection methods, respectively. The discrepancy between the PCA-based detection method and the others is likely due to the detection metric used for the PCA-based approach, which comes from previous literature.

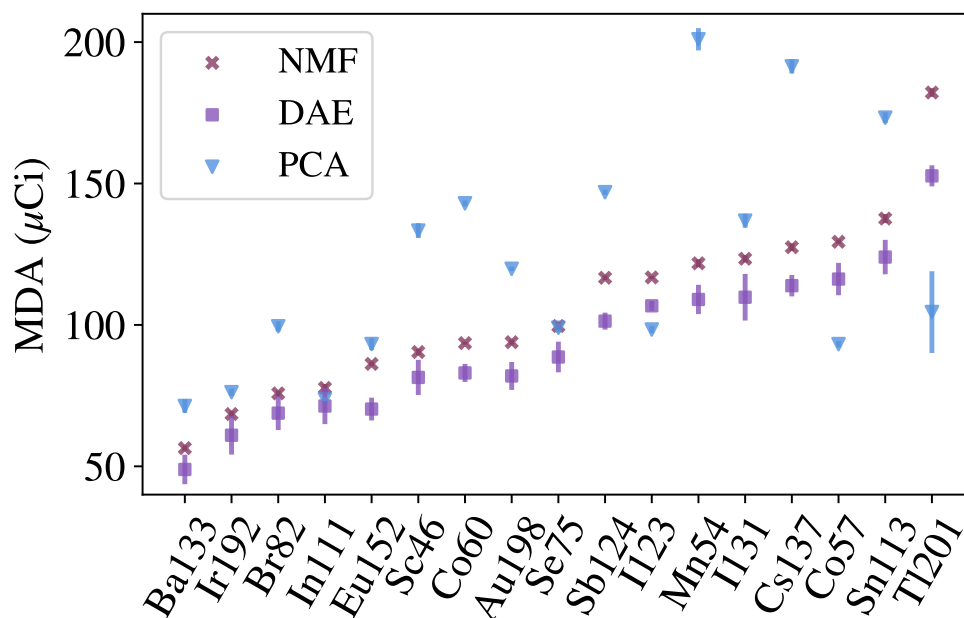


Figure 5.5: Comparison of MDA for the three detection methods across all 17 sources at a $1/8 \text{ hr}^{-1}$ FAR. Sources are sorted in ascending order MDA for the baseline NMF method. Each model was evaluated by injecting each source type across activities into each run of the background test set and computing the MDA. The background used, the test set, was separate from the training and validation background, and thus gives a sense of how well each model generalizes to unseen background data. Note that the discrepancy between the PCA-based method and the other two is likely due to the detection metric used for the PCA-based approach, which comes from the literature. The error bars shown were computed from Equation (3.18). Note that there is an overlap between DAE and PCA for ^{111}In and between NMF and PCA for ^{75}Se .

To assess timing performance, the average runtime per spectrum was computed for each method. Specifically, 100 background runs were randomly sampled, and the average runtime on a per-spectrum basis was computed over each run. The averages from all 100 runs were then averaged and presented here. The reason for averaging over runs is to provide a direct comparison with the RNN-based method in the next section, which is not evaluated on a single spectrum. Each model was evaluated on a 3.50 GHz Intel i7-5930K CPU. The average per-spectrum runtime for the NMF-, PCA-, and DAE-based detection methods were 0.37 ms, 0.05 ms, and 0.02 ms, respectively.

5.6.2 Identification

Values of λ for both feedforward and recurrent identification networks were randomly sampled between 10^{-3} and 10 using a log-uniform distribution. For the single convolutional layer used, the number of kernels was sampled between 8 and 64, and the number of output neurons of the first dense layer, or the recurrent layer in the case of RNNs, was sampled uniformly between 32 and 256. Feedforward models were trained using mini-batches of 256 spectra, and recurrent models were trained with mini-batches of 32 runs of spectra. For a given model, thresholds for each source were set using empirical values of outputs generated from background data to achieve an overall FAR of approximately $1/8 \text{ hr}^{-1}$. Due to simultaneously testing for multiple sources, a Bonferroni correction [92] was used to achieve the target FAR, resulting in an effective FAR for each source which is simply the target FAR divided by the number of sources (i.e., $1/(8 \times 17) \text{ hr}^{-1}$).

The optimization procedure was repeated 40 times for both types of models, resulting in the values $(\lambda, n_{\text{kernel}}, n_{\text{neurons}})$ of (0.49, 64, 66) for the feedforward networks, and (0.0013, 32, 115) for the RNNs. The distribution of MDA over the validation set for all of the 40 trials in the optimization routine is shown for both feedforward and recurrent models in Figure 5.6. This shows that, while there are outliers, the distribution of MDA for RNNs is generally lower than that for feedforward networks, indicating that RNNs often perform better at the same FAR. Figure 5.7 shows a comparison of the mean MDA for optimized feedforward and recurrent identification networks compared to the NMF-based identification benchmark. This figure indicates that the NMF-based detection method and feedforward network perform roughly the same at the same FAR, while the RNN-based identification method often provides an improvement, as expected. Specifically, the RNN was seen to provide a 17% improvement over the feedforward network. As with the detection methods, the average per-spectrum runtime was computed for the NMF-based identification, feedforward network, and recurrent network, yielding 9.54 ms, 0.07 ms, and 0.05 ms, respectively.

5.7 Conclusion

The goals of this chapter were to introduce spectral anomaly detection using autoencoders, establish a baseline of current state-of-the-art identification networks relative to simpler methods, and improve upon the current state-of-the-art using recurrent neural networks. In doing so, ANN-based detection showed a 12% and 23% improvement over the NMF- and PCA-based detection models, respectively, while the current state-of-the-art ANN-based identification was on par with the NMF-based identification method. Furthermore, the ANN-based methods showed a reduction in computational time. The improvement in detection performance and reduction in computation time make ANN-based detection a compelling candidate. Regarding identification, the reduction in runtime while achieving similar performance to NMF, an established benchmark method,




Figure 5.6: Histogram of mean MDA for RNN and feedforward (FF) ID networks evaluated on the validation set. Each model, 40 in total, was trained using a random value of the L2 regularization coefficient λ , number of kernels in the convolutional layer, and number of neurons in the first dense layer. This distribution shows a general trend of improvement when using recurrent layers. The models corresponding to the lowest mean MDA for both feedforward and recurrent networks are examined further on the test set.

makes ANN-based identification worth considering for practical applications as well. Lastly, ANN-based identification was seen to improve, on average, with the use of recurrent layers, meaning that relatively inexpensive performance improvements can be made by including temporal modeling.

Despite the encouraging results, the use of neural networks over linear models involves tradeoffs. Networks require significantly more overhead in the form of data preparation, model design, optimization, etc. Meanwhile, methods such as NMF generally require significantly less in this regard – NMF simply needs to be performed on a matrix of input spectra using a single hyperparameter k , the number of components. However, much of the overhead with networks comes at a fixed cost; the computational burden of networks at runtime is generally significantly lower than iterative methods (e.g., NMF). For this reason, the networks examined here may be more practical than NMF-based detection or identification in scenarios with limited computational resources, for example, in low-power applications.

Though the intention here was to evaluate network-based approaches in operationally-




Figure 5.7: Comparison of the three methods evaluated on the test set: NMF-based identification, a feedforward network (FF), and an RNN-based identification method. Sources are sorted in ascending order MDA for the baseline NMF method. A total FAR of approximately $1/8 \text{ hr}^{-1}$ across all sources was achieved by setting a threshold for each source individually set based on an effective FAR of $1/(8 \times 17) \text{ hr}^{-1}$. The RNN is seen to generally provide an improvement over its feedforward counterpart, though there are a few notable examples, such as ^{133}Ba . Note that there is an overlap of points between NMF and FF for ^{60}Co , ^{123}I , ^{131}I , and ^{192}Ir .

relevant conditions, the performance in many realistic scenarios remains unclear. For one, models must be able to generalize to new environments, meaning these methods should be assessed on real-world data with higher spectral variability. While efforts were made to provide variable background in generating the original simulated data, effects such as elevated radon levels following rain and gain drift were not included. Additionally, other scenarios of practical interest include having multiple sources in a run, either at the same location or at various points in a run, and examining the effect of shielding sources. Future efforts should then assess the impact of shielding and combinations of sources, similar to that of ref. [35], under operationally-relevant conditions.

Due to the number of parameters and operations, neural networks are generally not as interpretable, though this is an active area of research [93]. The nature of nuclear safety and security, the primary application of such algorithms, warrants tools and methods for introspection of networks to better assess behavior. For example, in the case of NMF,

an alarm for a ^{137}Cs source corresponds to an excess of counts associated with the ^{137}Cs template, which an operator can interpret and act on. However, it is not clear how the identification networks examined here could be interpreted, due to the number of interconnected parameters involved in making a decision – the operator must trust that the network is behaving correctly. Additional research is needed in the area of interpretability of spectral models, for example, understanding convolutional kernels as with ref. [81], or generating saliency maps which relate the most significant input features in determining a given network output. In summary, the competitive performance with established methods and additional improvements from temporal information makes these methods worth examining further.

Chapter 6

Concluding Remarks

6.1 Summary

The intention of this work was to develop and analyze methods for spectral gamma-ray anomaly detection and source identification, specifically to lay the foundation for advanced methods which are augmented with information from non-radiological sensors. While a variety of detection and identification methods have been introduced in the past, advances in sensor technology, portable computation, and data analysis methods have allowed for the possibility of more sophisticated approaches which take additional information about the local environment into account. Data-driven methods for spectral analysis provided a promising direction for detection and identification algorithms due to their ability to model distributions of radiological background, which varies with location and time. NMF is one such data-driven method, and in Chapter 4, it was shown to be capable of background and source modeling, anomaly detection, and source identification. In addition to providing improved detection sensitivity over benchmark methods from the literature, NMF results in interpretable models, capable of learning patterns such as downscattering and full-energy peaks. The interpretability of these models is particularly of interest when considering augmentation with non-radiological data. For instance, certain NMF background components can be related to physical features such as gamma-ray signatures from the sky (e.g., cosmic radiation), and the weights associated with these components can potentially be correlated with features from images collected simultaneously. By exploiting such correlations, algorithms may improve background estimates, which is essential in detecting weak sources.

While NMF has been shown to be a worthwhile approach for spectral analysis, there has not been a clear path to augmenting NMF with additional sensor, and recent approaches have relied on ad-hoc analyses to explore correlations [67]. As a result, neural networks were considered, as they are also capable of modeling entire spectra and can easily be extended to co-learn with additional complementary data sources (e.g., via multimodal machine learning [94]). One major limitation of neural networks applied

to gamma-ray spectroscopy was that previous work had not provided an evaluation of these models under operationally-relevant conditions, benchmarked against a known algorithm. A goal for this research was then to establish a baseline of performance for neural network-based detection and identification methods. Neural network-based detection was compared to NMF- and PCA-based detection, and the neural network showed a modest improvement for the experiments performed. Additionally, an identification method using neural networks was compared to NMF-based identification, and both achieved a similar MDA across sources. Identification using networks was extended by use of recurrent neural networks, which allowed for time dependence of modeling and further improved performance. Altogether, this work has introduced novel methods for detection and identification, and shown that these methods outperform commonly-used methods from the literature.

6.2 Future Work

While the approaches developed and assessed in this work are an advancement in spectral processing on their own, the next step, which was a significant factor in considering these methods, is to augment algorithms with contextual data. For example, a simple experiment could be to co-learn autoencoders for anomaly detection using both spectra and images as inputs. Specifically, training can be simultaneously performed on pairs of spectra and images, giving a shared hidden representation, and detection can be performed as before. By constraining background estimates based on learned features from images, such a model may have improved accuracy over a model trained only using spectra, yielding enhanced detection sensitivity.

As seen throughout this work, algorithm development and assessment relies on the ability to train and evaluate methods under realistic conditions. Real-world data require extensive measurements, suffer from sensor robustness issues (e.g., gain shift with temperature), and involve data cleaning and curation, each of which involves significant effort. Algorithm practitioners could then greatly benefit from publicly-available, curated data sources. An attempt at such large-scale gamma-ray data dissemination has been ongoing, with adoption from users in the community [95, 96]. This model of data collection and sharing should be expanded further, providing rich datasets for training and evaluation, and more importantly, to more effectively enable collaboration between researchers. Not only can spectral data can be distributed, but also data from contextual sensors taken simultaneously, allowing for exploration of data fusion approaches by algorithm developers.

In addition to reducing the burden of generating and curating datasets on algorithm developers, standardized datasets would be useful in benchmarking methods, similar to the common approach in computer vision benchmarks (e.g., ref. [97]). As seen with the field of computer vision, the research community as a whole benefits through this process, as the best-performing methods are known to all, and researchers can more ef-

fectively determine which concepts are worth examining further. An organized effort to data collection and dissemination has the potential to not only improve individual research directions, but the field of spectral detection as a whole.

Appendix A

Sigmoid Maximum Likelihood Estimation Derivation

The probability of detection as a function of x , $q(x; \mu, \sigma)$, is parameterized using a sigmoid function of two parameters μ and σ :

$$q(x; \mu, \sigma) = \frac{1}{1 + e^{-(x-\mu)/\sigma}}. \quad (\text{A.1})$$

Consider a set X of N source activities $X = \{x_1, x_2, \dots, x_N\}$. A source of type s with activity x_i is injected into n trials of gamma-ray background, each trial either being a single spectrum or a time series of gamma-ray spectra. At activity x_i , suppose k_i events are detected out of the n trials. The likelihood of detecting k_i events out of n trials for the probability of detection $q(x_i; \mu, \sigma)$, abbreviated as q_i is given by

$$p(k_i|n, q_i) = \binom{n}{k_i} q_i^{k_i} (1 - q_i)^{n-k_i}. \quad (\text{A.2})$$

The total likelihood over all samples in X with corresponding number of detections $K = \{k_1, k_2, \dots, k_N\}$ is given by

$$p(K|n, q) = \prod_{i=1}^N \binom{n}{k_i} q_i^{k_i} (1 - q_i)^{n-k_i}, \quad (\text{A.3})$$

meaning the negative log-likelihood is

$$-\ln p(K|n, q) = \sum_{i=1}^N -\ln \binom{n}{k_i} - k_i \ln q_i - (n - k_i) \ln(1 - q_i). \quad (\text{A.4})$$

Let $z_i = \frac{x_i - \mu}{\sigma}$. The maximum likelihood estimates of μ and σ are then given by

$$\hat{\mu}, \hat{\sigma} = \arg \min_{\mu, \sigma} \left\{ -\ln p(K|n, q) \right\} \quad (\text{A.5})$$

$$= \arg \min_{\mu, \sigma} \left\{ \sum_{i=1}^N -k_i \ln q_i - (n - k_i) \ln(1 - q_i) \right\} \quad (\text{A.6})$$

$$= \arg \min_{\mu, \sigma} \left\{ \sum_{i=1}^N k_i \ln(1 + e^{-z_i}) + (n - k_i)(z_i + \ln(1 + e^{-z_i})) \right\} \quad (\text{A.7})$$

$$= \arg \min_{\mu, \sigma} \left\{ \sum_{i=1}^N z_i(n - k_i) + n \ln(1 + e^{-z_i}) \right\} \quad (\text{A.8})$$

$$= \arg \min_{\mu, \sigma} \left\{ \sum_{i=1}^N \frac{1}{\sigma} (x_i - \mu)(n - k_i) + n \ln(1 + e^{-(x_i - \mu)/\sigma}) \right\} \quad (\text{A.9})$$

$$(\text{A.10})$$

This objective function can then be minimized using an optimization routine to yield the estimates $\hat{\mu}$ and $\hat{\sigma}$.

Appendix B

Probability of Detection Sampling

This appendix introduces three different methods that can be used for sampling source activities for computing the probability of detection for a given detection or identification algorithm.

B.1 Fixed Activity Sampling

In this simplest case, a fixed set of source activities \mathcal{A} are used. At each activity $A \in \mathcal{A}$, n trials are performed, resulting in k detections. In this approach, the range of activities \mathcal{A} must be known *a priori*, and are often initially determined through trial-and-error. That is, the set of activities \mathcal{A} which is expected to contain the MDA for a given algorithm and source scenario must be known. Using the source activities and detections, the probability of detection can be computed using the method described in Appendix A.

B.2 Adaptive Sampling

In some cases, fixed activity sampling for computing MDA is impractical. Specifically, if all of the points in \mathcal{A} are far below or far above the MDA, the detection algorithm may report 0% or 100% detection probability, respectively. This is particularly an issue when there is no prior expectation on the MDA, for example, when testing an algorithm under a new set of parameters.

To avoid this, a wide range of activities must be sampled in order to capture the sigmoid transition region, however, this may result in many unnecessary computations. An alternate approach is to automate finding the transition region by beginning with a set of activity boundaries $(A_{\text{low}}, A_{\text{high}})$ which is large enough such that it is likely to contain the MDA for the detection scenario of interest. Additionally, boundaries on the detection probability $(p_{\text{low}}, p_{\text{high}})$ are set (e.g., (0.01, 0.99)). A random activity $A_i \sim U(A_{\text{low}}, A_{\text{high}})$ is drawn, and the detection probability $\hat{p}(A_i)$ is estimated. If $\hat{p}(A_i) \leq p_{\text{low}}$, A_{low} is updated to A_i , and similarly, if $\hat{p}(A_i) \geq p_{\text{high}}$, A_{high} is updated to A_i . This procedure is repeated

for additional activities, and the effect is a smaller sampling region which includes the sigmoid transition. After N samples have been taken that fall in the region $(p_{\text{low}}, p_{\text{high}})$, the samples can be fit to a sigmoid function, and the MDA is computed.

B.3 Adaptive Bayesian MDA Estimation

An additional approach is to exploit the known functional form that probability of detection curves have – a sigmoid. The probability of measuring k events is given by a binomial distribution:

$$p(k|n, q) = \binom{n}{k} q^k (1 - q)^{n-k}, \quad (\text{B.1})$$

where q is the detection probability from equation (3.15). The binomial probability mass function $p(k|n, q)$ can then be written as a function of the sigmoid parameters as $p(k|n, x, \mu, \sigma)$.

An adaptive approach to sampling can be performed as follows. An initial estimate for the parameters μ_0 and σ_0 can be made, and using these initial estimates, prior probabilities $p(\mu_0)$ and $p(\sigma_0)$ can be specified. Viewing $p(k_0|n, x_0, \mu_0, \sigma_0)$ as a likelihood function, Bayes' theorem can be invoked, yielding the posterior probabilities for μ and σ . Specifically, Bayes' theorem is invoked sequentially yielding new estimates for μ and σ as new activity sampling points x yield new values k .

Suppose i iterations have been performed, resulting in estimates μ_i and σ_i . A new sampling point x_{i+1} is taken, resulting in k_{i+1} detected events. New estimates of the posterior distributions are computed as

$$p(\mu_{i+1}|k_{i+1}, n, x_{i+1}, \mu_i, \sigma_i) = \frac{p(k_{i+1}|n, x_{i+1}, \mu_i, \sigma_i)p(\mu_i)}{p(k_{i+1})} \quad (\text{B.2})$$

$$p(\sigma_{i+1}|k_{i+1}, n, x_{i+1}, \mu_{i+1}, \sigma_i) = \frac{p(k_{i+1}|n, x_{i+1}, \mu_{i+1}, \sigma_i)p(\sigma_i)}{p(k_{i+1})}. \quad (\text{B.3})$$

Note that the posterior distribution for σ_{i+1} uses the recent update for μ_{i+1} , and that the choice to compute the posterior of μ_{i+1} first was arbitrary. Also note that the denominator term $p(k_{i+1})$, which represents the marginal likelihood over the entire parameter space, is simply a normalizing constant and is disregarded. Instead, the product of the terms in the numerator are taken and their product is normalized to unity.

For simplicity, the priors $p(\mu)$ and $p(\sigma)$ are modeled using Gaussian distributions having mean and variance (μ, ω_μ^2) and $(\sigma, \omega_\sigma^2)$, respectively. The mean and variance of each prior is computed using the mean and variance of the respective posterior distribution on the previous iteration, or initial estimate in the case of the first iteration. A new sample x_{i+1} can be drawn using this $p(\mu_i)$, and as new samples are taken, the posterior distributions of μ and σ are updated, resulting in a better estimate of the true sigmoid curve. In summary, this approach is used to simultaneously find the sigmoid parameters and sample activities from regions that are likely to contain samples useful for improving the model.

Appendix C

Spectral Anomaly Detection Derivation

This appendix derives the PCA formulation used in this work, along with an anomaly detection metric that is used with PCA-based models.

C.1 Principal Component Analysis

Note that this section is largely based on a derivation in [53], and is included here to reflect the specific implementation used in this work. Consider a matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ of n samples, each a length- d vector. The true covariance matrix $\Sigma \in \mathbb{R}^{d \times d}$ of the data-generating process is estimated as the matrix \mathbf{S} [36]

$$\mathbf{S} = \text{cov}[\mathbf{X}] = \mathbb{E} \left[(\mathbf{X} - \mathbb{E}[\mathbf{X}])^\top (\mathbf{X} - \mathbb{E}[\mathbf{X}]) \right] = \frac{1}{n-1} \tilde{\mathbf{X}}^\top \tilde{\mathbf{X}},$$

where $\tilde{\mathbf{X}} = \mathbf{X} - \boldsymbol{\mu}$, and $\boldsymbol{\mu}$ is the mean spectrum estimated from \mathbf{X} .

Due to interactions between gamma rays and matter, described in Section 2.1, gamma-ray spectra generally contain continua with enhanced counts at lower energies. That is, there are generally significantly more counts in gamma-ray spectra at lower energies, and as a result, the variance in the number of counts is also enhanced at lower energies. Due to this high variability, PCA performed on the covariance matrix will return components that capture variability in these high-variance low-energy bins. To capture components from bins that vary together across the entire spectrum, PCA is performed on the correlation matrix, which is computed as

$$\mathbf{C} = \boldsymbol{\sigma}^{-1} \mathbf{S} \boldsymbol{\sigma}^{-1}$$

where

$$\boldsymbol{\sigma} = (\text{diag}(\mathbf{S}))^{\frac{1}{2}} \in \mathbb{R}^{d \times d} = \begin{bmatrix} \sigma_1 & 0 & 0 & \dots & 0 \\ 0 & \sigma_2 & 0 & \dots & 0 \\ 0 & 0 & \sigma_3 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \sigma_d \end{bmatrix}$$

Note that the resulting matrix \mathbf{C} has unit variance across each energy bin.

The eigen-decomposition is then performed on \mathbf{C} as

$$\mathbf{C} = \mathbf{V}\mathbf{Q}\mathbf{V}^\top.$$

The matrix $\mathbf{V} \in \mathbb{R}^{d \times d}$ is then an orthogonal basis where the rows are basis vectors of correlated bins. The resulting basis vectors are ordered according to the amount of variability in the data attributed to the vector. The first k basis vectors can then be kept, yielding a compact matrix $\mathbf{V}_k \in \mathbb{R}^{k \times d}$ containing the k most important basis vectors.

The coordinates of a new set of spectra $\mathbf{X}' \in \mathbb{R}^{m \times d}$ in the k -dimensional subspace can then be found by the transformation

$$\mathbf{A} = \left(\frac{\mathbf{X}' - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \right) \mathbf{V}_k^\top \in \mathbb{R}^{m \times k}.$$

Spectra can then be reconstructed as $\hat{\mathbf{X}}'$ by performing the inverse transformation, which yields an approximation of \mathbf{X}' :

$$\begin{aligned} \hat{\mathbf{X}}' &= \mathbf{A}\mathbf{V}_k\boldsymbol{\sigma} + \boldsymbol{\mu} \\ &= \left(\frac{\mathbf{X}' - \boldsymbol{\mu}}{\boldsymbol{\sigma}} \right) \mathbf{V}_k^\top \mathbf{V}_k \boldsymbol{\sigma} + \boldsymbol{\mu} \\ &= (\mathbf{X}' - \boldsymbol{\mu}) \boldsymbol{\sigma}^{-1} \mathbf{V}_k^\top \mathbf{V}_k \boldsymbol{\sigma} + \boldsymbol{\mu}. \end{aligned}$$

C.2 Spectral Anomaly Detection

The Spectral Anomaly Detection (SAD) method [53] uses the method in the previous section for generating a low-dimensional reconstruction of background gamma-ray spectra. In particular, SAD uses the L_2 norm between measured spectra \mathbf{X}' and the reconstructed spectra $\hat{\mathbf{X}}'$ as a measure of error:

$$\|\mathbf{X}' - \hat{\mathbf{X}}'\|_2 = \|\mathbf{X}' - (\mathbf{X}' - \boldsymbol{\mu}) \boldsymbol{\sigma}^{-1} \mathbf{V}_k^\top \mathbf{V}_k \boldsymbol{\sigma} - \boldsymbol{\mu}\|_2 \quad (\text{C.1})$$

$$= \|(\mathbf{X}' - \boldsymbol{\mu})(\mathbb{I} - \boldsymbol{\sigma}^{-1} \mathbf{V}_k^\top \mathbf{V}_k \boldsymbol{\sigma})\|_2 \quad (\text{C.2})$$

$$= \|(\mathbf{X}' - \boldsymbol{\mu})\mathbf{T}\|_2 \quad (\text{C.3})$$

where $\mathbf{T} = \mathbb{I} - \boldsymbol{\sigma}^{-1} \mathbf{V}_k^\top \mathbf{V}_k \boldsymbol{\sigma} \in \mathbb{R}^{d \times d}$ is a linear operator.

The quantity $\|(\mathbf{X}' - \boldsymbol{\mu})\mathbf{T}\|_2$ scales with the magnitude of spectra (i.e., the gross counts), and a normalization is applied to account for this scaling. The complete SAD metric is then

$$\text{SAD}(\mathbf{X}') = \frac{\|\mathbf{X}' - \hat{\mathbf{X}}'\|_2}{\sqrt{\|\mathbf{X}'\|_1}} = \frac{\|(\mathbf{X}' - \boldsymbol{\mu})\mathbf{T}\|_2}{\sqrt{\|\mathbf{X}'\|_1}}. \quad (\text{C.4})$$

Bibliography

- [1] Alan E. Waltar. "Nuclear Technology's Numerous Uses". In: *Issues in Science and Technology* 20.3 (2004), pp. 48–54. ISSN: 07485492, 19381557. URL: <http://www.jstor.org/stable/43312453>.
- [2] *IAEA Safety Glossary: 2018 Edition*. Vienna: International Atomic Energy Agency, 2019. ISBN: 978-92-0-104718-2. URL: <https://www.iaea.org/publications/11098/iaea-safety-glossary-2018-edition>.
- [3] *The Fukushima Daiichi Accident*. Vienna: International Atomic Energy Agency, 2015. ISBN: 978-92-0-107015-9. URL: <https://www.iaea.org/publications/10962/the-fukushima-daiichi-accident>.
- [4] Yukihsa Sanada and Tatsuo Torii. "Aerial radiation monitoring around the Fukushima Dai-ichi nuclear power plant using an unmanned helicopter". In: *Journal of Environmental Radioactivity* 139 (2015), pp. 294–299. ISSN: 0265-931X. DOI: <https://doi.org/10.1016/j.jenvrad.2014.06.027>.
- [5] Blair Thornton et al. "Distribution of local ¹³⁷Cs anomalies on the seafloor near the Fukushima Dai-ichi Nuclear Power Plant". In: *Marine Pollution Bulletin* 74.1 (2013), pp. 344–350. ISSN: 0025-326X. DOI: <https://doi.org/10.1016/j.marpolbul.2013.06.031>.
- [6] International Atomic Energy Agency. *Ukraine to Receive Mobile Radiation Detection Vehicle*. Apr. 2010. URL: <https://www.iaea.org/newscenter/news/ukraine-receive-mobile-radiation-detection-vehicle> (visited on 05/21/2020).
- [7] National Nuclear Security Administration. *NNSA to conduct aerial radiation measurements over Washington, D.C.* URL: <https://www.energy.gov/nnsa/articles/nnsa-conduct-aerial-radiation-measurements-over-washington-dc> (visited on 11/13/2020).
- [8] Kurt Guthe. "The Global Nuclear Detection Architecture and the Deterrence of Nuclear Terrorism". In: *Comparative Strategy* 33.5 (2014), pp. 424–450. DOI: 10.1080/01495933.2014.962961.

- [9] International Atomic Energy Agency. *List of Confirmed Incidents Involving HEU or Pu*. URL: <https://web.archive.org/web/20080724044058/https://www.iaea.org/NewsCenter/Features/RadSources/table1.html> (visited on 05/21/2020).
- [10] George M. Moore. *Lessons from a Mexican theft*. URL: <https://thebulletin.org/2013/12/lessons-from-a-mexican-theft/>.
- [11] Igor Jovanovic and Anna S. Erickson. *Active Interrogation in Nuclear Security*. Springer, 2018.
- [12] M. Askins et al. *The Physics and Nuclear Nonproliferation Goals of WATCHMAN: A WAter CHerenkov Monitor for ANtineutrinos*. 2015. arXiv:1502.01132 [physics.ins-det].
- [13] R. C. Runkle et al. "Analysis of spectroscopic radiation portal monitor data using principal components analysis". In: *IEEE Transactions on Nuclear Science* 53.3 (2006), pp. 1418–1423.
- [14] S. M. Robinson et al. "Time Series Evaluation of Radiation Portal Monitor Data for Point Source Detection". In: *IEEE Transactions on Nuclear Science* 56.6 (2009), pp. 3688–3693.
- [15] D. Hellfeld et al. "Real-Time Free-Moving Active Coded Mask 3D Gamma-Ray Imaging". In: *IEEE Transactions on Nuclear Science* 66.10 (2019), pp. 2252–2260.
- [16] Mark S. Bandstra et al. "RadMAP: The Radiological Multi-sensor Analysis Platform". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 840 (2016), pp. 59–68. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2016.09.040>.
- [17] Joseph C. Curtis et al. "Simulation and validation of the Mobile Urban Radiation Search (MURS) gamma-ray detector response". In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* August (Sept. 2018). ISSN: 01689002. DOI: 10.1016/j.nima.2018.08.087.
- [18] T. H. Y. Joshi et al. "Measurement of the Energy-Dependent Angular Response of the ARES Detector System and Application to Aerial Imaging". In: *IEEE Transactions on Nuclear Science* 64.7 (2017), pp. 1754–1760.
- [19] T. J. Aucott et al. "Routine Surveys for Gamma-Ray Background Characterization". In: *IEEE Transactions on Nuclear Science* 60.2 (2013), pp. 1147–1150.
- [20] Yanming Guo et al. "A review of semantic segmentation using deep neural networks". In: *International Journal of Multimedia Information Retrieval* 7.2 (2018), pp. 87–93. DOI: 10.1007/s13735-017-0141-z.
- [21] Kenneth S Krane. *Introductory nuclear physics*. Wiley, 1988.
- [22] Gordon Gilmore. *Practical gamma-ray spectrometry*. John Wiley & Sons, 2011.

- [23] Glenn F. Knoll. *Radiation Detection and Measurement*. John Wiley & Sons, 2012.
- [24] Edward C. Morse. *Analytical Methods for Nonproliferation*. Springer, 2016.
- [25] R.J. Livesay et al. "Rain-induced increase in background radiation detected by Radiation Portal Monitors". In: *Journal of Environmental Radioactivity* 137 (2014), pp. 137–141. ISSN: 0265-931X. DOI: <https://doi.org/10.1016/j.jenvrad.2014.07.010>.
- [26] Simon R. Cherry, James A. Sorenson, and Michael E. Phelps. *Physics in nuclear medicine*. Saunders, 2012.
- [27] National Research Council (US) and Institute of Medicine (US) Committee on State of the Science of Nuclear Medicine. *Advancing nuclear medicine through innovation*. National Academies Press, 2007.
- [28] D. Huang et al. "CoDetect: Financial Fraud Detection With Anomaly Feature Detection". In: *IEEE Access* 6 (2018), pp. 19161–19174.
- [29] Gilberto Fernandes et al. "A comprehensive survey on network anomaly detection". In: *Telecommunication Systems* 70.3 (Feb. 2018), pp. 447–489. DOI: [10.1007/s11235-018-0475-8](https://doi.org/10.1007/s11235-018-0475-8).
- [30] J. Liu et al. "Anomaly Detection in Manufacturing Systems Using Structured Neural Networks". In: *2018 13th World Congress on Intelligent Control and Automation (WCICA)*. 2018, pp. 175–180.
- [31] Prateek Tandon et al. "Detection of radioactive sources in urban scenes using Bayesian Aggregation of data from mobile spectrometers". In: *Information Systems* 57 (2016), pp. 195–206. ISSN: 0306-4379. DOI: <https://doi.org/10.1016/j.is.2015.10.006>.
- [32] K. Miller and A. Dubrawski. "Gamma-Ray Source Detection With Small Sensors". In: *IEEE Transactions on Nuclear Science* 65.4 (Apr. 2018), pp. 1047–1058. ISSN: 0018-9499. DOI: [10.1109/TNS.2018.2811049](https://doi.org/10.1109/TNS.2018.2811049).
- [33] Kenneth D. Jarman et al. "A comparison of simple algorithms for gamma-ray spectrometers in radioactive source search applications". In: *Applied Radiation and Isotopes* 66.3 (2008), pp. 362–371. ISSN: 0969-8043. DOI: <https://doi.org/10.1016/j.apradiso.2007.09.010>.
- [34] T. H. Joshi et al. "A Comparison of the Detection Sensitivity of the Poisson Clutter Split and Region of Interest Algorithms on the RadMAP Mobile System". In: *IEEE Transactions on Nuclear Science* 63.2 (Apr. 2016), pp. 1218–1226. ISSN: 0018-9499. DOI: [10.1109/TNS.2016.2537206](https://doi.org/10.1109/TNS.2016.2537206).
- [35] M. Kamuda, J. Stinnett, and C. J. Sullivan. "Automated Isotope Identification Algorithm Using Artificial Neural Networks". In: *IEEE Transactions on Nuclear Science* 64.7 (2017), pp. 1858–1864.

- [36] Larry Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer Texts in Statistics. New York: Springer, 2004.
- [37] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. 1st ed. Springer, 2007.
- [38] J. A. Nelder and R. W. M. Wedderburn. "Generalized Linear Models". In: *Journal of the Royal Statistical Society. Series A (General)* 135.3 (1972), pp. 370–384. ISSN: 00359238. URL: <http://www.jstor.org/stable/2344614>.
- [39] S. Agostinelli et al. "Geant4 – a simulation toolkit". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [40] Cihangir Celik et al. "A Directional Detector Response Function for Anisotropic Detectors". In: *Nuclear Science and Engineering* 193.12 (2019), pp. 1355–1370. DOI: 10.1080/00295639.2019.1631028.
- [41] *NuDat 2.8*. <https://www.nndc.bnl.gov/nudat2/>. [Online; accessed 21-May-2020].
- [42] A. A. Sonzogni. "NuDat 2.0: Nuclear Structure and Decay Data on the Internet". In: *AIP Conference Proceedings* 769.1 (2005), pp. 574–577. DOI: 10.1063/1.1945075.
- [43] Brian C. Plimley et al. *Becquerel (Bq) v0.1*. DOI: <https://doi.org/10.11578/dc.20180925.1>.
- [44] "Benchmarking the Gamma-Ray Response of the Radiological Multi-sensor Analysis Platform". M.S. Thesis. University of California, Berkeley, 2014.
- [45] E. I. Novikova et al. "Designing SWORD–SoftWare for Optimization of Radiation Detectors". In: *2006 IEEE Nuclear Science Symposium Conference Record*. Vol. 1. Oct. 2006, pp. 607–612. DOI: 10.1109/NSSMIC.2006.356228.
- [46] *Detecting Radiological Threats in Urban Areas*. <https://www.topcoder.com/challenges/30085346>. [Online; accessed 14-February-2020].
- [47] K. J. Bilton et al. "Non-negative Matrix Factorization of Gamma-Ray Spectra for Background Modeling, Detection, and Source Identification". In: *IEEE Transactions on Nuclear Science* 66.5 (2019), pp. 827–837. DOI: 10.1109/TNS.2019.2907267.
- [48] K. J. Bilton et al. "Modeling Shielded Gamma-ray Source Spectra using Non-negative Matrix Factorization". In: *2019 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*. 2019, pp. 1–7. DOI: 10.1109/NSS/MIC42101.2019.9059733.
- [49] S. Kaski and J. Peltonen. "Dimensionality Reduction for Data Visualization [Applications Corner]". In: *IEEE Signal Processing Magazine* 28.2 (2011), pp. 100–104.

- [50] Ian T. Jolliffe and Jorge Cadima. "Principal component analysis: a review and recent developments". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 374.2065 (2016). DOI: 10.1098/rsta.2015.0202.
- [51] Michael E. Tipping and Christopher M. Bishop. "Probabilistic Principal Component Analysis". In: *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 61.3 (1999), pp. 611–622. ISSN: 13697412, 14679868. URL: <http://www.jstor.org/stable/2680726>.
- [52] D. Boardman, M. Reinhard, and A. Flynn. "Principal Component Analysis of Gamma-Ray Spectra for Radiation Portal Monitors". In: *IEEE Transactions on Nuclear Science* 59.1 (2012), pp. 154–160.
- [53] Prateek Tandon. "Bayesian Aggregation of Evidence for Detection and Characterization of Patterns in Multiple Noisy Observations". English. Ph.D. United States – Pennsylvania: Carnegie Mellon University, 2015. URL: <https://search.proquest.com/docview/1774007802/abstract/463397BF9CAF4561PQ/1> (visited on 05/22/2018).
- [54] Michael Collins, Sanjoy Dasgupta, and Robert E. Schapire. "A generalization of principal component analysis to the exponential family". In: *Advances in Neural Information Processing Systems*. MIT Press, 2001.
- [55] P. Tandon et al. "Suppressing Background Radiation Using Poisson Principal Component Analysis". In: *arXiv:1605.08455 [physics, stat]* (May 2016). arXiv: 1605.08455. URL: <http://arxiv.org/abs/1605.08455> (visited on 03/07/2018).
- [56] Pentti Paatero and Unto Tapper. "Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values". en. In: *Environmetrics* 5.2 (June 1994), pp. 111–126. ISSN: 1099-095X. DOI: 10.1002/env.3170050203. (Visited on 05/23/2018).
- [57] Daniel D. Lee and H. Sebastian Seung. "Learning the parts of objects by non-negative matrix factorization". In: *Nature* 401 (Oct. 1999). DOI: 10.1038/44565;. URL: <http://dx.doi.org/10.1038/44565>.
- [58] S. Kullback and R. A. Leibler. "On Information and Sufficiency". In: *Ann. Math. Statist.* 22.1 (Mar. 1951), pp. 79–86. DOI: 10.1214/aoms/1177729694.
- [59] Daniel D. Lee and H. Sebastian Seung. "Algorithms for Non-negative Matrix Factorization". In: *Proceedings of the 13th International Conference on Neural Information Processing Systems*. NIPS'00. Denver, CO: MIT Press, 2000, pp. 535–541. URL: <http://dl.acm.org/citation.cfm?id=3008751.3008829>.
- [60] H. Akaike. "A new look at the statistical model identification". In: *IEEE Transactions on Automatic Control* 19.6 (Dec. 1974), pp. 716–723. ISSN: 0018-9286. DOI: 10.1109/TAC.1974.1100705.

- [61] C. Boutsidis and E. Gallopoulos. "SVD based initialization: A head start for nonnegative matrix factorization". In: *Pattern Recognition* 41.4 (2008), pp. 1350–1362. ISSN: 0031-3203. DOI: <https://doi.org/10.1016/j.patcog.2007.09.010>.
- [62] S. S. Wilks. "The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses". In: *The Annals of Mathematical Statistics* 9.1 (1938), pp. 60–62. ISSN: 00034851. URL: <http://www.jstor.org/stable/2957648>.
- [63] Steven G. Self and Kung-Yee Liang. "Asymptotic Properties of Maximum Likelihood Estimators and Likelihood Ratio Tests Under Nonstandard Conditions". In: *Journal of the American Statistical Association* 82.398 (1987), pp. 605–610. ISSN: 01621459. URL: <http://www.jstor.org/stable/2289471>.
- [64] R. J. Estep, C. W. McCluskey, and B. A. Sapp. "The multiple isotope material basis set (MIMBS) method for isotope identification with low-and medium-resolution gamma-ray detectors". In: *Journal of Radioanalytical and Nuclear Chemistry* 276.3 (June 2008), pp. 737–741. ISSN: 1588-2780. DOI: 10.1007/s10967-008-0626-6.
- [65] M. L. Koudelka and D. J. Dorsey. "A Modular NMF Matching Algorithm for Radiation Spectra". In: *2016 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. June 2016, pp. 284–289. DOI: 10.1109/CVPRW.2016.42.
- [66] K. J. Bilton et al. "Neural Network Approaches for Mobile Spectroscopic Gamma-ray Source Detection". In: *IEEE Transactions on Nuclear Science (under review)* (2020).
- [67] Mark S. Bandstra et al. "Attribution of gamma-ray background collected by a mobile detector system to its surroundings using panoramic video". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 954 (2020). Symposium on Radiation Measurements and Applications XVII, p. 161126. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.08.085>.
- [68] S. Khan et al. *A Guide to Convolutional Neural Networks for Computer Vision*. 2018.
- [69] Mark Kamuda, Jifu Zhao, and Kathryn Huff. "A comparison of machine learning methods for automated gamma-ray spectroscopy". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 954 (2020). Symposium on Radiation Measurements and Applications XVII, p. 161385. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2018.10.063>.
- [70] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. <http://www.deeplearningbook.org>. MIT Press, 2016.
- [71] Jeffrey L. Elman. "Finding Structure in Time". In: *Cognitive Science* 14 (1990), pp. 179–211.
- [72] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (1997), pp. 1735–1780.

- [73] Kyunghyun Cho et al. *Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation*. 2014. arXiv: 1406.1078 [cs.CL].
- [74] P. Olmos et al. "A new approach to automatic radiation spectrum analysis". In: *IEEE Transactions on Nuclear Science* 38.4 (1991), pp. 971–975.
- [75] P. Olmos et al. "Drift problems in the automatic analysis of gamma-ray spectra using associative memory algorithms". In: *IEEE Transactions on Nuclear Science* 41.3 (1994), pp. 637–641.
- [76] V. Pilato et al. "Application of neural networks to quantitative spectrometry analysis". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 422.1 (1999), pp. 423–427. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(98\)01110-3](https://doi.org/10.1016/S0168-9002(98)01110-3).
- [77] Eiji Yoshida et al. "Application of neural networks for the analysis of gamma-ray spectra measured with a Ge spectrometer". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 484.1 (2002), pp. 557–563. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(01\)01962-3](https://doi.org/10.1016/S0168-9002(01)01962-3).
- [78] M.E. Medhat. "Artificial intelligence methods applied for quantitative analysis of natural radioactive sources". In: *Annals of Nuclear Energy* 45 (2012), pp. 73–79. ISSN: 0306-4549. DOI: <https://doi.org/10.1016/j.anucene.2012.02.013>.
- [79] Liang Chen and Yi-Xiang Wei. "Nuclide identification algorithm based on K-L transform and neural networks". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 598.2 (2009), pp. 450–453. ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2008.09.035>.
- [80] Mark Kamuda and Clair J. Sullivan. "An automated isotope identification and quantification algorithm for isotope mixtures in low-resolution gamma-ray spectra". In: *Radiation Physics and Chemistry* 155 (2019). IRRMA-10, pp. 281–286. ISSN: 0969-806X. DOI: <https://doi.org/10.1016/j.radphyschem.2018.06.017>.
- [81] G. Daniel et al. "Automatic and Real-Time Identification of Radionuclides in Gamma-Ray Spectra: A New Method Based on Convolutional Neural Network Trained With Synthetic Data Set". In: *IEEE Transactions on Nuclear Science* 67.4 (2020), pp. 644–653.
- [82] Eric T. Moore et al. *An Application of CNNs to Time Sequenced One Dimensional Data in Radiation Detection*. 2019. arXiv: 1908.10887 [physics.app-ph].
- [83] Eric T. Moore et al. *Transfer Learning in Automated Gamma Spectral Identification*. 2020. arXiv: 2003.10524 [physics.data-an].
- [84] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *NIPS*. 2012.

- [85] Kaiming He et al. *Deep Residual Learning for Image Recognition*. 2015. arXiv: 1512.03385 [cs.CV].
- [86] Christian Szegedy et al. *Going Deeper with Convolutions*. 2014. arXiv: 1409.4842 [cs.CV].
- [87] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. "Learning Internal Representations by Error Propagation". In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*. Ed. by David E. Rumelhart and James L. McClelland. Cambridge, MA: MIT Press, 1986, pp. 318–362.
- [88] Yann LeCun et al. "Efficient BackProp". In: *Neural Networks: Tricks of the Trade*. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 1998. Chap. 2, p. 546. ISBN: 978-3-540-65311-0. DOI: 10.1007/3-540-49430-8_2.
- [89] Sergey Ioffe and Christian Szegedy. "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift". In: *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*. 2015.
- [90] Nitish Srivastava et al. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". In: *J. Mach. Learn. Res.* 15 (2014), pp. 1928–1958.
- [91] James Bergstra and Yoshua Bengio. "Random Search for Hyper-Parameter Optimization." In: *J. Mach. Learn. Res.* 13 (2012), pp. 281–305. URL: <http://dblp.uni-trier.de/db/journals/jmlr/jmlr13.html#BergstraB12>.
- [92] Eric W. Weisstein. *Bonferroni Correction*. URL: <https://mathworld.wolfram.com/BonferroniCorrection.html>.
- [93] Julius Adebayo et al. "Sanity Checks for Saliency Maps". In: *Proceedings of the 32nd International Conference on Neural Information Processing Systems*. NIPS'18. Montréal, Canada: Curran Associates Inc., 2018, pp. 9525–9536.
- [94] Tadas Baltrušaitis, Chaitanya Ahuja, and Louis-Philippe Morency. *Multimodal Machine Learning: A Survey and Taxonomy*. 2017. arXiv: 1705.09406 [cs.LG].
- [95] Shreyas Cholia et al. *Berkeley Nuclear Data Cloud (BDC) v2.3, Version 00*. Feb. 2018. URL: <https://www.osti.gov/servlets/purl/1434262>.
- [96] *Berkeley Data Cloud*. URL: <https://bdc.lbl.gov/> (visited on 11/01/2020).
- [97] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. *CIFAR-10 (Canadian Institute for Advanced Research)*. URL: <http://www.cs.toronto.edu/~kriz/cifar.html> (visited on 11/01/2020).