Sensor Management Problems of Nuclear Detection

Tamra Carpenter
Jerry Cheng
Fred Roberts
Minge Xie

DIMACS, Rutgers University, Piscataway, NJ 08854

Abstract

Terrorist nuclear attack is a potentially devastating threat to homeland security. It is increasingly important to have the capability to intercept illicit nuclear materials entering the country and to monitor for nuclear threats emerging from within. This article describes a variety of approaches to sensor management in a multi-institution project on nuclear detection, which is based at Rutgers University and supported by the US Department of Homeland Security Domestic Nuclear Detection Office in collaboration with the US National Science Foundation. These approaches revolve around formulating the related problems using precise mathematical language and then developing tools of the mathematical sciences to solve them. The article provides an overview and summary of the project. In so doing, it touches on four themes that have emerged as the areas of greatest emphasis in the project: 1) methods to exploit data from radiation sensors and shipping manifests for classification and decision making; 2) ways to optimize sequential decisions in layered inspection processes; 3) detection using a fleet of mobile radiation sensors; and 4) data sampling strategies for nuclear detection.

1. Introduction

The effective use of sensors for nuclear and radiological detection requires choosing the right type of sensor, putting it in the right place and activating it at the right times. It also involves interpreting the results of sensor alarms and making decisions that balance various types of risk and uncertainty based on those results. This article describes a variety of approaches to sensor management for nuclear detection that revolve around formulating the related problems using precise mathematical language and then developing tools of the mathematical sciences to solve them. It emphasizes a variety of approaches to sensor management in a multi-institution project on nuclear detection, which is based at Rutgers University and includes Princeton University and Texas State University-San Marcos.

The nuclear and radiological materials whose detection is of particular concern are radiation dispersion devices (RDDs) – more commonly known as dirty bombs – and special nuclear materials (particularly highly enriched uranium and weapons-grade plutonium) that could

provide the fissile material for a nuclear weapon. Throughout this paper, we will use the generic term "nuclear detection" to include detection of any radiation-emitting material of concern. RDDs could potentially contain a number of different radionuclides, some of which are used commercially, so the RDD threat is less specific than special nuclear materials and harder to differentiate from benign sources of radiation such as those from medical procedures or naturally occurring radioactive materials like the clay found in pottery and kitty litter. The need to distinguish true threats from commonly occurring benign sources and background sources of radiation is a particular challenge in nuclear detection [GEH03].

Nuclear detection arises in a variety of different contexts that pose overlapping, but sometimes quite different, research challenges. This paper aims to provide an overview of several problems arising in the following settings for nuclear and radiological detection:

Border crossings – At borders, vehicles move through radiation portal monitors (RPMs) that provide passive, non-intrusive screening for the presence of nuclear and radiological materials. In this setting, we have brief contact with all entering vehicles and can detain them for further inspection when alarms occur. Here, the emphasis is on preventing nuclear materials from entering the country, so detection is the main priority and false alarms are tolerated as a natural consequence [GEH03]. Typically, inspection at borders is a layered process and all alarming vehicles are subject to further scrutiny in subsequent layers which can include passing the vehicle through another RPM, screening with handheld radioactive isotope identification devices (RIIDs), or manual inspection of the vehicle. Each progressive level of scrutiny introduces additional delay and inspection cost. Since all RPM alarms are followed with further inspection, new methods can potentially reduce false alarm rates through enhanced analysis of sensor data at each layer and by optimizing the choice of which inspections to perform next.

Ports of Entry – At seaports and other ports of entry, we have huge numbers of shipping containers that must be screened, and this must be done in a way that mitigates risk without introducing excessive delays and the ensuing disruptions to commerce. As at borders, we have a layered inspection process that includes passive radiation monitoring and manual inspection of containers, but we also have more stringent testing capabilities using gamma radiography, as well as pre-port information on arriving ships and the cargo they contain. As containers arrive, we must decide which containers to inspect more carefully, and we need to do this without causing excessive disruption to port operations. For a particular detection technology, we have to identify the best method for assessing the risk of a container and perhaps the sequence of screening tests to apply.

Special events – At special events such as a major concert in a city park, a political rally, or a large festival or parade, there may be no existing infrastructure and possibly no restricted points of entry. In such cases, there is a need to locate a system of sensors to provide maximal protection. Here, we can consider the development of methods to determine where to locate sensors so as to optimize detection, as well as routing strategies to efficiently patrol an area or venue.

Urban settings – Major metropolitan areas present attractive targets, but cover large geographic areas that may be difficult to monitor and/or patrol. They present many of the same challenges

as special events, but on a grander scale – both in terms of geographic area and duration. Here we discuss methods to use either stationary or moving detectors (e.g., in vehicles) with fixed or random routes. At special events and in urban settings, radiological sources may be contained within moving vehicles or carried by people, introducing the added challenges of detecting devices and materials in transit, typically without the ability to choose whom to screen. If an event is detected, there will be a delay while response measures are enacted and this introduces the problem of identifying and tracking the source vehicle or person.

This paper touches on four themes that have emerged as the areas of greatest emphasis in several projects related to inspection and nuclear detection that we are involved in. The primary research thrusts we consider in this paper are: 1) methods to exploit data from radiation sensors and shipping manifests for classification and decision making; 2) ways to optimize sequential decisions in layered inspection processes; 3) detection using a fleet of mobile radiation sensors; and 4) data sampling strategies for nuclear detection.

2. Analysis of Manifest and Sensor Data

Detection of nuclear materials entering the US currently relies on two important sources of data. One is the radiation sensors that are deployed at all major border crossings and ports of entry to scan containers entering the US and the other is documents submitted to US Customs and Border Protection (CBP) prior to a shipping container entering the US. Challenges in our project have included obtaining and understanding the information that these sources provide; identifying how we might make greater use of these data throughout the inspection process; and building models to support enhanced decision-making based on these data.

2.1 Manifest Data. Customs information is collected at overseas points of embarkation using a variety of customs forms, including a ship's manifest and bill of lading. Recently, there has been increasing emphasis by US CBP on improving the quality of customs data resulting in new, more stringent requirements that accurate manifest information be submitted at least twenty-four hours before cargo is loaded onto a US-bound vessel. Prior to arrival at US ports, CBP does screening based on such data to determine whether the shipment poses a risk. Identifying mislabeled or anomalous shipments through scrutiny of manifest data is one step in a multi-layer inspection process for containers arriving at ports [WWB06]. Enhancing capabilities to extract information from such data may prove useful in screening for radioactive and nuclear materials.

Early in 2008, we obtained manifest data that provides information on cargo entering US ports over several days. The raw data include over 30,000 records, which were parsed to create a database for use by various teams on our project. Each manifest contains 120 attributes that include information about the shipper, consignee, notify party, and the shipment itself. Shipment data include size, weight, export codes (for hazardous materials, products covered by tariffs, etc.), and a physical description of the cargo manually entered by an inspector. Some of the manifest information is contained in free-form text while the rest is categorical or numeric. It is our observation that there is considerable leeway in the level of information provided, as well as little structure or commonality in the text fields. These issues present challenges and introduce uncertainty that must be dealt with when using this data for screening.

As noted in [WWB06], the Automated Targeting System (ATS) of US CBP is already using manifest information to classify containers as being either "trusted" or "untrusted", but it may have difficulty in correctly classifying a container transporting nuclear material via a trusted shipper. The work of McLay et al [MLN09] suggests that effective prescreening, such as that based on manifest data, can be an important component of cargo screening when there are limited screening resources.

The Canada Border Services Agency (CBSA) is also using information available in shipping documents to classify containers according to risk [HOZ] and to improve their inspection process [HCS09]. Like the US CBP, CBSA has an automated system that assigns risk scores to indicate the likelihood that a container entering the country has undesirable contents. In recent work, Hoshino et al [HOZ] have looked for ways to improve the existing system by explicitly taking into account the facts that 1) the problem is inherently unbalanced with only a very small fraction (roughly 2%) of containers being deemed dangerous; and 2) the likelihood of finding a dangerous container seems to vary with time. To deal with these issues they propose a two-stage approach that first fits a baseline classifier without considering time and then applies a time adjustment factor that results in substantially improved performance. In other work [HCS09], Hoshino and CBSA colleagues have developed user-friendly classification methods to predict the presence of fumigants to reduce the time and expense of chemical testing of every marine container referred for further inspection.

Our project includes several studies that leverage classification algorithms and apply them in the context of shipping manifest data. We describe some relevant activities in the remainder of this subsection.

Bayesian Binary Regression. Our project is leveraging anomaly detection methods developed for the intelligence community to look for anomalies and general trends in manifest data. We have begun applying Bayesian LASSO logistic regression using the Bayesian binary regression (BBR) software developed at DIMACS [GLM07, MGLF05] to help analyze manifest data. In particular, we used the following logistic regression model to analyze and discover the potential associations between the risk status (Y) and the origination/destination and contents of the cargo, as well as the history of the shipping company, etc. (covariates X's):

$$\log\left(\frac{P(Y_i=1)}{1-P(Y_i=1)}\right)=\sum_{j=1}^{m}\alpha_j X_{i,j}\ .$$

In this model, $Y_i = 1$ if the i^{th} container is selected for further inspection and $Y_i = 0$ if it is not. The $X_{i,j}$'s represent the values of covariates, as well as possible interactions among covariates, associated with the i^{th} container. Bayesian logistic regression finds the maximum *a posteriori* (MAP) estimate of the parameter vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_m)^t$ under a Laplace prior distribution. This approach can effectively deal with the large number of covariates/fields extracted from manifest data, as well as the sparsity of the data. From the manifest data, we identify the covariates and interactions among them with statistical significance. This information can help us build a predictive model to assign risk scores to incoming containers.

The response variable of risk status is not available in the manifest data. Nevertheless, we performed a variety of simulation studies based on the manifest data to determine the

effectiveness of Bayesian LASSO regression and the BBR software in such an application. Specifically, we selected a small set of covariates from the manifest data (such as port of origin, cargo contents, etc) and hypothesized a regression relationship between the risk status and the selected set of covariates. We assumed this relationship to be the "true" model and based on this assumed relationship, we simulated the risk status (Y) for each container. Now, pretending that we did not know the assumed relationship, we applied the BBR algorithm and LASSO regression using the simulated risk status and all covariates associated with the containers from the manifest data. Most of the time, the BBR and LASSO regression could identify the selected set of covariates as significant contributors to container risk status, and the predicted risk scores were consistent with simulated risk score from the assumed "true" model [CCX09]. This suggests that the proposed logistic model/BBR approach could indeed provide an effective tool for processing information in the manifest data.

Higher-order Naïve Bayes. In another area of research, project member Bill Pottenger and his students are applying a higher-order naïve Bayes (HONB) algorithm [GLP09] to classify shipments in the manifest data using the hazardous material export code as the class. In theoretical work, they are developing a novel approach to learning that exploits relationships between attribute/feature values across different shipment manifests. In empirical work, they are selecting nominal classes within the manifest data that result in the most useful models.

The underlying assumption in many traditional machine learning algorithms is that instances are independent and identically distributed (i.i.d.). Such models are called "first-order" because in general they only leverage relationships between attributes within instances (e.g., co-occurrence relationships), and do not leverage connections that link attributes from different instances. These critical independence assumptions that are made in traditional machine learning algorithms prevent them from going beyond instance boundaries to exploit latent "higher-order" relations between instances. Work in our project moves beyond instance boundaries to exploit the latent information captured in higher-order co-occurrence paths between instances within a dataset. The algorithms being developed leverage implicit co-occurrence relationships between attributes in different instances or manifests. Pottenger and his team believe that algorithms leveraging higher-order associations between different attributes of each shipment will allow for more precise identification of anomalous shipment data, especially when such algorithms are part of an online learning environment. The related work in the project assumes that descriptions of products such as "IKEA home furnishings" will be more likely to match with certain container types or ports of departure than will other products; thus, anomalies may be discoverable in manifests that do not observe similar associations. Higher-order naïve Bayes is especially useful considering the online nature of the manifest data, which implies sparsity during initial model learning.

Results obtained on benchmark corpora [GLP09] show that higher-order naïve Bayes generates more accurate models on sparse data than first order naïve Bayes classifiers, especially with small training sets. Extensive experiments on several data sets from different domains support this conclusion. However, it remains to be determined in future project work whether the classification models point to anomalous shipments that would be identified as "high risk cargo" by inspection domain experts.

2.2 Radiation Sensor Data. Detecting nuclear materials at borders and seaports relies on data from the radiation portal monitors that are deployed at all major ports to scan vehicles and containers entering the country [GEH03]. At present, there are roughly 1100 radiation portal monitors installed and they inspect approximately 90% of the containers and vehicles entering the country [W08]. Much of the nation's commercial life depends on the contents of these containers that are carried on the roughly 57,000 trucks, 2,500 aircraft, and 580 sea vessels entering the US each day. Given this volume of cargo, there are two competing priorities in inspection: 1) to process cargo quickly so as not to cause congestion and resulting disruptions to commerce and 2) to prevent entry of any illicit nuclear or radiological material. To meet these dual objectives, CBP has adopted a multi-layer approach for inspection, which consists of a "routine" inspection followed by a more stringent inspection of containers that were identified as suspicious during the routine inspection. Wein et al [WWB06] describe the current layered approach and consider how to optimize an 11-layer security system that includes shipper certification, container seals, the ATS system, passive and active radiation testing, and manual inspection to improve detection. Our project is developing approaches for making decisions during routine screening.

Statistical Learning. As trucks at border crossings move though portal radiation sensors, the portal captures the energy spectrum every tenth of a second across a range of channels going from low frequency to high frequency. It can be 256 channels or coarser bands consisting of frequency counts in only 5 non-overlapping, exhaustive bands corresponding to channels 0-5, 6-10, 11-40, 21-80 and 81-256, respectively [KR86]. Project member Siddhartha Dalal formulated a Bayesian learning approach for modeling the energy emitted by an unknown source and classifying it as belonging to one of K defined classes [DH09]. These classes would include radioactive materials of high concern, such as high energy Uranium, depleted Uranium, Plutonium, Cobalt-57, and Barium-133, as well as benign materials that may be common sources of emission such as medical waste or kitty litter.

Denote by $Z = (R_1, ..., R_5)$ the observed radiation counts in the 5 non-overlapping channels from the training set of the portal radiation sensors data. Dalal and Han [DH09] assumed that the total count $N = R_1 + R_2 + \ldots + R_5$ given a class $C = c$, $c = 1, 2, \ldots, K$, follows a Poisson distribution. Furthermore, they assumed that, given total count $N = n$ in class $C = c$, the observed radiation counts $Z = (R_1, ..., R_5)$ follows a multinomial distribution. Based on these Poisson process and multinomial models – both of which are conventional assumptions for this type of count data – they were able to derive the following classification rule from Bayes formula:

$$P\{C = c|Z = (r_1, ..., r_5)\} \propto P\{N = n|C = c\}\, P\{R_1 = r_1, ..., R_5 = r_5) \,|N = n, C = c\}P\{C = c\}.$$

Here, $(r_1, ..., r_5)$ are observed radiation counts in class c in the training data, $n = r_1 + r_2 + \ldots + r_5$ and $P\{C = c\}$ can be estimated by the fraction of containers of class c in the training set.

This classification model can be used to develop a scoring model that assigns a risk score to each new container. A potential risk scoring model [CCDX09] could be

$$s^* = \sum a_c P\{C = c|Z^* = (r_1^*, ..., r_5^*)\}$$

where: (r_1*, ..., r_5*) is the radiation sensor reading of a new container; a_c is the average risk score of the class c computed from the training data and classification model; and the sum $\sum$ is over all c = 1, …, K classes. We can use this predictive model to assign a risk score to each incoming container. This will give us a likelihood that the new container has nuclear or illicit material so further investigation can be conducted accordingly.

Machine Learning. In a second study using radiation sensor data, project members Bill Pottenger and Jason Perry have begun to apply machine learning techniques to analyze gamma-ray spectra generated by CZT-based handheld detectors to see whether they could distinguish non-threat sources of radiation from possible threat materials. One way to cast this problem as a machine learning problem is to train a set of classifiers to identify the presence of any gamma-emitting radioisotopes from a predetermined library, as above. However, another approach may be necessary to build a robust real-world solution for distinguishing threat-from non-threat isotopes. For instance, in security applications, very specific types of accuracy may yield practical advantages for reducing false alarms. One example would be a mechanism to provide a very high confidence level in detecting known non-threat isotopes such as Technetium-99m (Tc99m) – the most common medical isotope and one which would generally not present a threat. A system to accurately discern Tc99m as the radiation source could safely indicate when no further inquiry is necessary. At the same time, the system must be sensitive to a wide variety of other known and unknown radioisotopes, so that no potential sources of threat are missed. These, in turn, must be distinguished from fluctuations in the natural radioactive background, in order to minimize false alarms. This requires both an optimal framing of the machine learning problem and a very finely-tuned classification system which takes advantage of all available data.

In initial investigations using data obtained data from a CZT-based hand-held detector, Perry [Per09] formulated a three-class classification problem with classes corresponding to: 1) presence of Tc99m; 2) presence of other known and unknown isotopes; and 3) all natural background radiation conditions. Using Support Vector Machines (SVMs) for classification, his experiments showed that the single-isotope Tc99m class is distinguishable from the other classes of isotopes with near-perfect accuracy. However, separating the class with all other isotopes from the normal background noise and detector anomalies is much more difficult and requires further study.

Statistical Change Detection and Identification. Project members Savas Dayanik, Warren Powell and Kazutoshi Yamazaki have developed new online statistical change detection and identification rules to identify pattern changes in sensor readings that indicate the presence of either hazardous materials or a malfunctioning of the sensor [Po07, DPY08]. We envision these procedures operating in real time as vehicles or containers are scanned through portals or other sensing devices. The algorithms are designed to make diagnoses within a prescribed low level of false alarms and to work with sensing devices that have only a small amount of associated computational capability. In general, the method models a set of potential "disruptions" that include a variety of sensor failure modes and detection events corresponding to detection of different materials. The methods analyze sensor readings and both determine when to sound an alarm and identify the suspected alarm trigger (i.e. the failure mode or material detected).

A typical change detection and identification rule consists of a pair of an alarm delay (the difference between the time that some disruption first occurs and the time that the algorithm declares an alarm) and a diagnosis rule. Because observations are collected sequentially over time, optimal rules are typically the solutions of a dynamic program in a Bayesian framework [DGP08, DG09]. Unfortunately, optimal solutions are often not in closed form, and due to the curse of dimensionality of state space, their numerical implementations require large computing power and memory. Therefore, a classical dynamic programming approach to the change detection and identification problem does not easily lead to online optimal rules that can be run on devices with limited associated computational capability. However, it is possible to develop simple nearly optimal decision rules by combining dynamic programming and renewal theory for stochastic processes, which is the methodology adopted by Dayanik et al [DPY08].

Their methods [DPY08,Yam09] seek to find an alarm time and an identification rule such that the decision risk associated with any potential "disruption" is below a specified threshold (which is an adjustable model parameter) and the detection delay is minimized. The specific types of risk considered include the risk associated with a real nuclear event whose detection is missed or delayed, risks from investigating false alarms, and misdiagnoses of detected real disruptions. Precise solution of this problem would require solving a constrained stochastic optimization problem to find a rule that would concurrently minimize for each of the potential disruptions, and it is unclear whether this would even have a solution. Therefore, Dayanik et al [DPY08] studied the optimal asymptotic performance as the bounds on allowable decision risk converge to zero. Results show that, for small allowable decision risk, the minimum expected detection delay over all admissible rules can be attained by a common admissible rule that can exploit recursive equations to minimize the need for computational power that may be lacking with small sensing devices. The result is a simple, computable policy that determines when a signal change has occurred and determines the cause of the change (i.e. the sensor failure mode or material detected) in the presence of noisy readings from sensors that may fail due to aging or operational stress. This policy closely approximates an optimal policy [Yam09], but is much easier to compute, potentially allowing it to be used in real time.

2.3 Combining Data Sources. Together, sensor data and manifest data provide terabytes of data on millions of containers and their contents. While methods have emerged to analyze each set of data separately (including the methods that we have described), efforts to combine these data for more powerful capabilities for detecting illicit nuclear material are still relatively new. Methods for combining such data hold the promise of considerable improvement in detection. Recently, we obtained an additional month of manifest data that are coordinated with radiation detection data that we also hope to attain. Such data – linking radiation sources, manifest data, and radiation portal readings – would be used to provide a rich source of training data for building a classifier for specific ports of entry. Our future work will study how to combine manifest data with sensor data in our statistical and machine learning methods. For instance, we will study how to compare materials claimed on the manifest with those identified in classification using radiation sensor data to identify potential anomalies. In addition, we will investigate how we can use manifest information to inform the classification task itself in order to improve accuracy. For instance, such methods could include information from the manifest to determine which materials to consider when defining the K classes in Dalal and Han's model [DH09]. Our aim is to use these combined methods for more powerful decision-making capabilities during the

routine screening process, enabling us to more definitively identify suspicious containers and reduce delay of benign containers. Along these lines, researchers at Lawrence Livermore National Laboratory [LLNL07] have developed a Context-Aware Nuclear Evaluation System (CANES), which combines data from multiple types of sensors (such as RPMs and RIIDs) with context data that includes information on distance from source and type of conveyance and applies machine-learning algorithms for threat assessment.

3. Optimizing Sequential Decision-making Strategies for Inspection

In addition to analyzing the sensor data itself, another aspect of sensor management is deciding which tests to apply to incoming cargo and in which order to apply them in light of practical considerations such as budgets on inspection time and/or cost. To date, several researchers have studied paradigms for modeling and optimizing container inspection at ports [SS03, WWB06, BEK08, Ram08, CR09, ESX09, AMM06, MMR07, MMR09]. Rather than focusing on making a decision based on given sets of data (as in Section 2), the emphasis here is on determining the sequence of tests to apply to optimize the inspection process.

At ports of entry, we envision a stream of entities arriving for inspection and a decision maker having to decide how to inspect each one. This includes deciding which to subject to further inspection and which to pass through with only minimal levels of inspection. Viewed this way, the process becomes a sequential decision making problem. Sequential decision making is an old subject, but one that has become increasingly important as traditional methods for making sequential decisions fail to keep pace with problem scale. Enumerative algorithms for optimizing port-of-entry inspection rapidly come up against the combinatorial explosion caused by the many possible alternative inspection strategies. Moreover, methods must incorporate practical considerations – such as sensor error – which introduce uncertainty into the models. Work on these topics is being conducted as part of several other projects that are closely aligned with ours and have some overlapping participants. In particular there is another nuclear detection project based at Rutgers and led by Endre Boros and Paul Kantor, and there are several projects on port-of-entry inspection also based at Rutgers. These projects are developing approaches that bring into the analysis many of the complications – such as sensor error – that arise from practical considerations.

In the port of entry inspection projects (see [BEK08] for an overview), the project teams are building on the initial approach to the port-of-entry inspection problem taken by Stroud and Saeger [SS03], who studied a case that involved different potential tests (we will call them sensors) for deciding whether a cargo contains illicit material. Four such tests currently in use are evaluation of ships manifests, passive radiation signatures, radiographic images, and induced fission. All of the tests have costs associated with them, including the cost of a reading indicating illicit material when there is none (a false positive), the cost of a reading indicating there is no illicit material when there is (a false negative), time costs of using the sensor, delay costs of waiting for the sensor, and fixed cost of equipment, labor, etc. For each sensor the readings for cargo containing illicit material (positives) and readings for cargo not containing illicit material (negatives) are random variables. The model Stroud and Saeger created assigns an output of 0 (absence of illicit material) or 1 (presence of illicit material) for each sensor. In general, n sensors will yield a string (vector) of 0's and 1's of length n, and can be modeled with

a binary decision tree (BDT). Stroud and Saeger developed enumerative methods to find the binary decision tree of sensors that would minimize total cost of inspection. The problem becomes intractable already for n = 4 if one relies on brute force methods since the number of possible trees expands rapidly, but Stroud and Saeger were able to extend their method to n = 4 by making some assumptions about the types of decision functions captured by the BDT. However, their method is not feasible for higher values of n.

The project teams built on this initial approach, making the models and algorithms better suited to address inspection issues that might arise in practice. For instance, the heuristics developed will need to be able to scale up to perhaps 20 or more of sensors. Furthermore, there are several interdependent aspects of port-of-entry inspection that need to be explored in tandem [BEK08]. Some of them are:

- Developing simulation models of inspection stations as one part of an operating port. These models can be used to assess the efficiency and effectiveness of security field operations, aid decision makers in quantifying the trade-off between security goals and their attendant costs, provide feedback for devising improved operations, as well as to provide estimates for some of the cost parameters (such as delays) used in some of the optimization models.

- Studying the sensitivity of optimal and near optimal trees to the input parameters [AMM06]. As input parameters such as the costs of false positives and false negatives, the costs of delays, etc., are estimated with more or less accuracy, one wants solutions whose sensitivity to changes in these parameters is known and tolerable. Team studies led by Saket Anand, David Madigan, and Fred Roberts [AMM06] show that the optimal inspection strategy is remarkably insensitive to variations in the parameters needed to apply the Stroud-Saeger method. An important research challenge is to understand why.

- Developing new computational approaches that are inexpensive, scalable, and able to incorporate various cost factors with enough flexibility to include future technologies. Such approaches are based on efficient search heuristics [MMR07, MMR09, BEK08], linear programming [BFK09], and dynamic programming [GWB08] and are now able to address problems involving many more sensors in very little time. In related research, Concho and Ramirez-Marquez [Ram08, CR09] have used evolutionary algorithms to optimize a decision tree formulation of the inspection process. Their approach is based on the assumption that readings r_j by the jth sensor are normally distributed, with a different distribution depending on whether the container in question is "bad" or "good." Thresholds t_j are used to determine outcomes of inspections, with a container declared suspicious by the jth sensor if $r_j > t_j$. Here, the cost function used depends upon the number of sensors used and the cost of opening a container for manual inspection if needed, but does not take into account the cost of false positives or false negatives, which is a key feature of the work in [SS03], [AMM06], [MMR07], and [MMR09].

- Investigating the optimum threshold levels for sensor alarms so as to minimize overall cost as well as minimize the probability of not detecting hazardous material [AMM06, BEK08, ESX09, MMR07, MMR09]. Zhu et al [ZLY09], in work extending that of Elsayed et al [ESX09], consider sensor measurement error independently from the natural variation in the

container attribute values. They model situations when measurement errors exist (and are embedded) in the readings obtained by the inspection devices and use a threshold model to identify containers at risk for misclassification. They study optimization of container inspection policies if repeated inspection of at-risk containers is part of the process.

- Exploring use of sensors with many possible outputs categories. Boros, Kantor and their colleagues [BFK09], in a parallel nuclear detection project, used a large-scale linear programming model approach and considered more container classifications than just the bad or good. They demonstrated the value of a mixed strategy applied to a fraction of the containers. They then added budget constraints to the problem in [GWB08].

Several other authors have also considered ways to optimize and improve layered screening systems that include some of the above aspects. Wein et al [WWB06] consider several of the above issues in a detailed study on how to optimize the inspection strategy for detecting nuclear weapons (or their building blocks) at ports. In so doing, they develop operational models and make specific recommendations on which key uncertainties are most important to resolve, how to improve the existing screening process, as well as how to most effectively utilize new technologies. McLay et al [MLN09] develop a linear programming model for screening cargo for nuclear materials at ports of entry. Their approach defines a framework for determining alarms when there are limited screening resources. Jacobson et al [JKK06] look at baggage screening at airports and compare 100% screening with one type of screening device with screening with a second device when the first device says a bag is suspicious. They calculate costs and benefits of the two methods.

4. Managing Static and Mobile Sensors

In some cases, such as ports and border crossings, entering vehicles are funneled through checkpoints that provide natural locations for radiation sensors. Even in these cases, practical considerations arise because of differing sensor operating characteristics – different sensors have different capacities for inspection over a given time and vary in cost and performance. Wein et al [WLCF07] consider the spatial location of radiation portal monitors at overseas ports to improve detection (which depends on scan time) without creating bottlenecks that would create excessive congestion in the port. Jacobson et al [JMV05] consider the problem of deployment of baggage screening devices at airports, formulating it as an integer programming problem that takes into account various practical complications.

In less structured settings, such as urban environments, desirable locations are less obvious and need to be determined in other ways. In our project we have explored two different scenarios. The first is a variant of more traditional static sensor location problems that require locating sensors to respond to a set of uncertain events. In this case, we assume that we are placing sensors in a set of fixed locations to minimize the risk of missing a threat. Typically, an implicit assumption is that locations for these sensors must be chosen judiciously because they are too expensive to locate "densely" over the area to be covered. In this way, sensor placement problems are closely related to well-studied facility location problems in the optimization literature [KH79, KH79a, M90]. However, the sensor placement problem has sources of uncertainty that are not part of the traditional facility location problem. In sensor placement, it

makes sense to consider environments where events to be monitored occur with low probability. Thus, the locations that we need to "cover" have considerable uncertainty yielding stochastic versions of more traditional problems. These stochastic variants appear to be significantly more complex and are not yet well studied.

Dimitrov et al [DGM08] locate sensors along a transportation network using a stochastic interdiction model. Here the scenario is that a smuggler needs to get from a given origin to a given destination in the network and the "interdictor" needs to locate sensors to minimize the probability that the smuggler can reach the destination without detection. They envision this problem arising in border protection. Another complication that arises in sensor networks is sensor error. Neidhardt et al [NLK08] consider optimizing the positioning of error-prone sensors to monitor an area. Their placement strategy seeks to reduce error by of having areas "covered" by multiple sensors in an "equitable" fashion through use of a minimax objective.

Wein and Atkinson [WA07], Atkinson and Wein [AW08] and Atkinson, Cao and Wein [ACW08] develop a detection-interdiction model to assess the efficacy of deploying a ring of sensors to protect an urban area from attack with various types of nuclear devices. Their models assume that an adversary is attempting to drive an already-assembled nuclear device into an urban area to maximize expected damage, while a defender combines use of a ring of radiation sensors and a fleet of interdiction vehicles to prevent penetration into the city. Their studies consider sensor errors resulting in missed detections and false alarms that occupy interdiction vehicles.

The second scenario that we explore in our project assumes that sensors are cheap and mobile. Under this basic paradigm, we examine the utility of locating sensors on vehicles such as taxis and police cars in an urban setting. Here the problem is no longer locating sensors; it is developing the statistical capabilities to reconcile readings from multiple sensors that are moving and may be prone to errors. Our project to date has emphasized this second scenario, with a focus on managing a "fleet" of mobile sensors. We are examining the viability of fleet-based sensing and addressing related statistical challenges in detection. Hochbaum [H09] and Hochbaum and Fishbain [HF09] have considered a similar fleet-based surveillance scenario, while Neidhardt et al [NLK08] have considered a two-level sensing system that includes a static network of sensors augmented with a mobile pool of opportunistic sensors, such as those on cell phones.

4.1 Opportunistic Surveillance with Mobile Sensors

We envision "opportunistic sensing" as one possible paradigm for sensing with vehicles. In this case, we imagine vehicles (whether taxis, police cars, or some other "fleet") that contain radiation sensors, but their movement is determined by activities other than surveillance, such as routine taxi pickups or police patrolling. This paradigm features a network of mobile sensors operating relatively independently to provide surveillance of an area as an artifact of movement in performing other duties. Such networks can operate in tandem with smaller, more carefully designed, static networks to provide additional coverage and corroboration of alarms (such as in [NLK08]), or they can operate independently.

To illustrate the concept, we envision installing small radiation detection devices, communication capabilities (through cellular networks) and global positioning systems (GPS) in taxis, police vehicles, fire trucks, and/or public transit vehicles to provide surveillance in major urban areas. Such networks aim to leverage technological advances in sensors and positioning systems, miniaturization of devices for sensing and communication, and the pervasiveness of human activity in dense urban areas. Recent advances have made communication infrastructure nearly ubiquitous, while detection devices and positioning systems have become both economical and portable. Thus, large-scale deployment of a mobile sensor network is becoming feasible and affordable. The New York City police department is already using small sensors in vehicles and on officers [K09, Rig09] for radiological detection. The idea of using massive networks of mobile sensors has been adopted and tested by the Radiation Laboratory at Purdue University, where they use a network of cell phones with GPS capabilities to detect and track radiation [Pur08].

The movement and extensive coverage afforded by sensors in taxis is appealing because it could provide pervasive surveillance in dense urban areas, while devices placed on emergency response vehicles or police cars could offer greater control to investigate suspicious regions and allow the possibility of including more powerful (and expensive) sensing capabilities on some vehicles. When vehicles equipped with sensors move within a certain range of a nuclear source, the radiation energy from the source will trigger the sensing devices to send an alarm notification and a GPS position to a central command center over a wireless network. This basic sensing paradigm has many attractive features. First, the random movement and extensive coverage of the vehicles provides constant surveillance for nuclear materials. Second, the mobile sensors do not need to be of high accuracy, since the failure of a small portion of them will not significantly hamper the effectiveness of surveillance coverage because of the sensors' random movements. Next, the movements of the sensors will (in most cases) be difficult to predict by an adversary, and because of the number of sensors, difficult to tamper with. Finally, because the sensors are mounted on vehicles, there are fewer size constraints and power consumption requirements.

Such a mobile sensor network would likely be supplemented by stationary sensors to cover locations with sparse traffic, such as a large park in the city. The methods we have developed can easily be envisioned for such mixed networks by simply viewing the stationary sensors as parked vehicles. While our algorithm can be readily adapted to a variety of settings, we work under the following basic assumptions:

- Nuclear sensors and Global Position System (GPS) tracking devices are installed on a large number of vehicles.
- The sensors and GPS devices constantly send detection and location information to a central surveillance center.
- Real-time tracking signals can be geolocated on a map of the area under surveillance.
- Real-time analysis is done at the surveillance center using sophisticated statistical algorithms to identify potential locations of nuclear sources that appear as clusters of positive sensor readings.

Because sensors are not always 100% accurate, there will potentially be false alarms and missed detections. Statistical methodologies have proven to be effective tools for detecting true signals against random errors. Thus, a challenge that we began to address early in the project was that of

processing sensor network information to identify "positive clusters" among the sensor signals that are not due to either random chance or known background sources. Multi-cluster spatial classification methods are ideal for such tasks. Our research has explored two innovative multiple spatial clustering methods. The first method (due to Demattei et al [DMD06, DMD07]) is based on data transformation and a step regression model. It provides a formal statistical test of significance against background noise based on the premise that points within a "cluster" should be spatially closer to each other than positive signals outside the cluster that are due to random chance/error. The second approach is based on the recent Ph.D. thesis of Lynette Sun [Sun08]. By mimicking the process of typical sample data generation, Sun [Sun08] and Xie et al [XSN09] developed an intuitive procedure that introduces a latent modeling structure and uses formal likelihood inference to detect multiple clusters occurring simultaneously within a defined region or time window. They apply model selection techniques to determine the number of clusters, and develop likelihood inference and Expectation Maximization/Markov Chain Monte Carlo algorithms to estimate model parameters, detect clusters and identify cluster locations. Their new method differs from the classical scan statistic in that it can simultaneously detect multiple clusters of varying sizes. This work is readily applicable to identifying clusters of vehicles with "positive" sensor readings for radiation.

This latent model approach was adapted by team members Jerry Cheng and Minge Xie to the context of nuclear surveillance. The method is flexible and able to accommodate a variety of extensions that make it well suited to the nuclear detection problem. The key idea is to use statistical notions of clustering, where a "cluster" involves an unusually large number of events/alarms clumping within a small region of time, space, or locations in a contiguous sequence (suggesting a moving source). Our methods are using modified versions of the traditional statistical method using scan statistics. The idea is to scan the entire study area and try to locate region(s) with unusually high likelihood of incidence. For example, one would use the maximum number of cases in a fixed-size moving window or identify the diameter of the smallest window that contains a fixed number of cases. Early work in our project demonstrated the applicability of this method to detecting clusters of positive radiation sensor readings from taxis. Cheng and Xie also performed simulations for both spatial classification methods under scenarios that include stationary and moving sources. Results of these preliminary simulations suggested that the proposed approach can effectively filter noise and background radiation sources to detect nuclear materials placed in a metropolitan area. For some details of the approach, see [CXR09].

In the first phase of our project, we emphasized use of taxis in radiation detection. Our subsequent discussions with law enforcement suggested reluctance to depend on the private sector (e.g., taxis) in surveillance. As the project progressed, our emphasis therefore shifted from considering taxis as the primary type of sensing vehicle to police cars or a combination of taxis and police cars. This concept employs the police vehicles in a manner similar to our initial ideas about taxis, but it explores use of smaller fleets, with possibly less random movement, and perhaps higher-quality sensing equipment. A central focus of more recent work has been to compare taxi-based "coverage" to police car "coverage" through simulation studies. This line of investigation aims to determine how many police cars might be enough to get sufficient "coverage" of a region when the police cars contain sensing devices but, as with taxis, their movement is directed toward normal police activity, not radiation detection.

As part of our efforts, project members Jerry Cheng, Fred Roberts, and Minge Xie applied statistical power analysis to determine the number of vehicles required to provide adequate coverage for surveillance of a given network [CXR09]. They developed a model and carried out a large number of simulations to gain intuition and assess detection power under a variety of different assumptions.

The simulations performed so far follow the same basic paradigm but systematically vary one or more of the parameters of interest. The surveillance area in our testing consists of a 4000 ft by 10000 ft. area, roughly equal to the area of the roads and sidewalks of Mid/Downtown Manhattan. In this phase of the work, we have disregarded the street network and simply consider that a specified number of vehicles are randomly located in the area at a particular "snapshot" in time. At the next time period, the vehicles are again randomly located in the region to correspond with a new "snapshot" in time. The parameters that we may adjust from one experimental run to the next include: the number of vehicles; the effective range of a sensor; and the rates for false positives and false negatives. In these experimental runs, we consider a stationary radiation source, placed randomly in the surveillance area.

We conducted a large number of experiments using this basic framework. For example, in one model, we assumed the effective range for a detector to be150 ft., a false positive rate of 2%, and a false negative rate of 5%, and we varied the number of vehicles (i.e. sensors). (We realize that this range may be beyond most presently-used detectors, but wanted to concentrate on methodology and relative comparisons, and we are experimenting with shorter ranges as well, as noted below.) We then ran at least 200 simulations for each number of vehicles and determined whether the source was detected. For each number of vehicles, this gave us an estimate of the "power," which is defined to be the probability of detecting a source for a single random placement of the vehicles (i.e. a single time period). In this model, we found that 4000 vehicles were needed to get even 75% power. With 2000 vehicles, the power was about 30%. To give this some perspective, we note that the New York City Police Department has 3000+ vehicles in 76 precincts in 5 boroughs, but at any given time only about 500 to 750 would be in the streets of Mid/Downtown Manhattan.

Of course, in practice, we would monitor the alarms over a period of time, not just at a single instant. To reconcile readings from several "snapshots" over time we may wish to use some decision rule such as: detection if a majority of the times there is an alarm; detection if at least once there is an alarm. The number of time periods is another variable that needs to be considered. It is not hard to show that if the statistical power is sufficiently high and majority rule detection is used, then with sufficiently many time periods, the detection probability can be increased significantly. We are currently exploring various rules for detection over time.

We also conducted studies in which we vary the effective range of a sensor. Given current technology, the range of a sensor may actually be closer to 25ft. than it is to the 150 ft. assumed in the experiment that we just described. However, since our project is intended to look beyond today's capabilities, we wondered: what would happen if we had a better detector, say with an effective range of 250 ft.? In an experiment similar to the previous one but with a sensor range of 250 ft., 2000 vehicles yielded 93% power.

There are other aspects of our model that need to be modified when the sensing vehicles are police cars as opposed to taxis. In particular, the assumption of random movement is less appropriate for police cars, since they will tend to remain in their own region/precinct, and they won't move around as randomly or as frequently as taxis. We did a very simple study that attempted to make movement slightly more realistic by dividing the region into 20 equal-sized precincts. (There are 22 police precincts in Manhattan.) Next, we placed k police cars randomly within each precinct. When we assumed that the number of police vehicles in each precinct is 25 – making for a total of 500 police vehicles – and assumed that each detector has a 250 ft. range, then our simulations estimated the power at 35%. This is not very good and is not significantly different than when the same number of vehicles is allowed to roam throughout the region. In other studies, we varied other parameters such as the false positive and false negative rates. We have also investigated hybrid models that involve a mixture of police cars and taxis.

We are implementing a variety of extensions to make our models more realistic, including: more complex hybrid models of taxis and police vehicles with different movement models; hybrid models that include some stationary detectors; hybrid models with more powerful detectors in police vehicles; more realistic movement models; moving sources; multiple sources; fusing information from multiple time periods. We are especially interested in exploring hybrid models that include police cars and taxis that have different models for movement and possibly sensors of differing capability. In work with graduate student Tsvetan Asamov, we recently introduced a street network with more realistic models for movement of vehicles, and we plan to use this in future studies.

A mobile sensor nuclear threat detection problem is also studied in [H09]. Here, the goal is to identify a small area in the region of interest that has a high concentration of alarms. The paper separates the two goals of small area and high concentration of alarms, which can be conflicting, and introduces a weighting factor for balancing the contribution of the two goals. In contrast to our work to date, this work has a specific model of a region as a network with streets and assumes vehicles move along streets, the paper formulates the problem of finding an "optimal" area as a mathematical programming problem and presents a polynomial time algorithm for solving it. The study is extended in [HF09] with discussions of false alarms, simulations, and methods of aggregating results over time to improve the algorithm's performance.

4.2 Randomized Surveillance Routing

The previous "opportunistic" model considers a surveillance scenario in which we gain surveillance capability by exploiting the random movement of a large vehicle fleet. Another line of research is to consider the case in which our fleet is not large enough to provide sufficient surveillance coverage by purely undirected movement. In this case, the idea is to equip a certain number of vehicles with sensors and dedicate them to the task of performing surveillance within an area. Unlike the previous case, the movement of these vehicles will be prescribed by some "controller". This controller would like to find routes for the vehicles that are "efficient" in the sense that they cover the entire region quickly but also appear "unpredictable" to an adversary [New09]. In this case, we represent the region by a graph, where links correspond with streets and nodes correspond with locations. For each vehicle, we would like to create a patrol route

that begins and ends at a designated location (e.g., police headquarters) and cannot be predicted by an adversary; yet, taken together, these routes cover the entire region efficiently. These two properties of efficiency and unpredictability are seemingly at odds with each other. Suppose for the moment that we have just one patrolman or patrol vehicle. An extremely efficient route would be a traveling salesman tour of the graph. However, given that it is very efficient, an observer knows that once a location is visited, it will not be visited again. Moreover, if the vehicle drives the same route each day, an observer could predict exactly where it will be at any given time. On the other hand, the vehicle could be very unpredictable by moving totally at random. In this case, it may take a long time to cover the entire region, but an attacker would not be "safe" just because a node was recently patrolled.

Clearly there is a tradeoff between efficiency and unpredictability, and Alantha Newman, a researcher on our project is trying to formalize this tradeoff by developing formal definitions for "unpredictability" and then defining a route optimization problem for selecting efficient routes that satisfy these formal definitions [New09]. Another approach that does not apply a strict definition for unpredictability but does exploit notions of randomness in surveillance [PPT+07] is the basis for surveillance at LAX airport. Here the approach is to consider surveillance to be a Bayesian game against an unknown adversary. Nonetheless, the ideas are similar: to find a tradeoff between an efficient assignment of inspection stations each day and one that is unpredictable. The problem is formulated as a game between an inspector and an attacker and mixed strategies are used to find good solutions. The methods have recently been extended to the Pittsburgh airport and are also being used to randomly assign federal air marshals to some of the flights between the US and Europe.

5. Data Sampling Strategies for Sensor Data

Sometimes when we have limited time or budget for data collection, it is advantageous to adjust our data sampling strategy in response to previously collected data. The specific settings that we consider involve considerable uncertainty, where the underlying probability distributions are either unknown or changing through time. Although we do not know the underlying distributions, we nonetheless have the ability to collect information through measurements to help us learn about the environment. In this sense, we may view information collection as a sequential decision problem in which our objective is to learn about our environment. There are a wide variety of practical decision-making settings in which a decision maker has the ability to collect a finite amount of information before he or she must render a decision. Examples of applications in nuclear detection include determining when to subject people and containers to additional scrutiny, positioning sensors, and deciding when to introduce a new sensing technology.

A number of practical adaptations of dynamic programming [Be57] techniques exist for finding near-optimal solutions to these types of sequential decision problems. One such adaptation that project members Peter Frazier, Warren Powell and Savas Dayanik are exploring is the knowledge gradient policy which makes sampling decisions by maximizing the expected value of the sampled information according to a simple heuristic metric [FPD08]. The problems addressed typically involve three dimensions:

1) the decision about what information to measure or collect;

2) the information that is observed when a measurement is made;
3) the decision that is made after observing the new information.

We refer to the first decision as the measurement decision w. We represent knowledge about a problem using a vector μ^n which captures the distribution of belief about a set of parameters after n measurements. After a measurement decision is made, we make an observation (such as on the level of nuclear radiation) which was uncertain before the measurement. Finally, we seek to make an economic decision which we denote by x.

Letting μ^{n+1} represent knowledge after $n+1$ measurements (which is a random variable before we have made our last observation), measurement decisions will have the fundamental structure of

$$\max_w E\left\{\max_x F\left(x \mid \mu^{n+1}(w)\right)\right\}$$

We could avoid the measurement and take an action now that requires solving

$$\max_x F\left(x \mid \mu^n\right)$$

The value of a measurement is called the *knowledge gradient*, and is given by

$$\nu^{KG} = \max_w E\left\{\max_x F\left(x \mid \mu^{n+1}(w)\right)\right\} - \max_x F\left(x \mid \mu^n\right).$$

For the problem of deciding what to measure, we have been exploring a class of measurement policies we call knowledge gradient policies. With this strategy, we choose the measurement w that yields the largest value of ν_w^{KG} where

$$\nu_w^{KG} = E\left\{\max_x F\left(x \mid \mu^{n+1}(w)\right)\right\} - \max_x F\left(x \mid \mu^n\right).$$

This policy chooses the measurement that would be best if you were going to make a single measurement, but it has also been shown to be asymptotically optimal. Not surprisingly, it has been found to work better than any other competing approaches for intermediate measurement budgets. The knowledge gradient (KG) policy [FPD08] offers an easily implemented rule that tells us how to sample information on competing alternatives to learn which is best. The initial version of the knowledge gradient policy applies in settings where the measurements are independent. The team later developed the Correlated Knowledge Gradient (CKG) [FPD08a] for determining how to collect information when measurements are correlated, as would occur when detecting nuclear radiation. KG and CKG are provably optimal in certain special cases; have provable bounds on suboptimality in other cases; and are very easy to implement and use.

As the project has progressed, the team has considered applying the KG method in more complicated settings, such as collecting information on a graph [RP09], possibly for moving sensors around a street network. In sensor management decisions, our choice of what to measure (which changes our knowledge state) depends on where we are (our physical state). A measurement at one physical state can affect decisions at other physical states. The result is what we call "the information collecting shortest path problem". The shortest path problem is fundamental to a wide range of optimization problems. [RT09] For example, these might arise in the process of designing emergency response measures (e.g. how to evacuate New York City, or how to guide emergency response forces). Often, the state of these networks is uncertain (e.g. travel times on links may vary), and it is necessary to collect information which may involve the

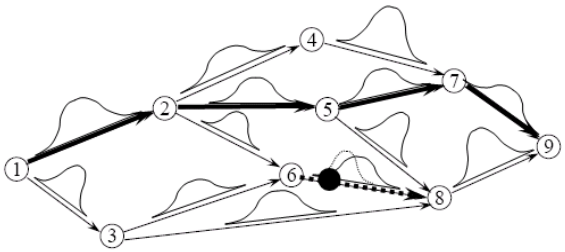

time-consuming process of dispatching people to collect it. Given this, we would like to collect the most valuable information first.

The immediate goal in this research is sensor management, where we have to move a physical sensor around the network to collect information. We found that an important stepping stone for this larger problem is the problem of determining which link in a network we should measure (without regard to the physical location of a sensor). For example, in the network depicted above, we show a probability density function on each link to describe our belief about the cost on that link. The dark path represents what we currently believe is the shortest path, while the dotted link is one that we might consider measuring. If we do measure this link, we may change our distribution of belief about the link based on the observed measurement. The question is: which link should we measure?

Project members Warren Powell and Ilya Rhyzov have explored this problem of information collection on a graph and have found that they can adapt an existing knowledge gradient method [FrPo08] to this new problem. In computational testing they have found that for networks where the paths are not too short, the knowledge gradient works consistently quite well relative to competing techniques [RP09].

Acknowledgements: We thank the members of our nuclear detection and port security project teams, especially those whose work is represented in this paper: James Abello, Saket Anand, Tsvetan Asamov, Endre Boros, Xueying Chen, Siddhartha Dalal, Savas Dayanik, Elsayed Elsayed, Peter Frazier, Emily Hogan, Paul Kantor, Mingyu Li, David Madigan, Sushil Mittal, Alantha Newman, Jason Perry, William Pottenger, Warren Powell, Ilya Rhyzov, Warren Scott, Kazutoshi Yamazaki, Christina Young, and Yada Zhu. We also gratefully acknowledge support from the National Science Foundation under grants SES 05-18543, DMS 09-15139, and CBET 07-36134, from NSA under grant H98230-08-1-0104, ONR under grant N00014-07-1-0299, and the US Department of Homeland Security Domestic Nuclear Detection Office under grant 2008-DN-077-ARI012-02.

References

[AMM06] S. Anand, D. Madigan, R. Mammone, S. Pathak and F. Roberts, "Experimental analysis of sequential decision making algorithms for port of entry inspection procedures," in S. Mehrotra, D. Zeng, H. Chen, B. Thuraisingham, and F-X Wang (eds.), *Intelligence and Security Informatics, Proceedings of ISI-2006*, Lecture Notes in Computer Science 3975, Springer-Verlag, New York, 2006, 319-330.

[AW08] M. Atkinson and L. Wein, "Spatial queueing analysis of an interdiction system to protect cities from a nuclear terrorist attack," *Operations Research* **56** (2008), 247-254.

[ACW08] M. Atkinson, Z. Cao and L. Wein, "Optimal stopping analysis of a radiation detection system to protect cities from a nuclear terrorist attack," *Risk Analysis* **28** (2008), 353-371.

[Be57] R. Bellman, *Dynamic Programming*, Princeton University Press, 1957.

[BEK08] E. Boros, E. Elsayed, P, Kantor, F. Roberts and M. Xie, "Optimization problems for port-of-entry detection systems," *Intelligence and Security Informatics: Techniques and Applications*, H. Chen and C. C. Yang (eds), Springer, 2008, 319-335.

[BFK09] E. Boros, L. Fedzhora, P. Kantor, K. Saeger and P. Stroud, "Large scale LP model for finding optimal container inspection strategies, *Naval Research Logistics Quarterly* **56** (2009), 389-486.

[CCDX09] X. Chen, J. Cheng, S. Dalal and M. Xie, "Enhancing inspection process in nuclear detection by combing information from different sources," working paper, DIMACS Center, Rutgers University.

[CCX09] X. Chen, J. Cheng and M. Xie, "A statistical approach for analyzing manifest data in pre-portal intelligence," working paper, DIMACS Center, Rutgers University.

[CXR09] J. Cheng, M. Xie, and F. Roberts, "Design and deployment of a mobile sensor network in surveillance of nuclear materials in metropolitan areas," in *Proceedings of the 15th ISSAT International Conference on Reliability and Quality in Design*, (2009).

[CR09] Concho, A., and Ramirez-Marquez, J.E., "An evolutionary algorithm for port-of-entry security optimization considering sensor threshold," *Reliability Engineering and System Safety*, in press.

[DH09] S. Dalal and B. Han, "Detection of nuclear material in containers entering the US: a Bayesian approach for analyzing radiation portal data," working paper, DIMACS Center, Rutgers University.

[DG09] S. Dayanik and C. Goulding, "Detection and identification of an unobservable change in the distribution of a Markov-modulated random sequence," to appear in *IEEE Transactions on Information Theory*.

[DGP08] S. Dayanik, C. Goulding and H. V. Poor, "Bayesian sequential change diagnosis," *Mathematics of Operations Research*, **33** (2008), 475-496.

[DPY08] S. Dayanik, W. Powell, and K. Yamazaki, "Asymptotic analysis of sequential change diagnosis problem," *Proceedings of the International Workshop on Applied Probability*, 2008.

[DMD07] C. Demattei, N. Molinari, and J-P. Daures, "Arbitrarily shaped multiple spatial cluster detection for case event data," *Computational Statistics and Data Analysis*, **51** (2007), 3931-3945.

[DMD06] C. Demattei, N. Molinari, and J-P. Daures, "SPATCLAS: An R package for arbitrarily shaped multiple spatial cluster detection for case event data," *Computer Methods and Programs in Biomedicine*, **84** (2006), 42-49.

[DGM08] N. B. Dimitrov, M. A. Gonzalez, D. P. Michalopoulos, D. P. Morton, M. V. Nehme, E. Popova, . A. Schneider and G. G. Thoreson, "Interdiction modeling for smuggled nuclear material," *Proceedings of the 49th Annual Meeting of the Institute of Nuclear Materials Management,* 2008.

[ESX09] E. Elsayed, C. Schroepfer, M. Xie, H. Zhang and Y. Zhu, "Port-of-entry inspection: sensor deployment policy and optimization," *IEEE Transactions on Automation Science and Engineering*, **6** (2009), 265-277.

[FrPo08] P. Frazier and W. B. Powell, "The knowledge gradient stopping rule for ranking and selection," *Proceedings of the Winter Simulation Conference*, 2008.

[FPD08] P. Frazier, W. Powell and S. Dayanik, "A knowledge gradient policy for sequential information collection," *SIAM J. on Control and Optimization*, **47** (2008), 2410-2439.

[FPD08a] P. Frazier, W. B. Powell, S. Dayanik, "The knowledge gradient policy for correlated rewards," *INFORMS Journal on Computing*, **21**:4 (2009), 585-598.

[GLP09] M. C. Ganiz, N. I. Lytkin, and W. M. Pottenger, "Leveraging higher order dependencies between features for text classification," *Machine Learning and Knowledge Discovery in Databases,* Buntine et al. (eds.), Lecture Notes in Computer Science, **5781** (2009), 375-390.

[GEH03] B.D. Geelhood, J. Ely, R. Hansen, R. Kouzes, J. Schweppe, and R. Warner, "Overview of portal monitoring at border crossings," *IEEE Nuclear Science Symposium Conference Record*, 1 (2003), 513-517.

[GLM07] A. Genkin, D. Lewis, and D. Madigan, "Large-scale Bayesian logistic regression for text categorization," *Technometrics*, **49** (2007), 291-304.

[GWB08] N. Goldberg, J. Word, E. Boros, and P. Kantor, "Optimal sequential inspection policies," to appear in *Annals of Operations Research*, also RUTCOR Research Report 14-2008, Rutgers University, and DIMACS Technical Report 2008-07, DIMACS Center, Rutgers University.)

[HF09] D. Hochbaum and B. Fishbain,"Nuclear threat detection with mobile distributed sensor networks," *Annals of Operations Research* (in press), DOI 10.1007/s10479-009-0643-z

[H09] D. Hochbaum, "The multi-sensor nuclear threat detection problem," J. W. Chinneck B. Kristjansson, and M. J. Saltzman (eds.), *Operations Research and Cyber-Infrastructure*, Operations Research/Computer Science Interfaces Series, Springer, Vol. 47, 2009, 389-399.

[HCS09] R. Hoshino, D Coughtry, S. Sivaraja, I. Volnyansky, S. Auer, and A. Trichtchenko, "Application and extension of cost curves to marine container inspection," *Annals Of Operations Research* (in press), DOI 10.1007/s10479-009-0669-2.

[HOZ] R. Hoshino, W. Oldford, and M. Zhu, "Two-stage approach for unbalanced classification with time-varying decision boundary: application to marine container inspection," working paper.

[JKK06] S. H. Jacobson, T. Karnani, J. E. Kobza, and L. Ritchie, "A cost-benefit analysis of alternative device configurations for aviation checked baggage security screenings," *Risk Analysis* **26** (2006), 297-310.

[JMV05] S. H. Jacobson, L. A. McLay, J. L. Virta, and J. E. Kobza, "Integer program models for the deployment of airport baggage screening security services," *Optimization and Engineering,* **6** (2005), :339-359.

[KR86] S. Kapoor and V. Ramamurthy, *Nuclear Radiation Detectors*, New Age Publishers, 1986.

[KH79] O. Kariv and S. Hakimi, "An algorithmic approach to network location problems. I: The p-centers," *SIAM J. of Applied Math*., **37** (1979), 513-538.

[KH79a] O. Kariv and S. Hakimi, "An algorithmic approach to network location problems. II: The p-medians," *SIAM J. of Applied Math*., **37** (1979), 539–560.

[K09] C. Kearney, "New York police expand dirty bomb security", Reuters News, http://www.reuters.com/article/domesticNews/idUSTRE56067720090702, July 1, 2009.

[LLNL07] Lawrence Livermore National Laboratory, Computing Directorate, Annual Report, pg. 38, 2007.

[MMR07] D. Madigan, S. Mittal and F. Roberts, "Sequential decision making algorithms for port of entry inspection: Overcoming computational challenges," in *Proceedings of the International Conference on Intelligence and Security Informatics*, 2007, 1-7.

[MMR09] D. Madigan, S. Mittal and F. Roberts, "Efficient sequential decision-making algorithms for container inspection operations," *Naval Research Logistics*, submitted.

[MGLF05] D. Madigan, A. Genkin, D.D. Lewis and D. Fradkin, Bayesian multinomial logistic regression for author identification, *Proceedings of the 25th International Workshop on Bayesian inference and Maximum Entropy Methods in Science and Engineering* (MaxEnt 05), 2005, 509-516.

[MLN09] L. McLay, J. Lloyd and E. Niman, "Interdicting nuclear material on cargo containers using knapsack problem models," *Annals of Operations Research* (in press), DOI 10.1007/s10479-009-0667-4.

[M90] B. Mirchandani, *Discrete Location Theory*, Wiley-Interscience, 1990.

[NLK08] A. Neidhardt, H. Luss and K. Krishnan, "Data fusion and optimal placement of fixed and mobile sensors," *Proceedings of the IEEE Sensors Applications Symposium* (SAS2008), 2008.

[New09] A. Newman, "Using globally random walks to patrol a network," working paper, DIMACS, Rutgers University, 2009.

[PPT+07] P. Paruchuri, J. Pearce, M. Tambe, F. Ordonez, and S. Karus, "An efficient heuristic approach for security against multiple adversaries," *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems,* Honolulu, Hawaii, 2007.

[Per09] J. Perry, "Clustering and machine learning for gamma ray spectroscopy," Rutgers University, Department of Computer Science, Class Project, May, 2009.

[Po07] W.B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, John Wiley and Sons, New York, 2007.

[Pur08] "Purdue University: Cell phone sensors detect radiation to thwart nuclear terrorism." Online: http://news.uns.purdue.edu/x/2008a/080122FischbachNuclear.html

[Ram08] J. Ramirez-Marquez, "Port-of-entry safety via the reliability optimization of container inspection strategy through and evolutionary approach," *Reliability Engineering & System Safety* **93** (2008), 1698-1709.

[RP09] I. Rhyzov and W. B. Powell, "Optimal learning on a graph," under review by *Operations Research*.

[Rig09] M. Riggio, "Status report on federal and local efforts to secure radiological sources," prepared statement of testimony before the United States House of Representatives Committee on Homeland Security, Subcommittee on Emerging Threats, Cybersecurity, and Science and Technology, September 14, 2009. http://homeland.house.gov/SiteDocuments/20090914150055-28907.pdf

[RT09] F.S. Roberts, F.S., and B. Tesman (2009) *Applied Combinatorics*, 2nd ed., Chapman&Hall/CRC, an imprint of Taylor&Francis.

[SS03] P.D. Stroud, and KJ. Saeger, "Enumeration of increasing Boolean expressions and alternative digraph implementations for diagnostic applications", *Proceedings Volume IV, Computer, Communication and Control Technologies: I*, (eds. H. Chu, J. Ferrer, T. Nguyen, Y. Yu), 2003, 328-333.

[Sun08] Q. K. Sun, "Statistical modeling and inference for multiple temporal or spatial cluster detection," Ph.D. thesis, Department of Statistics, Rutgers University, 2008.

[W08] D. Weier, "Radiation Detection for DHS Applications and the Radiation Portal Monitor Project," talk given at the DIMACS/DyDAn/LPS Workshop on Port Security/Safety, Inspection, Risk Analysis and Modeling, November, 2008.

[WA07] L. Wein and M. Atkinson, "The last line of defense: designing radiation detection-interdiction systems to protect cities from a nuclear terrorist attack," *IEEE Transactions on Nuclear Science* **54** (2007) 654-669.

[WLCF07] L. Wein, Y. Liu, Z. Cao, and S. Flynn, "The optimal spatiotemporal deployment of radiation portal monitors can improve nuclear detection at overseas ports," *Science and Global Security* **15** (2007) 211-233.

[WWB06] L. Wein, A. Wilkins, M. Bajeva, and S. Flynn, "Preventing the importation of illicit nuclear materials in shipping containers," *Risk Analysis*, **26** (2006), 1377-1393.

[XSN08] M. Xie, Q. Sun and J. Naus, "A latent model to detect multiple clusters of varying sizes," *Biometrics*, **65** (2009)1011-1020.

[Yam09] K. Yamazaki, "Essays on Sequential Analysis: Multi-Armed Bandit with Availability Constraints and Sequential Change Detection and Identification," Ph.D. thesis, Department of Operations Research and Financial Engineering, Princeton University, 2009.

[ZLY09] Y. Zhu, M. Li, C.M. Young, M. Xie, and E. Elsayed, "Port of entry inspection policies: Incorporation of measurement errors," *Annals of Operations Research*, tentatively accepted, 2009.