A Mobile Sensor Network for the Surveillance of Nuclear Materials in Metropolitan Areas

by

Jerry Cheng, Minge Xie and Rong Chen
Dept. of Statistics
Rutgers University
New Brunswick, New Jersey 08903

Fred Roberts
DIMACS
Rutgers University
New Brunswick, New Jersey 08903

ABSTRACT

Nuclear attacks are among the most devastating terrorist attacks, with severe losses of human lives as well as damage to infrastructure. It becomes increasingly vital to have sophisticated nuclear surveillance and detection systems deployed in major cities in the U.S. to deter such threats. In this paper, we outline a robust system of a mobile sensor network and develop statistical algorithms and models to provide consistent and pervasive surveillance of nuclear materials in major cities. Specifically, the network consists of a large number of vehicles, such as taxicabs and police cars, on which nuclear sensors and Global Position System (GPS) tracking devices are installed. Real time readings of the sensors are processed at a central surveillance center, where mathematical and statistical analyses are performed. We use simulations to evaluate the effectiveness and detection power of such a network.

Key Words: Nuclear Detection and Surveillance, Latent Models, Cluster Detection, Scan Statistics

1 Introduction

Threats to national and homeland security have become more dynamic and complex in the past decade due to global terrorism, increased opposition to U.S. interests, greater pursuit of nuclear power and expanded access by adversaries to sophisticated technologies and materials. Among all the threats, nuclear attacks are arguably the most devastating. They can cause severe losses and casualties in human lives as well as long term and large scale damage to infrastructure. As the result, there have been growing concerns regarding the prospect of transporting, storing and detonating nuclear materials or dirty bombs in the populous metropolitan areas. Thus it becomes increasingly vital to have sophisticated nuclear detection systems deployed in major cities. Proactive monitoring and detection via pervasive surveillance is crucial to detect and thwart the malicious attacks.

We propose in this paper a massive surveillance network of mobile sensors that are installed in vehicles such as taxicabs, buses, and police cars. In such a network, when vehicles with sensors move within a certain range of a nuclear source, the radiation energy from the source will trigger the sensor devices to send out wireless signals to a central command center along with the positions of the sensors. With the random movement and extensive coverage nature of the vehicles, this setup provides a constant surveillance of nuclear materials. In this massive network, the mobile sensors do not need to be of high accuracy, since the failure of a small portion of them will not significantly affect the effectiveness of the surveillance coverage due to sensors' random movements. Mounted on vehicles, the sensors have fewer size constraints and power consumption requirements. We use less sophisticated sensors which only report binary signals instead of the actual readings of radiation intensity in this study. A positive signal is generated when the intensity from a nuclear source exceeds a certain threshold. Due to the mobility of the sensors, regular inspection, maintenance and calibration can be conducted at a central location. Thus this setup will further reduce the cost. More importantly, it is almost impossible to tamper with such a network of devices.

A mobile sensor network is often supplemented by stationary sensors. In fact, in most cases such a supplement is necessary to cover locations with sparse or zero traffic, such as a large park in the city. Our algorithm can be directly applied to such a combination since the stationary sensors may be viewed as the vehicles that are not moving in our models. While our algorithm can be easily adapted to other similar networks, we will work with the following typical mobile sensor network designs.

- Nuclear sensors and Global Position System (GPS) tracking devices are installed on a large number of vehicles such as taxicabs, police vehicles, fire trucks, and buses.

- The sensors and GPS devices constantly send detection and location information to a central surveillance center. Real time tracking signals are marked onto a map of a metropolitan area under surveillance.

- Real time analysis is done at the surveillance center using sophisticated statistical algorithms for the detection of existence and potential locations of nuclear sources.

Based on our preliminary investigations of current sensor technology, it is feasible to manufacture portable nuclear sensors with high accuracy. For instance, the leading manufacturers such as Thermo Scientific and ICX Technology have produced portable nuclear sensors with long range (up to 250 feet) detection capabilities. Currently these sensors are still too expensive for a large scale deployment. However, with the rapid advancement in technologies, low-cost sensors with a medium range and a reasonable accuracy should be available in the near future.

Due to many attractive characteristics of sensor networks, there have been many studies and applications of the sensor networks in military and civil applications including surveillance, smart homes, remote environment monitoring. See [1] and [2] for a recent survey. Much of the research devotes to sensor placement, sensor reorganization and communications. In the area of radiation detection, the idea of using massive mobile sensors has been adopted and tested by the Radiation Laboratory at Purdue University[3]. They use a network of cell phones with GPS capabilities to detect and track radiation. The noise and false positive detection problems are tackled by setting and tuning the solid state devices. A multi-sensor nuclear threat detection problem was studied in [4] using a combinatorial network flow algorithm.

In this paper, we propose a mobile sensor network following the aforementioned design and use statistical algorithms to analyze the network. Since the sensor signals are not 100% accurate, there are always false alarms or missed detections. For example, a sensor might display positive readings when there is no such signal, or fail to detect a real signal nearby. From the viewpoint of statistical modeling, the occurrence of missed detections can be treated as random. Statistical methodologies are effective tools for detecting true signals against random errors. We consider probabilistic models for sensor reading and source detection. These models are generalized to include multiple sources with different aggregation rules. In our work so far, we do not have a specific model of vehicle movement, which will be a future research topic. We assume that vehicles randomly roam within the monitoring region at each time instance. If there are many vehicles with sufficiently random movements, this is a reasonable first approximation. Since we do not model the vehicle movement, our approach is robust against model misidentifications, although it may compromise some detection power compared to other methods with an accurate model of vehicle movement.

Our algorithm in this study is based on recently developed statistical methods for detecting multiple spatial clusters [5] [6] [7] [8]. In particular, our simulation studies use an algorithm from the latent modeling approach proposed by [7] [8], which mimics the process of typical sample data generation. The method introduces a latent modeling structure and uses likelihood inference to detect multiple clusters simultaneously in an entire region. It can filter out known and harmless sources efficiently and is suitable to analyze the signals from our mobile sensor framework.

As mentioned before, the proposed mobile sensor network typically consists of a large number of sensors. The detection capability of the system depends on the size of the network, as well as other parameters. Our discussions with law enforcement agencies reveal some reluctance to rely on private sectors (e.g., taxicabs) in surveillance. However, are there

enough police cars to get sufficient coverage in a region? How many vehicles are needed for sufficient coverage? How does the answer depend upon the range, the false positive and false negative rates of detectors? These are some of the questions that we investigate in the paper.

The rest of the paper is arranged as follows. Section 2.1 discusses the nuclear intensity and sensor reading models. Section 2.2 covers the detection model. Section 2.3 reviews and outlines developments of statistical methodologies to detecting multiple spatial clusters. Section 3 describes simulation studies on several practical scenarios. We use the simulations to estimate detection powers of the network with different sets of parameters. Section 4 concludes the paper with discussions and future research directions.

2 Models and Methodology

2.1 Nuclear Intensity and Sensor Reading Models

We consider a nuclear source in this paper as a small portable nuclear device transported by an individual via trucks or bags [9]. As the nuclear radiation starts from a source, the total energy stays as a constant due to the Conservation Law of Energy. For simplicity, we assume that radiation travels in spherical waves. Let $z(r)$ denote the intensity at distance r. The total energy remaining a constant for all r is $4\pi r^2 z(r)$, where $4\pi r^2$ is the surface area of the sphere with radius r. As the radius increases by a factor of k, the surface area of the sphere will increase by a factor of k^2. As a result, the radiation intensity z decreases by the inverse square of the distance r [10]: $z(r) = c/r^2$, where the constant c is a factor related to the total energy of the source. Since the nuclear detection device is triggered by radiation intensity, getting closer to the nuclear source will better the chance for detection. The ubiquitous nature of the mobile sensor network takes advantage of this property.

As previously mentioned, we assume that the sensors report binary signals. Let S denote the status of the sensor's reading with the value of 1 for a positive reading and 0 otherwise. We describe S with a threshold model:

$$S = \mathbf{1}_{\{z(r) \geq d\}} = \mathbf{1}_{\{c/r^2 \geq d\}} \tag{1}$$

where d is a threshold for detection and $\mathbf{1}_{\{.\}}$ is the indicator function. That is, if the intensity $z(r)$ at the sensor location is greater than the threshold d, the sensor will detect the source; otherwise the sensor reports a negative reading.

In practice there might be multiple nuclear sources, whose energy levels and positions will jointly determine the reading status of a sensor. In this paper, we assume that nuclear energies from difference sources are additive. For example, they are all within same spectrum of frequencies. Let Ω be the number of sources, c_ω be energy factor of the ωth source, r_ω be the distance from the sensor to this source. The aggregation of intensities from all sources at the senor location is: $z_{total} = \sum_{\omega=1}^{\Omega} c_\omega / r_\omega^2$. From the threshold model (1), the reading S can be determined by:

$$S = \mathbf{1}_{\{z_{total} \geq d\}} = \mathbf{1}_{\{\sum_{\omega=1}^{\Omega} c_\omega / r_\omega^2 \geq d\}} \tag{2}$$

2.2   Nuclear Detection Model

As with any detection device, nuclear sensors are not necessarily 100% accurate. The inaccuracy may stem from the variability in the manufacturing process, routine wear and tear, missing scheduled maintenance and calibrations, and undetected malfunctions. In addition, random traces of weak environmental nuclear signals can also trigger false alerts. For example, a person who just went through a radioactive therapy or a bag of cat litter can trigger positive alarms. We regard such sources as *trivial* sources as they are weak and last a very short period of time. Furthermore the wireless signals from the mobile sensor to the control center may incur transmission errors.

We use the two parameters *sensitivity* and *specificity* to assess the average performance of a sensor device. In the context of nuclear detection, sensitivity, denoted as η, presents the probability of detecting nuclear sources where there are indeed such materials. Specificity, denoted as ζ, is the probability of not detecting any nuclear materials where there in fact do not exist any. Let D be the binary indicator of a sensor detecting a true nuclear source, D equal to 1 for the positive detection and 0 otherwise. We have $\eta = \mathbf{P}(D = 1|S = 1)$ and $\zeta = \mathbf{P}(D = 0|S = 0)$.

The quality control characteristics of a sensor, false negative rate (FNR) and false positive rate (FPR), can be expressed in η and ζ as: $FNR = \mathbf{P}(D = 0|S = 1) = 1 - \eta$ and $FPR = \mathbf{P}(D = 1|S = 0) = 1 - \zeta$. Then the probability of detecting a nuclear source is:

$$\mathbf{P}(D = 1) = \mathbf{P}(D = 1|S = 1)\mathbf{P}(S = 1) + \mathbf{P}(D = 1|S = 0)\mathbf{P}(S = 0)$$
$$= (1 - \zeta) + (\zeta + \eta - 1)\mathbf{P}(S = 1). \tag{3}$$

Under the perfect scenario where both η and ζ are 1, the detection D is the same as the reading S. In practice, extremely accurate sensors are not necessary. However, the effectiveness of the detection methods does depend on the accuracy.

The threshold model (1) can be expressed as $S = \mathbf{1}_{\{A\}}$, where $A = \{r \leq (c/d)^{1/2}\}$ is a sphere, or a circle on a 2-dimensional map, centered at the nuclear source and with radius $R = (c/d)^{1/2}$. The ratio of the probabilities of a positive reading inside and outside the set A is $P(D = 1|A)/P(D = 1|\bar{A}) = P(D = 1|S = 1)/P(D = 1|S = 0) = \eta/(1 - \zeta)$. In the case when both FNR and FPR are less than 25%, for instance, we have the ratio greater than 3. That is, the sensor is 3 times more likely to report a positive signal ($D = 1$) inside A than inside $\bar{A}$ with moderate accuracy. This type of statement matches the definition of a spatial cluster in the statistical literature, in which the clusters are defined as areas within which an incident of interest is more likely to happen (i.e., with a higher probability of happening per squared unit) than outside these areas. In our setting, an incident of interest is an alert signal with $D = 1$. Thus, detecting spheres or circles like A is equivalent to the cluster detection problem in statistics.

2.3   Statistical Methods for Detecting Spatial Clusters

A traditional statistical method to detect a cluster of events in spatial data is via *Scan Statistics* [11] [12] [13] [14]. The most commonly used scan statistic is the maximum number

of cases in a fixed size moving-window that scans through the study area. The test based on this scan statistic has been shown to be a generalized likelihood ratio test for a uniform null against a false alternative. A related scan statistic is the diameter of the smallest window that contains a fixed number of cases. Other scan statistics and related likelihood based tests for localized temporal or spatial clustering have been developed, often using a range of fixed window sizes or a range of fixed number of cases [15] [16] [17] [18]. Scan statistics methods have also been developed under the Bayesian framework [19] [20] [21] [22] [23].

Scan statistics procedures have been successful in detecting a single significant cluster, and they also have had some success in detecting multiple clusters of fixed sizes. But difficulties arise for detecting multiple clusters of varying sizes. In recent years, there have several attempts to overcome this difficulty. A well known approach is a stepwise regression model together with model selection procedures to locate and determine the number of unusually high clustering regions [5] [6]. These approaches rely on a weighted least square formulation, although the response variable (gaps between incidents) is typically non-Gaussian. Recently, Xie, Sun and Naus [7] developed a latent cluster model for temporal data which allows the use of the standard likelihood inference for detecting multiple clusters. Sun [8] extended the temporal cluster detection to spatial data and developed a spatial cluster detection method to simultaneously detect multiple clusters of varying sizes, as well as a significant single cluster. These approaches are based on likelihood inference and they are more efficient in detecting clusters of varying sizes than the weighted least squares approaches developed in [5] [6]. We use the likelihood inference based method in our study and outline below the main points of the approach.

We first assume that there are k non-overlapping clusters with the centers and radii as latent random variables: $\mathbf{O} = (\mathbf{o_1}, \mathbf{o_2}, ..., \mathbf{o_k})$, $\mathbf{r} = (r_1, r_2, ..., r_k)$. Denote the j^{th} cluster by A_j. We assume the observations $\mathbf{y} = (y_1, y_2, ..., y_n)$ are i.i.d. samples from a piecewise uniform density function:

$$f_\theta(\mathbf{y}|\mathbf{O}, \mathbf{r}, k) = \begin{cases} c\alpha_i & \text{if } \mathbf{y} \in A_i \\ c & \text{if } \mathbf{y} \notin \bigcup_{i=1}^{k} A_i \end{cases}$$

where c is a normalizing constant, θ is the collection of all unknown parameters, α_i is the relative density in cluster i.

If we want to know whether there are any significant clusters in the data, we can test a hypothesis $H_0: \alpha_1 = ... = \alpha_k = 1$ versus H_1: at least one $\alpha_j \neq 1$. In order to do this, we need to calculate the observed likelihood:

$$f_\theta(\mathbf{y}, k) = \int ... \int f_\theta(\mathbf{y}, \mathbf{O}, \mathbf{r}, \delta = k) d\mathbf{O} d\mathbf{r}$$

$$= \int ... \int f_\theta(\mathbf{y}|\mathbf{O}, \mathbf{r}, k) f_\theta(\mathbf{O}, \mathbf{r}|k) P_\lambda(\delta = k) d\mathbf{r}$$

where δ is a random variable such that $\{\delta = k\}$ is the event that k non-overlapping clusters occur in the region.

Since the integration is difficult to compute directly, we use an Expectation-Maximization (EM) algorithm (Dempster et al [24]) to solve the estimation problem where we treat

$(\mathbf{y}, \mathbf{O}, \mathbf{r}, \delta = k)$ as the complete responses and $(\mathbf{y}, \delta = k)$ as the observed ones. When a cluster is significant (i.e. $\alpha_j \neq 1$), the cluster region is determined by the conditional distribution of $\mathbf{o_j}$ and r_j given $(\mathbf{y}, \delta = k)$.

In reality, we don't know the value for the number of clusters. We use a standard model selection process, either the AIC (Aikaike Information) or BIC (Bayesian Information) Criterion, to determine the k value.

2.4   A design question on number of vehicles

Typically, the larger the number of vehicles, the higher the statistical power of detection. Here, the statistical power can be interpreted as the probability of detecting a true nuclear source by the network. The required number of vehicles in the surveillance network can be quantified by statistical power analysis. We have developed a model and carried out a large number of computer tests to assess power of detection under different assumptions. For example, we have studied a surveillance network that covers an area of 4000 feet by 10000 feet, roughly equal to the area of the roads and sidewalks of Mid/Downtown Manhattan. In this phase of our research, we are disregarding the street network. We fix key parameters such as effective range, false positive and false negative rates for the sensors.

3   Simulation Study

The key of processing sensor network information is to identify clusters from real signals mingled with random noises. In this section we conduct simulation studies to demonstrate that the proposed network and method can effectively detect single and multiple nuclear sources.

3.1   Detecting multiple nuclear sources

Let us first consider a study window of a rectangular area $(0, 100) \times (0, 300)$. We assume that there are two unknown and possibly malicious sources, one at position (33,225) and the other at (66, 60). The choices of the positions are for illustrative purposes and they do not impose any restrictions in the study.

We simulate 5000 points uniformly in the study window. Each point represents a vehicle mounted with a nuclear sensor at a given time point. The models (1) and (3) determine whether each point will be turned on to positive or not. In (1), we assume the energy factor c=20 for the two possible sources, and the energy threshold d=0.2. As a result, the nuclear source will not trigger the sensor if its distance is more than 10 units away. In (3), we assume all the sensors have η of 0.990 and ζ of 0.995. For each point i, we calculate the distance to the stationary source r_i^s and the distance to the moving sources r_i^m. Now the reading probability $\mathbf{P}(S_i = 1)$ in (2) is $I\{c/(r_i^s)^2 + c/(r_i^m)^2 \geq d\}$, and the detection probability $\mathbf{P}(D_i = 1)$ in (3) is $0.005 + 0.985\mathbf{P}(S_i = 1)$ from which the binary D_i is generated. We keep

the point if D_i is 1 and delete it otherwise. With all the points, we apply the Latent Model Clustering software by Sun [8].


Figure 1: Cluster detection. The left plot displays positive signals from the sensors via simulation. The right plot gives the detected clusters representing two nuclear sources.

Figure (1) has the simulated sensors with positive readings and detected clusters. The "+" in the left plot marks the locations of positive signals, and the dotted circles/ellipsoids in the right plot mark the clusters detected by the proposed method. They exactly point to the true locations of the nuclear sources.

3.2   Design parameters in the mobile sensor network

In this simulation, we aim to design a mobile sensor network in an area of the similar size of Midtown and Downtown Manhattan. The study region is set henceforth to a 40×100 rectangular, with one unit representing 100 feet in real distance. Thus, the rectangular has roughly the same area as totality of the streets and side walks south of the Central Park in New York City. We first consider a network consisting of taxicabs, where the vehicles or sensors roam randomly within the entire study region.

The set of the network parameters is designed as: 1) sensor detection range: 250 feet (scaled to 2.5 units) versus 150 feet (scaled to 1.5 units) (both higher than currently practical); 2) error rates (false positive rate and false negative rate): (2%, 5%) versus (5%, 10%); 3) number of sensors: from 1500 to 4000 with increments of 500. The numbers are picked in line with the 13,000 taxicabs in New York City; 4) network type: taxicab type versus police car type. At each set of network parameters, we repeat 200 simulations and compute how many times (in percentage) the proposed algorithm can correctly detect a randomly placed nuclear source as an estimate of the statical power of detection $P(D = 1|S = 1)$.

(a) Different detection ranges   (b) Different error rates
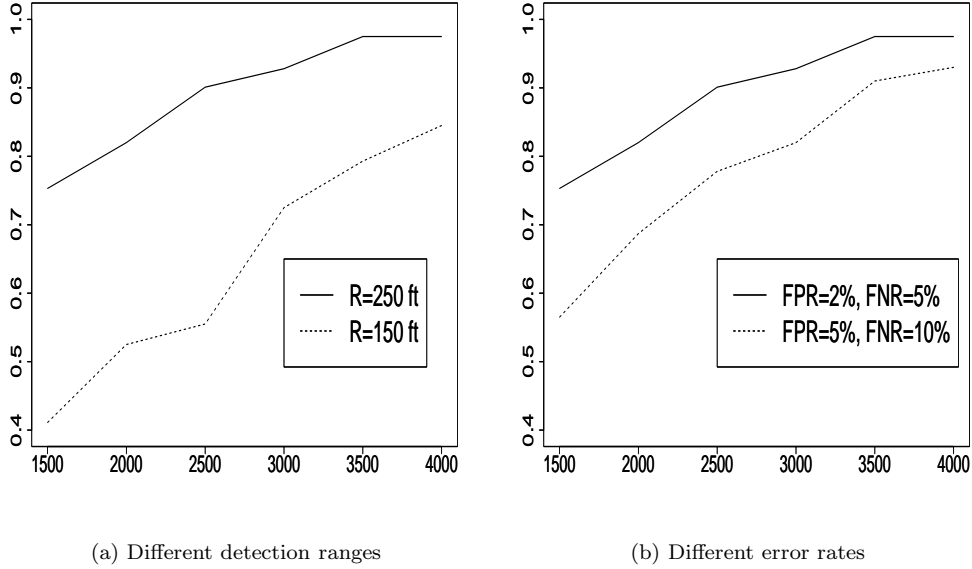
Figure 2: Detection Power Comparison: (a) Two different sensor ranges with the same error rates (2%, 5%). (b) Two different sets of error rates with the same sensor range of 250 ft. The horizontal axis is the number of sensors.

We start the study on the network of taxicabs with error rates (2%, 5%) to compare the effectiveness of the long-ranged (250 feet) against the short-ranged (150 feet) sensors. Figure 2(a) plots the two power curves and it shows that we need more than 3,000 long-ranged sensors to achieve around 95% detection power, while the short-ranged sensors yield about 20% less power with the same number of sensors.

Then we continue on the tax cab type network with fixed sensor detection range of 250 feet, and study the difference between the better quality sensors (error rates of 2% and 5%) and the lower quality sensors (error rates of 5% and 10%). The two power curves are plotted in Figure 2(b), which reveals that inferior sensors lag about 20% in the detection power at the sensor number of 3,000.

We also consider a network that consists of police vehicles. Note that the police patrol cars usually are limited into the boundaries of their precincts. Thus, in our simulation we consider the case that the sensors are moving within subregions to which they belong. The New York City Police Department has 3000+ vehicles in 76 precincts in 5 boroughs, 22 in Manhattan. Perhaps 500 to 750 are in the streets of the Manhattan borough at one time. In our police car model, the study region is divided into 20 identical subregions, each with the dimension of 10×20. A subregion in our simulation represents roughly a precinct. Within each of the 20 subregions, we simulate 25 police cars with the long-ranged (250 feet) and better quality (error rates of 2% and 5%) sensors and place them at random locations. With a randomly generated source in the whole study region, this setup achieves about 35% in detection. It suggests that the number of police cars in Manhattan would not be sufficient

to detect nuclear sources with a high power.

4 Discussion and Future Work

This paper outlines a robust mobile sensor network and a statistical algorithm to provide consistent and pervasive surveillance for nuclear or biological materials in major cities. Simulation studies suggest the proposed network and method can effectively detect nuclear signals placed in a spatial region. Although we only illustrate our approach at a fixed time point, we can collect and analyze such information at consecutive short time intervals, and it can be extended to provide dynamic surveillance for nuclear or biological materials in a metropolitan area.

Since static sensor networks provide a complementary detection capability, we need to study the network with a combination of both static and mobile sensors for optimal results. In addition, the algorithm that we have discussed is computationally intensive. In practice the signal points can be much more than what we have in the simulations. So improving efficiency and speed of the algorithms is the key to enable us to dynamically monitor and detect nuclear materials.

References

[1] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayiric, "Survey on sensor networks," IEEE Communications Magazine, pp. 102-114, August 2002.

[2] I. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayiric, "Wireless sensor networks: a survey," Computer Networks, vol. 38, pp. 393-422, 2002.

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

[4] D. Hochbaum, "The multi-sensor nuclear threat detection problem," in Proceedings of the Eleventh INFORMS Computing Society (ICS) Conference, 2008, in press.

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

[6] 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, vol. 84, pp. 42-49, 2006.

[7] M. Xie, Q. Sun, and J. Naus, "A latent model to detect multiple temporal clusters," Biometrics, 2009, in press.

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

[9] "Federal Emergency Management Agency (FEMA): Are You Ready?" (2008). [Online]. Available: http://www.fema.gov/areyouready/nuclear_blast.shtm

[10] L. Wein, A. Wilkins, M. Baveja, and S. Flynn, "Preventing the impor- tation of illicit nuclear materials in shipping containers," Risk Analysis, vol. 26, 2006.

[11] J. Glaz and N. Balakrishnan (Eds.). "Scan Statistics and Applications." Boston: Birkhauser, 1999.

[12] J. Glaz, J. Naus, and S. Wallenstein, "Scan Statistics and Applications." New York: Springer, 2001.

[13] N. Balakrishnan and M. Koutras, "Runs and Scans with Applications." John Wiley and Sons, 2001.

[14] J.C. Fu and W.Y Lou, "Distribution Theory of Runs and Patterns and its Applications: A Finite Markov Chain Imbedding Approach." World Scientific, 2003.

[15] M. Kulldorff and N. Nagarwalla, "Spatial disease clusters: Detection and infection," Statistics in Medicine, vol. 14, pp. 799-810, 1995.

[16] J. Naus, and S. Wallenstein, "Multiple window and cluster size scan procedures, Methodology and Computing in Applied Probability," vol. 6, pp. 389 400, 2004.

[17] A. Dembo and S. Karlin, "Poisson approximation for r-scan processes," Ann. Appl. Probab, vol. 2, pp. 329357, 1992.

[18] X. Su, S. Wallenstein, and D. Bishop, "Non-overlapping clusters: Approximation distribution and application to molecular biology," Biometrics, vol. 57, pp. 420-426, 2001.

[19] R. Gangnon and M. Clayton, "Bayesian detection and modeling of spatial disease maps," Biometrics, vol. 56, pp. 922-935, 2000.

[20] A. Lawson, "Markov chain monte carlo methods for putative pollution source problems," Environmetal Epidemiology, vol. 14, pp. 2473-2486, 1995.

[21] R. Gangnon and M. Clayton, "A hierarchical model for spatially clus- tered disease rates," Statistics in Medicine, vol. 22, pp. 3213-3228, 2003.

[22] L. Knorr-Held and G. RaBer, "Bayesian detection of clusters and discontinuities in disease maps," Biometrics, vol. 56, pp. 13-21, 2000.

[23] D. Denison and C. Holmes, "Bayesian partitioning for estimating disease risk," Biometrics, vol. 57, pp. 143-149, 2001.

[24] A. Dempster, N. Laird, and D. Rubin, "Maximum likelihood from incomplete data via the em algorithm," Journal of the Royal Statistical Society. Series B (Methodological), vol. 39(1), pp. 1-38, 1977.