

Bayesian Aggregation of Evidence for Detection and Characterization of Patterns in Multiple Noisy Observations

Prateek Tandon

July 2015

Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Artur Dubrawski
Manuela Veloso
Paul Scerri
Simon Labov (LLNL)

CMU-RI-TR-15-23

Keywords: Robotics, Bayesian Aggregation, Retrospective Learning, Mobile Sensors, Nuclear Threat Detection

Abstract

Effective use of Machine Learning to support extracting maximal information from limited sensor data is one of the important research challenges in robotic sensing. This thesis develops techniques for detecting and characterizing patterns in noisy sensor data. Our Bayesian Aggregation (BA) algorithmic framework can leverage data fusion from multiple low Signal-To-Noise Ratio (SNR) sensor observations to boost the capability to detect and characterize the properties of a signal generating source or process of interest.

We illustrate our research with application to the nuclear threat detection domain. Developed algorithms are applied to the problem of processing the large amounts of gamma ray spectroscopy data that can be produced in real-time by mobile radiation sensors. The thesis experimentally shows BA's capability to boost sensor performance in detecting radiation sources of interest, even if the source is faint, partially-occluded, or enveloped in the noisy and variable radiation background characteristic of urban scenes.

In addition, BA provides simultaneous inference of source parameters such as the source intensity or source type while detecting it. The thesis demonstrates this capability and also develops techniques to efficiently optimize these parameters over large possible setting spaces. Methods developed in this thesis are demonstrated both in simulation and in a radiation-sensing backpack that applies robotic localization techniques to enable indoor surveillance of radiation sources.

The thesis further improves the BA algorithm's capability to be robust under various detection scenarios. First, we augment BA with appropriate statistical models to improve estimation of signal components in low photon count detection, where the sensor may receive limited photon counts from either source and/or background. Second, we develop methods for online sensor reliability monitoring to create algorithms that are resilient to possible sensor faults in a data pipeline containing one or multiple sensors.

Finally, we develop Retrospective BA, a variant of BA that allows reinterpretation of past sensor data in light of new information about percepts. These Retrospective capabilities include the use of Hidden Markov Models in BA to allow automatic correction of a sensor pipeline when sensor malfunction may occur, an Anomaly-Match search strategy to efficiently optimize source hypotheses, and prototyping of a Multi-Modal Augmented PCA to more flexibly model background and nuisance source fluctuations in a dynamic environment.

*To my parents and cute little brothers and sisters too
who keep copying everything I do
and to our puppy of new
Chewbie Doobie Doo*

...And also to my iCreate robots.



These iCreates are ones I worked on at University of Southern California (USC) in undergrad. After I graduated, (unknown to me at the time) several were shipped to Carnegie Mellon University (CMU) for use here. I found them in storage gathering dust. Their model is obsolete, but they still boot up, so I updated their PlayerStage configuration to ROS. They have made cameos in my project videos on YouTube. Average readers may not understand sentimentality about roombas, but, after a Robotics PhD, you too would appreciate a pack of consistently-performing, fault-tolerant roombas!

Acknowledgements

I would like to thank my advisor, Prof. Artur Dubrawski, for everything he has done guiding me through the trials and travails of the PhD. It has been quite a journey, and I am very glad he has not given up on me (yet). In all seriousness, though, thank you, Artur, for reviewing my late night poster drafts and always being prompt with papers and deadlines. Thank you for caring for your students no matter the circumstance.

Thank you also to my PhD Committee members: Prof. Manuela Veloso, Prof. Paul Scerri, and Dr. Simon Labov. You have served as additional guides on my PhD journey. Manuela, I really enjoyed TAing Graduate AI. Paul, I still have your robots so remind me to return them. Simon, thank you for all the physics expertise which made my thesis even possible.

Thank you to my Auton Lab coauthors Prof. Jeff Schneider, Dr. Peter Huggins, Dr. Adam Zagorecki, Vladimir Ermakov, and Rob MacLachlan. Thank you to my undergraduate student coauthors Aashish Jindia, Parag Srivastava, and Jay Jin for help with the research. In addition, thank you to the numerous researchers, postdocs, PhD Students, research scientists, software engineers, system admins, and staff members at the Auton Lab that have collaborated with me over the years on aspects related to my PhD work. Thank you to Mike Baysek, Karen Chen, Dr. Madalina Fiterau, Dr. Roman Garnett, Dr. Mathieu Guillame-Bert, Yamuna Krishnamurthy, Yifei Ma, Dr. Kyle Miller, Prof. Barnabas Poczos, Dr. Predrag Punosevac, Dougal J. Sutherland, Saswati Ray, Jarod Wang, and Dr. Liang Xiong. It has been a pleasure to work with all of you.

Thank you to my collaborators for this research at Lawrence Livermore National Laboratory. Thank you to Dr. Karl Nelson and Dave Trombino for extensive guidance related to NaI sensors, calibration algorithms, and spectral anomaly detection. Our backpack system would not have been possible without your support.

I would like to extend the thanks to additional CMU professors, students, and staff who have made my life easier along the way. These are people who have been influential professors, been on committees that have advised me, or TAed me for important class projects. Thank you to Prof. Chris Atkeson, Prof. Drew Bagnell, Prof. Matt Mason, Prof. Reid Simmons, Dr. Debadeepta Dey, Dr. Mark Palatucci, and Dr. Mihail Pivtoraiko. Thank you to CMU staff members Suzanne Lyons Muth and Karen Widmaier for logistical support and ensuring this thesis goes out on time.

Thank you to the extensive (and most definitely incomplete) list of CMU Robotics friends and colleagues who I have gotten to know at RoboOrg events: Matt Barnes, Joydeep Biswas, Nathan Brooks, Ellen Cappel, Andrew Chambers, Jennifer Cross, Tony Dear, Felix Duvallet, Michael Furlong, Allie Del Giorno, Sam Horvath, Humphrey Hu, Breelyn Kane, Heather Knight, Shervin Javdani, Heather Jones, Karthik Lakshmanan, Robbie Paolini, Alex Schepelmann, Alex Styler, Michael Taylor, Prasanna Velagapudi, Jiuguang Wang, Nathan Wood, Ada Zhang, Yin Zhang. Thank you to the RI Table Tennis crew (of late), Martin Li and Wenhao Luo.

Thank you to the Academic Research Initiative (ARI) and Defense Threat Reduction Agency (DTRA) programs through which I am funded to complete this research. Thank you to the appropriate coordinating offices for hosting us at our regular conference meetings.

Finally, I ought ultimately acknowledge the professors at the University of Southern California (USC) who inspired me to spend five years completing a PhD thesis. Thank you to Prof. Milind Tambe, Prof. Michael Arbib, and Prof. Michael Crowley from whose early inspiration provided me the springboard to become a researcher.

Contents

1	Introduction	1
1.1	Problem Statement	1
1.2	Nuclear Threat Detection Domain	2
1.2.1	Technical Challenges in Nuclear Threat Detection	2
1.2.2	Related Work	3
1.2.3	Innovations of Bayesian Aggregation	4
2	Bayesian Aggregation Method	7
2.1	Stages of Bayesian Aggregation	7
2.2	Related Work	9
2.3	Benchmark methods	10
3	Measurement SNR Estimation	13
3.1	Anomaly Detection Approaches	14
3.1.1	Principal Component Analysis	14
3.1.2	Spectral Anomaly Detector	16
3.1.3	Poisson Principal Component Analysis	16
3.2	Match Filtering Approaches	17
3.2.1	Correlation Match Filter	17
3.2.2	Energy Windowing Regression	17
4	Sensor Fusion with Bayesian Aggregation	19
4.1	Algorithm Description	19
4.2	Experimental Setup	22
4.3	Results and Discussion	23
4.3.1	Incorporation of Source Intensity and Type Hypotheses in BA	24
4.4	Inference of Source Properties	28
4.5	Conclusions	29
5	Adaptive Post-processing of Hypothesis Scores	31
5.1	Adaptive Grid Bayesian Aggregation	32
5.1.1	Algorithm	32
5.1.2	Experiments and Results	33
5.1.3	Efficient Maintenance of Hypothesis Space	34

5.2	Branch and Bound for Handling Dynamic Occlusions	35
5.2.1	Algorithm	36
5.2.2	Experiments and Results	37
6	Using Bayesian Aggregation to Guide Data Collection and Search Efforts	39
6.1	Use of BA Information in Single Agent Source Search	39
6.1.1	Single Agent Algorithm	39
6.1.2	Experiments and Results	41
6.2	Application of BA to Multi-Agent Radioactive Source Search	42
6.2.1	Algorithm	43
6.2.2	Results	45
6.3	Discussion and Conclusions	47
7	Low Photon Count Signal and Noise Component Estimation	49
7.1	Introduction	49
7.1.1	Typical Measurement SNR Estimation	50
7.1.2	Poisson Modeling	50
7.1.3	Poisson PCA	51
7.1.4	Zero-Inflated Regression Poisson Match Filter	51
7.2	Experiments and Discussion	52
7.2.1	Benchmarking Poisson PCA vs. Gaussian PCA	52
7.2.2	Benchmarking Poisson Match Filters vs. Gaussian Match Filters	56
7.2.3	Incorporation into BA Framework	57
7.3	Conclusion	61
8	Using BA for Online Sensor Reliability Monitoring	63
8.1	Introduction	63
8.2	Using BA for Sensor Characterization	64
8.3	Sensor Fault Monitoring	65
8.3.1	Simulation of Sensor Failures	65
8.3.2	Diagnosing Sensor Failures in Background Modeling	67
8.3.3	Incorporating Sensor Failure Hypotheses into BA	70
9	Retrospective Learning and Applications to Radiation Sensing	73
9.1	Introduction	73
9.1.1	Retrospective Learning	73
9.1.2	Multiple Bayesian Hypotheses Approach	75
9.2	Anomaly Match Bayesian Aggregation (AM-BA) Approach	79
9.2.1	Application to Radiation Source Detection Problem	80
9.2.2	Experiments with Clustering Source and Background	82
9.2.3	Discussion and Conclusion	85
9.3	Multi-Modal Augmented Principal Component Analysis	85
9.3.1	Augmented PCA Method	85
9.3.2	Application of Augmented PCA to Radiation Source Detection	86

10 Conclusion	93
10.1 Thesis Contributions	93
10.2 “User Guide” for Bayesian Aggregation	95
10.3 Generalizability of BA Method	100
10.4 Recommended Future Research Directions	100
10.5 Final Remarks	102
A	103
A.1 Data Sets for Mobile Threat Detection Problem	103
A.1.1 Sacramento Data Set	104
A.1.2 RadMap and SUNS Data Sets	105
A.1.3 Pittsburgh Data	105
A.2 Robotic Radiation Sensing Backpack Data Platform	106
A.2.1 Backpack System Architecture	106
A.2.2 Experiments with Indoor Localization	107
A.3 Sensor Calibration and Energy Windowing	108
A.3.1 Energy Selection Heuristic for Match Filters	114
A.4 Radiation Point Source Simulation Methods	115
A.5 Principal Component Analysis Derivation	116
A.6 Spectral Anomaly Detector Derivation	117
A.7 Poisson Principal Component Analysis Derivation	118
A.7.1 General EPCA	118
A.7.2 PPCA Derivation	118
A.7.3 Poisson PCA Final Algorithm	120
A.7.4 Zero-Inflated Poisson Regression Match Filter	121
Bibliography	125

List of Figures

2.1	Variance Reduction Pipeline. The branches in the figure represent different popular options for the pipeline components.	7
3.1	PCA Scree Plot	14
3.2	First Three Principal Components of Sacramento Background Data.	15
3.3	PCA Example	15
3.4	Example of Regression Match Filter	17
4.1	BA Alternate and Null Sensor Models	20
4.2	Example data processing through BA. Figure 4.2a shows input data: a raw trajectory with an injected point source affecting measurements. Figure 4.2b shows the output: a resulting BA threat probability map. The red region in Figure 4.2b indicates the correctly detected radiation source.	22
4.3	Setup for experimental evaluation	22
4.4	ROC curve results for algorithms.	23
4.5	Intensity Modeling ROC Results	25
4.6	Intensity Modeling ROC Results	26
4.7	Source templates injected into background.	27
4.8	Source Type Modeling ROC Results	27
5.1	Example of Adaptive Grid BA. The algorithm grows the grid adaptively around the detected point source to better estimate the true source location.	32
5.2	Expansion Rules for Adaptive Grid	33
5.3	Adaptive Grid Source Localization Result	33
5.4	Build and Query Time Comparison for Tree Data Structures	35
5.5	Figure shows an example of a static occlusion scenario. Sensor observations in green are occluded by a wedge with particular start and end angles and receive no injection, while observations in red are unoccluded. Dynamic occlusions, which occur at particular points in time, can further complicate analysis of a wedge depending on the time of the observation. In a dynamic occlusion scenario, some of the green measurements may actually be red, necessitating a search for the source in both space and time dimensions simultaneously.	36
5.6	ROC Detection Performance without and with occlusions	37
6.1	Example Entropy Reduction Curve for Exhaustive Algorithm	41

6.2	Source Detection Capability of Exhaustive ($T=3$) and Greedy Strategies	41
6.3	Effect of Setting Prior Probabilities on Threat Detection Performance	42
6.4	Entropy of Non-Source vs. Source-Like Vehicle	43
6.5	Illustration of Entropy Propagation Algorithm Steps	45
6.6	Results of Small-Scale Simulation	46
7.1	SKL comparison of methods as a function of distance to the source.	53
7.2	SKL comparison of methods when detecting a source with shorter measurement lifetime.	54
7.3	SKL comparison of methods when detecting a faint source.	55
7.4	Experiments comparing match filter estimators.	56
7.5	Experiments with Weighted Mean Method.	57
7.6	Example BA sensor models for Spectral Anomaly Detector and Poisson PCA. . .	58
7.7	SKL comparison of methods when detecting a source with shorter measurement lifetime.	59
7.8	Detection results for using ZIP Regression Match Filter information in BA. . . .	60
8.1	Empirical Models of Different Size Sensors.	64
8.2	Effect of Sensor Model Misspecification.	64
8.3	Simulation of Sensor Dropout and Stalling.	66
8.4	Mean Spectrum Comparison for Normal Spectra and Spectrum Corrupted by Gain Drift.	66
8.5	Mean Spectrum Comparison for Normal Spectrum and Spectrum Corrupted with Crystal Problems.	67
8.6	Sensor Failures in PCA Reconstruction Residual.	68
8.7	Sensor Failures in Canonical Variable Space.	69
8.8	Sensor Failure in BA Hypothesis Space.	70
8.9	Sensor Failure Diagnosis Results.	71
8.10	Sensor Failure Diagnosis Results: Lowered Source Intensity.	72
9.1	ROC Results	77
9.2	AUC Results as a Function of Data Corruption Percentage $T\%$	78
9.3	Histogram of Forward Probabilities	79
9.4	Null and Alternate Hypothesis (Source Agnostic) Anomaly Models	81
9.5		81
9.6	Results of AM-BA Experiments.	82
9.7	Clustering of Sacramento Background Spectra	83
9.8	Clustering of Sacramento Background Spectra	84
9.9	Operational Clustering using PCA Residuals	84
9.10	PCA Scores for Background Data	87
9.11	Rotation of Principal Components due to Background Changes	88
9.12	Rotation of Principal Components due to Nuisance Sources	88
9.13	PCA Scores for Background Data	89
9.14	Cross-Validating over Source Library	90

9.15	Cross-Validating over Nuisance Library	91
9.16	Cross Validation Experiments with Stronger Nuisance Source	91
10.1	Trajectory data containing point source affecting sensor observations.	95
10.2	Easy source detection scenario with strong signal in individual measurement. . .	96
10.3	Source detection scenario requiring aggregation of observations to detect source.	96
10.4	Comparison of WC and BA on data with a distributed nuisance source.	97
10.5	Match filter BA improves detection capability.	97
10.6	ZIP BA improves detecting sources in sparse, low count spectra.	98
10.7	Augmented PCA can help lower false alarm rate.	98
10.8	Fault-Tolerant BA can account for possible sensor failures in data.	99
A.1	Sacramento Data Set	105
A.2	System Diagram	106
A.3	Screenshots of Android Application Capabilities	107
A.4	Wi-Fi Localization Accuracy	108
A.5	Estimated Detection and Localization Capability of Wi-Fi Localizer	108
A.6	24 Hours of Collected Background and Source Isotope data	109
A.7	DGE Gain Computation Steps for Example Thorium Feature	110
A.8	Spectra Before and After Gain Drift Calibration with DGE	111
A.9	Constructed Spline Result	113
A.10	Mapping of 1024 channel-space of sensor to energy (“EnergyInstrument” profile)	113
A.11	The sensor receives a new background spectrum in 1024 channel space (Figure A.11a). The DGE procedure aligns the new spectrum with reference spectra (Figure A.11b). Calibration produces a 128-dimensional energy-space vector (Figure A.11c).	114
A.12	Snapshot of Source Simulator Capabilities	115

List of Tables

4.1	Intensity Multiplier to Count Rate Conversion	25
4.2	Confusion matrix for intensity inference: BA-Marg / BA-Max.	28
4.3	Confusion matrix for source type inference: BA-Marg / BA-Max.	28
5.1	Comparison of Grid Sizes	34
7.1	Correlation Comparison of Spectral and Poisson PCA SNR Scores	59
8.1	Confusion Matrix for Diagnosing Sensor Failure Modes	68
8.2	Accuracy for Diagnosing Sensor Failure Severity	69
A.1	Radiation Data Sets	103
A.2	Key Data Properties	103
A.3	Channel-Energy Features	112

Chapter 1

Introduction

1.1 Problem Statement

Analyzing and extracting maximal information from limited sensor data is a key research challenge for robotic sensing. This thesis develops Machine Learning techniques for aggregating evidence from multiple low Signal-to-Noise Ratio (SNR) observations to boost detectability and property characterization of a source or process producing the signals.

Bayesian Aggregation (BA), the name of our method, is an algorithmic framework for leveraging data fusion from multiple sensor observations to boost the capability to detect and characterize the properties of the signal generating source. Different stages of the BA algorithm such as background noise estimation, nonparametric modeling of SNR expectation, and Bayesian data fusion across multiple sensors or viewpoints help account for different aspects of the sensor measurements.

The BA framework supports building highly detailed empirical sensor models that can scale to analyze large amounts of sensor data in real time to identify and characterize signals of interest. The data-analytic technique is applicable to domains involving both single or multiple mobile sensors used to make noisy observations about the environment where the goal of the sensing is to understand and characterize the properties of the underlying signals in the data.

The thesis develops methods to boost detection of sources of interest, even if the source is partially-occluded or enveloped in background noise. Methods can simultaneously infer properties of sources such as intensity or type and efficiently optimize estimates of such parameters.

The thesis also addresses critical challenges to robustifying BA to sensor data with less reliable information. For instance, in the case of low signal detection, the sensor may receive low signal from the source and/or background. In this case, BA can be augmented with appropriate statistical models to improve performance. When scaling BA to more sensors, a key challenge is that sensors may dynamically malfunction in the field. BA can be made robust to possible sensor failures in the data by incorporating explicit hypotheses for sensor failure in the BA hypothesis space. By testing for the possibility of sensor malfunction, BA can enable more fault-tolerant source detection.

We also develop Retrospective BA, a variant of BA that can reinterpret sensor data based on new data clues about those percepts. Sensor failures, background fluctuations, and benign

nuisance isotopes, if not accounted for, can cause a detector system to sound false alarms. Retrospective BA utilizes various constructs such as Hidden Markov Models, a joint Anomaly-Match hypothesis search strategy, and Multi-Modal Augmented Principal Component Analysis to automate re-analysis of data in the presence of such occurrences.

Developed techniques are demonstrated with application to the nuclear threat detection domain. We demonstrate BA's capability both in simulation and in a robotic radiation-sensing backpack system used to detect radiation sources.

1.2 Nuclear Threat Detection Domain

From preventing the proliferation of nuclear weapons around the globe to curbing untold tragedy from a terrorist group detonating a nuclear weapon, a nation's capability to effectively monitor nuclear materials around the world is increasingly key to safety in the 21st century. Since their invention, the destructive capability of nuclear weapons is unrivaled by any other technology known to humans. Whether the threats from nuclear materials are international or domestic, from large, well-organized and funded state actors or from rogue independent militia groups, one thing is paramount: Whether in the wrong hands or even in the malfunctioning "right" hands, nuclear weapons and catastrophe can go hand in hand.

The threats from nuclear weapons are only increasing. Countries or terrorist groups acquiring nuclear weapons in an unchecked or unregulated manner is a scary situation for the international community at large [58]. Furthermore, the smuggling of such weapons (encased in engineered shielding material) onto U.S. soil is a very real threat, and such warheads could cause massive harm if detonated in an urban area. Tactical nuclear weapons, highly destructive yet portable warheads that can be easily transported around an environment make up most of the nuclear arsenal that has ever been produced [47]. Dirty bombs built from radioactive sources [2], radioactive emanation from stolen medical or industrial-use isotopes [50], or health risks from a malfunctioning nuclear power plant [3] are other frightening but very real threat scenarios with recent precedents.

The increasing demand for nuclear monitoring of such scenarios compels the development of newer and better technologies for threat detection. While chokepoint monitoring can be effective if the source goes through the choke point, if the source is able to be smuggled through other routes, the challenge becomes one of wide-area detection.

For these broad area monitoring problems, mobile radiation detectors are a promising technology. Typical mobile radiation detectors are highly portable radiation sensing units, deployed in a moving vehicle for outdoor broader area radiation monitoring or carried by pedestrians for indoor surveillance. Their aim is to aid law enforcement officers in identifying the possible sources of radiation in a geographic area and alert if a possible threat is detected.

1.2.1 Technical Challenges in Nuclear Threat Detection

Advances in hardware design have allowed for the ability to collect significant amounts of radiation spectrum data in real time. One of the fundamental challenges is to automate the analysis of the large amounts of collected sensor data. Some of the key challenges are:

1. Modeling and characterizing typical fluctuations in radiation background and detecting anomalous behavior
2. Providing accurate detection and localization of radioactive sources of different types under various environmental conditions
3. Suppressing the impact of background fluctuations and environment nuisance sources on the accurate detection of threats
4. Planning sensor placements, routes, and trajectories to find static and mobile sources in various cluttered urban scenes.

In general, well developed algorithms can augment the sensitive detection capabilities of the hardware but maintain low false detection rates. Variability in the background radiation rates and nuisance sources in the environment can cause false alarms for mobile radiation detectors. The remedy is to account for the expectable variation in background and common potential nuisances in modeling so that a detector can tell a truly threatening radioactive source of interest from a benign one.

The algorithms can also boost the capability to detect radioactive sources in low SNR and highly nonuniform settings that may result from incidental anisotropy on the radioactive source. Incidental anisotropy is directional occlusion of the source by virtue of the environment. Obstacles such as trees, passing cars, and buildings can cause static and/or dynamic obstruction in the path from the sensor to the source, affecting the measured SNR by the detector. The goal is to detect threatening radiation sources even in the presence of partial occlusion of the source.

A significant useful capability provided by Bayesian algorithms is the capability to infer properties of the radioactive source such as the source intensity (also referred to as the source count rate multiplier) or the source (e.g. isotope) type. It is useful for a law enforcement officer using a mobile radiation detector to know such parameters of the source so that an appropriate response can be made.

Finally, as the system gathers observations about the region, it can characterize its own uncertainty about where radioactive sources do and do not exist in the environment. This uncertainty analysis can gate trajectory route planning for future locations for the sensor to explore. The algorithms thus not only provide passive information for the operator to use, but can provide active advice for where the system operator should search next for possible radioactive sources of interest.

1.2.2 Related Work

The problem of detecting a radioactive point source from observations of radiation spectra has been previously studied by many research efforts.

A well-known early study in radiation source detection looked at the case of detecting a fully isotropic source using a mobile spectrometer [77]. In their mathematical model, they identified that background photon count rates are typically unknown. Thus, characterizing and suppressing the background is an important challenge in finding the source. Interestingly, they found that background and source photon counts both scale proportionally with the surface area of the sensor, making it difficult to gain from a larger detector.

To improve robustness and performance on actual collected sensor data, algorithms in the literature leveraged statistical models that could capture variability and imperfections in data in an attempt to improve detection of radioactive sources in real world settings. These formulations were based both on frequentist and Bayesian approaches and resulted in better signal separation models for signal and noise components on actual collected sensor data.

One of the popular approaches is K-Sigma, which models collected photon counts under a Poisson likelihood model that takes into account the distribution of total gross photon counts in collected radiation spectra [9]. The K specifies a detection parameter such that a spectrum is flagged as containing a source if the total counts in the spectrum is larger than K standard deviations away from the total counts in a mean background spectrum.

Building upon the anomaly detection theme, the Spectral Anomaly Detector algorithm [45] employs Principal Component Analysis (PCA) to capture major directions of background fluctuation and variation. The Spectral Anomaly Detector can be used to capture (and suppress) the key principal linear directions of variance in background spectra containing multiple energy bins. Projecting a new radiation observation onto a learned basis for background (and subtracting out projections in these directions) results in a spectral anomaly score, which can be used to decide whether the observation exhibits source-like behavior or is more background-like.

Both vanilla K-Sigma and the Spectral Anomaly Detector score individual sensor spectrum observations. Using Bayesian techniques, algorithms could be extended to account for fluctuations in signal and noise across multiple correlated sensor observations via Bayesian data fusion [43, 44]. For instance, particle filters [55] are a popular approach to aggregating multiple observations to detect or track a moving target in a Bayesian framework. They are immensely popular in areas such as robotic localization and tracking [19]. Particle filtering methods have been developed for tracking a radioactive source [10, 51, 60, 61].

Aside from the Bayesian data fusion approach, the Weighted Combining (WC) method [46] (also known as the “ $\frac{1}{r^2}$ back projection” method) has been popular for fusing evidence from multiple observations to detect sources. The WC method uses $\frac{1}{r^2}$ weighting (where r is distance to a source) on measurements to flag a source location. The method is presumed to be optimal for flagging the locations in the environment that maximize the estimated SNR at the locations, given the source is isotropic. The algorithm maintains a geographic background “map” and a geographic source “map” which are iteratively updated. Given a new measurement, WC estimates the signal and noise components of the measurement and adds these estimates to its running estimates of signal and noise at geographic locations in its maps. The geographic location with the highest SNR score, after aggregation, is predicted to be the source location.

The Self-Weighted Combining (SWC) method [46] is an extension to WC that uses the difference between estimated source counts and estimated background counts as the weighting when fusing measurements instead of the $\frac{1}{r^2}$ weighting. This method is attractive, since it has been demonstrated to account for occlusions but at a slight reduction in overall sensitivity.

1.2.3 Innovations of Bayesian Aggregation

As we will see, our BA approach builds upon many of the existing works, but provides many improvements not yet explored in the literature.

First, it is one of the premier Bayesian methods to fully utilize spectral information in empirical modeling of data likelihoods from real-world background data instead of just total photon counts in a spectrum. Using non-parametric density estimation techniques and appropriate measurement scoring schemes, BA can effectively suppress background radiation and non-threatening radiation emanated by nuisance sources without making *a priori* parametric assumptions about the distribution of local background and nuisances [67]. These capabilities help provide robust signal separation when compared to other methods of aggregating evidence.

Second, our framework enables not only source detection and localization from multiple spatially-correlated observations but also inference of characteristics of the source such as the source intensity multiplier or source type. The inferential capabilities of BA allow more than just determination of source parameters [40, 59, 76]. BA allows simultaneous tracking of multiple, multi-modal hypotheses about source parameters in a joint hypothesis space. We develop scalable branch and bound and pruning mechanisms for searching (at multiple resolutions) over source parameters such as source location, intensity, isotope type, as well as static and dynamic occlusion scenarios. These computational innovations enable efficient evaluation of many possible source configuration scenarios that may manifest in the real world.

BA resembles the framework underlying particle filtering [55] and sequential importance sampling [51] and, given sufficiently extensive posterior sampling, offers equivalent theoretical detection performance as particle filters. However, BA is designed to be substantially more computationally scalable by using fast data structures and a manageably complex space of source parameter hypotheses. By using data structures such as kd-trees and R-trees to speed up computation of hypotheses, BA can enable rapid evaluation of many real world scenarios.

The statistical and computation innovations developed in this thesis enable our BA framework to build upon prior cited work.

Chapter 2

Bayesian Aggregation Method

2.1 Stages of Bayesian Aggregation

The problem of identifying a source of interest from cluttered, varying background noise can be viewed in the language of variance reduction. The ultimate goal is to identify the true signal due to the source of interest, removing or otherwise accounting for the background fluctuations in the process.

The BA method [63]¹ for detecting a source from multiple sensor observations utilizes a set of stages to account for different aspects of variance in the data.




Figure 2.1: Variance Reduction Pipeline. The branches in the figure represent different popular options for the pipeline components.

¹Tandon, Prateek, Artur Dubrawski, Jeff Schneider, Adam Zagorecki, Simon Labov, and Karl Nelson. Machine Learning for Effective Nuclear Search and Broad Area Monitoring. ARI Annual Review 2011.

The first stage of variance reduction involves sensor calibration and feature selection. During this stage, possible sources of temporal gain drift in sensor spectra are filtered out via calibration. Many sensors suffer from different types of calibration errors, and this stage in the BA pipeline can help account for such challenges. In the radiation domain, spectrometers are often affected by gain drift. Gain drift is systematic shift in where photon counts fall in the energy spectrum due to changing temperature, humidity, lighting and other temporal effects. The net effect is that photon counts may appear in energy bins adjacent to their true energy values. To handle these drift effects, sensors are often characterized and calibrated with respect to their gain drift properties. For example, the Dynamic Gain Estimation algorithm is often used to calibrate for gain drift in our lab’s hardware data collection system (see appendix A.3 for details).

Once calibrated, feature selection is employed to help tune the sensor to particular features of interest in sensor observations. For example, in radiation data, energy windowing is often employed. Energy windowing or tuning the sensor to particular ranges or windows of “interesting” bins in the energy spectrum can boost the expected signal from radioactive sources that may have key well-known peak features in these areas. Energy windowing is useful if the operator is looking for particular radioactive source types of interest, and detection of such sources of interest can be greatly improved by this type of feature selection. An example algorithm for performing energy windowing automatically, given source templates, is described in appendix A.3.1.

The goal of the SNR estimation stage of BA (described in Chapter 3) is to decompose an observed sensor spectrum into the contribution of a source component (if any) and the contribution of typical background variation in the data. Though background is noise, there is often structure in it. Environments often contain typical background variation that is a function of temperature, time, humidity, etc. Such variation can often be characterized by statistical algorithms and subtracted out from the measured spectra. SNR estimation is often done by finding principal directions of systematic variation in the data and/or by matching directions of variance to those expected by a human observer. For example, on the radiation data, Principal Component Analysis (PCA) and Energy Windowing Regression are often employed to estimate background and source components in an observed spectrum.

In the sensor modeling stage of BA (described in Chapter 4), expectation of variance in source signal observed by the sensor is modeled non-parametrically using empirical models. For example, on the radiation data, the empirical models allow discounting spurious false positives in the data by encoding a data-driven expectation that radiation source signal follows particular exposure laws on the detector surface. Since models are empirical, they can flexibly account for sophisticated statistical distributions of variance. Since models are nonparametric, they can flexibly account for nonlinearity or other non-uniformity in the data that does not follow any known parametric distribution.

Once a satisfactory model of the sensor is built, multiple sensor observations can be aggregated to increase detectability of a radiation source, even if the source of interest is very weak or partially occluded. BA uses conditional independence assumptions and Bayesian update rules to aggregate multiple observations in a computationally efficient manner. In the radiation domain, it is expected that, for an isotropic (or even partially occluded) radioactive source, multiple observations of the source at different locations in its proximity will be correlated with respect to the presence of the source. Mid-level fusion of multiple sensor observations can thus possibly aid detectability of the radioactive sources.

Efficient maintenance of the source hypothesis space can be a nontrivial computational problem. The cartesian product of possible source locations, combined with possible source intensities and anisotropy configurations, can be quite expansive and deciding which hypothesis to test next is a key challenge. The hypothesis space can be dynamically post-processed and adapted in several ways (e.g. as described in Chapter 5). While a rectangular grid spanning the collected data is often used as a starting point, the hypothesis space can be refined via various post-processing techniques. For instance, the hypothesis space can be analyzed at multiple resolutions and grown adaptively as necessary. Also, iterative Retrospective Learning algorithms are hypothesized to be able to aid detectability of a radioactive source.

Once a hypothesis space is maintained efficiently, it can be used to gate robotic route-planning efforts for both single and multi-agent systems (described in Chapter 6). The current hypothesis space of probable source locations can be used by one or many mobile detectors to decide where the sensors should go next to maximally reduce the uncertainty of the world. Active planning and subsequent collection of sensor observations at locations in the environment that maximize reduction of uncertainty in where source signal exists (and where it does not) serves as another way to reduce variability in the data. For example, in the radiation domain, techniques can be developed to gate exploration and exploitation behaviors for a team of law enforcement vehicles to find a radioactive source in traffic.

2.2 Related Work

Particle filtering [19] is a related approach to maintaining a space of hypotheses. In a sense, the vanilla version of BA can be thought of as a particle filter with static particles spaced as a uniform grid in the region of observation. The fact that BA employs a base level grid of static particles allows it to avoid loss of diversity issues involved with particle respawning [55]. Furthermore, by using a base number of uniformly spaced particles, BA can be used in the multi-source and anomaly detection settings where there might be several real or nuisance sources on the map. BA can robustly account for all such hypotheses, whereas a Particle Filter would likely struggle with multi-source behavior.

BA uses special data representation structures such as KD-trees and R-trees to speed up computation of particles actually affected by new data. In the nuclear threat detection problem, particles are only affected locally by spatially proximal measurements. Use of tree data structures speed up the particle computation several-fold by avoiding unnecessary updating of particles not affected by new geo-local measurements.

BA is scalable and inherently parallelizable in ways that vanilla particle filtering is not. First, BA makes the assumption that data updates to a particle are independent from each other. Second, BA makes the assumption that particle updates are independent of each other. The net result is that the BA approach can scale to cases with many heterogeneous sensors and multitudes of sensor data. BA is readily implementable in a Map-Reduce framework (such as Hadoop) commonly developed on cloud computing grids in industry.

Finally, BA does not suffer from the kidnapped robot problem because there are always particles in all regions. The uniform grid means that no sensor resetting (commonly used with particle filters) is necessary in cases where the radioactive source may momentary teleport or

disappears due to hardware failures in the sensor. BA can readily incorporate models of hardware limitation, and the grid of hypothesis locations ensures that the algorithm doesn't abruptly stop working due to hardware failure (a particle filter might lose all of its particles in this scenario).

2.3 Benchmark methods

Developed BA methods are compared to other existing methods to detect a radioactive source by sensor fusion that are commonly used in the radiation domain. These methods provide good benchmarks to our BA approach for aggregating multiple observations to detect a source, and are used to showcase its potential practical utility in an operational setting.

The Weighted-Combining (WC) method [46] uses idealistic $\frac{1}{r^2}$ distance weighting on measurements to flag a source location. The method is presumed to be optimal for flagging the locations in the environment that maximize the estimated Signal-to-Noise Ratio (SNR). The algorithm maintains a background "map" and a source "map" which are iteratively updated. Given a new radiation measurement vector M , the scalar total sum of photon counts in the spectrum, Y , and an estimate of its background counts B , the number of source photon counts is estimated as $S = Y - B$. The Weighted Combining Method of scoring source location hypotheses estimates the SNR at locations (x,y) on the map, aggregating evidence using the following update rules:

$$\begin{aligned} Map_S(x, y) &= Map_S(x, y) + \frac{S}{R(x,y)^2} \\ Map_B(x, y) &= Map_B(x, y) + \frac{B}{R(x,y)^4} \\ SNR(x, y) &= \frac{Map_S(x,y)}{\sqrt{Map_B(x,y)}} \end{aligned} \tag{2.1}$$

The method fuses evidence from multiple measurements to detect a source, taking into account expected exposure to the source due to R^2 sensing laws. There are two key limitations in applying WC to threat detection, however. First, the WC method does not account for momentary nuisance (but benign) sources in the data that can cause false alarms. Imagine if a banana truck drove by the sensor and caused a momentary SNR spike, or if the sensor drove by a radiology department in a hospital. In both cases, the WC method would flag the source location as containing threats, even though the SNR spikes are really caused by characteristic nuisance sources existing in the urban scene being surveyed. Accounting for typical sources of variation in urban scenes is a key challenge in threat detection. Second, in threat detection, source counts can be within the tolerance of background. Maximizing the raw SNR metric may not be enough to detect the source, and distributional models of SNR may be necessary.

The Self-Weighted Combining (SWC) method is an extension to WC that uses the difference between estimated source counts and estimated background counts as the weighting when fusing measurements instead of the idealistic $\frac{1}{r^2}$ weighting. Under idealized conditions, this weighting strategy should help account for occlusions that would cause WC to fail. The challenge in making the SWC method effective, however, is that estimation of source and background components is prone to noise. Variability and imperfection in the estimation procedure can harm SWC's capability to be effective. It is possible that BA trained estimates can boost the capability of SWC to detect threats.

Methods are compared on data sets publicly available in the radiation community (described in appendix A.1). The thesis commonly uses either the Sacramento or SUNS data for experiments. The Sacramento data set contains real field data collected over a period of five consecutive days by a van driving mostly in down-town Sacramento equipped with a double NaI planar detector. The data consists of GPS trajectories of the van as well as the observed spectrum at each environment location the van traversed. The SUNS data set is a simulated data stream (albeit based on real high-resolution data) for pedestrian-wearable spectrometers. The data set consists of several pedestrian trajectories as well as the observed sensor reading at each location from the pedestrian's sensor.

Chapter 3

Measurement SNR Estimation

Signals received by sensors are typically presumed to be composed of multiple additive components that contribute to the overall sensor measurement. The problem of Measurement SNR Estimation is to estimate the signal and noise components of a sensor measurement to form its Signal-To-Noise Ratio (SNR) score.

Signal processing defines the Signal-to-Noise Ratio of a sensor measurement M as:

$$SNR(M) = \frac{\mu_{sig}}{\sigma_{bg}} \quad (3.1)$$

In the radiation domain, a radiation spectrum M is assumed to be made up of the additive contribution of background radiation and the contribution of a signal generating source. The scalar total sum of photon counts, Y , in the spectrum is presumed to given by:

$$Y = B + S \quad (3.2)$$

where B is the scalar sum of photon counts in the background component and S is the scalar sum of photon counts in the source component. Note that S can be zero if there is no source contribution. The goal of SNR estimation is to estimate the quantities B and S in M .

The signal is given by the source signal estimate, and the noise term is formed as the square root of the background signal estimate. Since noise in radiation spectroscopy is drawn from a Poisson distribution, a square root relationship is often used for the background noise term. Thus, we obtain the common convention and definition of SNR in the radiation detection literature [32]:

$$SNR(M) = \frac{S}{\sqrt{B}} \quad (3.3)$$

Two fundamental ways to estimate the SNR of a measurement are anomaly detection and match filtering. In anomaly detection, the goal is to identify deviation from typical background variation using a learned background model, regardless of what type of source may be causing the deviation. In contrast, match filtering assumes a particular source template or design that the algorithm is trying to “match” against. By allowing the operator to specify a source template, match filtering can boost estimation of background and source components in relation to that particular source.

3.1 Anomaly Detection Approaches

The anomaly detection approach treats the problem of estimating source signal in a measurement as that of finding misfit to typical, expected data patterns. The approach makes no assumption about the type of source being sought. Instead, a model of typical background variation is constructed, and deviation of subsequent observations from the model is measured. The component of the observed spectrum that fits the background model serves as the explanation for the background, and the total remaining counts in energy bins serves as the estimate of source signal.

3.1.1 Principal Component Analysis

Principal Component Analysis (PCA) is an algorithm for explaining variation in data [7]. PCA finds a set of mutually orthogonal linear directions in potentially high dimensional data that maximally explain its variance. The projection onto the found “principal components” reveals the fit of the data to the general trend. The reconstruction error reveals anomalous behavior from general variation. Complete derivation of PCA mathematics is included in appendix A.5 for the reader’s convenience.

Many times in data analysis, the directions of variability do not have an easy physical interpretation when plotted. For radiation spectroscopy data, however, PCA’s results have clear interpretation that help both scientific understanding of data trends and engineering of filtering mechanisms to perform background subtraction.

PCA, when applied to radiation spectra collected in typical urban scenes, captures the principal directions of variation in the main spectral shapes of background. A large percentage of the variability in the data can be accounted for by a few top principal components. The scree plot in Figure 3.1 shows how much variability is explained for different choices of number of principal components in the Sacramento data set. For the Sacramento data (and similar data sets), $N = 5$ is a good empirical choice for the number of principal components [45].




Figure 3.1: PCA Scree Plot

The PCA algorithm has clear, scientific interpretation on the radiation data. The first three principal components (PCs) are shown for the Sacramento data in Figure 3.2.




Figure 3.2: First Three Principal Components of Sacramento Background Data.

PC 1 correlates with mean counts. PCs 2-3 show systematic “buldges.” These are caused by temporal fluctuations due to temperature, weather, and locality of the sensor. These systematic variation effects are captured by PCA and can be subtracted out to find source signal.




Figure 3.3: PCA Example

The matrix of top k (we choose $k = 5$) PCs, P_k , can be used to estimate SNR of new measurements. The data projection of a new measured spectrum M on the learned PCA basis is given by $T_k = MP_k$. The total sum of squared reconstruction error over all bins serves as the source estimate, $S = \|M - T_k P_k^T\|_2$. The spectral anomaly detector treats the background estimate, B , as proportional to the total counts in the spectrum and simply uses it as a normalizer. The SNR score for the measurement is $SNR(M) = \frac{S}{\sqrt{B}}$.

Figure 3.3 illustrates using PCA to estimate background and source components. The observation (blue) is the original measurement. The observation is projected onto the background basis learned by PCA (red). The projection amount is calculated, and the remainder in the bins after subtracting out the projection (green) serves as the source estimate.

3.1.2 Spectral Anomaly Detector

The spectral anomaly detector is a PCA variant that is popular in the radiation sensing literature [45]. By conditioning PCA based on properties of radiation data, the spectral anomaly detector is one of the standard benchmarks currently available for SNR estimation in the domain.

The spectral anomaly detection method is a null-space anomaly-detector that performs Gaussian PCA-type background subtraction. It essentially removes the best-fit linear combination of “principal spectra” learned from scaled correlations between energy bins across training measurements. Unlike standard PCA, however, Singular Value Decomposition (SVD) is performed on the correlation matrix as opposed to the covariance matrix. A general rule of thumb with PCA is to use the covariance matrix version when the variable scales are similar and the correlation matrix version when the variables exist on different scales. The use of the correlation matrix standardizes the photon count data, ensuring the variables exist on the same scale. In addition, a Poisson noise model is used to create the SNR score instead of a Gaussian one. For a thorough explanation of the algorithm’s mathematics, the reader is invited to read appendix A.6.

3.1.3 Poisson Principal Component Analysis

One of the contributions of the thesis is to apply Poisson Principal Component Analysis (Poisson PCA) to the nuclear threat detection problem space. Standard PCA projections are computed assuming the data is Gaussian distributed. However, spectroscopy data is often better approximated using a Poisson process. Poisson PCA is an extension of PCA that allows Poisson-type loss instead of Gaussian sum-of-squares error used by traditional PCA [11]. Poisson PCA gives a space of typical background spectra which are always nonnegative (unlike regular PCA). Furthermore, accumulations of measurement deviations in unlikely energy bins are less likely to over fit. For a full formal description of the mathematics, see appendix A.7.

The Gaussian approximation to a true Poisson distribution is inadequate for situations where photon count rates from the background or the source may be low. This situation occurs for many real-world scenarios. The radioactive source being pursued may be faint or occluded, measurement time may have been limited, and/or sensors used for data collection may have been limited in sensitivity. In these cases, Poisson PCA can be advantageous in modeling the spectroscopy data over typical Gaussian PCA methods. Discussion of such scenarios and experimental results is provided in Chapter 7.

3.2 Match Filtering Approaches

Match filtering contrasts with anomaly detection approaches in that a source of a particular type is sought. The algorithm has access to a source template which it can use to “match” against. Because of this, the algorithm can be made more sensitive to energy regions and subspaces of the data in which the source signal is likely to be detected.

3.2.1 Correlation Match Filter

The classic match filter is the correlation based match filter. The filter measures the correlation between a new sensor observation vector R and the source template T :

$$\text{corr}(R, T) = \frac{\text{cov}(R, T)}{\sigma_R \sigma_T} = \frac{E[(R - \mu_R)(T - \mu_t)]}{\sigma_R \sigma_T} \quad (3.4)$$

A slightly more advanced match filter is one that uses a PCA model to subtract out the top principal components of variation before correlating with the source template.

3.2.2 Energy Windowing Regression

More powerful match filtering approaches can be derived using regression techniques. One key approach is Energy Windowing Regression [46].

Energy Windowing Regression divides the overall set of energy levels of a spectrum into two disjoint sets: a set of “feature-selected” energy bins where the source is likely to appear (denoted y) and a set of “non-feature-selected” energy bins where the source is unlikely to appear (denoted x). The background levels in the set y are estimated from the photon counts in the set x , which is presumed to be mostly background. Energy Windowing Regression uses a regression model to predict the sum of background counts in the feature-selected bins from the background levels in the non-feature selected bins. Figure 3.4 shows an example division of an energy spectrum into x and y subsets.




Figure 3.4: Example of Regression Match Filter

To find x and y given a source template and a mean background spectrum, the Energy Selection Heuristic is generally employed as a preprocessing step. Details of this heuristic for energy selection are provided in appendix A.3.1. This algorithm can be used to find the sets x and y for a particular source design. The heuristic ensures that $\text{length}(x) + \text{length}(y) = (\text{length of original spectrum})$.

Given these sets, a variety of regression estimators can subsequently be used to facilitate the prediction of counts in y from the counts in x . The most simple is the Simple Linear Regression estimator in which the sum of counts in the non-feature selected bins is regressed to predict the sum of counts in the feature-selected bins. This is a one-dimensional linear regression where only the sum of counts in x is used to predict y , regardless of the energy levels of bins in x .

A more expressive model is the Least Squares Estimator that uses not just the sum of counts but the full data matrix X , taking into account the number of photon counts that fell in different energy levels of the non-feature-selected energy bins (x):

$$\hat{B} = (X^T X)^{-1} X^T y_{data} \quad (3.5)$$

Here X is a $N \times |x|$ matrix of training samples where the data features are photon counts in non-feature-selected energy bins and y_{data} is the $N \times 1$ vector of ground truth total counts in feature-selected energy bins of training samples. Once trained, the $\hat{B}$ regression matrix (in this case, a $|x| \times 1$ vector) can be used for prediction.

Given a new observed spectrum vector, M , define M_W as the $|x| \times 1$ vector of photon counts of feature-selected energy bins in M and $M_{\bar{W}}$ as the $|y| \times 1$ vector of photon counts of non-feature-selected energy bins in M . Note that $\text{length}(M_W) + \text{length}(M_{\bar{W}}) = \text{length}(M) = 128$. The scalar background estimate is given by $B = \hat{B}^T M_{\bar{W}}$, the regression prediction of the background amount in the feature-selected bins from the background in the non-feature-selected energy bins. The scalar source estimate is given by $S = |M_W|_1 - B$ where $|M_W|_1$ is the scalar sum of total photon counts in the feature selected energy bins. The SNR score is formed as $SNR(M) = \frac{S}{\sqrt{B}}$.

To add regularization, sometimes a Ridge Regression estimator can be used in lieu of standard Least Squares:

$$\hat{B} = W_t (X (X^T \bar{W}_t X - \lambda I)^{-1} X^T) \quad (3.6)$$

W_t is the 1×128 indicator vector of feature selected bins, and $\bar{W}_t = \text{diag}(1 - W_t)$. The ridge regression process can help condition the regression in the case of large amounts of data points or features. This is true especially with tuning of the λ parameter which allows for regularization, sometimes yielding more compact models.

Chapter 4

Sensor Fusion with Bayesian Aggregation

After training a measurement SNR estimator, a sensor model can be built that takes into account the variability of estimates. The SNR estimation process is naturally imperfect, and the variation in the process can be taken into account by building probabilistic models of Signal-to-Noise Ratio (SNR) score distributions. BA allows for building a probabilistic, nonparametric sensor model robust to different types of variation in estimation. It is also used to fuse sensor measurements together in a computationally tractable fashion to detect sources of interest in a noisy environment.

4.1 Algorithm Description

BA uses a “validation” subregion of field data to assemble distributions of measurement scores for positive observations (source-injected background data) and negative observations (assumed pure background data). The measurement score distribution for negative data forms the null distribution (“ H_0 ”), while the measurement score distribution obtained for the positive observations serves as the alternate distribution (“ H_1 ”). The scoring function used for sensor measurements is a domain design choice. Distributions can be one or higher-dimensional and parametrized by different property variables. Here we describe BA sensor models developed for the radiation domain.

In primer models in the radiation domain, each distribution is 2-dimensional, parameterized by the SNR measurement score and a total source exposure statistic computed from sensor velocity, source intensity, and relative locations of the source and the sensor. The source exposure statistic is a multiplier for the Poisson parameters of injected source counts, taking into account all these factors. For a detector of negligible volume approximated as a point, the source exposure statistic can be defined as proportional to $\int dt/R^2$, where t is time, R is the distance between source and detector, and the integral is computed over the duration of the measurement as the detector moves.

The use of the point detector model allows for a simple exposure statistic that BA models can be parametrized by and that can be easily computed. Note that the use of the point detector exposure is a simplification. Real radiation detector systems can be large and integration over the area of detector surface would be necessary to improve the detector model (as well as incorporating

the dynamics of radiation sources).

Non-parametric density estimators (i.e. Histogram estimators, Kernel Density estimators [29]) can be used to assemble distributions for the null and alternate distribution. Non-parametric estimation helps account for the highly variable and complex empirical distribution of expected SNR as a function of exposure, affected by uncertainties in location, geometry, spectra of nearby nuisance sources, occlusion effects, and environmental conditions.

The SNR scores and total exposure values for positive observations are computed for each observation in validation data, and density estimation is applied to obtain a two-dimensional matrix of probability density values, with axes given by SNR score and total exposure statistic. For each value of the exposure statistic, the corresponding row of the matrix gives the estimated SNR score distribution. The density matrix is the basis for our estimate of $P(D|H_1)$, the probability of observing a set of given measurements D under the alternate hypothesis that a source of a particular intensity and type is present at a particular location. Similarly, a density matrix for negative observations produces an estimate of $P(D|H_0)$, the probability of observing given measurements under the null hypothesis that there is no source in sight.

Figure 4.1 shows example alternate and null hypothesis models obtained from field data for representative anomaly detector and match filter SNR scores.




Figure 4.1: BA Alternate and Null Sensor Models

The alternate models (Figure 4.1a and Figure 4.1c) show that as exposure increases, so does expected SNR (though there is copious variation in how it does). The null models (Figure 4.1b and Figure 4.1d) are flat with respect to exposure since there are no sources of interest to be exposed to. There is structure to these distributions (such as the intricate double tail behavior on the match model) that is not easy to capture directly using parametric models commonly used for spectroscopy. However, with sufficient data support, BA can capture these details in its nonparametric models.

The next step of BA is to spatially combine evidence as it is collected and to model probabilities of various source hypotheses across the environment. Alternate source hypotheses H_1 state that a particular location in the environment contains a source of a particular intensity and type, and null hypotheses H_0 states that no source with those parameters is present. For a given terrain, the scene can be covered with a set of hypothetical source locations (e.g. distributed over a regular planar grid). As new measurements are collected and added to the overall data D , BA maintains and updates estimates of the probabilities $P(H_1|D)$ for each source hypothesis and each null hypothesis $P(H_0|D)$.

The exposure and SNR are estimated for individual measurements D_j . These serve to index into the density estimation function and look up the probabilities $P(D_j|H_0)$ and $P(D_j|H_1)$. For a particular hypothesis $\mathbf{H}$ (either a H_1 or H_0), BA is motivated by the following equation for posterior probabilities $P(\mathbf{H}|D)$, assuming conditional independence of measurements:

$$P(\mathbf{H}|D) \propto P(\mathbf{H}) \prod_{D_j \in D} P(D_j|\mathbf{H}) \quad (4.1)$$

where $P(\mathbf{H})$ is the prior probability (belief) assigned to $\mathbf{H}$. For sufficiently low values of source exposure statistic, $P(D_j|H_1) \simeq P(D_j|H_0)$. This motivates the following algebraic manipulation when evaluating the Likelihood Ratio between H_1 and H_0 :

$$\begin{aligned} \frac{P(H_1|D)}{P(H_0|D)} &\simeq \frac{P(H_1)}{P(H_0)} \prod_{D_i \in D} \frac{P(D_i|H_1)}{P(D_i|H_0)} \\ &\simeq \frac{P(H_1)}{P(H_0)} \prod_{D_k \in D \wedge D_k \in D'} \frac{P(D_k|H_1)}{P(D_k|H_0)} \prod_{D_j \in D \wedge D_j \notin D'} \frac{P(D_j|H_1)}{P(D_j|H_0)} \\ &\simeq \frac{P(H_1)}{P(H_0)} \prod_{D_k \in D'} \frac{P(D_k|H_1)}{P(D_k|H_0)} \end{aligned} \quad (4.2)$$

where $D' \subset D$ is the subset of measurements whose value of source exposure statistic is sufficiently high for a source location. The algebraic manipulation shows that, for a new measurement, only hypothetical source locations within a particular distance of it (e.g. 20m) need to have their posterior scores updated as the measurements with low exposures cancel from the likelihood ratio. For each new measurement, kd-trees [42] efficiently find the set of hypothesized source locations locally affected by the measurement. This enables efficient and more scalable computation of Bayesian probabilities than cited previous Bayesian methods.

Figure 4.2 shows an example result of BA. Figure 4.2a shows physical locations where measurements were made and the location of an injected point source. Figure 4.2b shows the mapping of threat probabilities by BA. It shows correct detection and localization of the injected synthetic source.




Figure 4.2: Example data processing through BA. Figure 4.2a shows input data: a raw trajectory with an injected point source affecting measurements. Figure 4.2b shows the output: a resulting BA threat probability map. The red region in Figure 4.2b indicates the correctly detected radiation source.

4.2 Experimental Setup

Using the point source simulator (described in appendix A.4) and a geographic subset of the Sacramento data, an isotropic source with a fixed intensity was repeatedly injected at random locations in the subregion off the road. Source locations were restricted to be at least 10m away from the centers of the roads to simulate realistic road-side source conditions and to avoid creating trivial cases where a source was very close to a measurement taken on the road. Figure 4.3 shows an example of our evaluation setup




Figure 4.3: Setup for experimental evaluation

A true positive is defined as the top scoring grid point within 40m of the true injected source

when injection of source counts is performed. A false positive is defined as the top scoring grid point in the same 40m disc around the source location when injection is not performed. Detection success is measured with Receiver Operating Characteristic (ROC) curves using the true positives and false positives collected over a set of source location simulations. The 40m distance was used to avoid road-side degeneracy in localizing a source due to symmetry effects. On a perfectly straight road, the trajectory does not give information about which side of the road the source may be on. Thus, if the source is injected 10-20m off the road, the top scoring hypothesis for the source is degenerate to a maximum of 40m around the trajectory.

Different versions of BA were benchmarked against each other as well as in comparison to an alternate method of aggregation currently fielded called the Weighted Combining (WC) method (described in section 2.3). Accuracy of inference of true parameters of the source is summarized in confusion matrices, parameterized with a particular setting of the false positive rate (FPR). In our experiments, we used $\text{FPR} = 0.01\%$, and therefore the numbers reported in confusion matrices reflect the fraction of correct inferences over the total number of experiments at $\text{FPR} = 0.01\%$. $\text{FPR} = 0.01\%$ is a commonly used metric since the beginning of the ROC curve ($\text{FPR} = 0.00\%$) is prone to noise and less stable than the $\text{FPR} = 0.01\%$ mark. In our set of at least 1000 false positives, $\text{FPR} = 0.01\%$ was always a stable cutoff for comparison. This cutoff is equivalent to the average interval between consecutive false alerts being 2hrs 46min 40sec of system operation.

4.3 Results and Discussion

10,000 instances of the same test area of the city, each with a single, simulated, fixed intensity source (source count rate = 255 counts/sec, randomly placed in each instance 10-20m from the trajectory), were created. The WC and BA were compared with match filter and anomaly detection SNR estimators. Figure 4.4 plots ROC curve performance of algorithms using the previously described evaluation scheme. Each test instance provides one true positive (when injected) and one false positive (when not injected) example.




Figure 4.4: ROC curve results for algorithms.

In our experiments, match filtering outperforms anomaly detection with regularity simply

because it is provided with the exact knowledge of the types of sought sources. The ROC curve indicates that BA outperforms the WC method no matter which SNR estimator is used. BA captures the variability in SNR score as a function of expected exposure non-parametrically for each estimator. It can thus outperform the WC method which makes a Gaussian assumption when combining the SNR scores. The WC method cannot capture the more intricate structures in the distribution of SNR scores which BA can.

Furthermore, the WC method is affected by occasional nuisance sources in the data which can increase the WC source estimate and degrade the performance of this method. The WC method is susceptible to false positives resulting from inaccuracies in source amount estimation. The BA algorithms, however, use distributional models of SNR scores (robustly estimated with large amounts of data) that also consider the exposure statistic to enable dismissal of momentary spikes in the SNR if the expected empirical trend for source exposure is not consistently followed.

4.3.1 Incorporation of Source Intensity and Type Hypotheses in BA

The source location hypothesis space of BA can be extended to incorporate source intensity and source type variations. Several methods of including source intensity and source type information in BA were benchmarked.

The simplest method is BA-Specific which trains a separate model for each discrete setting of a parameter. If there are four expected source intensities, $\{I_1, \dots, I_4\}$, there are four distinct alternate hypothesis BA models $\{M_{I_1}, \dots, M_{I_4}\}$, each trained on data for one specific intensity parameter setting. BA-Specific methods assume $P(H_1|D) = P(H_1|D, M_{I_j})$. BA-Agnostic trains a single model learned from data containing all possible discrete settings of the parameter. Thus, there is only a single BA-Agnostic model, $M_{I_1, \dots, I_4}$, learned from a data set that agnostically combines data from all intensities. BA-Agnostic assumes $P(H_1|D) = P(H_1|D, M_{I_1, \dots, I_4})$.

Posterior distribution information from an ensemble of BA-Specific models can also be used to detect a source. BA-Max maximizes over the posterior intensity hypothesis distribution to score an alternate hypothesis:

$$P(H_1|D) = \max_j P(H_1|D, M_{I_j}) \quad (4.3)$$

BA-Marg, in contrast, marginalizes over the posterior intensity hypothesis distribution to score an alternate hypothesis:


$$P(H_1|D) = \sum_j P(H_1|D, M_{I_j}) \quad (4.4)$$

For intensity modeling experiments, we used four different settings of the intensity multiplier, increasing in a geometric series, to simulate increasingly pronounced sources. Sources are well in the tolerance of background. At the lowest intensity setting, source stand-off count rate at the distance of 10m is 150 photons per second, while mean background count rate observed in this subregion of the city without synthetic injections is 1,263 per second with standard deviation of 267 photon counts per second. All conversions to source count rate for the four chosen intensity multipliers are shown in Table 4.1.


The match filter estimator was used for all experiments. In the energy bins selected by the match filter estimator, the stand-off count rate at 10m is shown in the “Source Window Count Rate” column of Table 4.1. The mean background count rate in the source window was 125 per second with standard deviation of 31 counts per second.

Table 4.1: Intensity Multiplier to Count Rate Conversion


Intensity	Source Count Rate	Source Window Count Rate
I1	150 counts/sec	31 counts/sec
I2	196 counts/sec	77 counts/sec
I3	255 counts/sec	101 counts/sec
I4	331 counts/sec	131 counts/sec




(a) Results for SRC=i1



(b) Results for SRC=i2



(c) Results for SRC=i3



(d) Results for SRC=i4

Figure 4.5: Intensity Modeling ROC Results

Four sets of 1,000 simulated source-injected worlds were prepared, each set having a different

source intensity. Figures 4.6a-4.6d show ROC detection results for the four intensity settings.

Since the WC method does not specifically model source intensity, the BA methods trained with the correct intensity model outperform it easily. BA-Marg performs about as well as BA-specific models trained on the true source intensity – for all four intensity settings. If intensity is unknown *a priori*, BA-Marg is the theoretically optimal detection algorithm.

We extended this experiment to the case where the intensity of the source was treated as a real number, and the four specific intensity models were created over real ranges of possible source intensity, rather than single known values. The four intensity ranges used were {100-150 counts/sec, 150-196 counts/sec, 196-255 counts/sec, 255-331 counts/sec}. The code can work with arbitrary intensity ranges. The test source intensities were generated at random in the full range of 100 counts/sec-331 counts/sec, so no algorithm tested the exact intensity of the source. Figure 4.6 shows the ROC performance of BA-Marg (marginalizing over the BA-specific models for intensity ranges) in comparison to WC in this case of real-valued, completely unknown source intensity. Results are summarized in the four different setting ranges of true source intensity.




Figure 4.6: Intensity Modeling ROC Results

We also performed experiments with source type variations. Three different source templates (shown in Figure 4.7) were chosen for injection.







Figure 4.7: Source templates injected into background.



(a) Source Type 1 ROC Results



(b) Source Type 2 ROC Results



(c) Source Type 3 ROC Results

Figure 4.8: Source Type Modeling ROC Results

Figures 4.8a-4.8c show ROC detection trends in detecting the three different source types. BA-Marg and BA-Max have little loss when compared to the BA-Specific model trained for the true source type. BA-Marg is the method of choice since marginalization over the intensity and

source type variations provides not only the theoretically optimal detection power but promising empirical results as well.

4.4 Inference of Source Properties

A major capability of BA not existing in previous methods in the literature is the capability to simultaneously infer source properties while detecting it [67]¹. In addition to helping provide robust detection and localization of a radioactive point source, the posterior probabilities from the BA models can be used to infer the source intensity and source type parameters for the source. To evaluate accuracy of such inference capability, we compared the parameter inferred at the hypothesized source location at 0.01% false positive rate for each of the evaluated synthetic environments.

Table 4.2: Confusion matrix for intensity inference: BA-Marg / BA-Max.

True Intensity →	I1	I2	I3	I4
Correct Inference	35.4% / 36.2%	20.7% / 21.2%	21.4% / 21.5%	46.8% / 47.0%
Incorrect Inference	23.4% / 24.3%	47.4% / 48.2%	55.3% / 57.2%	38.7% / 39.1%
False Detection First	41.2% / 39.5%	31.9% / 30.6%	23.3% / 21.3%	14.5% / 13.9%

Table 4.2 shows intensity inference results for our two methods of inferring a parameter - maximizing over the BA-specific posterior hypothesis distribution (BA-Max) and marginalizing over the posterior hypothesis distribution (BA-marg). BA-Max slightly outperforms BA-Marg in intensity inference, though differences are not significant. Intensity is a quite confusable parameter since there is a fair bit of location leeway in detecting a source within a 40m radius. A detection algorithm can model a source as either being a strong source far away or a weak source close by and still succeed in detecting it, though yield incorrect identification of source intensity.

Table 4.3: Confusion matrix for source type inference: BA-Marg / BA-Max.

True Source Type →	Source 1	Source 2	Source 3
Correct Inference	57.8% / 57.6%	77.5% / 77.5%	87.0% / 86.3%
Incorrect Inference	2.50% / 2.50%	0.80% / 0.80%	5.40% / 5.20%
False Detection First	39.7% / 39.9%	21.7% / 21.7%	7.60% / 8.50%

¹Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Simultaneous Detection of Radioactive Sources and Inference of their Properties. IEEE Nuclear Science Symposium 2013.

Table 4.3 shows source type inference results for BA-max and BA-marg. Results are not significantly different between algorithms, though BA-Marg slightly outperforms BA-max on the task. Both algorithms produce accuracies within the 58-87% range, showing BA’s capability to tell different source types apart robustly.

4.5 Conclusions

Our results indicate that the BA algorithm can boost the detection of radioactive sources as well as simultaneously help determine their characteristics such as location, intensity, and source type.

With regards to detection, BA can boost the performance of anomaly detector and match filter estimators commonly used by the radiation sensing community. BA captures the distribution of expected SNR scores as a function of exposure without making parametric assumptions about expected distribution. This approach allows for the dismissal of background nuisance sources that might otherwise thwart the capability of such SNR estimators in practice.

With regards to determining properties of detected sources simultaneously with its detection, we have shown how intensity and source type information can be incorporated into BA models. Our experimentation with various methods suggests that our method of choice, BA-Marg, can simultaneously provide robust detection of sources of different intensities and source types (as evaluated by ROC curves) while inferring their characteristics (as evaluated by accuracy on confusion matrices).

BA’s algorithmic enhancement of detection capability enables mobile radiation detection systems to provide more sensitive and precise nuclear threat detection. Systems using BA will be able to more accurately classify threats from non-threats while lowering the false alarm rates. BA’s source parameter inference capabilities allow law enforcement officers to have real-time knowledge of the likely properties of threatening (or non-threatening) radiation sources discovered in the field, allowing for informed responses.

Chapter 5

Adaptive Post-processing of Hypothesis Scores

BA can be powerful with a fixed, descriptive, and well-chosen hypothesis space. BA can be made more powerful via methods of post-processing BA scores under multiple resolutions and with different parameter settings. In this chapter, we develop algorithms for adaptively post-processing BA hypothesis scores to improve system performance.

In the radiation domain, BA scores can be post-processed under multiple resolutions and different parameter settings to try to improve detection of sources and inference of their properties. The generalized radiation source search problem involves searching over different parameter settings (such as location, intensity, and source type) for best explaining the presence of a radioactive point source (or lack thereof) in data collected by a mobile spectrometer. In this chapter, we build post-processing techniques for detecting a source, taking into account the following parameters:

1. (x,y) geographic location of radioactive point source
2. angular wedges of possible static anisotropy (e.g. partial occlusion of source).
3. contiguous temporal subsets of measurements during which a dynamic occlusion can occur.

To improve estimates of the source location parameter, a fixed hypothesis grid of BA scores can be expanded iteratively using adaptive grid approaches. The expansion need only occur in geographic areas of the world predicted to possibly contain sources. Data structures (such as R-trees) can be used to maintain an adaptive resolution of grid point hypotheses in a computationally efficient manner.

Furthermore, the space of possible static and dynamic occlusions can be searched efficiently via appropriate pruning strategies. This study develops and experimentally verifies a branch-and-bound method for searching over static and dynamic occlusion scenarios.

5.1 Adaptive Grid Bayesian Aggregation

The adaptive grid variant of BA extends the static grid version to allow adaptively growing the grid hypothesis space, where necessary, to test promising source locations [31]¹. The key is to build expansion rules to dynamically adjust local resolution of source parameters (e.g. source location, intensity, etc) for BA hypotheses so that resolution is increased only in regions likely containing sources. Expanding unnecessarily adds computational burden that must be avoided. The strategies developed for efficient expansion of search nodes are similar in spirit to Gradient Descent and Nelder-Mead optimization.

Figure 5.1 shows an example of an adaptive grid BA in action.




Figure 5.1: Example of Adaptive Grid BA. The algorithm grows the grid adaptively around the detected point source to better estimate the true source location.

5.1.1 Algorithm

A BA grid (at a particular static resolution) can be represented as a set of squares where each square has 4 grid points as its vertices. Adaptive Grid BA adds additional BA hypothesis grid points in square regions of the static grid where significant variation exists between the scores of its vertices. Grid resolution is thus iteratively increased for squares where the min and max BA scores of the square's vertices varies more than threshold.

If the minimum BA score (out of the square's four vertices) is called m and the maximum BA score on the square is called M , the following decision rule is used to decide whether to increase the square's resolution:

$$m + C * (M - m) \geq T \quad (5.1)$$

where C and T are parameters. T is a cutoff threshold for statistical significance. In the radiation threat detection problem, T is generally set to the BA score at the 0.01% False Positive Rate. C is often set to 1.

¹Huggins, Peter, Prateek Tandon, Artur Dubrawski, Simon Labov, and Karl Nelson. Dynamic Placement of Sensors for Rapid Characterization of Radiation Threat. DTRA Review July 2013.

When expanding a square, unnecessary expansion can be avoided by increasing grid resolution only in the direction of the score increase. Figure 5.2 show the expansion rules (geometrically) for the two possible cases: expansion along the square edge and expansion along the square diagonal.




Figure 5.2: Expansion Rules for Adaptive Grid

5.1.2 Experiments and Results

1,000 test blocks were prepared, each with an injected point source. The Adaptive Grid BA evaluated source location hypotheses at an initial grid resolution of 4m and used the developed expansion strategy to increase grid resolution to 2m and 1m in areas likely to contain sources. Figure 5.3 shows the results of using Adaptive Grid BA to localize a radioactive source in comparison to BA methods that use only fixed grid resolutions of 4m, 2m, and 1m respectively.




Figure 5.3: Adaptive Grid Source Localization Result

The results show that the Adaptive Grid BA algorithm can get roughly the same localization accuracy as the 1-m fixed grid algorithm. It can do so, however, at a fraction of the computational cost. Table 5.1 compares the grid size of the tested Adaptive Grid BA hypothesis space in comparison to grid sizes at fixed world resolutions.

Table 5.1: Comparison of Grid Sizes

Grid Resolution	Number of Grid Points	% vs. 1m Gridding
4-m	11,501	6.25%
2-m	46,507	25%
1-m	184,207	100%
Adaptive Grid	57,027	31%

The Adaptive Grid BA only uses 31% of the hypothesis space of the 1-m grid resolution but obtains comparable localization accuracy. Adaptive Grid BA effectively obtains the boost of running BA at a finer grid resolution without incurring the computational cost of the full grid search.

5.1.3 Efficient Maintenance of Hypothesis Space

BA assumes that source location hypotheses far from new measurements do not affect the probabilities at those locations significantly. When the BA system processes new measurements, the source location hypothesis spatially close to the measurement are queried from a database to have their BA scores updated. When the BA hypothesis grid is large, locating source hypothesis locations can be a potentially time-consuming database operation.

In vanilla BA, kd-trees are used to store source location hypotheses. When new measurements arrive, the kd-trees facilitate efficient geographic search in $O(\log n)$ time to find the affected grid points requiring update. Kd-trees function well for vanilla BA because vanilla BA uses a static grid. In this case, the kd-tree structure needs only one build operation and does not need to ever add new grid points. An adaptive grid algorithm, however, expands and contracts grid resolution dynamically and thus requires constant addition and maintenance of new grid location hypotheses. Kd-trees are limited in that they cannot efficiently add new hypothetical source location grid points. Adding a new point to a vanilla kd-tree requires the full time to build the structure in $O(n \log n)$.

R-trees [42] have comparable worst-case build and access times compared to kd-trees, and have been observed to be faster in practice. Furthermore, R-Trees can add new grid points efficiently. An implementation of R-trees has been compared to the previous kd-tree implementation [31]², using the source location hypothesis space used by BA on our benchmark Sacramento data.

²Huggins, Peter, Prateek Tandon, Artur Dubrawski, Simon Labov, and Karl Nelson. Dynamic Placement of Sensors for Rapid Characterization of Radiation Threat. DTRA Review July 2013.

Figure 5.4 shows runtime benchmarks for building and querying grid points with the R-tree software library in comparison to the previously used kd-tree software library.




Figure 5.4: Build and Query Time Comparison for Tree Data Structures

As Figure 5.4a and Figure 5.4b show, the R-tree software library provides a speed up for maintaining the source hypothesis grid. The R-tree library is faster both in building the tree structure and querying for spatially proximal grid points when a new sensor observation is received. Running times for R-tree queries are not only faster but also have less variance than the kd-tree queries (which can have significant benchmark performance differences from run to run). Having a stable algorithmic runtime is important in real-time source detection since occasional algorithmic delays may also delay real-time processing in a time critical threat detection scenario.

R-trees also allow fast dynamic insertions and deletions of points, whereas kd-trees need to be periodically rebuilt when points are added or deleted (causing delays in a real-time system). Thus R-trees are a necessary replacement of kd-trees in Adaptive Grid BA to facilitate dynamic generation and pruning of source hypotheses. R-Trees are the data structure of choice when working with Adaptive Grid BA.

5.2 Branch and Bound for Handling Dynamic Occlusions

Previous BA variants do not consider source occlusions. In this section, we extend BA to deal with cases of dynamic and static occlusion during observation of a radioactive source [31]³. Static occlusions, parameterized by wedge start and end angles, can be applied to radioactive sources to simulate partial shielding of source. An example static occlusion wedge is shown in Figure 5.5. Dynamic occlusions, further parameterized by start and ending time of occlusion

³Huggins, Peter, Prateek Tandon, Artur Dubrawski, Simon Labov, and Karl Nelson. Dynamic Placement of Sensors for Rapid Characterization of Radiation Threat. DTRA Review July 2013.

simulate momentary shielding of a source in time (for instance, by a truck passing between sensor and source and temporarily blocking view of the source).




Figure 5.5: Figure shows an example of a static occlusion scenario. Sensor observations in green are occluded by a wedge with particular start and end angles and receive no injection, while observations in red are unoccluded. Dynamic occlusions, which occur at particular points in time, can further complicate analysis of a wedge depending on the time of the observation. In a dynamic occlusion scenario, some of the green measurements may actually be red, necessitating a search for the source in both space and time dimensions simultaneously.

5.2.1 Algorithm

Searching over all possible wedges and time periods naively gives an $O(n^4)$ algorithm. This naive algorithm is much too slow for practical use. Branch and bound techniques can be used to speed up the search using the following strategy. At each time step, the current top scoring BA hypothesis is maintained. The search is pruned using the following bounds:

1. If the sum of all positive BA scores at a grid point is less than best BA score found so far, then the wedges or time intervals associated with the grid point need not be evaluated. Move on to the next grid point.
2. If the sum of all positive BA scores at a grid point plus the sum of all negative BA scores outside of the current wedge is less than the best BA score found so far, then the time intervals of the wedge need not be searched. Move on to the next wedge.

These bounds substantially reduce the search space but still find the optimal solution. It was not possible to evaluate the full empirical runtime of the naive $O(n^4)$ algorithm simply because the runtime exceeded the PhD student's available computational resources. Thus, this branch and bound formulation was strictly necessary to finish experiments in finite time for this thesis.

5.2.2 Experiments and Results

1,000 test scenarios were created, each with an injected point source. Two data sets were prepared from this underlying data: one with and one without dynamic occlusions around the point source. Three BA methods were benchmarked on the data: Isotropy BA that did not model occlusions, Static Anisotropy BA which modeled static occlusions, and Dynamic Anisotropy BA which modeled dynamic occlusions. Figure 5.6 shows the detection performance of the different methods on the two data sets.




Figure 5.6: ROC Detection Performance without and with occlusions

As Figure 5.6b shows, using a BA that uses branch and bound search over dynamic occlusion scenarios can give significant performance boost to the source detection problem when dynamic occlusions are truly present in the data. Figure 5.6a shows that when occlusions are not present in the data, the performance of all methods does not differ significantly. With the branch and bound approach, BA with the dynamic occlusion hypotheses runs 300x faster than a real time measurement cycle on a single CPU in a MATLAB prototype.

Chapter 6

Using Bayesian Aggregation to Guide Data Collection and Search Efforts

The BA algorithm allows for the maintenance and testing of many, multi-modal hypotheses about the world. This capability also allows for quantification of uncertainty with respect the world, and can be used to guide further data collection and search efforts aimed at efficiently reducing the uncertainty. In this section, we provide preliminary examples of how the BA hypothesis space and derived information can be plugged into route-planning efforts for mobile sensors. Developed methods are demonstrated on the radiation source search problem.

In the single agent source search problem, a single mobile sensor must find a static or mobile source. In the multi-agent case, a team of sensors is tasked with the same objective. The multi-agent case is generally considered computationally harder since the joint space of possible team actions, which can be quite large, must be searched over to find the optimal plan. In this chapter, we prototype example algorithms for both the single and multi-agent route-planning problems that can use BA to detect a static or mobile radioactive source in an urban scene.

6.1 Use of BA Information in Single Agent Source Search

BA information can be used in a single agent route planning framework to allow a single mobile radiation detector to find a static or mobile radioactive source [63]¹. In our test scenario, a radiation source exists somewhere in the environment, and the agent has to decide where to collect sensor spectra next in the environment to find it.

6.1.1 Single Agent Algorithm

In the radiation source search problem, the key is to cover as much important ground as possible. Having only a single mobile sensor means that, if there is uncertainty in the world about where sources may be, the uncertainty must be reduced quickly so that the true source may be detected efficiently.

¹Tandon, Prateek, Artur Dubrawski, Jeff Schneider, Adam Zagorecki, Simon Labov, and Karl Nelson. Machine Learning for Effective Nuclear Search and Broad Area Monitoring. ARI Annual Review 2011.

Entropy is a common way of quantifying uncertainty or ambiguity about the world with respect to multiple, competing Bayesian hypotheses. Reduction in entropy (also known as information gain) is adopted as our metric.

For a radiation source search space with k possibilities (where a possibility is a full description of a source with its parameters), there exist posterior probabilities $P(H_{0,k}|D)$ and $P(H_{1,k}|D)$. For each possibility k , the probabilities $P(H_{0,k}|D)$ and $P(H_{1,k}|D)$ form a probability distribution which sums to one. The total entropy of the current hypothesis space can be computed as:

$$E_{total} = \sum_{k \in Grid} E_k = - \sum_{k \in Grid} [P(H_{0,k}) \log_2 P(H_{0,k}) + P(H_{1,k}) \log_2 P(H_{1,k})] \quad (6.1)$$

The total entropy metric quantifies how much uncertainty exists in the world in relation to the null and alternate hypotheses for each possibility. When evaluating trajectories that a mobile detector should take, the goal is to find trajectories that maximize the total reduction in world entropy (or “information gain”) so that the source can be found most efficiently.

Two ways of estimating the information gain of a particular trajectory are tested. The first method is to sum the current entropies along the trajectory. This is a computationally efficient way to estimate the information gain, but it assumes that a pass through the trajectory will zero out all entropy on nearby nodes entirely. This assumption leads to only an approximate solution in practice. A more realistic way to estimate the expected information gain of traversing a trajectory is to use Monte Carlo simulation. The process is to:

1. Pick N consecutive points in the training data at random where N is the number of points in the trajectory.
2. Simulate traversing the trajectory in hypothetical simulation and estimate the information gained by the hypothetical traversal.
3. Report the expected information gain from the distribution produced by multiple simulation trials.

The Monte Carlo simulation approach helps provide more realistic estimates for the information gain. The Monte Carlo algorithm can, however, require many simulation trials to estimate information gain well.

BA allows incorporation of prior knowledge (either from secondary sensing modalities or other sources of information) into alternate and null hypotheses. Assigning a high alternate prior probability to a particular source location can be used to bias algorithms to further investigate that location. A high null prior probability encodes that a source was not found, so there may not be more need to search a region. Setting the priors on hypotheses can be used to bias algorithms towards source pursuit or exploration behaviors. By modulating the prior probabilities, the algorithm can automatically produce such behaviors.

The single agent algorithm also allows setting of the time horizon of forward lookahead as a parameter. Two time horizon parameter settings are tested. The first algorithm is “Greedy” at each time step, computing information gain using only immediate trajectories. The second algorithm type is “Exhaustive” and evaluates all possible trajectories T time steps ahead. Though

exhaustive planning requires more computational runtime than greedy planning, exhaustive planning can make more optimal decisions in efficiently searching for the radioactive source.

Figure 6.1 shows an example entropy reduction curve for an experiment running the Exhaustive algorithm with $T = 3$.




Figure 6.1: Example Entropy Reduction Curve for Exhaustive Algorithm

6.1.2 Experiments and Results

Algorithms were benchmarked on a data set of 1000 simulated worlds, each with a single injected radioactive point source. The key metric of success was time to detect the source. Activity Monitoring Operating Characteristic (AMOC) curves conveniently summarize the trade-offs between number of detections and the number of observations required to make the detection.

Figure 6.2 shows an AMOC comparison of Exhaustive ($T = 3$) with the greedy algorithm.




Figure 6.2: Source Detection Capability of Exhaustive ($T=3$) and Greedy Strategies

The results indicate that the Exhaustive strategy outperforms the Greedy strategy in number of observations required to locate the radioactive point source. This indicates that the Exhaustive

algorithm is making more optimal routing decisions than the Greedy algorithm. It can thus be more efficient in detecting radioactive threats.

Figure 6.3 shows the results of varying the prior parameter settings of source location hypotheses on average number of detections.




Figure 6.3: Effect of Setting Prior Probabilities on Threat Detection Performance

Setting a low threat prior probability encodes the assumption that the environment is mostly threat-less. This means the vehicle need not explore much once it has found something. Figure 6.3 shows that setting a low threat prior probability requires fewer observations to find the source because, once the vehicle finds a suspicious-looking area, it exploits its current knowledge to continually monitor the possible detection. In comparison, uniform priors encourage more exploratory behavior. With this setting, the algorithm will cover more geographic space to map radiation threats on the map more comprehensively (though it may miss a detection by not following up in monitoring a previously flagged location).

The experiments showcase the capability of BA algorithm information in helping plan routes for a single mobile sensor for subsequent data collection. The entropy metric, easily computed from the BA posterior probabilities, helps the operator decide where to search next for a radioactive source. The prior probabilities of BA hypotheses allow encoding additional secondary information into the search process.

6.2 Application of BA to Multi-Agent Radioactive Source Search

The single-agent approach in the previous section can be extended to allow a team of spectrometer-carrying mobile sensors to survey a city block for a mobile radiation source. In our scenario, there are N traffic vehicles to survey and M mobile sensors. One of the traffic vehicles is carrying a radiation source. The goal is to identify the mobile source as quickly as possible using the team of mobile sensors. The locations of all vehicles are presumed to be known at all times.

This problem setting is, in theory, significantly harder than the previously studied single agent problem for a few reasons. First, to achieve optimality, the large space of team plans must be efficiently searched at each time step to find the team plan that, when effectively executed, will minimize the time required to find the source. Second, the radiation source is now mobile and can be substantially more difficult to track than a static radiation source. Finally, the configuration


space of possible team vehicle location deployment and possible source locations is large and grows exponentially with the dynamic time evolution of the system. The main route planning issue becomes one of role allocation in a team: How should the surveillance vehicles be assigned to survey the vehicles in traffic so that mobile radiation source can be found efficiently?

6.2.1 Algorithm

A graph representation was built on top of the segmented city block to support reasoning over routes. The graph nodes represent (latitude, longitude) locations in the world. Edges represent actions vehicles can take. Actions on the graph allow vehicles to choose a speed of travel (from speed 1 to 5) or to stay at their current location (speed=0).

The BA detector algorithm maintains, for each traffic vehicle, the probability distribution vector $[P(H_1), P(H_0)]$. $P(H_1)$ is the alternate hypothesis probability that the traffic vehicle is a mobile source, while $P(H_0)$ is the null hypothesis probability that the car is not a mobile source. As sensor readings are taken within 20 meters proximity to a traffic vehicle on the road, each reading is scored as source-like or background-like by BA sensor models, and evidence is spatially aggregated.

At any given time step, the entropy of a traffic vehicle distribution can be calculated. The entropy metric quantifies uncertainty with respect to the planner's belief in each traffic vehicle's state as a mobile source. With enough sensor measurements in its vicinity, a traffic vehicle's entropy will approach zero as the detection algorithm decides whether the car is truly the mobile source or not. However, if a car's state is uncertain, the entropy will remain high, warranting additional observations. Figure 6.4a shows that the entropy of a standard vehicle in traffic is quickly zeroed out as a function of the number of measurements. However, the entropy of a mobile source shown in Figure 6.4b typically continues to be high. More observations will allow the entropy to drop to zero, and the algorithm will converge to the conclusion the car is a mobile source.



(a) Entropy of Typical Non-Source Vehicle

(b) Entropy of Suspicious Source-Like Vehicle

Figure 6.4: Entropy of Non-Source vs. Source-Like Vehicle

The key problem of interest is role allocation: to effectively allocate M mobile surveillance vehicles to N mobile surveillance targets on the road. Naively, the size of this space is too large to search exhaustively. Making different assumptions about the role allocation cost function can, however, aid scalability.

If one assumes a bipartite matching, one can cast the problem within an Assignment Problem framework and solve via specialized methods such as the Hungarian Algorithm [38]. In general, one can employ linear programming, mixed-integer linear programming, and binary integer programming to help solve the problem. For more details of our exploration of such techniques to tackle larger role allocation problems for source search, see [62] as we do not cover such cases here. We instead focus on the use of BA information in the planning frameworks.

Two different approaches to role allocation were tested in this study: search with point-scoring heuristics and search with an entropy propagation heuristic. Both approaches use the same general architectural strategy. In each approach, a global planner allocates targets from 1 to N for each law enforcement vehicle. A local planner then plans paths for each law enforcement vehicle to reach its target. Currently, local route planning is done via Dijkstra’s algorithm [12].

Point-scoring heuristics create a score for each car, and then compute a sample statistic over all cars to use a heuristic value. Point-scoring heuristics for assigning the nearest target in Manhattan (L1) distance and assigning the maximum entropy target were tested. These heuristics, though simple, can perform well for small numbers of surveillance targets. As the number of surveillance targets increases, however, point-scoring functions will not perform as well, necessitating design of a more advanced heuristic.

Inspired by value iteration, our idea was to use Entropy Propagation to propagate entropy in space and time to estimate expected reduction in entropy along a path. The idea is that even if an individual mobile sensor is technically allocated to investigate one vehicle, it may be able to obtain source detection information about many vehicles by traveling an appropriately planned route to its surveillance target. Optimizing the reduction of entropy along a surveillance path is thus the important criterion of the Entropy Propagation algorithm.

The Entropy Propagation algorithm is as follows:

1. For each target vehicle t , take its point entropy and propagate out based on a motion model.
2. Add all car entropy maps together to create overall entropy map.
3. For each law enforcement vehicle v and target vehicle t , there is a shortest path $p = \text{computePath}(v, t)$.
4. The score $S(v, t)$ of the assignment of v to t is the expected entropy reduced along p .
5. The best allocation of law enforcement vehicles to targets is the one that maximizes expected entropy reduction over the paths of the entire team.

Figure 6.6 illustrates an example of this process. Figure 6.5a shows the point entropy of a particular traffic vehicle, Figure 6.5b shows the entropy propagated in time due to the traffic vehicle’s motion model, and Figure 6.5c shows the combined entropy map of all surveillance targets. This entropy cost map nicely encodes expected reduction of entropy in time. By planning paths in the expected entropy cost map, one can hope to reduce real entropy in the world.




Figure 6.5: Illustration of Entropy Propagation Algorithm Steps

6.2.2 Results

Performance of heuristics was tested in a simulated setting with 4 non-source surveillance targets in traffic, 3 surveillance vehicles and 1 traffic vehicle carrying a radiation source (and thus being the mobile radiation source). This small-scale scenario models surveillance of a residential area. 30 simulation scenarios were run for 1000 time steps. Spawn locations and destinations of traffic on the road were generated from a uniform random distribution. The system dynamics described by the graph representation were applied to allow the cars to choose actions to stay or move. The source was made strong enough and 1000 time steps was long enough so that the source was caught by all algorithms. The algorithms differ in the time it takes to catch the mobile source.

Greedy and exhaustive versions of algorithms were compared. Greedy algorithms iterate through the list of law enforcement vehicles and assign the currently unassigned highest scoring target. Exhaustive algorithms score the team assignment vector overall and search over all possible assignments. Since our number of surveillance targets is small, exhaustive search in this

manner was feasible.

Several metrics were used to understand the results of simulations (shown in Figure 6.6).




Figure 6.6: Results of Small-Scale Simulation

The key metric is Time to Detect (TTD) which measures, as a function of detector threshold, the number of simulation time steps it takes to find the source. The detector threshold is the threshold on $P(H_1)$ to flag a car as a mobile source. The mean TTD is the time at which the mobile source is found is averaged over the 30 simulations, at various settings of the threshold. Figure 6.6a demonstrates that the entropy propagation methods dominate other methods in TTD. Statistical significance is verified by drawing nonparametric 95% bootstrap confidence intervals. The mean TTD is tightly within an envelope significantly outperforming the other algorithms, showing that the entropy propagation methods have the best TTD even when the data are slightly different.

A secondary metric involves the analysis of first true positives as a function of detector threshold. A true positive is defined as flagging the mobile source. Flagging a car that is not a mobile

source counts as a false positive. The first true positive metric measures the number of times (out of 30 runs), at a given threshold, the mobile source is flagged before other cars. Figure 6.6b shows that the exhaustive entropy propagation method flags fewer false positives than other algorithms in finding the source.

A final metric is the ROC curve shown in Figure 6.6c, which compares true positive rate to false positive rate. The entropy propagation methods outperform other methods at various settings of the false positive rate. Two important characteristics of ROC curves are where they start and how efficiently they reach a true positive rate of 1. The starting point states the true positive rate that can be achieved if the detector flags no spurious points. The efficiency of the ROC curve in reaching a perfect true positive rate denotes the false positive rate with which a perfect detection rate can be achieved. Both Greedy and Exhaustive entropy propagation methods perform well on both of these parts of the ROC.

6.3 Discussion and Conclusions

These algorithms illustrate the capability of BA information to be used in both single and multi agent route planners to find a radioactive source. The computed entropy maps from BA can help guide where subsequent sensor measurements should be taken by a team of mobile sensors, helping the team decide how to allocate its resources to facilitate the overall surveillance plan. In turn, new information from sensors also helps boost the accuracy of the subsequent BA posteriors that are obtained from the new data. Thus, BA helps the operators make decisions, and the operators helps BA become increasingly accurate in its predictions thereby creating a closed-loop system.

Chapter 7

Low Photon Count Signal and Noise Component Estimation

7.1 Introduction

The problem of Low Photon Count Signal and Noise Component Estimation is to succeed in detecting a radioactive source even when the effective photon count rate received by a sensor is especially low for either the background or source components. The low photon count rate scenario may occur for a variety of reasons due to properties of the source and/or of the sensor(s) used.

An important class of cases is where the photon count rate from the source may be low relative to the background, making estimation of source counts difficult. The radioactive source being searched for may be very weak, occluded by signal-attenuating obstacles (e.g. be inside a building), or be observable only far away in standoff distance from the sensor. If the source is deliberately engineered to be shielded, it may be directionally attenuated, leading to a loss in photon counts detected by the sensor. In this class of cases, the photon counts in the source component are expected to be low, though the background count rate may be as usual.

Another set of cases is where the photon count rate from both background and source are low. This can occur due to limitations of the sensor data collection process and/or hardware. For instance, a sensor may only have been able to briefly survey an area, collecting only a handful of photon counts from both background and source due to lack of data collection time. Additionally, sensors may also be limited in their intrinsic material detection capability and sensitivity. Sensors can vary in size, shape, and materials which affects their detection capabilities. Smaller sensors can be more cost-effective to deploy, but they will receive fewer counts from both the background and the source than larger sensors due to their smaller volume.

In all of these cases, the constraint of fewer photon counts in radiation spectra makes the detection problem harder. Fewer photon counts in either a source or background component generally means that each component is harder to estimate. An algorithm has less data to reason about the delimitation of signal components and is subject to greater possible estimation error. The goal is to develop algorithms that can extract the most useful information possible out of available photon count sensor data to succeed in detecting the source.

7.1.1 Typical Measurement SNR Estimation

Measurement SNR estimators (introduced in Chapter 3) are statistical estimators that seek to estimate the number of photon counts in a radiation spectrum due to background (B) and the number of photon counts due to the source (S). By estimating these two quantities, an SNR estimator can thus calculate its estimate of the Signal-To-Noise Ratio (SNR) of the spectrum as:

$$SNR = \frac{S}{\sqrt{B}} \quad (7.1)$$

Previously, we discussed typical anomaly detection and match filtering approaches to SNR estimation. The goal of anomaly detection is to flag spectra that are distinct from typical background. Anomaly detection typically assumes no knowledge of a source template and simply uses a model of expected background to measure the background (typical variation) and source (unusual deviation) components of a measurement. Match filtering, in contrast, allows the user to specify a specific source template of interest to match against spectrum observations. Both methods allow estimation of background and source components. Principal Components Analysis (PCA) and the Spectral Anomaly Detector [45] are common approaches to anomaly detection, while Energy Window Regression [46] is commonly used for match filtering on radiation data.

7.1.2 Poisson Modeling

spectroscopy data is often presumed to be created with a Poisson Process. However, the available anomaly detection and match filtering methods typically rely on Gaussian models of the data. In anomaly detection, both the vanilla PCA and Spectral Anomaly Detector methods assume a Gaussian distribution of the data and primarily find linear directions of maximal variance. Likewise, Energy Windowing Regression is typically based on a Linear Regression (also Gaussian) model of the data.

When the photon count rate for both source and background components is sufficiently high, a Poisson distribution can be approximated well by a Gaussian distribution. In these cases, the Gaussian modeling assumption does not invoke much loss. When the photon count rate for either or both components is low, however, the Gaussian approximation ceases to work as a sufficient model of the data.

The thesis develops original techniques to augment the estimation process with Poisson models to boost performance. Our study experiments with using Poisson Principal Component Analysis (Poisson PCA) [11] instead of standard Gaussian PCA to estimate source SNR using anomaly detection. Similarly, we experiment with using Zero-Inflated Poisson (ZIP) Regression [75] instead of Linear (Gaussian) Regression when estimating SNR with a match filter. Our results suggest that these Poisson-based methods can mine additional useful information from the data, unavailable in a pure Gaussian model, to help detect potentially harmful sources of radiation.

In addition, we use our Bayesian Aggregation (BA) algorithm to further improve detection performance. BA leverages spatial aggregation of multiple correlated radiation observations to robustly test source location hypotheses [67]. By aggregating evidence statistically over multiple

observations, detection performance is made more robust against possible spurious noise in the data that can cause false alarm.

7.1.3 Poisson PCA

Poisson PCA is an extension of PCA that allows a Poisson-type loss function instead of the sum of squared error loss function used by traditional PCA [11]. The sum of squared error loss function in traditional PCA imposes a Gaussian model on the data. Poisson PCA uses a Poisson-based loss function instead.

Standard PCA finds key linear directions of variation in the data and allows for negative components. Poisson PCA, in contrast, gives a space of typical background spectra which are always non-negative. Accumulations of measurement deviations in unlikely energy bins are much less likely to be over-fitted. We use the Poisson PCA formulation in [53] for our experiments.

7.1.4 Zero-Inflated Regression Poisson Match Filter

The standard Energy Windowing Regression match filter uses the Least Squares Estimator to derive source and background components, imposing a Linear (Gaussian) Regression model on the data. One can use a Poisson Regression model for the expectation instead:

$$E(y|x) = e^{\lambda^T x} \quad (7.2)$$

where y is the predicted sum of counts in the energy window, x is counts in the predictor energy bins, and λ is the Poisson mean. Using Poisson Regression instead of Linear Regression in match filtering can help improve the model.

One hurdle is that extremely low photon count data can also be sparse and contain excess zeros that occur by chance in no-signal cases. These excess zeros, if not accounted for, may lead to over-dispersion in the estimates of the λ parameters in the Poisson model when signal is expected. Since the mean and variance are given by the same parameter in a Poisson distribution, fitting a Poisson distribution to data with very sparse number of counts may not lead to a good model of the data. An effective numerical trick is to use a Zero-Inflated Poisson (ZIP) model:

$$P(y_j = 0) = \pi + (1 - \pi)e^{-\lambda} \quad (7.3a)$$

$$P(y_j = h_i) = (1 - \pi) \frac{\lambda^{h_i} e^{-\lambda}}{h_i!}, h_i \geq 1 \quad (7.3b)$$

The ZIP model on the radiation data uses a two-step hierarchal approach. Logistic regression is used to first classify the presence of non-zero counts from predictor energy bins. If the spectrum is predicted to have non-zero background counts in the source window, then Poisson Regression is run to predict the amount. The separate probability densities for the zero count and non-zero count cases help prevent over-dispersion when estimating the mean parameter of the Poisson distribution for the case of expected non-zero source signal.

7.2 Experiments and Discussion

Our experiments consist of benchmarking different anomaly detection methods with respect to each other as well as comparing different match filter methods. Match filters typically outperform anomaly detectors when the expected source type is known. In all of our experiments, we compare methods mostly using a single shielded fissile material source type. It is generally established that match filters will outperform anomaly detectors when a particular source type is expected [67]. Thus, this comparison was not particularly interesting and is omitted. There are, however, interesting significant differences that emerge internally within the category of anomaly detection methods and internally within the category of match filter methods.

7.2.1 Benchmarking Poisson PCA vs. Gaussian PCA

We benchmarked the detection capability of Poisson PCA with respect to other anomaly detectors in the literature such as the vanilla Gaussian PCA and Spectral Anomaly Detector methods. A training set of 13,278 un-injected background measurements was subsampled from our urban data set. All methods built null models for the training data, storing their top N principal components.

A testing set of 13,015 background spectra was subsampled from the data. The un-injected data was used as a negative example set. The point source simulator was used to inject synthetic point sources into the testing data to form a positive example set containing source-injected measurements at different distances to the source, from 1 to 20m.

All methods assigned scores to all positive and negative examples based on their learned models from training. Distributions of positive and negative scores were assembled, and discriminative capability was compared between methods. The key success metric used was the Symmetric Kullback-Leibler (SKL) divergence between positive and negative point score distributions:

$$KL(P||Q) = \sum_i P(i) \ln \frac{P(i)}{Q(i)} \quad (7.4a)$$

$$SKL(P||Q) = \frac{1}{2}[KL(P||Q) + KL(Q||P)] \quad (7.4b)$$

where P indicates the distribution of positive scores and Q indicates the distribution of negative scores for a particular method. Typically, KL divergence can be used to non-parametrically compare distributions. Standard KL divergence, however, is not symmetric, so the definition of the SKL divergence helps to fix this.

The SKL formula estimates the difference (average log-odds ratio) between two distributions. A higher SKL indicates that the distribution of positive scores is increasingly distinct from the negative score distribution, and thus source is more detectable against background. The top-performing method for the low count detection problem should produce the greatest divergence between the positive and negative score distributions.

Detecting a far away source

Each anomaly detection method was allocated some number, from $N = 2$ -5 of their top Principal Components (PCs). SKL performance was compared at 1-20m standoff distances to a source.




Figure 7.1: SKL comparison of methods as a function of distance to the source.

Figure 7.1a-7.1d compare the Gaussian PCA, Spectral Anomaly Detector, and Poisson PCA methods for different numbers of PCs. Figure 7.1e shows the maximum SKL performance plotted over the best combination of PCs at each distance for each method.

The experiment dealt with the case of low photon count rate from the source but normal background photon count rates. The source count rate was 65 counts/sec at 10m standoff from the source, while the background count rate was 1263 ± 267 counts/sec. The source is thus well within the tolerance of the background, though the background rate is not considered low.

The SKL metric peaks near the source for all methods as all methods succeed in close-range detection. It falls off for all methods as standoff distance is increased to the source, but falls off slower for Poisson PCA. The results suggest that Poisson PCA outperforms other methods in detecting far away sources at distances where source-originating photon counts are very low.

Detecting a source with shorter lifetime

We also performed an experiment where measurements had reduced sensor lifetime. The counts in the measurement (for both background and source) were scaled to simulate halved sensor measurement intervals. Figure 7.7 shows the maximum SKL performance. The results indicate that Poisson PCA can help detect sources at shorter measurement interval. The net effect is that sensors can travel faster, covering more area while still providing useful levels of source detection.




Figure 7.2: SKL comparison of methods when detecting a source with shorter measurement lifetime.

Detecting a faint source

The first experiment was repeated with a weaker source (i.e. lowered count rate) injected into the background data. For this experiment, the overall counts from the source were scaled down by half from the first experiment. The source count rate was 33 counts/sec at 10m standoff, and the background count rate remained unchanged.




Figure 7.3: SKL comparison of methods when detecting a faint source.

The results indicate that Poisson PCA outperforms other methods in detecting lower intensity sources. Interestingly, there is even more possible utility to Poisson PCA for detecting sources at large standoff distances to the sensor. In the first experiment we saw Poisson PCA outperform Gaussian methods at about 13m standoff to the source, while with the lowered intensity, the performance gain can come at even 9m standoff from the fainter source. More benefit manifests in the top four and five principal components (Figure 7.3c-7.3d).

7.2.2 Benchmarking Poisson Match Filters vs. Gaussian Match Filters

An experiment was designed to test different match filter SNR estimators in their capability to estimate source and background counts from spectra that were low in both background and source counts. The average background count rate was reduced to only 22 counts/sec with the source count rate being even less, making it a very challenging detection scenario.




Figure 7.4: Experiments comparing match filter estimators.

We compared three different match filters: the vanilla Linear (Gaussian) estimator, a Poisson regression method, and the Zero-Inflated Poisson (ZIP) method. Each estimator was trained on a set of 1,527 radiation observations consisting of background radiation to match against a shielded fissile materials source template. The estimators were subsequently tested on a set of 23,389 radiation background observations with source-injected counts at a range of different distances and exposures to a point source. Each estimator estimated the background and source count amounts in each of the observations, and the errors were plotted.

Figure 7.4a-7.4b show mean absolute error in estimation of source and background photon

counts for each method. Poisson Regression has a slight gain over standard Linear (Gaussian) Regression, while ZIP Regression has a major advantage in the match filter. Figure 7.4c shows mean absolute error in estimation of Signal-to-Noise Ratio (SNR). Both Poisson and ZIP Regression outperform Linear (Gaussian) Regression in reducing SNR estimation error.

Interestingly, the Gaussian and ZIP Regression Match Filter estimation errors were only weakly correlated. The background component SNR estimation errors exhibited a linear correlation value of only 0.5519, while the source component SNR estimation errors exhibited only a 0.2745 linear correlation value. This indicates that the two methods are extracting different information from spectra to flag sources.




Figure 7.5: Experiments with Weighted Mean Method.

Figure 7.5a plots the signed error of the Gaussian and ZIP models. The direction of the signed error is opposite. Gaussian models tend to systematically overestimate background components while ZIP models systematically appear to overestimate the source component. Using this observation, we created a mixed method that would fuse the Gaussian and ZIP component estimates in a Weighted Mean. The idea was to regress the Gaussian and ZIP components against the true amounts to attempt to cancel out signed estimation error as much as possible. This method had some success in lowering the SNR estimation error even further as shown in Figure 7.5b.

7.2.3 Incorporation into BA Framework

To further boost performance detection for these challenging signal estimation problems, we leveraged the Bayesian Aggregation (BA) algorithm [67]. BA is a framework for spatially aggregating multiple noisy, low Signal-to-Noise Ratio (SNR) observations to score a geographic location as possibly containing a point source.

Training BA consists of building empirical sensor models based on collected field data (though threats are often simulated). Distributions of SNR scores are assembled as a function of exposure on the detector’s surface. Two sensor models are built: one for the “null” hypothesis (H_0) that is assembled from raw background radiation data and one for the “alternate” hypothesis (H_1)

model that is assembled from data containing injection of a particular source intensity and type. These empirical models establish a testable expectation of what new measurements are supposed to look like if they contain (or do not contain) source signal.




Figure 7.6: Example BA sensor models for Spectral Anomaly Detector and Poisson PCA.

Figure 7.6a-Figure 7.6b show example BA sensor models using the Spectral Anomaly Detector SNR estimator. Figure 7.6c-Figure 7.6d show example BA sensor models using Poisson PCA as the underlying spectrum SNR estimator.

Use of Poisson PCA information in BA

We experimented with the use of Poisson PCA information in the BA framework. We first built BA sensor models using SNR scores estimated using Poisson PCA.

As can be seen in Figure 7.6d, the SNR trend as a function of exposure for Poisson PCA for the alternate hypothesis can differ in shape from the expected SNR trend for a top-performing

Gaussian PCA method like the Spectral Anomaly Detector (shown in Figure 7.6b). The differences in trend indicate that Poisson PCA and the Spectral (Gaussian) PCA may be extracting different information from the same spectra in flagging spectra as containing source-injection or not.

This conjecture is further supported by a correlation analysis of Poisson PCA and Spectral Anomaly Detector scores.

Table 7.1: Correlation Comparison of Spectral and Poisson PCA SNR Scores

Data Set	Spearman Rank Correlation	95% Confidence Intervals in Rank Correlation
Null Scores	0.1621	[0.1202,0.1862]
Alternate Scores	0.1410	[0.1018,0.1671]

As can be seen in Table 7.1, Spectral and Poisson PCA scores are only weakly correlated for both the null and alternate SNR score distributions. This indicates the methods are detecting different features of source signals in the same spectra. The correlation analysis suggested using Poisson PCA and Spectral Anomaly Detector methods in a “Combined Aggregation” scheme. The Combined Aggregation method uses a weighted combination of Poisson PCA and Spectral Anomaly Detector BA scores after BA is applied to each method.

1,000 bootstrap replicants of a geographic subset of the data were prepared. Each subset of data was injected with exactly one radioactive point source using our point source simulator and a shielded fissile materials source template. The source count rate was 65 counts/sec at 10m standoff from the source, while the background count rate was 1263 ± 267 counts/sec. A 2m grid of hypothetical source locations was overlaid on the trajectory and both methods scored each grid point as possibly being the source location. The region was scored by both methods with and without source injection.




Figure 7.7: SKL comparison of methods when detecting a source with shorter measurement lifetime.

Figure 7.7 shows the detection accuracy for these methods with respect to the scores of the top 30 false positives. As can be seen, the Combined Aggregation method that used information from both Poisson PCA and Spectral Anomaly Detector (“Combined Aggregation”) outperformed the method that just used information from the Spectral Anomaly Detector in BA (“Spectral Aggregation”). The optimal weights for the Combined Aggregation were [0.85,0.15] on Spectral and Poisson PCA scores, respectively. The result suggests that using Poisson PCA information in BA can help improve threat detection capability.

Note that Poisson PCA can lose to Gaussian PCA, in practice, due to challenges of overfitting and over-dispersion. Since the mean and variance in a Poisson method are given by the same parameter, the fit to a Poisson is highly sensitive to that parameter. Figuring out how to condition Poisson PCA well enough on the data to meet its expected theoretical performance is a subject of future work. Note that ZIP handles this problem of over-dispersion better than the current implementation of Poisson PCA.

Incorporation of ZIP Regression Match Filter into BA

We performed a similar experiment benchmarking the ZIP Regression Match Filter in BA against other match filters in BA. We compared the Linear (Gaussian) Regression and ZIP Regression match filters both on their own and in the BA framework. We also experimented with marginalizing over the posterior scores of all of these methods in a combined scheme (BA-Marg).

1,000 simulated worlds were prepared, each with an injected shielded fissile material point source. The background had an average count rate of 22 counts/sec with the source count rate being fewer than 10 counts/sec at 10m standoff. A 2m grid was overlaid onto the world, and all methods scored the grid points as being the likely location of the source.

Figure 7.8 shows the Receiver Operating Characteristic (ROC) detection results of the methods that indicate true positive rate as a function of the false positive rate.




Figure 7.8: Detection results for using ZIP Regression Match Filter information in BA.

The ZIP Regression Match Filter BA outperforms the Linear (Gaussian) Regression Match Filter BA. BA-Marg provides a slight advantage and is technically the method of choice. The results indicate that using ZIP Regression information in BA can lead to a more powerful match

filter method than simply using the Linear (Gaussian) Regression Match Filter when the photon counts in spectra are low.

7.3 Conclusion

Our study experimented with augmenting the Bayesian Aggregation (BA) framework with Poisson-based SNR estimators to aid analysis of gamma-ray radiation spectra where the photon count rates of either or both the source and background are relatively low.

Typically, Gaussian assumptions suffice for estimating the background and source components of a radiation spectrum to calculate its SNR when the spectrum contains a sufficient number of photon counts in its additive components. When the photon counts in the spectrum from either the background or source are relatively low, however, the Gaussian approximation to the Poisson distribution is not effective. In these cases, SNR estimators such as the Poisson PCA anomaly detector or ZIP Regression match filter may be useful in improving source detection capabilities of radiation threat detection systems.

The results of our experiments suggest that the use of Poisson models in SNR estimation can help improve threat detection when photon counts from the source (and also potentially from the background) can be relatively low. Our study illustrated the use of such Poisson models for a variety of real world scenarios where this may occur. For instance, we showed Poisson PCA can help detect a radioactive source that is faint or far away or when a sensor moves fast by it. ZIP Regression can help detect a source when smaller, less sensitive sensors are used for data collection. In all of these tested cases, the Poisson-based SNR estimators appear to improve detection capability. Exploiting the spatial correlation between measurements in BA can further improve performance for these challenging source finding problems.

We hope that our results will enable more sensitive and accurate threat detection as well as possibly provide new functionalities at the absolute boundaries of detectability with spectrometers. For example, if mobile radiation detectors need fewer photon counts to locate a radioactive source, they may be able to move faster, monitoring wider areas in the same interval of time. Similarly, if sources can be detected at fainter SNRs and at larger standoff, there would be substantial impact on the sensitivity of threat detection systems. The net result is increasing the safety of the population at risk.

Chapter 8

Using BA for Online Sensor Reliability Monitoring

8.1 Introduction

The BA framework allows sensor fusion over multiple sensing modalities. In this chapter, we develop a fail-safe version of BA to enable online monitoring of sensor calibration properties and evaluate reliability of collected data from multiple sensor streams.

BA allows multiple heterogenous sensors to produce data for analysis. The sensors can be different in their inherent properties, calibration settings, and in the dimensions of data they produce. Sensors are likely to produce information at different levels of reliability for testing different aspects of a Bayesian hypothesis space. For example, radiation sensors can have different sizes, be made of different materials, be tuned to different parts of the energy spectrum, and provide different energy information and levels of granularity of data. Some sensors may be better for detecting different types of sources than others. Sensors may have different response to noise. These sum of these considerations mean that the informativeness of spectra for a source detection problem must be taken into account. Characterizing the initial properties of sensors is an important primary step.

Even with the best possible initial calibration, sensor properties may change in the field. Sensors may malfunction in the field without warning. For example, radiation sensors are known to be susceptible to electronic drifts, data drops, or hardware problems. Detecting sensor failures when they occur is a key step. Understanding the severity of a sensor failure and down-weighting sensors that are no longer providing meaningful data are important considerations in ensuring robust system performance.

This chapter provides:

1. Application of BA to characterizing properties of multiple sizes of sensors.
2. Description and methods for simulating common radiation spectrometer failure modes.
3. Methodology for diagnosing sensor failures from spectral measurements.
4. Incorporation of corruption hypotheses into BA.

8.2 Using BA for Sensor Characterization

BA can be applied to characterize individual sensors empirically in terms of the data they produce. Sensors with different sizes and intrinsic calibration parameters produce different distributions of data features in BA. BA can mine important aspects of these data features. In addition, the Bayesian update rules allow for seamless fusion of information between sensors, taking into account the variability of individual sensors.

The BA framework conveniently allows a user to specify different probability models for different sensors. Figure 8.1a shows an empirical model for a small sensor, while Figure 8.1b shows an empirical model of a larger sensor.




Figure 8.1: Empirical Models of Different Size Sensors.

The probability models for the two sensor sizes can be substantially different. As one can see in Figure 8.2, a loss in detection performance is expected if the probability distribution for a sensor is not well-estimated.




Figure 8.2: Effect of Sensor Model Misspecification.

8.3 Sensor Fault Monitoring

Sensors can further be characterized in the reliability of the data they are producing, and particularly if they are producing faulty data. BA can be used to diagnose when a sensor failure is occurring. During this study, we built a version of the SNR estimation phase in BA that provides dynamic measurement-level detection of sensor faults both via building models for individual sensors and by cross-validating sensor observations against each other.

Several principles came into play. First, many types of sensor failures manifest as deviation from a typical background model. When such deviation is observed, it is evidence to suggest a malfunction is occurring. That deviation can be captured both in the background modeling phase of BA as well as in the sensor empirical probability models. BA allows cross-validating a sensor's data with itself in time using Retrospective approaches (see Chapter 9). Second, sensor data from multiple sensors can be cross-correlated to detect a faulty sensor. Sensors that are proximal to each other are expected to measure similar observations of the environment. If sensor readings differ, it is evidence to suggest one of the sensors is malfunctioning.

We will first start by discussing common failure modes with radiation sensors and simulation methods for those failures.

8.3.1 Simulation of Sensor Failures

Sensor Dropout and Sensor Stalling

Sensor dropout occurs when a sensor (due to hardware malfunction) fails to produce data during a particular time period the data production is expected. For radiation spectrometers, it is often expected that spectra are produced by the hardware every N seconds (in the data used in our experiments, $N = 2$ seconds is often used as a baseline). However, when sensor dropout occurs, there may be blocks of N seconds where data is not produced. The system thus drops the data, and some radiation spectra are blatantly missing. When a spectra is missing, zero photon counts are reported for the spectra when photon counts are binned in time.

A related condition is sensor stalling. Sensor stalling occurs when a sensor fails to produce the correct data at a given time, and instead lumps data associated with the current time step with the next one. In spectroscopy, sometimes counts from one time block fall into the next adjacent time block, spuriously creating a higher number of photon counts in the spectrum. Sensor stalling differs from sensor dropout in the sense that sensor dropout loses data from a time step completely whereas sensor stalling moves the data into the next time step.

Figure 8.3 shows a histogram of the sum of gross counts in spectra for the different cases. First, we see the distribution of gross counts for a normal functioning sensor. This distribution is compared with sensor failure cases. The sensor dropout gross counts distribution shows a sensor with a 5% probability of spuriously dropping a reading. This manifests in several spectra having zero counts (shown by the bar at zero). The sensor stalling case shows a sensor with a 5% probability of stalling. Like in sensor dropout, there is a bar at zero counts. There is also a fatter distribution tail on the right hand side of several larger than usual (in gross counts) measurements.




Figure 8.3: Simulation of Sensor Dropout and Stalling.

Both sensor dropout and stalling can be handled via effective timestamp processing, flagging of suspicious measurements, and monitoring of regular background levels.

Gain Drift

Gain drift in a sensor can occur when environmental conditions (such as temperature near the crystal) change, leading to a shift in the location of expected photo-peaks in the energy space. A gain drift causes photon counts to fall in energy bins adjacent to where they are expected to occur. Figure 8.4 shows a comparison of the expected mean spectra for a normal spectrum and one corrupted by a light gain shift.




Figure 8.4: Mean Spectrum Comparison for Normal Spectra and Spectrum Corrupted by Gain Drift.

Sensor Crystal Problems

Sometimes sensor crystals can break in the field, but still produce reasonable data. A way to simulate crystal fractures is to apply a gain to a spectrum to create a gain-shifted version of

that spectrum. The gain-shifted version of the spectrum is subsequently added to the original spectrum. The new spectrum is normalized to the same total counts as the original spectrum.

Figure 8.5 shows a comparison of the expected mean spectra for a normal spectrum as well as one corrupted with a crystal fracture.




Figure 8.5: Mean Spectrum Comparison for Normal Spectrum and Spectrum Corrupted with Crystal Problems.

The crystal problem failure mode is different from gain drift since the simulation of crystal problems exhibit a specific type of nonlinear drift. The nonlinear nature can be difficult to capture using typical parametric models. The nonparametric modeling capabilities of BA can likely aid diagnosing of crystal fractures when such sophisticated drift occurs. BA can be augmented with hypotheses for crystal fractures with explicit simulation and modeling of such occurrences.

8.3.2 Diagnosing Sensor Failures in Background Modeling

Many types of sensor failures can be diagnosed via anomaly detection approaches. If sensor failure changes the data, one can model typical data and observe perturbations from typical behavior. For spectroscopy, the deviation of the spectrum shape from a failing sensor can be computed from a typical background model. We developed two approaches for performing this type of anomaly detection based on Principal Component Analysis (PCA) and Canonical Correlation Analysis (CCA).

Multi-Sensor Principal Component Analysis Method

Training Principal Component Analysis (PCA) on data containing typical background can build a basis for expected normal functioning of the spectrometer. Sensor failures can then be diagnosed based on the PCA reconstruction error (or linear residual) when a faulty observation is projected onto the learned basis for typical functioning. While changes in actual radiation background can also change the sensor data, the residual is likely to be different. Sensor failures can have characteristic residual shapes as shown in Figure 8.6




Figure 8.6: Sensor Failures in PCA Reconstruction Residual.

Figure 8.6 shows the typical residual for two good sensors and two faulty sensors. The residuals of the two good sensors generally match each other closely, while the faulty sensors have larger reconstruction error from background. Also, the residuals for different failure modes are different, so it may be possible to tell what type of sensor failure is occurring based on just the PCA reconstruction error.

The different residual shapes were leveraged in a simple diagnosis method for sensor failures. The PCA residuals from new measurements were correlated with precomputed expected residual templates for different sensor failure modes (as well as with the working case). A sensor was flagged as containing a possible type of failure if it exhibited maximum correlation with that particular residual template (in comparison to all residual templates in the database). Using just this very simple method, we were able to diagnose sensor failure with high accuracy. Gain drift and broken crystal simulations are often parameterized by a gain parameter g . Table 8.1 shows the confusion matrix for different sensor failure modes with data corrupted with a medium corruption gain parameter.

Table 8.1: Confusion Matrix for Diagnosing Sensor Failure Modes

	Working Sen- sor Model	Sensor with Gain Drift Model	Sensor with Broken Crystal Model
Working Sensor Data	93.05%	4.027%	2.92%
Sensor with Gain Drift Data	1.045%	89.89%	9.07%
Sensor with Gain Drift Data	0.86%	9.19%	89.95%

The experiment can be extended to maintaining a residual template for a range of possible gains for each type of sensor failure. This can help infer the severity of the corruption from gain drift or crystal fracture. An experiment was conducted for five different levels of the gain parameter with the sensor failure modes for gain drift and broken crystal. For these 10 possible levels of failure, the accuracies are reported in Table 8.2.

Table 8.2: Accuracy for Diagnosing Sensor Failure Severity

Sensor Failure Type	Gain	Accuracy
Gain Drift	$g=1.5$	62.19%
	$g=2.0$	65.26%
	$g=2.5$	63.45%
	$g=3.0$	63.45%
	$g=3.5$	55.55%
Broken Crystal	$g=1.5$	54.32%
	$g=2.0$	39.90%
	$g=2.5$	42.11%
	$g=3.0$	42.54%
	$g=3.5$	50.12%

The simple method observes above chance accuracy for all failure modes and performs well in diagnosing gain drift. Broken crystal can be confused with gain drift, leading to a drop in performance. Broken crystal diagnosis can be augmented with BA distributional analysis to improve accuracy as we will see.

Multi-Sensor Canonical Correlation Analysis Method

Another approach to diagnosing sensor failures is to use Canonical Correlation Analysis (CCA) on data from multiple related sensor streams. CCA takes in two sensor data streams and learns multiple pairs of bases (or “canonical variables”). One basis in the pair is learned for each sensor such that projections of new sensor data onto respective bases maximize directions of correlation. By looking at the canonical variable space, it can be seen that with larger number of canonical variables used in the model, the correlation between projections onto respective bases of combinations of good sensors and bad sensors differs.




Figure 8.7: Sensor Failures in Canonical Variable Space.

Figure 8.7 shows an example comparison of sensor projections. Correlations between sensor

projections between respective bases drops off as more canonical variables are used. In fact, by canonical pair $N = 5$, a clean separation emerges in the correlation of faulty sensor versus normal sensor and normal sensor versus normal sensor for both cases of faulty sensors (gain drift and crystal problems). It is no likely accident that by $N = 5$ we start seeing a clear separation between the cases. This is also the number of principal components often chosen for PCA-based methods in the spectroscopy literature [45]. These examples show how both PCA and CCA can be used to diagnose sensor failures in the background modeling phase of BA itself.

8.3.3 Incorporating Sensor Failure Hypotheses into BA

The sensor modeling phase of BA flexibly allows incorporation of sensor failure mode hypotheses. One can have a sensor failure hypothesis for different types of sensor failure modes and test the hypothesis that a failure is occurring. The BA algorithm should be able to detect whether a sudden change in received photons is due to geographic changes in background, a detected radioactive source of interest, or a crystal malfunction.

When sensor failures impact the data, the use of the anomaly detection linear estimators to separate background and source signals leads to a different trend in the estimated Signal-to-Noise Ratio (SNR) than normal. This difference can be statistically modeled as a function of exposure on the detector's surface. The null hypothesis is the case that there is no source present, and the alternate hypothesis is the case a source is present. Figure 8.8 shows the log likelihood difference of BA models learned from empirical data with and without sensor corruption for both the null and alternate hypotheses.




Figure 8.8: Sensor Failure in BA Hypothesis Space.

Differences in model log likelihood can be used to tell the four possibilities apart: $\{\text{sensor working, sensor not working}\} \times \{\text{source, no source}\}$. If a measurement falls within the red regions of the alternate or null models, the more likely the sensor crystal is to be working. If a measurement falls within the blue regions of either plot, the more likely the sensor crystal is


malfunctioning. BA thus allows convenient encoding of different sensor failure hypotheses in robustifying statistical testing of the null and alternate hypothesis.

Experiments

We performed experiments with the corruption-resilient BA models in simulated threat detection scenarios. 1,000 bootstrap replicants of a geographic subset of the data were prepared. Each test block was injected with a single fissile materials source of a particular intensity.

For initial experiments, we injected a relatively intense Source Count Rate (SCR). The source had a count rate of 653 counts/sec at 10m standoff. Note that the source is still well in the tolerance of background since the background count rate was 1263 ± 267 counts/sec. Also, the closest visibility of the source was 10m, so the reported source count rate represents the strongest possible source signal available to algorithms.

To make the scenario more challenging, the detection scenario data was modified with failing sensor simulation of a crystal fracture. Methods were compared on data both with and without corruption. Three methods were benchmarked: BA without corruption hypotheses, BA with corruption hypotheses, and BA-Marg which marginalizes over the posterior probabilities generated from both sub-hypotheses.



(a) Sensor Data with Corruption

(b) Sensor Data without Corruption

Figure 8.9: Sensor Failure Diagnosis Results.

Figure 8.9 shows that the BA with corruption hypotheses does better when the data has corruption, while the BA without corruption hypotheses does better when the data is normal. Marginalizing over the posterior probabilities (“BA-Marg”) does about as well as the top-performing strategy for each case making it the method of choice.

The experiment was repeated with lowered source count rates. The background count rate remained the same, but the source count rate was lowered to 522 counts/sec at 10m standoff and 392 counts/sec at 10m standoff respectively. Figure 8.10 shows the similar results.




Figure 8.10: Sensor Failure Diagnosis Results: Lowered Source Intensity.

As the source becomes less intense, the utility of using the corruption models diminishes. The interesting thing is that BA without any corruption hypotheses can be relatively resilient to moderate levels of data corruption on its own. This is an implicit strength of the empirical variational models captured by the BA algorithm.

Once a failure is detected, we can also attempt to retrospectively correct the pipeline for the sensor failure. In the next section, we formalize Retrospective Learning, an algorithmic framework that allows reinterpretation of past percepts in light of new clues about them.

Chapter 9

Retrospective Learning and Applications to Radiation Sensing

9.1 Introduction

We propose Retrospective Learning as a new paradigm in robot sensor data analysis that allows robots to be retrospective about the data in relation to new information about percepts. In the particular circumstance of BA, it enables automated reevaluation of Bayesian hypotheses about the world in light of new understanding of its past percepts and reevaluation of those past percepts with respect to changes in top hypotheses.

9.1.1 Retrospective Learning

Retrospective Learning builds models that can be retrospective in relation to the data, analyzing data under multiple resolutions, varying granularity of detail, parameter settings, and in light of temporally evolving hypotheses about the world. A robotic agent using Retrospective Learning can reinterpret past percepts in order to obtain a better understanding of current hypotheses about the world and use current hypotheses to gain a better understanding of its past percepts. The goal is specifically to apply retrospective capability in sensor data analysis (as opposed to planning or controls).

From investigative detective work where evidence is reinterpreted in light of new clues to solve a mystery to revisiting mistakes on homework to better prepare for a test, humans use retrospective capabilities all the time to better understand and test hypotheses about their current world. Retrospection is a complex process whose inner workings are a subject of continuous research in psychology [23, 30, 36]. However, a key aspect of retrospective analysis that we can capture algorithmically for use in sensor data analysis is a two-step iteration:

1. Evaluating current high-level hypotheses about percepts in light of new understanding of those percepts
2. Evaluating understanding of percepts in light of current top high-level hypotheses

Robots do not yet use this type of retrospective analysis to understand their environment. The current dominant paradigm in robot sensor data analysis is to process a sensor data stream

only once, on the fly, to update models. Once filtering (such as in a Kalman or Particle Filter) processes a percept, the percept is typically discarded and never revisited. This design decision is made often to save memory and/or achieve the computational constraints of real-time sensing and control. However, as memory becomes more abundant and less expensive and computer hardware becomes faster and more powerful, building retrospective analysis capabilities for robots on sensor data becomes more feasible.

Retrospective Learning aims to design algorithms that utilize past sensor data in smarter ways than simply discarding it. Retrospective Learning uses both past data and current hypotheses to extract more insight into the data set to attempt to gain a better estimate of hypotheses and data parameters than if the agent were not to reanalyze the data. In this way, an agent employing Retrospective Learning could try to obtain a better account of the sequence of sensed events that led it to conclude on a particular hypothesis, as well as refine interpretation of sensed events given particular important or dominant hypotheses.

Retrospection could be one mechanism for robots to reason within a dynamic world with evolving conditions. New data may come to the robotic agent that might cause re-evaluation of top hypotheses and re-filtering of past percepts. In the radiation sensing domain, for example, imagine that the latest percept tells the robot the sensor is likely to be suffering gain drift. A Retrospective Learning algorithm could go back and reinterpret the past percepts in light of this information. It could try to mine the percepts with the gain drift hypothesis and attempt to figure out the source of the gain drift, when it started corrupting the data, and so forth. Another example of where Retrospective Learning could be useful is camera data. Imagine that the latest camera data indicates lighting conditions are changing. Could an algorithm be retrospective about how it analyzes data before and after such environmental change?

The paradigm of Retrospective Learning has been prototyped in the indoor navigation domain with Inertial Measurement Units (IMU) [20]. Normally, when a particle filter operates on such data to track a user's location in an indoor space, the particle filter discards preexisting particle history to save memory. However, keeping the particle history in memory and using retrospective analysis of it has shown benefit to tracking a user's location. With retrospective analysis, the particle histories can be used to test secondary and tertiary particle trajectory hypotheses to reconstruct the particle filter behavior under other (and perhaps more feasible) interpretations of the sensor data if large numbers of particles died. This capability provides a more robust particle estimate in the case of errors in hypotheses and, in general, provides a type of statistical data conditioning that ensures smoother particle trajectories in retrospect.

This chapter prototypes different types of Retrospective capabilities in BA. Three avenues are explored:

1. Use of multiple Bayesian hypotheses and a Hidden Markov Model (HMM) in BA to efficiently recompute hypothesis scores based on new sensor data clues with the Forward-Backward algorithm.
2. Use of a dual Anomaly Match BA (AM-BA) strategy that uses anomaly detection to prune a large hypothesis space before testing potentially computationally intensive specific hypotheses.
3. Development of a multi-modal Augmented PCA to allow more flexibility in modeling intricate variances in BA.

Developed algorithms are benchmarked on the radiation threat detection domain. The goal is to boost real-world performance in the problem of using radiation sensors to detect and characterize radioactive sources. Retrospective Learning approaches are demonstrated to be practically useful in robustifying analysis of spectroscopy data.

9.1.2 Multiple Bayesian Hypotheses Approach

The BA hypothesis space provides a natural way of grounding Retrospective capabilities. Different world scenarios can easily be represented as hypotheses in the Bayesian hypothesis space. In addition, the Bayesian update rules of BA provide a formal way of testing a variety of hypotheses about the world. BA thus readily allows the maintenance of an ensemble of alternative world models that can be rapidly tested to find the hypotheses that fit the data well. It should be possible to dynamically reestimate local past hypothesis history if major world changes are suggested from current percepts.

By having an ensemble of Bayesian hypotheses, one can identify models to fit particular environments as well as dynamically retrain any number of models (in retrospect) to match and explain world events. For instance, one can have different hypotheses for different background types, different sensor function modes, etc. These models can be tested to provide a fine granularity of variation modeling to maximally fit many possibilities to explain world data. Care must be taken to avoid overfitting.

As hypotheses are added to the Bayesian hypothesis space of BA, however, more cases need to be tested. Computational considerations require the development of efficient decision rules to facilitate model switching and recompilation of probabilities.

Hidden Markov Models

Hidden Markov Models (HMMs) [49] provide a natural machinery for retrospectively on a set of incoming sensor observations using different world models. The system is modeled with hidden world states whose classification can be reinterpreted after the fact. HMMs are defined with the following components:

- Set of States $S = \{s_1, \dots, s_n\}$
- Set of Observations $\{o_1, \dots, o_n\}$
- Transition Probability Matrix: $P(s^{t+1}|s^t)$
- Sensor Model: $P(o|s)$
- Initial State Probabilities: $[P(s^0 = s_1), P(s^0 = s_2), \dots, P(s^0 = s_n)]$

For each measurement, the HMM maintains the probability of the system being in a particular state at that particular time step. These probabilities can be re-estimated given new measurements (close in time) that indicate a different sensor state. Though the HMM model assumes independence of observations given knowledge of the current state, the sequence of hidden states that produced the observation data (which the HMM infers) can be reinterpreted to redraw observation probabilities.

The classical Forward-Backward algorithm provides this retrospective capability. The two steps to the Forward-Backward algorithm are:

1. The forward pass on the data which computes the probability of the state based on the observation sequence.
2. The backward pass computes the probability of subsequent observations given the current state.

These two passes are combined to determine the posterior probabilities of state at a given time. The backward pass allows incorporation of new observations in updating knowledge of the individual states that produced those observations.

Application to Sensor Fault Diagnosis

We developed a Retrospective version of BA, based on an HMM formulation, to efficiently reinterpret data when sensor failure is detected. During real time operation, a sensor may first produce good data for some period of time and then malfunction producing corrupt data (upon a breakdown point). The experiments in the last chapter developed measurement-level classifiers for telling whether a data point was from a working sensor or a malfunctioning sensor. However, the algorithms provide no rules for how to reinterpret past percepts in the presence of a sensor failure.

The Retrospective version of BA can model switch between the working model and failure model when failure is detected. It can identify the point of initial sensor breakdown efficiently and correct the data pipeline from that time forward. While the BA algorithm may first test the sensor data with the set of $\{H_1, H_0\}$ for the “sensor working” hypothesis, measurements can be flagged as possibly erroneous for post-processing by failure hypotheses. The probabilities can be later replaced by probabilities drawn from the failure model using the HMM Forward-Backward algorithm. Our sensor state monitoring HMM has the following formulation:

- States = {Sensor Working (W), Sensor Malfunctioning ($\bar{W}$)}
- Observations = Radiation Spectra $\{R^1, \dots, R^n\}$

$$\text{State Transition Matrix} = \begin{matrix} & \begin{matrix} W & \bar{W} \end{matrix} \\ \begin{matrix} W \\ \bar{W} \end{matrix} & \begin{bmatrix} 1 - \alpha & \alpha \\ 0.0 & 1.0 \end{bmatrix} \end{matrix}$$

The state transition matrix encodes the prior probability α of a sensor switching from working to failing. Once a sensor fails, it remains in the failure state.

- Sensor Model = $P(R^i | s^i)$. We fit a logistic model to the threshold on the ROC that maximized classification accuracy for our simple residual template classifier. The logistic model nicely maps the correlation scores of the classifier into probabilities of failure (numbers in the range [0,1]) such that a higher probability indicates a greater likelihood of failure. The sensor working probability is obtained via take one minus that probability.

Experiments

A simulation was run where the sensor was working normally for the first epoch of data collection and then suffered a broken crystal. The sensor subsequently malfunctioned, producing erroneous data for the remaining $T\%$ of the time. Several variants of BA were compared:

1. No Corruption Model: BA that uses only a model for a working sensor.
2. Corruption Model: BA that uses only a model for a corrupt sensor.
3. HMM Model Switch: BA that uses the HMM model to determine failure state. Based on the classification decoding of the hidden state from the HMM, BA draws appropriate model probabilities from the corruption and no corruption models to test observations as supporting H_1 vs. H_0 .
4. HMM Throw Away Data: A BA algorithm that uses the HMM model to classify the state of the sensor as working vs. failing. This algorithm only uses data that is not flagged as containing possible failure by the classifier to make predictions for grid points.
5. BA-Marg: Marginalizing over posterior probabilities.

Methods were benchmarked on data with substantial corruption ($T = 75\%$) as well as no corruption for two different source count rates (SCR). Figure 9.1 shows the ROC results.




Figure 9.1: ROC Results

Several observations can be drawn from the results. In general, there is a significant drop in performance for algorithms when corruption is present in the data. Without corruption, sources at these particular count rates have 90 – 98% detectability, but corrupt data brings the performance dramatically down even when the sources are readily visible. Corruption creates many spurious false positives that hinder performance in world regions that may not contain the source. Throwing away data does not appear to be a viable solution. Data with corruption still contains substantial source detection information and throwing away data loses valuable detection power.

The retrospective model switching capability helps bring back performance when corruption is present in the data. When the source is strong and data corruption is especially troublesome, model switching outperforms using just the BA models with or without corruption hypotheses because these models alone exhibit model mismatch on this data sets which exhibits examples of both working and corruption behavior. Using just the BA with corruption hypotheses, for instance, has model mismatch on the 25% of the data that is normal.

We also looked at the effect of varying the $T\%$ parameter. This translates into having different percentage of the total data amount exhibiting corruption. The scenario is that for $(100 - T)\%$, the data is good and then the sensor breaks so that $T\%$ of the data exhibits malfunction. Figure 9.2 shows the Area Under the ROC Curve (AUC) performance for algorithms at different values of T and source count rate (SCR).




Figure 9.2: AUC Results as a Function of Data Corruption Percentage $T\%$

As corruption percentage in the data increases, all methods tend to deteriorate in performance. HMM Model Switch tends to be the top performing method that deteriorates the least and is the method of choice. HMM Model Switch (with uses the forward-backward algorithm) outperforms simply marginalizing over posterior hypothesis probabilities (BA-Marg) or just doing a forward pass (Forward). The reason is that forward probabilities, which both BA-Marg and Forward are based on, can be noisy and specify erroneous state transitions. Figure 9.3 shows the histogram of unsmoothed probabilities, which can have wide spread.




Figure 9.3: Histogram of Forward Probabilities

Retrospective smoothing based on the Forward-Backward algorithm incorporates the constraint that a sensor crystal produces bad data only after a crystal crack, and robustly dismisses virtually all spurious state transitions that would previously exist in the unsmoothed forward probabilities. A very clean histogram emerges after retrospective smoothing, much more accurately identifying the failure point in each experiment.

Robustifying sensor algorithms to possible sensor failures is important in avoiding false alarms. The Retrospective HMM-based BA can enable resilient detection of radioactive sources under the possibility of sensor malfunction to provide more robust threat detection.

9.2 Anomaly Match Bayesian Aggregation (AM-BA) Approach

A desirable property for Retrospective Learning is efficient search over the set of possibilities of when and to what degree retrospection should occur. Much of this is deciding how computational resources should be allocated to different areas of the search space over parameters of an optimization problem. We hypothesized in the thesis proposal that Active Learning tools could be useful for formalizing such processes. Our earlier adaptive grid framework is a form of Active Learning that adapts geographic grid resolution to areas that are more likely to contain a source of interest than others. Areas that are less likely to contain the source are ignored. In this section, we extend such ideas to save even more computation in efficient retrospection.

For many search problems, there tends to be a tradeoff between sensitivity and specificity. One can test very general hypotheses about the world but risk missing very specific cases. Similarly, one can test very specific, specialized hypotheses about the world but miss out on a more general property of data. A concrete way to think about this tradeoff is with anomaly detectors and match filters in the radiation domain. Anomaly detectors make fewer assumptions about the data (i.e. are more general detectors) than match filters but are less sensitive for specific hypotheses. Match filters test very specific hypotheses about sources but might miss a source they are not designed to test.

The idea of Anomaly Match Bayesian Aggregation (AM-BA) is to use general anomaly data and specific match information in a joint strategy. Anomalies are first mapped and used to prune

the search space of the specific (which can be subsequently be precisely match filtered). Thus, fewer overall search queries have to be made with specific constrained search after the pruning. This controls efficiently where retrospective computational processes are allocated.

The AM-BA strategy is to:

1. Run anomaly detector BA on the grid of hypotheses
2. Threshold the anomaly BA scores at a particular threshold level T
3. Run Match Filters only on the grid points that are above the threshold level T

With the general (anomaly) pruning, less specific (match) searches have to be run to find the source. We will demonstrate application of this idea to the radiation source detection problem to facilitate efficient search and determination of source parameters from data.

Note that the threshold T can be chosen specifically for the search problem. While we do not provide a general strategy for choosing this parameter, we show an example way of doing so for the radiation domain.

9.2.1 Application to Radiation Source Detection Problem

The problem of finding a single isotropic radioactive point source in data is to identify the source location as well as the source parameters (e.g. count rate or intensity and source type). A source search scene may have a large number of possible source location hypotheses. If one has to search over all possible source parameters for a large number of source locations, the computational requirements can quickly become unscalable. To address this challenge, we propose to be smarter about where source parameter tuning is optimized using AM-BA.

In the radiation domain, anomaly detectors and match filters are SNR estimators commonly used on radiation observations to separate estimated source components from typical background fluctuation. Anomaly detectors assume no knowledge of the source, simply finding anomalies to typical background fluctuation. Match filters, in contrast, use knowledge of a source template or design in estimating components. Both types of estimators are useful in BA for estimating and assembling distributions of the Signal-to-Noise Ratio (SNR) of measurements. Both estimators can be used in unison in an AM-BA strategy for finding the radiation point source.

By first running the anomaly detector BA on the hypothesis grid of source locations, anomalies are identified on the map. The anomaly map quickly prunes, in geography, where sources may not exist. The remaining grid points can subsequently be optimized over other parameters (e.g. type, intensity, etc) via match filtering. The number of grid points flagged as possibly containing sources by the anomaly detector is substantially less than the full size of the grid, resulting in computationally efficient parameter optimization. There is a tradeoff, however, between pruning the number of grid points and detection performance. Pruning the grid with too high a threshold for anomaly score may result in loss of detection performance. The criterion for choosing T for the search problem is thus to identify an empirical threshold for the anomaly scores that results in maximum computational savings without loss of threat detection capability.

Experiments

Using a library of 899 fissile material source templates, an anomaly detector BA was trained with data containing injection of **all** source templates. Since the alternate model is trained with all source templates in the library, it is agnostic to any particular source template. The null and alternate hypothesis models are shown in Figure 9.4a and Figure 9.4b respectively.




Figure 9.4: Null and Alternate Hypothesis (Source Agnostic) Anomaly Models

In testing the learned models, we simulated 42,000 bootstrap replicants of data with injection of sources from the library. Each data replicant had one injected point source. Injected source templates from the library had a range of correlation with the mean raw background spectrum in the data. Typically, sources highly correlated with mean background are harder to detect than sources very distinct from it. AM-BA was run on these data replicants, and detection performance and grid savings were reported. The experiment was repeated for four different source intensities / count rates. Injected sources are typically within the tolerance of background count rate (1263 ± 267 counts/sec). Source count rates (SCR) are given at 10 m standoff.




Figure 9.5

Figure 9.5a shows that, as the threshold on anomaly score increases, more grid points are pruned. Figure 9.5b shows that as the T increases, detection performance accuracy can drop after a point. For each source type, we identified the Area Under the ROC Curve (AUC) operating point where detection performance remains optimal but maximum number of grid points are pruned.




Figure 9.6: Results of AM-BA Experiments.

Figure 9.6a shows the maximal AUC detection performance possible for different source templates with respect to correlation with mean background. Figure 9.6b shows the maximal grid point percentage pruned with the anomaly BA. In most cases, nearly 70–75% of the computation can be saved without loss of the AUC of detection performance for each source type.

9.2.2 Experiments with Clustering Source and Background

To further speed up search over the source library, hierarchical clustering was used to structure the source library into a dendrogram so sources very different in shape from the estimated signals need not be searched. We experimented with both clustering on the vanilla source template library as well as operational clustering where the clustering algorithm has access to the anomaly residuals from background. Agglomerative and Divisive Clustering show promising data representation.

Multi-dimensional Scaling (MDS) was used to select and visualize distance functions for the clustering. Since radiation data is often presumed Poisson, one of the more promising distance functions was based on the Poisson Log Likelihood distance between radiation vectors (assuming independence of energy bins).

The Probability Density Function (PDF) for a Poisson distribution is given by:

$$f(x|\lambda) = \frac{\lambda^x}{x!} e^{-\lambda} \quad (9.1)$$

Computing the Poisson likelihood between a radiation spectrum a and a reference spectrum b (both N -dimensional) assuming conditional independence between energy bins means the probabilities are multiplicative:

$$f(a|b) = \prod_{i=1}^N f(a_i|b_i) \quad (9.2)$$


To avoid underflow, however, the distance computation is implemented as a sum of log probabilities.

$$\log f(a|b) = \sum_{i=1}^N \log f(a_i|b_i) \quad (9.3)$$

Since the Poisson PDF is not initially symmetric, the distance must be computed both directions in the sum to symmetrize. Thus, the final used Poisson Log Likelihood Distance $d(a, b)$ is given by:

$$d(a, b) = \sum_{i=1}^N \log f(a_i|b_i) + \sum_{j=1}^N \log f(b_j|a_j) \quad (9.4)$$

The results of running MDS on the distance matrix of Poisson Log Likelihoods for un-injected background spectra in the Sacramento data is shown in Figure 9.7. K-Means Clustering is run on the 2D data to find significant data clusters.



(a) Emergent Background Clusters

(b) Interpretation in Space and Time

Figure 9.7: Clustering of Sacramento Background Spectra

As Figure 9.7a shows, background points cluster nicely into four distinct clusters in the Sacramento data. These clusters correspond to significant spatio-temporal variation in the data. The variation is strongly dominated by time. As can be seen in Figure 9.7b, the different clusters mostly correspond to different segments of time that data was collected during (even if in the same general geographic location).

Running MDS and K-Means Clustering on source and background points jointly, as shown in Figure 9.8, can reveal their joint arrangement.




Figure 9.8: Clustering of Sacramento Background Spectra

Using the Poisson Likelihood distance function (results shown in Figure 9.8a), the sources appear as a tail on the distribution of background points. An operator may find the insight that the tail overlaps more with certain types of background spectra than others useful. This information may help adapt optimization of search parameters in real time to detect sources. The structure is even more apparent when using Spearman Correlation as the distance function between observations as shown in Figure 9.8b.

We also experimented with a form of operational clustering (shown in Figure 9.9) that clusters PCA residuals from source-injected data as opposed to raw radiation observations. Since an operator is likely to have PCA residuals available to them from spectra, this form of clustering helps us understand how source-injected measurements (for different sources) really differ from the background.




Figure 9.9: Operational Clustering using PCA Residuals

The analysis reveals the spread in detectability of different types of sources. Some sources are easier to detect since they create more distinct residuals from background. Other sources are

more difficult to detect using anomaly detection. The spread of the detectability of the sources can further help optimize where specific match filters should be run.

9.2.3 Discussion and Conclusion

Efficient source parameter optimization is key for nuclear threat monitoring. This thesis develops AM-BA to help save computation in searching the parameters of a space of possible source locations. The method helps better decide when and where Retrospection should occur on the data to improve results while avoiding computational burden. The method can likely be made more efficient through learning more about the structure of the underlying (i.e. via MDS and clustering) data.

Though this thesis mainly demonstrates these Retrospective Learning ideas on the radiation source search problem, the ideas are expected to be useful in many domains. Deciding where an agent should Retrospect is an important problem across many fields, and how that Retrospection should occur is a continual research challenge.

9.3 Multi-Modal Augmented Principal Component Analysis

A key requirement for Retrospective Learning is flexibility in modeling. New components may need to be added on the fly for data analysis and old components removed. Retrospective Learning hopes to enable the creation of highly dynamic models that can explain lots of variance in percepts, and be able to efficiently reinterpret that variance in light of new data.

One of the key conventional approaches to variance modeling is Principal Component Analysis (PCA). Though vanilla PCA is useful on its own for explaining data variation, it can be made more useful with appropriate encoding of additional variation components that can be dynamically estimated from data.

In this section, we develop a Multi-Modal Augmented PCA that allows incorporation of multiple covariance components in a PCA model. The method is benchmarked on the radiation source finding problem to enable modeling of different types of background and nuisances sources that may exist in an urban scene, and to be robust when confronted with such dynamic conditions.

9.3.1 Augmented PCA Method

PCA can be trained using a wide range of covariance matrices. The MATLAB function PCA-COV, for instance, allows a user to specify a covariance matrix for Singular Value Decomposition (SVD) to be run on. Since SVD can be run on any arbitrary matrix, the covariance matrix input to this function can be virtually anything.

We experiment with covariance matrices that are a weighted combination of a set of pre-specified covariance components $\{U_1, U_2, \dots, U_N\}$. These covariance matrices have the following form:

$$\text{cov}(A) = \sum_i^N u_i \text{cov}(U_i) \quad (9.5)$$

$$\sum_i^N u_i = 1 \quad (9.6)$$

The overall covariance matrix is a weighted combination of covariance components and the set of weights $\{u_i\}$ sum to 1. The covariance components and weights are often chosen based on the domain and problem. We demonstrate use of the Augmented PCA method on the radiation domain.

9.3.2 Application of Augmented PCA to Radiation Source Detection

In the radiation domain, PCA is a common tool for decomposing a radiation spectrum into source and background components. In the vanilla case, PCA is trained on background data and learns a basis for characteristic background fluctuation. The top principal components capture typical background variation that can be subtracted out to find anomalous (potentially source-like) signal in subsequent sensor readings. In this study, we experiment with augmenting PCA with multiple covariance components to allow more modeling flexibility in specifying expected source and background components.

Real world background radiation landscapes can be extremely variable, requiring flexible PCA models. Urban scenes can contain many distinct types of background radiation. In addition, such scenes may contain different types of benign nuisance sources that, if not accounted for, can cause false alarms when detecting a threat. The goal is to develop a PCA that is less sensitive to changes in typical radiation background and the presence of nuisances to provide robust detection of true threats.

Augmented PCA has previously been used to augment a PCA model with different expected types of sources to create a PCA that is more sensitive to particular types of sources in a source library [31]. In our work, we use a Augmented PCA of the following form:

$$\text{cov}(A) = \sum_i^N \alpha_i \text{cov}(B_i) + \sum_{j=1}^M \beta_j \text{cov}(S_j) \quad (9.7)$$

$$\sum_i^N \alpha_i + \sum_j^M \beta_j = 1 \quad (9.8)$$

where the $\{B_i\}$ are a set of different types of characteristic background covariances expected in the data and $\{S_j\}$ are a set of covariances of potential nuisance sources in the environment. The sets of weights on the covariance components, $\{\alpha_i\}$ and $\{\beta_j\}$, jointly sum to 1.


Experiments: Multiple PCA Covariance Components

An experiment was run with two simulated background types, B_1 and B_2 . Background type B_2 had a gain-shifted Poisson mean in comparison to B_1 . The following Augmented PCA model was prepared:

$$\text{cov}(A) = (w)\text{cov}(B_1) + (1 - w)\text{cov}(B_2) \quad (9.9)$$

$$0 \leq w \leq 1 \quad (9.10)$$

Two testing sets of background data were prepared – one each for B_1 and B_1 . Each data set only contained one type of background. Distributions of PCA reconstruction error scores were assembled for different choices of the weight w . Figure 9.10 shows the results for some key choices of w .



(a) PCA Scores for Background Data B_1

(b) PCA Scores for Background Data B_2

Figure 9.10: PCA Scores for Background Data

Figure 9.10a shows the PCA reconstruction error distributions for background type B_1 , while Figure 9.10b shows the same thing for background type B_2 . When $w = 0$, all the weight is on $\text{cov}(B_2)$ and thus the background data from B_2 has least reconstruction error for that weight setting. Similarly, when $w = 1$, all the weight is on $\text{cov}(B_1)$ and thus the background from from B_1 is well modeled. Choosing an intermediate w between $[0, 1]$ generally finds a model in the middle of the two extremes. Note that distribution relations between reconstruction errors are not expected to be symmetric. B_2 may be easier to model with B_1 basis than B_1 with the B_2 basis, leading to smaller reconstruction errors for the first case.

One can also assemble the distribution of minimum reconstruction errors for each point over the entire ensemble of PCA models maintained at different choices of w . We call this method “Min-Ensemble” and is shown to be very close to the $w = 0$ and $w = 1$ for our simple two radiation background example. When more components are added to the Augmented PCA model, the results aren’t as simple and the “Min-Ensemble” becomes more useful. From now on in the

text, when we refer to “Augmented PCA”, we refer to the Min-Ensemble over all weighted PCA models for a specified discrete setting of weights.

One can interpret the setting of the weight w as rotating the top principal components from some original state (we use $w = 0$ as the reference). Figure 9.11 shows the angular rotation different principal components (PCs) undergo as w increases. When w gets larger, the concentration of $cov(B_1)$ increases in the Augmented PCA model and the concentration of $cov(B_2)$ decreases.




Figure 9.11: Rotation of Principal Components due to Background Changes

Interestingly, the spectral shape fluctuations manifest most in the fourth PC. The first couple PCs are still seemingly dominated by the mean spectrum and the (unchanged) temporal variation in the data.

A similar result manifests when we inject a nuisance in B_1 and replace the B_2 in the augmented PCA model with a S_1 nuisance source covariance and rerun the experiment.




Figure 9.12: Rotation of Principal Components due to Nuisance Sources

Figure 9.12 shows that the nuisance source really causes the fifth and fourth PCs to rotate most. The earlier PCs are less affected by the addition of the nuisance component in the PCA

model. This follows intuition that the variation introduced is less significant than the natural temporal variation in the data, though it can still affect results substantially and needs to be accounted for.

Experiments: Full Augmented PCA

A data set containing both background types was prepared. An Augmented PCA model of the following form was used:

$$\text{cov}(A) = (\alpha)\text{cov}(B_1) + (\beta)\text{cov}(B_2) + (1 - \alpha - \beta)\text{cov}(S) \quad (9.11)$$

$$\alpha + \beta = 1 \quad (9.12)$$

where $\text{cov}(S)$ is the covariance of a nuisance source. A cross-validation experiment was performed over the sources in the threat library. A source from the library was selected to be a true injected source, and another source was chosen to be an injected nuisance source. The goal of algorithms was to succeed in detecting the actual source while not creating false alarm due to the nuisance.

To make the cross-validation most interesting, the source library was first clustered into key source archetypes. Multi-Dimensional Scaling (MDS) was run on the source library to visualize source distances in two dimensions, and K-Means Clustering was run on that space. Figure 9.13 shows the results.





Figure 9.13: PCA Scores for Background Data

Figure 9.13a shows the three identified source clusters in the library. The clusters capture different source shapes in the library that are at different levels of correlation from mean background in the data set. Figure 9.13b shows that cluster 1 contains sources that are highly correlated with mean background, sources in cluster 3 exhibit mid-range correlation, and cluster 2 sources exhibit low correlation with mean background.

For the cross-validation, a subset of 100 sources were uniform randomly selected from the threat library. When the source was injected as the threat in the data, the source was cross-validated against other sources in the 100 source subset that were not in its cluster that act as injected nuisances. This helps provide good test examples to see whether nuisances can really getting subtracted out by Augmented PCA.

Figure 9.14 shows the results of a cross-validation experiment that compares the Augmented PCA model to vanilla PCA methods that only model one background component (B_1 or B_2 but not both and without any nuisance source modeling).



(a) AUC Performance across Source Library

(b) Distribution of AUC Scores

Figure 9.14: Cross-Validating over Source Library

Figure 9.14a shows the Area under the ROC Curve (AUC) results for fixing a particular injected nuisance template (with particular correlation with mean background) and cross-validating over the source library. The cross-validation for each source is run over sources not in the same cluster.

As the correlation of a nuisance with mean background increases, nuisances are easier to subtract out since PCA is already designed to subtract out projections that look like background. As the correlation of a nuisance with mean background decreases, there is more variation in results, as the result depends on exactly where the similarity between the nuisance and background occurs in the energy space.

Throughout the possible injected nuisances, even those weakly correlated with background, Augmented PCA is able to be resilient to the nuisances and detect the true threat when it is injected into the data. Figure 9.14b plots the distribution of AUC scores for different methods. Augmented PCA has a peak at higher level of AUC then other vanilla PCA methods that just model single background components for either B_1 or B_2 .

The result helps explain what types of nuisances can cause problems for detecting all types of sources. This helps warn an operator of what nuisances to watch out for (statistically speaking) and helps gate the sensitivity of models. An operator using the algorithm can better select parameters to detect different types of sources in a match filtering (or other) framework.

One can also run the cross-validation experiment in the opposite direction: fix a source template and vary the nuisance. This helps quantify how susceptible algorithms are to confusing particular source templates as nuisances. Figure 9.15 shows the results.




Figure 9.15: Cross-Validating over Nuisance Library

There is much variation, but Figure 9.15a shows many sources that be detected without succumbing to nuisances with Augmented PCA. Figure 9.15b shows that Augmented PCA provides a better distribution of AUC scores for sources than other methods.

Both cross-validation experiments were also repeated for a stronger injected nuisance source (Figure 9.16).




Figure 9.16: Cross Validation Experiments with Stronger Nuisance Source

As expected, performance drops somewhat. However, as Figure 9.16b shows, Augmented

PCA can still provide high AUC scores for detecting sources even in the presence of the stronger nuisance.

Discussion and Conclusion

Our experiments show that by using a covariance model in PCA that is a linear combination of covariance components, PCA can be augmented with additional information to improve background suppression and threat detection. The resulting model can be trained to be less sensitive to known nuisances as well as to known systematic shifts in background radiation variability, yielding lower false detection rates at comparable levels of threat detectability.

The use of Multi-Modal Augmented PCA for background modeling is a key component of Retrospective Learning and Retrospective BA. The model allows incorporation of new knowledge (i.e. covariance components) to reanalyze older data. This form of retrospection hopes to obtain better explanation of the data than was previously available.

Chapter 10

Conclusion

In this section, the contributions and future research directions of the thesis are discussed.

10.1 Thesis Contributions

The thesis has several contributions, many of which have been published in formal literature in refereed papers and journals and/or been presented as posters or talks at conferences. The following is a list of published contributions:

- We have proposed the Bayesian Aggregation (BA) data pipeline for helping detect and characterize patterns from a signal-generating source or process from multiple noisy sensor observations. Key to the framework is the aggregation of evidence, the fusion of multiple correlated observations, to help separate signal from noise and help localize the source ¹. The five stages of the pipeline are sensor calibration, measurement SNR estimation, sensor fusion, postprocessing, and route planning. The pipeline allows for many stages of flexibility in implementation and analysis of data. All stages of the pipeline have been prototyped.
- We have demonstrated the critical capability of BA to simultaneously infer the properties of a source while detecting it ². We prototyped the capability to use Bayesian methods to simultaneously track multi-model hypotheses of sources that allow inferring intensity (i.e. count rate) and source isotope type of detected sources, a new capability not previously existing in algorithms in the threat detection literature ³.
- As part of the BA pipeline capabilities, we have demonstrated the capability to efficiently

¹Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Source Location via Bayesian Aggregation of Evidence with Mobile Sensors. Symposium on Radiation Measurements and Application 2012.

²Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Simultaneous Detection of Radioactive Sources and Inference of their Properties. IEEE Nuclear Science Symposium 2013.

³Tandon, P., Huggins, P., Maclachlan, R., Dubrawski, A., Nelson, K., and Labov, S. (2015). Detection of Radioactive Sources in Urban Scenes Using Bayesian Aggregation of Data from Mobile Spectrometers. Information Systems: Special Issue on Mining Urban Data.

search over source location hypotheses using an adaptive grid ⁴. The method also enables efficient search over possible static and dynamic occlusion world scenarios and Bayesian hypotheses using a Branch and Bound algorithm.

- As part of the BA pipeline capabilities, we have prototyped the use of BA information in typical single-agent and multi-agent route planning strategies for mobile sensors. In the single-agent case ⁵, we employed Active Learning to plan routes for a single mobile sensor to find a single static radiation source. In the multi-agent case ⁶, we prototyped the capability for BA information to be used in guiding a team of mobile sensors in searching for a single mobile source in traffic.
- In addition to experiments with BA in simulation, we have prototyped a Robotic Radiation Sensing Backpack ⁷ that serves as a data collection platform. The backpack contains a NaI spectrometer and employs inconspicuous localization technologies such as Wi-Fi localization and inertial navigation to help detect radiation sources in GPS-denied settings. The backpack system has been successfully demonstrated in indoor inspection applications for the U.S. Domestic Nuclear Detection Office (DNDO).
- We have extended the capabilities of BA to operate on sensor data that has low photon counts from background and/or source. For the case where photon counts from background are normal but source counts are low, we have used Poisson PCA ⁸ to help find the source. This can be useful if the source is faint, far away, or sensor measurement time was quick. For the case where photon counts from both background and source are low, we have used Zero-Inflated Poisson (ZIP) Regression Match Filtering ⁹ to help detect the source. This method can be useful for modeling data from smaller sensors.
- The capabilities of BA have been further extended to account for possible sensor faults in the data. We have developed a fault-tolerant version of BA that can detect sensor malfunction and maintain hypotheses for sensor failure ¹⁰.
- We have developed Retrospective Learning capabilities in BA that allow BA to reinterpret past percepts based on new data clues. One mechanism of retrospection is a Hidden Markov Model (HMM) that allows smoothing of past state estimates using the Forward-Backward algorithm. In the radiation domain, this mechanism helps to correct a data pipeline in the presence of sensor faults in the data ¹⁰.

⁴Huggins, Peter, Prateek Tandon, Artur Dubrawski, Simon Labov, and Karl Nelson. Dynamic Placement of Sensors for Rapid Characterization of Radiation Threat. DTRA Review July 2013.

⁵Tandon, Prateek, Artur Dubrawski, Jeff Schneider, Adam Zagorecki, Simon Labov, and Karl Nelson. Machine Learning for Effective Nuclear Search and Broad Area Monitoring. ARI Annual Review 2011.

⁶Tandon, Prateek. Multi-agent Planning for Mobile Radiation Source Tracking and Active City-wide Surveillance. Graduate Artificial Intelligence Project (published in ARI Annual Report). Spring 2013.

⁷Tandon, Prateek, Vladimir Ermakov, Aashish Jindia, Artur Dubrawski, Simon Labov, and Karl Nelson. Portable Radiation Monitoring Platform for Effective Nuclear Search. ARI Annual Review 2012.

⁸Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Suppressing Background Radiation using Poisson Principal Component Analysis. IEEE Nuclear Science Symposium 2012.

⁹Tandon, Prateek, Artur Dubrawski, Peter Huggins, Robert Maclachlan, Karl Nelson, and Simon Labov. Poisson Match Filter for Low Photon Count Data from Portable Spectrometers. ARI Annual Review 2014.

¹⁰ Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Fault-Tolerant Source Detection using Bayesian Sensor Reliability Models. IEEE Nuclear Science Symposium 2015.

- Another mechanism is the use of an Anomaly-Match (AM-BA) strategy to first map anomalies, and then use anomaly information to gate search with specific match filters. In the radiation domain, this technique can be useful for pruning a grid of source location hypotheses using anomaly information before running computationally intensive search over the space of match filter parameters ¹¹.
- An additional source of Retrospection in Retrospective BA comes from an Augmented PCA which allows flexible, on-the-fly modeling of variance components in BA. In the radiation domain, Augmented PCA allows for incorporating multiple variance components for different types of background shapes and potential nuisance sources in the data ¹². This allows BA to be more resilient against possible nuisances in the data that might otherwise cause false alarm.

10.2 “User Guide” for Bayesian Aggregation

In this section, we provide a step-by-step explanation of when and how Bayesian Aggregation (BA) and its subcomponents can improve detection of a source in sensor data. The presented methods hierarchy is backwards-compatible. More advanced methods will perform on simpler cases. Advanced methods can, however, handle more complex real-world scenarios.

Suppose we have a data set of sensor observations containing a source. We have no prior knowledge of the presence of source, nor knowledge of any of its parameters. All that is given to us is a set of sensor observations made along a discrete trajectory (shown in Figure 10.2), and our goal is to algorithmically inspect the data and detect point sources in it.




Figure 10.1: Trajectory data containing point source affecting sensor observations.

The simplest case is where the source is very strong and detectable using a scoring metric that operates only on individual observations.

¹¹Tandon, Prateek, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Anomaly-Match Bayesian Aggregation (AM-BA) for Efficient Radiation Source Search. IEEE Nuclear Science Symposium 2015.

¹²Tandon, Prateek, Artur Dubrawski, Peter Huggins, Karl Nelson, and Simon Labov. Multi-Modal Principal Component Analysis for Robust Threat Detection. ARI Annual Review 2015.




Figure 10.2: Easy source detection scenario with strong signal in individual measurement.

In this very simple scenario, any of the sensor observations has sufficient source signal (well above the mean background) to tell that the source is there. A detector algorithm that scores only individual observations is sufficient for this trivial case. If, however, the source signal is in the tolerance of the background variance, this naive approach will not work.

The next detection strategy one can employ is to try to aggregate multiple observations to detect a source. One of the common approaches in the literature is the Weighted Combining (WC) Method [46]. The WC method aggregates the source signal and background estimates for nearby observations to compute a total estimated Signal-To-Noise Ratio (SNR) score for a hypothetical source location. WC will work to detect many sources by computing aggregate SNR scores over a grid of hypothetical source locations. Figure 10.3 shows use of WC to aggregate SNR for one example grid point.




Figure 10.3: Source detection scenario requiring aggregation of observations to detect source.

Since the SNR estimation process is imperfect, and signal and noise estimates may have non-Gaussian variance, nonparametric Bayesian estimators can improve performance by modeling empirical variability more flexibly. If there are distributed nuisance sources in the data (that

cause false positives), if the source is occluded, and/or the source is especially weak, the use of nonparametric BA sensor models have been shown to further boost source detectability.




Figure 10.4: Comparison of WC and BA on data with a distributed nuisance source.

Figure 10.4a shows the detection heat map of WC in the presence of a natural distributed nuisance source in the data, while Figure 10.4b shows the BA performance. WC misses the real source, flagging the distributed nuisance as the hottest spot. BA robustly estimates SNR using nonparametric models of expected SNR as a function of exposure. BA thus dismisses the distributed nuisance source as not following the expected SNR-exposure trends of point sources while finding the true point source. BA additionally supports inferring properties of the injected threat, such as the source intensity or source type.

If the type of source being sought is roughly known, using a “source-aware” BA may help. When a match filter estimator (for the correct source type template) is used in lieu of anomaly detector estimators in BA, the match filter shows a much stronger detection score (shown in Figure 10.5 for our test scenario).




Figure 10.5: Match filter BA improves detection capability.

The source detection problem may become more challenging due to a variety of scenarios. For instance, the signal components in the data may be sparse or low count. Using a less sensitive sensor or measuring signal from a weak source may create this challenge. In this case, the use of Poisson estimators (in lieu of Gaussian ones) can improve detection capability. Figure 10.6a shows an example sparse, low count spectrum collected by a small spectrometer. Figure 10.6b shows example ROC detection improvement that Zero-Inflated Poisson (ZIP) BA provides over a standard Linear Regression based match filter in low count settings.




Figure 10.6: ZIP BA improves detecting sources in sparse, low count spectra.

Another key challenge is the potential presence of benign nuisance sources or noise fluctuations in the sensor data. In the radiation domain, false positives may emerge from common nuisances (such as radiation emanating near hospital radiology departments) and/or changes in background modes. Figure 10.7a shows some common data nuisances.





Figure 10.7: Augmented PCA can help lower false alarm rate.

If one knows the source type being sought, a source-aware match filtering BA can help dismiss signal components from sources not being sought (such as the nuisances). If however, the source type being sought is unknown, the anomaly detector BA can be informed of nuisance covariances via an Augmented PCA model.

Figure 10.7b compares the distribution of PCA residuals for typical background and background injected with nuisance sources. If the background data is typical, the vanilla PCA approach will keep PCA residuals low. However, if there are nuisance fluctuations in the data, the vanilla PCA will create higher PCA residuals, potentially leading to false positives. Using an Augmented PCA model, the residuals for nuisance fluctuation will be lower (and performance on regular background is preserved).

Sensor failures can also affect data analysis. If one or multiple of the sensors in the data processing pipeline stop working correctly, false positives can be created from the corrupt data. Incorporating BA sensor models for possible failure cases can help make your sensor data analysis more fault tolerant.



(a) Probability mass shift with sensor failure data.

(b) Improvement of fault tolerant BA.

Figure 10.8: Fault-Tolerant BA can account for possible sensor failures in data.

We can see that when sensor failures are prominently affecting the data, the expected probability model of the sensor changes (Figure 10.8a). If not accounted for, sensor corruption can lead to a drop in performance. Performance can be recovered by using the fault tolerant BA. Figure 10.8b shows an example ROC improvement of the fault tolerant BA when compared to the original version when sensor failures are occurring in the data.

The maintenance of a large number of source hypotheses and possible source parameters can be a computational challenge. To improve scalability of BA algorithms, one can use adaptive grid approaches (outlined in Chapter 5) and Retrospective Learning approaches (outlined in Chapter 9) to improve search over the parameter space. Once one can squeezed the most juice out of current data, one can also use the route planning strategies (outlined in Chapter 6) for single or multiple sensors.

10.3 Generalizability of BA Method

The research has developed a Bayesian Aggregation (BA) data processing software pipeline that can help identify sources of signal in very noisy, high dimensional data. The application domain we have demonstrated our methods on is the problem of detecting radioactive threats. However, the software pipeline we have prototyped and developed can be useful to other problem spaces. We list some possible scenarios for future application (beyond the scope of this thesis, but still plausible from the scope of this thesis). These potential applications of BA illustrate the generalizability of BA to other problem spaces.

- **Medical Diagnosis Domain:** Using one or multiple sensors on patients to monitor patient health is a possible future application area. Such multi-modal sensor data is also very noisy and often requires advanced machine learning techniques such as ours to make sense of. BA could support building empirical models of sensors in this domain and aggregate multiple observations to improve classification of patient health.
- **Robotic Boats Domain:** The problem of sensing obstacles robustly on a river using robotic boats is a challenge for navigating the robot boats. BA could be used to create an aggregate view of possible upcoming obstacles, using sensor percepts from all boats to help the team navigate possible river obstacles successfully.
- **Sound Localization Domain:** Localizing the source of a sound signal with a microphone can be another useful application area. BA null models could be created from data with background noise in an environment. Data could be collected at different distances to a sound source in the environment, and alternate models could be built from this sound-source injected data. The method could then be used to localize a similar sound source in a novel similar setting.
- **Neural Data Analysis Domain:** EEG, EMG, FMRI, neural recording, and other types of neural data is also very noisy, and consists of many weak, low-count component signals. Many of these data types are often also thought to be influenced by generating processes that are likely to be Poisson. Our Poisson modeling techniques, in conjunction with BA, may help improve analysis and understanding of these very difficult any noisy data streams. Aggregating multiple observations may also boost recognition of user intent from these data streams.

10.4 Recommended Future Research Directions

There are several directions of potential future research:

- **Testing of a full BA system.** This thesis mostly developed BA and its subcomponents as independent systems. Creating and testing a full system with all the subcomponents is a natural next step. In order to do this, we would need to create a large test set of possible world scenarios – sources injected with various intensities, types, in the presence of data nuisances, with sensor failures, and so forth. We would create a BA algorithm that could account for all the different world possibilities and benchmark it against other methods.

Testing this full BA system was beyond the scope of this thesis and is a recommended direction for future work.

- **Further robustification and application of techniques to the real world.** The techniques developed for aggregating evidence to boost SNR from multiple observations of a source as well as, especially, the techniques for modeling low count Poisson-distributed data with Poisson PCA and ZIP match filter likely have many other real world applications. For instance, neural population data can also be sparse and Poisson-distributed. By applying the methods to additional real world application domains, the techniques can be further refined and developed.
- **Further development of statistical techniques.** One interesting statistical extension to the ZIP model might be the use of a generalized inflation model that inflates other parts of a spectrum, rather than just zeros. Using a hierarchical model for inflating multiple possible densities may lead to an interesting and useful method.
- **Deeper physics-based sensor failure simulation.** For the sensor failure modeling, one interesting extension might be to obtain a more intricate model of crystal fractures to make detection more accurate. For instance, crystal fractures that are parameterized by where, geometrically, on the crystal surface fractures occur can help better simulate (and detect) such sensor failures. BA is likely useful since these types of sensor failures can be highly nonparametric and nonlinear.
- **Literature connections for Retrospective Learning.** It would be interesting to explore the connections of Retrospective Learning to the similar Machine Learning paradigm of Online Learning. Additionally, the Retrospection mechanisms that we have developed using statistical principles can be refined with information from neuroscience or cognitive psychology perspectives.
- **Non-Markovian Retrospective Learning.** Currently we have developed a HMM-based retrospection for sensor failure analysis that makes the Markov assumption on the state. We have preliminary thoughts about extending to general dynamical systems that may not make the Markov assumption.
- **Further experiments with AM-BA.** The anomaly-match strategy can be tested with additional types of anomaly detectors and match filters. The methods can also be tested on other domains where anomaly information can be used to prune search in the general space for subsequent specific testing of hypotheses by match filters.
- **Improve scalability of Augmented PCA** One of the current limitations of the Multi-Modal Augmented PCA is the requirement for maintenance of an ensemble of PCA models. While we only experimented with a limited number of PCA models and did not experience computational challenges, we hope to improve the algorithm so that it does not scale in runtime with the number of PCA covariances we want to maintain. To improve algorithmic complexity, it is likely models can be parameterized by the rotation angles of their PCs. This may help save computation by indexing the models in a way that queries only relevant models as necessary, rather than the entire ensemble, for a new data point.

10.5 Final Remarks

This thesis has prototyped the Bayesian Aggregation of evidence to detect and characterize the properties of a signal-generating source or process from multiple noisy observations. Developed methods have been applied to the problem of detecting nuclear threats using mobile spectrometers. The results hope to advance the overall state-of-the-art in analysis of robotic sensor data using machine learning.

Appendix A

A.1 Data Sets for Mobile Threat Detection Problem

Many data sets exist to benchmark approaches to the mobile threat detection problem. Table A.1 references key data sets available for experiments in the radiation sensing community and their high level statuses. Table A.2 summarizes key data properties of collected data.

Table A.1: Radiation Data Sets

Data Set	Indoors or Outdoors?	Collection Agency	Public
Sacramento Data Set	Outdoors	Lawrence Livermore National Laboratory	Not Public
RadMap Data Set	Outdoors	Lawrence Berkeley National Laboratory	Public
SUNS Data Set	Outdoors	Lawrence Berkeley National Laboratory	Public
Pittsburgh Data Set	Indoors	Auton Lab, Carnegie Mellon University	Not Public

Table A.2: Key Data Properties


Data Set	Order of Data Available	Mean / Std Photon Count Rate (counts/sec)	Localization Modality
Sacramento Data Set	<1 GB	608 / 172	GPS
RadMap Data Set	Many TB	7313 / 810	GPS
SUNS Data Set	Many TB	32 / 37	GPS
Pittsburgh Data Set	Many GB	Ongoing	Wi-Fi and Inertial Navigation

The data sets break down across several critical dimensions:


1. **Was data collected indoors or outdoors?** A key difference between data sets is whether data was collected in an indoor or outdoor environment. Certainly, photon count rates and collected radiation spectra shape differ significantly between indoor and outdoor environments since different localities have different ambient isotopes and radiation (nuisance) sources. Furthermore, radiation behaves differently in open versus cluttered environments. In outdoor environments, radiation spectra, like many other signals as Wi-Fi, tend to follow a $\frac{1}{r^2}$ attenuation law with respect to the path from the radioactive source to the spectrometer (assuming perfect visibility of the source). Indoor environments tend to be highly cluttered spaces with many opportunities for source occlusion, so the radiation spectra are affected by significant attenuation and source shielding effects. This makes separating source signal from background signal in indoor environments harder. In general, indoor surveillance of radiation sources is considered to be a more challenging problem than outdoor surveillance.
2. **Who collected the data?** Was the data collected by a university lab, government lab, industrial lab, corporation, or some other entity?
3. **Is the data set publicly available?** Some data sets are closed, and not available to the public. Others are open-source and freely available for experiments by all.
4. **What is the order of data available?** Data sets contain different numbers of collected observations. Data sets are available that are on the order of several MB, GB, and TB, containing spectroscopy and secondary data records.
5. **What are the count statistics of spectral measurements?** The data sets differ significantly in their reported photon count rates. The photon count rate of a sensing system is the average number of photon counts that hit a sensor during a measurement interval. The photon count rate is typically reported in counts / second. Data collection for different sets used different types, sizes, and numbers of spectrometers. This is important since difference techniques and tweaks to algorithms work better for different numbers of photon counts in spectra.
6. **What user localization modality was used?** Each of the data sets contains an estimate of sensor position in the world for each sensor observation. For the outdoor environment data sets, Global Positioning System (GPS) was used to estimate location. Wi-Fi localization and inertial based navigation are popular approaches for indoor settings.

A.1.1 Sacramento Data Set

The Sacramento data set contains real field data collected over a period of five consecutive days by a van driving mostly in downtown Sacramento equipped with a double 4×16 NaI planar detector. Data contains roughly 70,000 spectral measurements, reflective of background and any existing nuisance sources, collected at intervals of about 1 second each while the vehicle was in motion. Annotation data recorded for each measurement include timestamp, longitude and latitude coordinates obtained from the GPS receiver, and the current speed of the vehicle. Figure A.1a visualizes the GPS trajectories of the vehicle on Google Maps. Figure A.1b shows an example recorded spectrum.



(a) GPS Trajectory of Vehicle



(b) Example Measured Radiation Spectrum

Figure A.1: Sacramento Data Set

A.1.2 RadMap and SUNS Data Sets

The RadMap data set is a data set collected by Lawrence Berkeley Lab. The data is collected via a trunk equipped with arrays of spectrometers. This allows very precise determination of radiation artifacts on different sides of the van. Spectra aggregated across the multiple sensors have well estimated spectral shape with large numbers of received photon counts per second.

The SUNS data set (a variant of the RadMap data provided by Lawrence Berkeley National Lab and subsampled by LLNL and CMU) simulates data collected by low-cost spectrometers. The data is similar to the Sacramento data set, though has some significant differences that present key modeling challenges.

The data consists of several pedestrian trajectories where each pedestrian is equipped with their own sensor. The data is thus simulated to be from multiple mobile sensors at different geographic locations in an environment, rather than from a single sensor stream. The sensors can be of potentially different types, and each sensor has its unique calibration parameters that impact its ability to identify radioactive sources. Characterizing the sensors' properties both statically and dynamically (in the field) are important challenges when dealing with the SUNs data. Finally, the sensor data is very low in photon counts. The pedestrian-mounted sensors are much smaller in size and are less sensitive than the much more powerful gamma ray spectrometers used to collect the Sacramento data. The low count nature of the data makes the background and source component estimation problem more difficult. Our techniques for dealing with the low count sensor modeling cases are described in Chapter 7.

A.1.3 Pittsburgh Data

The Pittsburgh data set is an indoor data set of mobile spectroscopy data currently being collected at Carnegie Mellon University. Data is mostly collected in the atrium of Newell Simon Hall (NSH) by the platform described in the next section. Localization of the user is provided by a

combination of Wi-Fi localization and inertial navigation (i.e. IMU-based dead-reckoning). In the next section, we describe initial development efforts and results of our platform.

A.2 Robotic Radiation Sensing Backpack Data Platform

We have built a portable in-house platform for collecting our own radiation data. The Robotic Radiation Sensing Backpack Platform[64]¹ is a spectrometer system contained completely within a backpack that can be comfortably worn by an operator. The system is designed to be inconspicuous, for use in surveying public areas, and contains no external protruding sensors.

The system supports data collection both in indoor and outdoor venues, applying techniques from the CMU Robotics Institute, in conjunction with existing spectrometer hardware to enable effective data collection and threat detection. The backpack contains a NaI spectrometer, a small laptop, cooling system, and the entire system is controllable via an Android smartphone application. The spectrometer collects radiation spectra, while secondary hardware provides the user location. GPS provides outdoor location. Location indoors is estimated with inertial (IMU-based) navigation and cellphone Wi-Fi localization with nearby Wi-Fi access points.

Spectroscopy and location datum are time-aligned and then broadcast over cellular networks for analysis on servers in the cloud. The BA algorithm facilitates spatial aggregation of observations to provide a probability map of detected sources. The Android application provides an intuitive user interface for the user to visualize the data and monitor the state of the backpack system.

A.2.1 Backpack System Architecture




Figure A.2: System Diagram

¹Tandon, Prateek, Vladimir Ermakov, Aashish Jindia, Artur Dubrawski, Simon Labov, and Karl Nelson. Portable Radiation Monitoring Platform for Effective Nuclear Search. ARI Annual Review 2012.

Figure A.2 shows the system architectural layout. The laptop is connected to the Digibase of the spectrometer using USB. The laptop polls the spectrometer, and received data are stored in a local MySQL database. Localization data is cached locally on the laptop and then sent to the server. Calibration and Bayesian Aggregation (BA) algorithms running on the server process incoming data feeds and produce posterior source detection probability maps in real time. Secure web services make server data available to mobile devices for visualization.

Processed data from the server is retrieved by a mobile application at regular intervals to render a constantly updating heat map of possible threats (Figure A.3a). The Android Application provides an interface to observe the live histogram of radiation counts at different time windows, allowing users to assess threats in real time (Figure A.3b).




Figure A.3: Screenshots of Android Application Capabilities

A.2.2 Experiments with Indoor Localization

The most recent revision of the backpack system utilizes mostly inertial navigation using IMU-based dead-reckoning to provide the user's indoor location. The capabilities of this approach to indoor localization were evaluated, and resulting threat detection capabilities were showcased in an indoor source detection scenario for the Domestic Nuclear Detection Office (DNDO) [20].

Previously, we experimented with Wi-Fi localization in the backpack system. A custom Wi-Fi localizer was trained using a fingerprinting algorithm for RSSI signal strengths from wifi access points inside NSH at CMU. RSSI signal strength data from 100 building accessing points was collected for a discrete set of locations. SKL-based feature selection was employed to select a stable and reliable set of access point features to use to predict each location. Such feature selection helped boost the accuracy of the localization as shown in Figure A.4 on a test set of discrete locations in the building. Localizer evaluation on a testing set of locations shows a median positioning error of 3.81 m.





Figure A.4: Wi-Fi Localization Accuracy

The robustness of BA was tested with respect to location uncertainty. 1000 worlds were simulated, each with a point source. Location error was simulated as Gaussian perturbations in trajectory locations with mean and std based on tolerances of our localization algorithm.



(a) Detection with Location Uncertainty

(b) Accuracy of Source Localization

Figure A.5: Estimated Detection and Localization Capability of Wi-Fi Localizer

Our results in Figure A.5 indicate that BA can be resilient to location uncertainty. As Figure A.5a shows, ROC curve detection performance remains accurate even with errors in positioning of collected measurements. As Figure A.5b, it is still possible to localize radioactive point sources in the environment with high accuracy even with modest location uncertainty of the user.

A.3 Sensor Calibration and Energy Windowing

Spectrometer software polls spectra in 1024 channel space. However, spectral anomaly detection is often performed in a space of 128 quadratically-binned energy levels which emphasize key

energies. A calibration process is required to convert between channel space to energy space, as well as to correct for data nuisances like gain drift.

Figure A.6 shows the first step of calibration: collecting data sets of sensor readings for back-ground and also in the presence of harmless radioactive test sources (called “check sources”).




Figure A.6: 24 Hours of Collected Background and Source Isotope data

Six total data sets of radiation spectra were collected - five for check sources and one for characteristic background. Each data set was collected by leaving the sensor running for 24 hours in the presence of a check source (or just background). Available check source isotopes were Cesium-137, Cobalt-57, Cobalt-60, Cadmium-109, and Barium-133. Photon counts were binned by the hour for each data set. This resulted in 24 spectra per data set, with a total of 144 collected spectra. The collected 24 spectra for each case (check source or background) are shown in Figure A.6.

There is visible drift in the channels over time as the 24 curves in each data set do not perfectly align. Gain drift in sensor readings is expected with changes in environmental temperature, lighting, humidity, etc. A sensor can be calibrated for such gain drift using the Dynamic Gain Estimation (DGE) algorithm. The DGE aligns spectra overtime with respect to known physical peak features that are expected to exist in all spectra. For example, Thorium is a known physical feature that occurs at around channel 750-800 in background and all check sources in 1024-channel space. It is a feature that observations can be aligned towards to calibrate out gain drift.

Given a set of collected radiation spectra $\mathbf{R} = \{R_1, \dots, R_N\}$ (for us $N = 144$) and a set of M physical peak features $\mathbf{F} = \{f_1, \dots, f_M\}$, DGE computes a corrective “gain” for each spectrum to rebin and align it with all the others. A key step is to compute a matrix of sub-gains g_{R_i, f_j} from which the overall gain $G(R_i)$ for each measurement is computed. The algorithmic steps involved in computing a sub-gain for a radiation spectrum are shown in Figure A.7 for an example Thorium peak feature.




Figure A.7: DGE Gain Computation Steps for Example Thorium Feature

The first step of the DGE algorithm is to compute a window around the feature of interest (shown in blue) in the spectrum. A window start $W_s(R_i, f_j)$ and window end $W_E(R_i, f_j)$ is specified in the energy space for where the feature is likely to exist. In our example, the chosen range is roughly [760,820] in the 1024-channel space. To robustify the window selection, the window can be algorithmically re-centered and resized around the max value.

The next step is to fit a line to the end points of the feature points in the window. In our example, the fit line is shown in green. Since the ends of the curve may be subject to spurious noise in channels, five points on each endpoint are used to seed a robust linear fit. The Gaussian

component (shown in red) of the signal is extracted via subtracting the original curve from the line.

Finally, the centroid of the Gaussian component is computed (shown as the red dot). The centroid is the sub-gain estimate, g_{R_i, f_j} , for this spectrum and feature and is defined as:

$$g_{R_i, f_j} = \frac{\sum_{k=W_S(R_i, f_j)}^{W_E(R_i, f_j)} p_{R_i}(k) \times k}{\sum_{k=W_S(R_i, f_j)}^{W_E(R_i, f_j)} p_{R_i}(k)} \quad (\text{A.1})$$

where k is an energy channel in the energy window of size $|W_E(R_i, f_j) - W_S(R_i, f_j)|$ and $p(k)$ is sum of counts in the Gaussian component at energy channel k . Using this procedure, we obtain the full 2D matrix of sub-gains $\mathbf{g} = g_{R_i, f_j}$. The next step is to fuse the sub-gain estimates to obtain an overall gain $G(R_i)$ for a measurement. The overall gain $G(R_i)$ is defined as the arithmetic mean of the individual sub-gain estimates:

$$G(R_i) = \frac{\sum_{k=1}^M \frac{g_{R_i, f_m}}{\text{var}(g_{f_m})}}{\sum_{k=1}^M \frac{1}{\text{var}(g_{f_m})}} \quad (\text{A.2})$$

where $\text{var}(g_{f_m})$ is the variance of the distribution of the gain values of feature f_m for all radiation spectra in the collection.

For our implementation, all $N = 144$ spectra were aligned with respect to a single Thorium peak around channel 750 in the background. The results are shown in Figure A.8.




Figure A.8: Spectra Before and After Gain Drift Calibration with DGE

As Figure A.8 shows, the DGE algorithm greatly reduces the misalignment variation that exists in the raw data. The output (Figure A.8b) of DGE shows cleaner, more aligned spectra.

Once spectra are aligned, a spline mapping 1024 channel-space to 128 energy-space can be fitted to the extracted spectral physical peak features. To robustly estimate the spline, we use a combination of both empirical estimation from data and well-known reference information in radiation physics.

Table A.3 shows a mapping of the channels of the physical peak features f_i to their well-known energies created using the Sodium-Iodine Catalog (NAICAT) [1]. NAICAT is a popular radiation physics isotope reference guide freely available on the web.

Table A.3: Channel-Energy Features

Spectra	Feature Name	Approx. Feature Channel Location	Energy	Source
Background	Thorium peak	750	2614	NAICAT, pg 85
Background	K40 peak	450	1460.83	NAICAT, pg 85
Cs-137	Thorium peak	800	2614	NAICAT, pg 255
Cs-137	K40 peak	450	1460.83	NAICAT, pg 255
Cs-137	662 KeV peak	200	661.657	NAICAT, pg 255
Cs-137	x-ray peak	10	31.8	NAICAT, pg 255
Co-57	Thorium peak	700	2614	NAICAT, pg 115
Co-57	K40 peak	450	1460.83	NAICAT, pg 115
Co-57	136 KeV peak	50	136	NAICAT, pg 115
Co-60	Thorium peak	800	2614	NAICAT, pg 120
Co-60	K40 peak	450	1460.83	NAICAT, pg 120
Co-60	1170 KeV peak	350	1170	NAICAT, pg 120
Co-60	1132 KeV peak	400	1132	NAICAT, pg 120
Cd-109	Thorium peak	800	2614	NAICAT, pg 277
Cd-109	K40 peak	450	1460.83	NAICAT, pg 277
Ba-133	Thorium peak	800	2614	NAICAT, pg 362
Ba-133	K40 peak	450	1460.83	NAICAT, pg 362
Ba-133	81 KeV peak	30	81	NAICAT, pg 362
Ba-133	356 KeV peak	120	356	NAICAT, pg 362

In constructing the spline to map channels to energies, we use the energies for peak features specified in Table A.3. Instead of using the channel values for peak features in the table, however, we use the estimated channel locations in the aligned empirical spectra. For bounding the endpoint behavior of the spline, it is helpful to include a point linear extrapolated out twice from the last two available data points as well the starting point (1,0). These guidelines help ensure a well-behaved spline shape.

The produced calibration spline from energy to channels is shown below in Figure A.9a. The spline has a characteristic shape that can be compared to a reference “ground truth” spline (Figure A.9b). The left-most plot is similar to the ground truth shape, validating the correctness of our constructed spline.




Figure A.9: Constructed Spline Result

The inverse spline from channels to energy was constructed via linear interpolation. The inverse map allows visualization of how the 1024 channels map to energy (Figure A.10).




Figure A.10: Mapping of 1024 channel-space of sensor to energy (“EnergyInstrument” profile)

Given a spectrum R_{New} in 1024 channel-space, converting to energy-space is a straightforward algorithm:

1. Align R_{New} with respect to the reference spectra using the DGE algorithm.
2. Use the rebin function as follows:

$$rebinnedSpectra = rebin(R_{New}, EnergyInstrument, gain * EnergyFeatures) \quad (A.3)$$

where $EnergyInstrument$ is the inverse spline mapping and $EnergyFeatures$ is the quadratic binning of energies required for the spectral anomaly detector.

The net effect is to remap and reduce the 1024 energies of the instrument to the 128 energy features useful for spectral anomaly detection. An example of this process is shown in Figure A.11.




Figure A.11: The sensor receives a new background spectrum in 1024 channel space (Figure A.11a). The DGE procedure aligns the new spectrum with reference spectra (Figure A.11b). Calibration produces a 128-dimensional energy-space vector (Figure A.11c).

A.3.1 Energy Selection Heuristic for Match Filters

If the source template is known, the sensor can be tuned towards energies where source signal is likely to manifest. Using an energy bin selection heuristic, the energy bins that maximize the Signal-to-Noise Ratio (SNR) between a given source template and mean background can be identified efficiently. We use the following Linear-time Subset Selection algorithm as a heuristic for finding promising energy levels that exhibit high SNR for a known source template.


Linear-time Subset Selection algorithm:

1. Compute SNR between the source template and mean background vector for each of the full set of energy bins.
2. Rank the energy bins in decreasing order of expected SNR.
3. Step through the sorted list ($i=1:128$) and compute the current SNR as the sum of the SNRs from $1:i$.
4. Terminate search when the SNRs of bins $1:i$ ceases to increase.

The output is a set of selected energy bins (y) that exhibit high SNR between the source template and the mean background.

A.4 Radiation Point Source Simulation Methods

All collected base data sets are presumed to be mostly background radiation with the occasional nuisance source. A high-fidelity source simulator is used to inject user-supplied synthetic radiation profiles into collected background radiation data to simulate the presence of point sources. Given a geographic subregion, the simulator chooses a random location which is within a prescribed distance to at least one measurement, and adds simulated source injection counts to the pre-existing background measurements. The simulator models injection counts, taking into account relative distance to the source, velocity of the sensor-carrying vehicle, duration of the measurement interval, and the Poisson distribution of photon counts. Figure A.12 shows the simulator capabilities.



(a) Example Source Template Injection into Background Spectrum (b) Example Injection of Point Source (Geographic View)

Figure A.12: Snapshot of Source Simulator Capabilities

Figure A.12a shows a typical background radiation spectrum measurement, a fissile materials source template, and the additive signal that results with injection. Figure A.12b shows the effect of our injection procedure geographically on a sample geographic region. Our simulator also

provides the capability to specify source intensity, apply various anisotropic configurations, and inject various source types.

A.5 Principal Component Analysis Derivation

Principal component analysis is a popular technique for explaining variability in a data set in terms of its top principal directions of variation. The derivation shown here is adapted from [7].

The input to Principal Component Analysis (PCA) is a set of observations, $\{x_n\}$ where x_n is a Euclidian variable vector with dimensionality D and $n = 1, \dots, N$. Each data point can be projected onto a scalar value $u_1^T x_n$. The mean of the projected data is $u_1^T \bar{x}$ where $\bar{x} = \frac{1}{N} \sum_{i=1}^N x_n$, the sample mean. The variance of the projected data is given by $S = \frac{1}{N} \sum_{n=1}^N (x_n - \bar{x})(x_n - \bar{x})^T$.

In order to maximize the projected variance $u_1^T S u_1$ with respect to u_1 we need to constrain u_1 to prevent $\|u_1\| \rightarrow \infty$. Thus, we use $u_1^T u_1 = 1$ as the normalization constraint. Our optimization problem is thus:

$$\begin{aligned} \max \quad & u_1^T S u_1 \\ \text{subject to} \quad & u_1^T u_1 = 1. \end{aligned}$$

The optimization can be performed via forming the Lagrangian $L(u_1, S, \lambda)$ and maximize with respect to u_1 by setting $\frac{\partial L}{\partial u_1} = 0$.

$$\begin{aligned} L(u_1, S, \lambda) &= u_1^T S u_1 + \lambda(1 - u_1^T u_1) \\ \frac{\partial L}{\partial u_1} &= S u_1 - \lambda u_1 = 0 \\ S u_1 &= \lambda u_1 \end{aligned} \tag{A.4}$$

The result says that u_1 must be an eigenvector of S . If we left multiply the equation, then we obtain:

$$u_1^T S u_1 = \lambda_1 \tag{A.5}$$

This result states the variance will be maximum where we set u_1 equal to the eigenvector having the largest eigenvalue. This eigenvector is commonly known as the “first principal component.” It can be shown that finding the next principal components can be achieved via an iterative process. Rigorous proof of this requires induction, and is too horrible to consider so we’ll omit it.

The algorithm involves first doing mean normalization and feature scaling

$$\begin{aligned} u_j &= \frac{1}{N} \sum_{i=1}^N x_i(j) \\ x_j &= \frac{x_j - u_j}{s_j} \end{aligned} \tag{A.6}$$

where s_j is a scale factor. Note that MATLAB does zero meaning for training data but NOT feature scaling by default. One can use MATLAB's flags to do various types of normalization.

The covariance matrix Σ is computed as

$$\Sigma = \frac{1}{N} \sum_{i=1}^N x_i x_i^T \quad (\text{A.7})$$

Then the eigenvectors of Σ is computed via Singular Value Decomposition (SVD). Some top components are kept from the matrix containing the eigenvectors of Σ .

A.6 Spectral Anomaly Detector Derivation

The spectral anomaly detector is a popular PCA variant used for modeling radiation background [45]. The input to the training procedure is an $N \times D$ matrix d of N D -dimensional observations and a number PC of principal components to keep. First, filtering is applied in energy to the data, and the data is smoothed via a 10s rolling window. Second, a special covariance matrix Σ that retains 0.01 of the mean is computed.

$$\Sigma = \frac{dd^T}{D} - 0.99mm^T \quad (\text{A.8})$$

where $m_j = \sum_{i=1}^N d_i(j)$. This is different from standard PCA in that standard PCA usually zero means the data. Next, we form the design matrix A and its inverse A^{-1} .

$$\begin{aligned} A &= \text{diag}\left(\frac{1}{\sqrt{\text{diag}(\Sigma) + 1}}\right) \\ A^{-1} &= \text{diag}(\sqrt{\text{diag}(\Sigma) + 1}) \end{aligned} \quad (\text{A.9})$$

Afterwards, the correlation matrix C is created using

$$C = A\Sigma A \quad (\text{A.10})$$

and SVD is performed on the correlation matrix to obtain the matrix of eigenvectors U . Finally, the null-space transform is applied to create our basis matrix T

$$T = I_{D \times D} - A^{-1}U_{PC}U_{PC}^T A \quad (\text{A.11})$$

where U_{PC} keeps the top PC principal components of U . The final T matrix will be $D \times D$.

To evaluate the spectral score of a new set of x_{test} observations (a $Y \times D$ matrix), the following procedure is used. The observations are once again filtered in energy. Then, we take the L_2 norm of the residuals to estimate the source signal component.

$$\xi = ||Tx_{test}^T||_2 \quad (\text{A.12})$$

The poisson noise, σ , is estimated via the sum of counts in the measurement

$$\sigma = \sum_{j=1}^D x_{test,j} \quad (\text{A.13})$$

The final spectral score is defined as the Signal-to-Noise Ratio (SNR):

$$SNR = \frac{\xi}{\sigma} \quad (\text{A.14})$$

A.7 Poisson Principal Component Analysis Derivation

Poisson PCA is a specific form of the more general Exponential PCA (EPCA) model. The EPCA model will be described first, and then the derivation of the Poisson PCA algorithm (adapted from [11] and [53]) will be shown.

A.7.1 General EPCA

An Exponential PCA model represents data b in terms of a link function f , a basis matrix U , and a low dimensional weight vector $\tilde{b}$. The use of a link function generalizes PCA to nonlinear models.

$$b \approx f(U\tilde{b}) \quad (\text{A.15})$$

A Bregman associated with the strictly convex function F is a type of distance formula between two points, p and q .

$$D_F^q(p||q) = F(p) - F(q) - \langle \nabla F(q), p - q \rangle \quad (\text{A.16})$$

Intuitively, the formula is the difference between the value of F at point p and the value of the first-order Taylor expansion of F around point q when evaluated at point p . Additional properties of the divergence are that it is guaranteed to be convex, linear, and nonnegative.

The EPCA cost function is a Bregman divergence between the high dimensional representation b and low dimensional representation of the data $U\tilde{b}$

$$B_{F^*}(b||U\tilde{b}) = F(U\tilde{b}) - bU\tilde{b} + F^*(b) \quad (\text{A.17})$$

where F is a convex function whose derivative is f and F^* is the convex dual of F .

A.7.2 PPCA Derivation

For the Poisson distribution, the chosen link function is

$$f(U\tilde{b}) = e^{U\tilde{b}} \quad (\text{A.18})$$

which corresponds to a Poisson error model for each component of the reconstructed data. This link function constraints the low dimensional representation, $e^{U\tilde{b}}$, to be positive. The variance of each belief component is proportional to the expected value. Thus, we have:

$$\begin{aligned} F(U\tilde{b}) &= \sum e^{U\tilde{b}} \\ F^*(b) &= b \ln b - \sum b \\ B_{F^*}(b||U\tilde{b}) &= \sum e^{U\tilde{b}} - bU\tilde{b} + b \ln b - \sum b = UKL(b||U\tilde{b}) \end{aligned} \quad (\text{A.19})$$

The first equation corresponds to integrating f to form F . The second equation forms F^* , the convex dual of F . The final result states that the divergence is equivalent to minimizing the Unnormalized Kullback-Leibler Divergence (UKL) between the original data and its reconstruction.

Thus, we can rewrite our cost function as:

$$L(B, U, \tilde{B}) = \sum_{i=1}^{|B|} e^{U\tilde{b}_i} - b_i U \tilde{b}_i \quad (\text{A.20})$$

The derivatives with respect to U and $\tilde{B}$ are:

$$\begin{aligned} \frac{\partial L}{\partial U} &= (e^{U\tilde{B}} - B) \tilde{B}^T \\ \frac{\partial L}{\partial \tilde{B}} &= U^T (e^{U\tilde{B}} - B) \end{aligned} \quad (\text{A.21})$$

Set

$$q(\tilde{B}_j) = U^T (e^{U\tilde{B}_j} - B_j) \quad (\text{A.22})$$

and the linearize about $\tilde{B}_j$ to find the roots of q .

$$\begin{aligned} \tilde{B}_j^{new} &= \tilde{B}_j - \frac{q(\tilde{B}_j)}{q'(\tilde{B}_j)} \\ \tilde{B}_j^{new} - \tilde{B}_j &= -\frac{U^T (e^{U\tilde{B}_j} - B_j)}{q'(\tilde{B}_j)} \end{aligned} \quad (\text{A.23})$$

Given a real-valued function f , its derivative f' , and a guess for its root x_n , we can iteratively compute the next guess for the root via Newton's method:

$$\begin{aligned} x_{n+1} &= x_n - \frac{f(x_n)}{f'(x_n)} \\ x_{n+1} - x_n &= -\frac{f(x_n)}{f'(x_n)} \end{aligned} \quad (\text{A.24})$$

An expression for $\frac{\partial q}{\partial \tilde{B}_j}$ can be found via

$$\frac{\partial q}{\partial \tilde{B}_j} = \frac{\partial}{\partial \tilde{B}_j} (U^T (e^{U\tilde{B}_j} - B_j)) = U^T D_j U \quad (\text{A.25})$$

where $D_j = \text{diag}(e^{U\tilde{B}_j})$.

So now we have:

$$\begin{aligned} \tilde{B}_j^{new} - \tilde{B}_j &= -\frac{U^T (e^{U\tilde{B}_j} - B_j)}{U^T D_j U} \\ (U^T D_j U)(\tilde{B}_j^{new} - \tilde{B}_j) &= U^T (B_j - e^{U\tilde{B}_j}) \\ \tilde{B}_j^{new} &= \frac{U^T D_j (U\tilde{B}_j + D_j^{-1}(B_j - e^{U\tilde{B}_j}))}{U^T D_j U} \end{aligned} \quad (\text{A.26})$$

This last equation corresponds to a weighted least squares problem that is solvable by standard linear algebra techniques. To ensure numerical stability, we add a regularization term to the divisor:

$$\tilde{B}_j^{new} = \frac{U^T D_j (U\tilde{B}_j + D_j^{-1}(B_j - e^{U\tilde{B}_j}))}{U^T D_j U + C_1 I_l} \quad (\text{A.27})$$

where I_l is the $l \times l$ identity matrix and C_1 is a regularization constant. The recommended value for C_1 is 10^{-5} .

Similarly, the result can be derived for U_i^{new} :

$$U_i^{new} = \frac{(U_i \tilde{B} + (B_i - e^{U_i \tilde{B}}) D_i^{-1}) D_i \tilde{B}^T}{(\tilde{B} D_i \tilde{B}^T + C_2 I_l)} \quad (\text{A.28})$$

where C_2 is the regularization constant (again, the recommended value is 10^{-5}).

A.7.3 Poisson PCA Final Algorithm

The results in the last section lead to the following algorithm:

1. Collect sample of data B into matrix.
2. Fix an estimate for $\tilde{B}$ and U randomly.
3. Do
 - For each column of $B_j \in \tilde{B}$, compute $\tilde{B}_j^{new}$ using current estimate of U .
 - For each column of $U_i \in U$, compute U_i^{new} using current estimate of $\tilde{B}$.
 - while $L(B, U, \tilde{B}) > \epsilon$
4. Reconstruct original data via $x = e^{U\tilde{B}}$. Poisson PCA Score is minimal $L(B, U, \tilde{B})$ at convergence.

A.7.4 Zero-Inflated Poisson Regression Match Filter

The Zero-Inflated Poisson (ZIP) model is a way of robustly fitting a Poisson Regression to sparse Poisson data with possible excess zeros.

On the radiation data, our ZIP model uses a two-step hierarchal approach. Logistic regression is used to first classify the presence of non-zero counts from predictor energy bins. If the spectrum is predicted to have non-zero background counts in the source window, then Poisson Regression is run to predict the amount. The separate probability densities for the zero count and non-zero count cases help prevent over-dispersion when estimating the mean parameter of the Poisson distribution for the case of expected non-zero source signal.

Here are the equations for the ZIP model:

$$P(y_j = 0) = \pi + (1 - \pi)e^{-\lambda} \quad (\text{A.29a})$$

$$P(y_j = h_i) = (1 - \pi) \frac{\lambda^{h_i} e^{-\lambda}}{h_i!}, h_i \geq 1 \quad (\text{A.29b})$$

We can use a ZIP model for the match filter estimator instead of the vanilla Gaussian Linear Regression. The following section provides code for training a ZIP match filter in MATLAB.

MATLAB Code for ZIP Regression

```

function ZIP_MODEL = trainZIP(background_train,W_t)

%inputs: background_train is a Nx128 matrix of
%pure background radiation data.
%W_t is a 128x1 vector specifying the energy window.
%W_t==1 in energy bins in the window and
%W_t==0 in energy bins outside the energy window

%% train zip model

%assemble sum of counts inside and outside window for each data point.
preds = sum(background_train(:,W_t==0),2);
windows = sum(background_train(:,W_t==1),2);

%train a logistic regression model to classify whether we get zero or
%nonzero sum of counts inside the window from the sum of counts outside
%the window.
windowed_nominal = windows>0;
%train logistic regression to classify as nonzero or not
[B,dev,stats] = mnrfits(preds,windowed_nominal+1);

%fit poisson regression to data points where we expect non-zero counts
%inside the window.
b_poiss = glmfit(background_train(windows>0,W_t==0),...
    windows(windows>0),'poisson');

%store variables and return structure
ZIP_MODEL.B = B;
ZIP_MODEL.b_poiss = b_poiss;
ZIP_MODEL.W_t = W_t;

```



```

function scores = testZIP(trajRad,ZIP_MODEL)
%inputs: trajRad is a Nx128 matrix of test spectra data points
%ZIP_MODEL is the model structure created in the trainZIP function

%assemble sum of counts inside and outside window
fsel = sum(trajRad(:,MODEL.W_t==0),2);
Y_i = sum(trajRad(:,MODEL.W_t==1),2);

%use logistic regression model to classify whether we get a non-zero
%prediction for counts inside the window from counts outside the window.
nonzero_pred = mnrvl(MODEL.B,fsel);
[unused,zero_indic]=max(nonzero_pred');

%if we do expect background counts, use the poisson regression to predict
%the number of background counts inside the window.
%If we don't expect any background counts, set the background estimate
%to 0.
B_i = glmval(MODEL.b_poiss,trajRad(:,MODEL.W_t==0),'log');
B_i(zero_indic==1) = 0;

%just make sure we don't ever get negative counts and that we
%never predict more counts than we have in our original
%spectrum.
B_i(B_i<0)=0;
B_i(B_i > Y_i) = Y_i(B_i > Y_i);

%Y_i-B_i is our estimate of source signal. It is basically
%      sum of source counts =
%      (sum of all counts inside spectrum)
%      -
%      (sum of estimated background counts)
%use our source estimate to compute the SNR = source / sqrt(background).
%To make sure that our code doesn't blow up for dividing by 0, use
%sqrt(background + 1) in the denominator in practice.
scores = (Y_i-B_i) ./ sqrt(B_i+1);

%just a sanity check to make sure our SNR scores are never negative.
%It can happen sometimes with a regression so threshold at 0.
scores(scores<0) = 0;

```


Bibliography

- [1] NAICAT. URL <http://www.inl.gov/gammaray/catalogs/pdf/naicat.pdf>. A.3
- [2] Nrc: Fact sheet on dirty bombs. URL <http://www.nrc.gov/reading-rm/doc-collections/fact-sheets/fs-dirty-bombs.html>. 1.2
- [3] Fukushima accident. URL <http://www.world-nuclear.org/info/Safety-and-Security/Safety-of-Plants/Fukushima-Accident/>. 1.2
- [4] Zoë; Abrams, Ashish Goel, and Serge Plotkin. Set k-cover algorithms for energy efficient monitoring in wireless sensor networks. In *IPSN '04: Proceedings of the third international symposium on Information processing in sensor networks*, pages 424–432, New York, NY, USA, 2004. ACM Press. ISBN 1-58113-846-6. doi: <http://doi.acm.org/10.1145/984622.984684>. URL <http://portal.acm.org/citation.cfm?id=984684&dl=GUIDE&coll=GUIDE&CFID=15151515&CFTOKEN=6184618>.
- [5] M A Batalin, G S Sukhatme, and M Hattig. Mobile robot navigation using a sensor network. *IEEE International Conference on Robotics and Automation 2004 Proceedings ICRA 04 2004*, 1(3):636—641 Vol.1, 2004. URL <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1307220>.
- [6] Maxim Batalin, Gaurav S Sukhatme, Myron Hattig, et al. Mobile robot navigation using a sensor network. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 1, pages 636–641. IEEE, 2004.
- [7] Christopher M Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New York, 2006. 3.1.1, A.5
- [8] Konstantin N Borozdin, Alexei V Klimenko, William C Priedhorsky, Nicolas Hengartner, Charles C Alexander, R Andres Cortez, Herbert G Tanner, and Xanthi Papageorgiou. Optimized strategies for smart nuclear search. In *Nuclear Science Symposium Conference Record, 2006. IEEE*, volume 2, pages 926–928. IEEE, 2006.
- [9] K Mani Chandy, Julian Bunn, and A Liu. Models and algorithms for radiation detection. In *Modeling and Simulation Workshop for Homeland Security*, pages 1–6, 2010. 1.2.2
- [10] Zhe Chen. Bayesian filtering: From kalman filters to particle filters, and beyond. *Statistics*, 182(1):1–69, 2003. 1.2.2
- [11] Michael Collins, Sanjoy Dasgupta, and Robert E Schapire. A generalization of principal

- components analysis to the exponential family. In *Advances in neural information processing systems*, pages 617–624, 2001. 3.1.3, 7.1.2, 7.1.3, A.7
- [12] Thomas H Cormen, Charles E Leiserson, Ronald L Rivest, Clifford Stein, et al. *Introduction to algorithms*, volume 2. MIT press Cambridge, 2001. 6.2.1
 - [13] J Cortes. Area-constrained coverage optimization by robotic sensor networks. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 1018–1023, 2008. doi: 10.1109/CDC.2008.4738883.
 - [14] Joshua Reich Department and Joshua Reich. Robot-Sensor Networks for Search and Rescue. In *In Proc. IEEE Intl Workshop on Safety, Security and Rescue Robotics*, 2006.
 - [15] Amol Deshpande, Carlos Guestrin, Samuel R Madden, Joseph M Hellerstein, and Wei Hong. Model-driven data acquisition in sensor networks. In *Proceedings of the Thirtieth international conference on Very large data bases-Volume 30*, pages 588–599. VLDB Endowment, 2004.
 - [16] Amol Deshpande, Samir Khuller, Azarakhsh Malekian, and Mohammed Toossi. Energy Efficient Monitoring in Sensor Networks. In Eduardo Laber, Claudson Bornstein, Loana Nogueira, and Luerbio Faria, editors, *LATIN 2008: Theoretical Informatics*, volume 4957 of *Lecture Notes in Computer Science*, pages 436–448. Springer Berlin / Heidelberg, 2008. ISBN 978-3-540-78772-3. URL http://dx.doi.org/10.1007/978-3-540-78773-0_38.
 - [17] Ma Di and Er Meng Joo. A survey of machine learning in Wireless Sensor networks From networking and application perspectives. In *Information, Communications Signal Processing, 2007 6th International Conference on*, pages 1–5, 2007. doi: 10.1109/ICICS.2007.4449882.
 - [18] Ma Di and Er Meng Joo. A survey of machine learning in wireless sensor networks from networking and application perspectives. In *Information, Communications & Signal Processing, 2007 6th International Conference on*, pages 1–5. IEEE, 2007.
 - [19] Arnaud Doucet and Adam M Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. *Handbook of Nonlinear Filtering*, 12:656–704, 2009. 1.2.2, 2.2
 - [20] Artur Dubrawski, Simon Labov, Jeff Schneider, Karl Nelson, Peter Huggins, Prateek Tandon, Rob MacLachlan, and Jay Jin. Presentation: Machine learning for effective nuclear search and broad area monitoring. ARI Annual Review, 2014. 9.1.1, A.2.2
 - [21] Stefan Funke, Alex Kesselman, Fabian Kuhn, Zvi Lotker, and Michael Segal. Improved approximation algorithms for connected sensor cover. *Wireless networks*, 13(2):153–164, 2007.
 - [22] Deepak Ganesan, Răzvan Cristescu, and Baltasar Beferull-Lozano. Power-efficient sensor placement and transmission structure for data gathering under distortion constraints. *ACM Transactions on Sensor Networks (TOSN)*, 2(2):155–181, 2006.
 - [23] Stefano Ghirlanda. Retrospective revaluation as simple associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 31(1):107, 2005. 9.1.1
 - [24] H González-Banos. A randomized art-gallery algorithm for sensor placement. In *Proceed-*

- ings of the seventeenth annual symposium on Computational geometry*, SCG '01, pages 232–240, New York, NY, USA, 2001. ACM. ISBN 1-58113-357-X. doi: 10.1145/378583.378674. URL <http://doi.acm.org/10.1145/378583.378674>.
- [25] C Guestrin, P Bodik, R Thibaux, M Paskin, and S Madden. Distributed regression: an efficient framework for modeling sensor network data. In *Information Processing in Sensor Networks, 2004. IPSN 2004. Third International Symposium on*, pages 1–10, April 2004. doi: 10.1109/IPSN.2004.1307317.
 - [26] Carlos Guestrin, Peter Bodik, Romain Thibaux, Mark Paskin, and Samuet Madden. Distributed regression: an efficient framework for modeling sensor network data. In *Information Processing in Sensor Networks, 2004. IPSN 2004. Third International Symposium on*, pages 1–10. IEEE, 2004.
 - [27] H Gupta, Zongheng Zhou, S R Das, and Q Gu. Connected sensor cover: self-organization of sensor networks for efficient query execution. *Networking, IEEE/ACM Transactions on*, 14(1):55–67, 2006. ISSN 1063-6692. doi: 10.1109/TNET.2005.863478.
 - [28] Himanshu Gupta, Zongheng Zhou, Samir R Das, and Quinyi Gu. Connected sensor cover: self-organization of sensor networks for efficient query execution. *Networking, IEEE/ACM Transactions on*, 14(1):55–67, 2006.
 - [29] Trevor Hastie, Robert Tibshirani, Jerome Friedman, T Hastie, J Friedman, and R Tibshirani. *The elements of statistical learning*, volume 2. Springer, 2009. 4.1
 - [30] Jan De Houwer and Tom Beckers. Higher-order retrospective revaluation in human causal learning. *The Quarterly Journal of Experimental Psychology: Section B*, 55(2):137–151, 2002. 9.1.1
 - [31] Peter Huggins, Prateek Tandon, Artur Dubrawski, Simon Labov, and Karl Nelson. Poster abstract: Dynamic placement of sensors for rapid characterization of radiation threat. DTRA Review, 2013. 5.1, 5.1.3, 5.2, 9.3.2
 - [32] Glenn F Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010. 3
 - [33] Andreas Krause, Carlos Guestrin, Anupam Gupta, and Jon Kleinberg. Near-optimal sensor placements: maximizing information while minimizing communication cost. In *Proceedings of the 5th international conference on Information processing in sensor networks*, IPSN '06, pages 2–10, New York, NY, USA, 2006. ACM. ISBN 1-59593-334-4. doi: 10.1145/1127777.1127782. URL <http://doi.acm.org/10.1145/1127777.1127782>.
 - [34] Andreas Krause, Ajit Singh, and Carlos Guestrin. Near-Optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *J. Mach. Learn. Res.*, 9:235–284, June 2008. ISSN 1532-4435. URL <http://dl.acm.org/citation.cfm?id=1390681.1390689>.
 - [35] Andreas Krause, Ram Rajagopal, Anupam Gupta, and Carlos Guestrin. Simultaneous Optimization of Sensor Placements and Balanced Schedules. *IEEE Transactions on Automatic Control*, 56(10):2390–2405, October 2011.
 - [36] John K Kruschke. Locally bayesian learning with applications to retrospective revaluation

and highlighting. *Psychological review*, 113(4):677, 2006. 9.1.1

- [37] John K Kruschke. *Doing Bayesian data analysis : a tutorial with R and BUGS*. Academic Press, Burlington, MA, 2011. ISBN 9780123814852 0123814855. URL http://www.worldcat.org/search?qt=worldcat_org_all&q=0123814855.
- [38] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 6.2.1
- [39] Anthony LaMarca, Waylon Brunette, David Koizumi, Matthew Lease, Stefan B Sigurdsson, Kevin Sikorski, Dieter Fox, and Gaetano Borriello. Making sensor networks practical with robots. In *Pervasive Computing*, pages 152–166. Springer, 2002.
- [40] A.H. Liu, J.J. Bunn, and K.M. Chandy. Sensor networks for the detection and tracking of radiation and other threats in cities. In *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pages 1–12, April 2011. 1.2.3
- [41] Guoqiang Mao, Barış Fidan, and Brian DO Anderson. Wireless sensor network localization techniques. *Computer networks*, 51(10):2529–2553, 2007.
- [42] Dinesh P. Mehta and Sartaj Sahni. *Handbook Of Data Structures And Applications (Chapman & Hall/Crc Computer and Information Science Series.)*. Chapman & Hall/CRC, 2004. ISBN 1584884355. 4.1, 5.1.3
- [43] M. Morelande, B. Ristic, and A. Gunatilaka. Detection and parameter estimation of multiple radioactive sources. In *Information Fusion, 2007 10th International Conference on*, pages 1–7, 2007. doi: 10.1109/ICIF.2007.4408094. 1.2.2
- [44] M.R. Morelande and B. Ristic. Radiological source detection and localisation using bayesian techniques. *Signal Processing, IEEE Transactions on*, 57(11):4220–4231, 2009. ISSN 1053-587X. doi: 10.1109/TSP.2009.2026618. 1.2.2
- [45] Karl Nelson and Simon Labov. Detection and alarming with sords unimaged data: Background data analysis. *Lawrence Livermore National Lab Technical Report*, 2009. 1.2.2, 3.1.1, 3.1.2, 7.1.1, 8.3.2, A.6
- [46] Karl Nelson and Simon Labov. Aggregation of Mobile Data. *Lawrence Livermore National Lab Technical Report*, 2012. 1.2.2, 2.3, 3.2.2, 7.1.1, 10.2
- [47] Tom Nichols, Douglas Stuart, and Jeffrey D McCausland. Tactical nuclear weapons and nato. Technical report, DTIC Document, 2012. 1.2
- [48] David Portugal and Rui P Rocha. On the performance and scalability of multi-robot patrolling algorithms. In *Safety, Security, and Rescue Robotics (SSRR), 2011 IEEE International Symposium on*, pages 50–55. IEEE, 2011.
- [49] Lawrence Rabiner and Biing-Hwang Juang. An introduction to hidden markov models. *ASSP Magazine, IEEE*, 3(1):4–16, 1986. 9.1.2
- [50] Tucker Reals. Radioactive material stolen in mexico. April 2015. URL <http://googl/BJMaeI/>. 1.2
- [51] Branko Ristic and Ajith Gunatilaka. Information driven localisation of a radiological point source. *Information Fusion*, 9(2):317 – 326, 2008. ISSN 1566-2535. doi:

- <http://dx.doi.org/10.1016/j.inffus.2007.06.004>. URL <http://www.sciencedirect.com/science/article/pii/S1566253507000577>. 1.2.2, 1.2.3
- [52] Branko Ristic, Mark Morelande, and Ajith Gunatilaka. A controlled search for radioactive point sources. In *Information Fusion, 2008 11th International Conference on*, pages 1–5. IEEE, 2008.
 - [53] Nicholas Roy, Geoffrey J Gordon, and Sebastian Thrun. Finding approximate pomdp solutions through belief compression. *J. Artif. Intell. Res.(JAIR)*, 23:1–40, 2005. 7.1.3, A.7
 - [54] V. S. Kumar Samparathi and Harsh K. Verma. Outlier Detection of Data in Wireless Sensor Networks Using Kernel Density Estimation. *International Journal of Computer Applications*, 5(6):28–32, August 2010. ISSN 09758887. doi: 10.5120/924-1302. URL <http://www.ijcaonline.org/volume5/number7/pxc3871302.pdf>.
 - [55] Thomas B Schön, Fredrik Gustafsson, and Rickard Karlsson. The particle filter in practice, 2009. 1.2.2, 1.2.3, 2.2
 - [56] Burr Settles. Active learning literature survey. *University of Wisconsin, Madison*, 52:55–66, 2010.
 - [57] Stefan Sigurdsson, Anthony Lamarca, Waylon Brunette, David Koizumi, Matthew Lease, Stefan B Sigurdsson, Kevin Sikorski, Dieter Fox, and Gaetano Borriello. Making Sensor Networks Practical with Robots. In *In Pervasive Computing. First International Conference, Pervasive 2002. Proceedings (Lecture Notes in Computer Science Vol.2414)*. 2002, pages 152–166. Springer-Verlag, 2002.
 - [58] Alexander Smith. Hiroshima 70th anniversary: What to know about nuclear weapons in 2015. URL <http://goo.gl/g2Zt8j/>. 1.2
 - [59] Jacob Stinnett. Bayesian algorithms for automated isotope identification. 2014. 1.2.3
 - [60] A. Sundaresan, P.K. Varshney, and N.S.V. Rao. Distributed detection of a nuclear radioactive source using fusion of correlated decisions. In *Information Fusion, 2007 10th International Conference on*, pages 1–7, July 2007. doi: 10.1109/ICIF.2007.4408143. 1.2.2
 - [61] Ashok Sundaresan, Pramod K Varshney, and Nageswara SV Rao. Distributed detection of a nuclear radioactive source based on a hierarchical source model. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on*, pages 2901–2904. IEEE, 2009. 1.2.2
 - [62] Prateek Tandon. Multi-agent Planning for Mobile Radiation Source Tracking and Active City-wide Surveillance. 2012. doi: 10.13140/RG.2.1.4213.0089. URL https://www.researchgate.net/publication/277775260_Multi-agent_Planning_for_Mobile_Radiation_Source_Tracking_and_Active_City-wide_Surveillance. 6.2.1
 - [63] Prateek Tandon, Artur Dubrawski, Jeff Schneider, Adam Zagorecki, Simon Labov, and Karl Nelson. Poster abstract: Portable radiation monitoring platform for effective nuclear search. *ARI Review*, 2012. 2.1, 6.1
 - [64] Prateek Tandon, Vladimir Ermakov, Aashish Jindia, Artur Dubrawski, Simon Labov, and Karl Nelson. Poster abstract: Portable radiation monitoring platform for effective nuclear

search. ARI Review, 2012. A.2

- [65] Prateek Tandon, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Poster abstract: Suppressing background radiation using poisson principal component analysis. IEEE Nuclear Science Symposium, 2012.
- [66] Prateek Tandon, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Poster abstract: Source location via bayesian aggregation of evidence with mobile spectrometers. Symposium on Radiation Measurements and Application, 2012.
- [67] Prateek Tandon, Peter Huggins, Artur Dubrawski, Simon Labov, and Karl Nelson. Poster abstract: Simultaneous detection of radioactive sources and inference of their properties. IEEE Nuclear Science Symposium, 2013. 1.2.3, 4.4, 7.1.2, 7.2, 7.2.3
- [68] Prateek Tandon, Artur Dubrawski, Peter Huggins, Robert Maclachlan, Karl Nelson, and Simon Labov. Poster abstract: Poisson match filter for low photon count data from portable spectrometers. ARI Review, 2014.
- [69] Matthew E Taylor, Manish Jain, Prateek Tandon, and Milind Tambe. Using dcops to balance exploration and exploitation in time-critical domains. In *Proc. Workshop on Distributed Constraints Reasoning (DCR09), at IJCAI*, volume 9, 2009.
- [70] Di Tian and Nicolas D Georganas. A node scheduling scheme for energy conservation in large wireless sensor networks. *Wireless Communications and Mobile Computing*, 3(2): 271–290, 2003. ISSN 1530-8677. doi: 10.1002/wcm.116. URL <http://dx.doi.org/10.1002/wcm.116>.
- [71] Marcos AM Vieira, Matthew E Taylor, Prateek Tandon, Manish Jain, Ramesh Govindan, Gaurav S Sukhatme, and Milind Tambe. Mitigating multi-path fading in a mobile mesh network. *Ad Hoc Networks*, 11(4):1510–1521, 2013.
- [72] Meritxell Vinyals, Juan A Rodriguez-Aguilar, and Jesus Cerquides. A survey on sensor networks from a multiagent perspective. *The Computer Journal*, page bxq018, 2010.
- [73] Ping Wang and Ting Wang. Adaptive Routing for Sensor Networks using Reinforcement Learning. In *Computer and Information Technology, 2006. CIT '06. The Sixth IEEE International Conference on*, page 219, 2006. doi: 10.1109/CIT.2006.34.
- [74] Yang Xiao and Yanping Zhang. Surveillance and tracking system with collaboration of robots, sensor nodes, and rfid tags. In *Computer Communications and Networks, 2009. ICCCN 2009. Proceedings of 18th International Conference on*, pages 1–6. IEEE, 2009.
- [75] M Xie, B He, and TN Goh. Zero-inflated poisson model in statistical process control. *Computational statistics & data analysis*, 38(2):191–201, 2001. 7.1.2
- [76] Eugene Yee, Ian Hoffman, and Kurt Ungar. Bayesian inference for source reconstruction: A real-world application. *International Scholarly Research Notices*, 2014, 2014. 1.2.3
- [77] KP Ziock and WH Goldstein. The lost source, varying backgrounds and why bigger may not be better. In *Unattended Radiation Sensor Systems for Remote Applications*, volume 632, pages 60–70. AIP Publishing, 2002. 1.2.2