Applied Stochastic Models in Business and Industry

Special Issue Paper

Received 23 October 2016, Revised 13 September 2017, Accepted 14 September 2017 Published online 17 November 2017 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/asmb.2293

A dynamic fusion system for fast nuclear source detection and localization with mobile sensor networks

Aude Grelaud[a], Priyam Mitra[b], Minge Xie[b] and Rong Chen[b*†]

This paper proposes a dynamic system, with an associated fusion learning inference procedure, to perform real-time detection and localization of nuclear sources using a network of mobile sensors. This is motivated by the need for a reliable detection system in order to prevent nuclear attacks in major cities such as New York City. The approach advocated here installs a large number of relatively inexpensive (and perhaps relatively less accurate) nuclear source detection sensors and GPS devices in taxis and police vehicles moving in the city. Sensor readings and GPS information are sent to a control center at a high frequency, where the information is immediately processed and fused with the earlier signals. We develop a real-time detection and localization method aimed at detecting the presence of a nuclear source and estimating its location and power. We adopt a Bayesian framework to perform the fusion learning and use a sequential Monte Carlo algorithm to estimate the parameters of the model and to perform real-time localization. A simulation study is provided to assess the performance of the method for both stationary and moving sources. The results provide guidance and recommendations for an actual implementation of such a surveillance system. Copyright © 2017 John Wiley & Sons, Ltd.

Keywords: SMC algorithm; Bayesian model choice; nuclear detection; real-time localization

1. Introduction

In the last decade, protecting human lives and infrastructures from terrorist attacks has been one of the priorities of governments around the world. Nuclear attacks with dirty bombs and other portable nuclear devices are one of the most devastating potential attacks, especially when happening in populous metropolitan areas. The Domestic Nuclear Detection Office of the USA, which is affiliated with the US Department of Homeland Security (DHS), was created in 2005 in order to increase the nation's capacity to detect the presence of unauthorized nuclear material and to prevent them from entering the country. The US Defense Threat Reduction Agency, in partnership with the US National Science Foundation and the US defense industry, has supported research programs focusing on the development of mathematical and statistical tools aimed at detecting nuclear weapons, chemical attacks, and related threats. In 2014, Defense Advanced Research Projects Agency launched the SIGMA program that aimed to revolutionize detection and deterrent capabilities for countering nuclear terrorism. A key component of SIGMA involves developing novel approaches to achieve low-cost and high-efficiency radiation detectors. The program also seeks to leverage existing infrastructure to help the deployment of sensors for developing game-changing detection and deterrent systems [1]. In summary, effective methods to detect the presence of nuclear material in large cities is in urgent need. In this paper, nuclear detection refers to inferring the possible existence of radiation-emitting materials of concern, in particular, radiation dispersion devices, also known as dirty bombs, as well as special nuclear materials including highly enriched uranium and weapons. The approach also provides an estimation of the location of the flagged potential nuclear source, which we refer to as localization.

One approach to protecting a major city is to place a small number of powerful but expensive sensors at different sensitive locations, such as major entry points of a city [2]. These sensors have high accuracy, but because of their high cost, it may be exceedingly expensive to provide coverage of the entire city. As a result, the surveillance system may have inherent

[a] *CREST-ENSAI, Bruz, France*
[b] *Rutgers University, Piscataway, NJ, U.S.A.*
**Correspondence to: Rong Chen, Rutgers University, Piscataway, NJ, U.S.A.*
[†] *E-mail: rongchen@stat.rutgers.edu*

blind spots, which can be exploited by potential terrorists. A different approach is to form a fixed sensor network on a grid of monitoring points with a large number of relatively inexpensive (hence low accuracy) sensors. There have been studies using such a fixed network for detection and monitoring of nuclear material, including mobile sources [3–9]. Small-scale experiments have been conducted to test such systems [3]. Fixed sensor network has its drawbacks, however, in that it is difficult to maintain and can be tampered.

Here, we study the use of a mobile sensor network for nuclear material surveillance. This idea was proposed by researchers from the Command, Control, and Interoperability Center for Advanced Data Analysis, a DHS Center of Excellence, in a project meeting of principle investigators hosted by DHS in 2009. It was subsequently reported in Carpenter *et al*. [10] and was also adopted by Hochbaum and Fishbain [11] who used a graph theory to study the problem. Similar ideas have been studied in Borozdin *et al*. [12] and Klimenko *et al*. [13]. In fact, Defense Advanced Research Projects Agency SIGMA program tested this approach by deploying sensors in emergency vehicles in Washington, D.C., in the early 2017 [14].

We emphasize here that the term 'detection' as used in this paper refers to a probabilistic decision rule to determine the existence of a nuclear source based on inaccurate sensor observations. We do not study physical detection of nuclear sources, as in sensor design or sensor signal analysis. Our basic assumption is that we have a mobile network equipped with a type of inexpensive sensors that provides inaccurate binary detection signals.

Such a network relies on taxicabs, buses, and police and emergency vehicles with nuclear material detection sensors installed in them. The sensors send a positive signal when they are within a certain distance from a nuclear source, and a negative signal otherwise, although false signals are possible [15]. These vehicles, hereafter referred to as taxicabs for simplicity, travel through most parts of the big city, and it is possible to track their exact location with GPS at all times. Even though the nuclear sensors deployed are not of high accuracy (hence inexpensive), having multiple and moving sensors in the neighboring area compensates for this weakness. A control center receives and analyzes the binary signals from these sensors at high frequency (say, every 30 s) and triggers an alarm if a nuclear source is detected through a fusion learning inference procedure. Analysis of binary and noisy detection signals is performed in real time. The consideration of false-positive and false-negative signals mitigates the weakness of using inexpensive sensors because the detection algorithm relies on the combination of all the signals. Another advantage is that, unlike a fixed sensor network, it is very difficult to tamper with such a mobile network. The sensors can be routinely checked at taxi depots or during annual vehicle maintenance checks with low maintenance cost. The same approach can be applied in other contexts; for instance, mobile sensors can be installed on military personnel in a battle field to detect chemical or biological agents. It can perhaps also be used to analyze data from connected networks of autonomous vehicles for preventing collisions or locating a nearby available vehicle on demand. A mobile network can also possibly have blind spots in areas that cannot be reached by any vehicles, although this may be mitigated by installing some fixed sensors at these locations. Cheng *et al*. [28] provides additional discussions on this type of mobile network.

In practice, the data can be displayed on a map, each sensor having a color indicating its signal and the presence of nuclear source producing a cluster of positive signals on the map. From a statistical point of view, the aim becomes assessing the likelihood of the presence of a cluster in a certain area. This kind of problem arises in epidemics in locating the first cases of a disease given the locations of all cases [16, 17]. These approaches are based on the analysis of data from one particular time and under the assumptions that incidence rate is low and cases are uniformly distributed within the neighboring area of the source. Commonly used statistical inference procedures are based on scan statistics [18–20] and their generalizations [21], where the methods utilize a moving window or box to scan through entire study area, seeking regions with unusually high rates of incidences. It can also be analyzed in a Bayesian framework based on scan statistics methods [22–26]. Some recent applications include the use of a Bayesian approach for real-time online disease surveillance. This approach detects peak locations of the intensity function of a point process [17] and fits a hierarchical model with latent variables associated with an Monte Carlo expectation maximization algorithm for detection of multiple clusters [27].

The nuclear surveillance problem is similar to the cluster detection problem earlier but has several key differences that prevent direct application of the existing methods. First, the presence of a nuclear source leads to an excess of positive signals, but the proportion of positive signals – playing the role of the incidence rate – is not small enough compared with those observed in epidemics. Second, the negative signals also contain some relevant information; therefore, it is necessary to analyze both positive and negative signals. Because of the moving nature of the sensors, it is difficult to treat negative signals as background noises.

Cheng *et al*. [28] studied nuclear detection procedure under the mobile sensor network we adopt here but used a static approach rather than a dynamic approach. Specifically, Cheng *et al*. [28] assumes a model with latent variables, and the estimation is performed with a Monte Carlo expectation maximization algorithm, extending that in Xie *et al*. [27]. If the source is not moving, a single fixed location is to be estimated, and combining observed at different time periods to improve detection and localization accuracy does not require any change in their methodology. However, as the source is more likely to be moving (being carried by a walking person, or in a moving vehicle), simple data combination is not

appropriate. Cheng *et al*. [28] proposed to estimate the trajectory of the source as the concatenation of single-time period location estimates. The path resulting from this method may contain quite large jumps between two consecutive time slots because the analysis is performed independently over time.

In this paper, we propose an approach that can accommodate both static and moving sources. In contrast to Cheng *et al*. [28], we perform real-time detection and localization, jointly utilizing the information observed at current and past time points. In order to connect the source locations at consecutive time periods, our approach assumes a simple motion model that mimics real traffic on a grid and covers both static and moving sources. Such an assumption significantly enhances detection and localization accuracy over static methods. We define a state–space model and propose Bayesian fusion learning inference of the parameters and hidden states. Detection is then based on Bayes factors for comparing the models both with and without a source. Such a Bayesian detection scheme with a distributed network has been studied in the literature (e.g., Varshney [29]). Our approach can be viewed as a special case under the general Bayesian distributed detection framework. Because of the dynamic setup, we utilize sequential Monte Carlo (SMC) method [30–32], which has been used extensively in many applications including source tracking under different settings [30, 31] and wireless communications signal processing [33]. SMC provides an estimation of the hidden states as well as the marginal likelihood, allowing for model comparison. The results of the development will help provide guidance and recommendations for actual implementation of such a surveillance system.

Our approach is similar to that in Rao *et al*. [34] in which a simple particle filter is used for localization of a mobile nuclear source with a fixed sensor network. There are several differences between ours and that in Rao *et al*. [34]. First, in this paper, we promote a novel mobile sensor network for nuclear surveillance of a major city using inaccurate binary signals. Second, we incorporate a motion model for the source to fully utilize the underlying motion dynamics, and we use a more sophisticated SMC method. Third, we use multiple panels and Bayesian decision approaches for the detection of the existence of a nuclear source. Sun *et al*. [35] proposed a control framework using an active sensor network to track a potential nuclear source.

The rest of the paper is organized as follows. Section 2 presents the state–space model used to analyze the data generated by the mobile sensor network. The estimation procedure based on SMC algorithm is introduced in Section 3. The detection and localization system based on model selection and fusion learning procedure is presented in Section 4. Section 5 contains a simulation study to assess the performance of the methodology and its several implementations. Section 6 concludes.

2. A dynamic system for mobile detection and localization network

2.1. State–space model

A state–space model consists of two equations: the state equation

$$x_t \sim h(\cdot|x_{t-1}), \tag{1}$$

which reflects the underlying dynamics of the unobserved states, and the observation equation

$$\mathbf{z}_t \sim f(\cdot|x_t), \tag{2}$$

which provides linkage between the observations and the states. Here, x_t is a sequence of underlying state that evolves according to (1), and $\mathbf{z}_t$ is an observation that reveals some characteristics of the underlying state x_t through the observation equation (2).

In our framework, the unobserved state variable x_t corresponds to the two-dimensional source location at time t. The conditional distribution $h(\cdot|\cdot)$ that governs the motion of the source from the location at time $t-1$ to that at time t is referred to as the source motion model. Details of the motion model will be specified later. The observations, $\mathbf{z}_t$, collected at time t, consist of the GPS locations, $\mathbf{y}_t = (y_t^{(1)}, \cdots, y_t^{(n_t)})$, of all the sensors in the system and the binary signals the sensors send to the control center $\mathbf{d}_t = (d_t^{(1)}, \cdots, d_t^{(n_t)})$. Here, n_t is the number of active sensors. The locations, $\mathbf{y}_t$, are assumed to be accurate, but the signals received by the control center could possibly be false.

This model assumes the presence of one source. It is compared with the baseline model, where no source is present, for the purpose of detecting the presence of a source. Details of the comparison are presented in Section 3.2. The extension to multiple source detection is discussed in Section 6.

State–space models have been widely used in many applications, including source tracking [29, 30, 36–43]. However, our approach is different from standard source tracking problems, mainly because in our case, both negative and positive binary signals are observed and used. The use of state–space formation allows the data collected at time t to be used for the estimation of the location of the source x_t with the information observed in the past, through an updating of the previous estimation of x_{t-1}, instead of performing an independent analysis at each time point. A source's motion possesses

strong continuity and memory, and combining observed information from consecutive times can improve the detection and localization.

2.2. Observations and sensor signal model

Here, we assume the source is a portable device emitting nuclear radiations. Its total energy E is unknown but remains constant. According to the Conservation Law of Energy [44], we assume that radiations travel in spherical waves; hence, the intensity of the signal at a distance l from the source is $E/4\pi l^2$. Note that a dense urban environment could present shielding effects and radiation may not travel in spherical waves [11]. It is possible to let these parameters dependent upon the distance and the location of the sensors in order to take into account shielding effects. This would require detailed knowledge of the entire city's geographical and topological property, which may be mapped with a detailed survey.

A nuclear sensor measures the level of nuclear energy through a complex system based on the recording of a spectrum of gamma ray emissions [45]. If the intensity exceeds a threshold, the sensor sends a positive signal; otherwise, the signal remains negative. In practice, the sensor measures the radiation intensity, but the measurement is not accurate for low-cost sensors, like the ones used here. The industry standard approach is thus to design sensors sending only a binary signal, controlling its sensitivity and specificity [46]. In statistics, there exists theory supporting the use of a dichotomized binary signal over a corresponding inaccurate measurement [47]. The range, r, is defined as the distance between the source and the sensor when the intensity is equal to the threshold. It depends on the quality of the sensor as well as on the unknown total energy of the source E. From now on, we use the range parameter r in the model instead of the total energy E. Following Cheng *et al*. [28], the detection status s of a sensor is binary and defined as

$$s = \mathbf{1}_{\{l \le r\}}, \tag{3}$$

where l is the distance between the sensor and the source. This represents the reading of a perfectly accurate sensor. We first treat the range parameter, r, as known. Its estimation, treated as a model selection problem, will be considered later in Section 3.2.

The signal received at the control center, denoted as d, may be inaccurate and may differ from the true status, s. The relationship between d and s is given by sensitivity, η, and specificity, γ, where

$$\gamma = \mathbb{P}(d = 0 | s = 0), \quad \eta = \mathbb{P}(d = 1 | s = 1).$$

Following Cheng *et al*. [28], the sensitivity and specificity are treated as known constant values, either given by the sensor manufacturer or through a preliminary study in a laboratory. In reality, they may depend on distance between the sensor and the source, especially when the source is outside the manufacturer-provided effective range. But because the detection signal is binary, the impact of the constant assumption as an approximation is limited, according to our preliminary study and field experts. Transmission problems may also cause the signal to be different from the status, and this error is included in η and γ. When background energy such as those emitted from a hospital is present, sensitivity and specificity may be adjusted to conform with the background. Indeed, based on our conversations with field experts, the parameters η and γ can be calibrated for practical use by placing known sources at various locations, while the manufacturers or laboratories provide the baseline measures.

Figure 1 shows an example of data in a focused region of 6×6 city blocks. The source and its detectable range is shown as the shaded circle. The circle points shown are true-negative signals, and the filled-grey points are false-negative signals in the detection region. The filled black dots are positive signals – the ones inside the shaded circle are true positive signals, and the ones outside are false-positive ones. Note that the figure is static. Our proposed approach aggregates information through time to enhance the information, because the true-positive signals tend to be present inside the area of detection but the false-positive signals outside the detection area (in an area of similar size) are random and are more likely to disappear over time.

The sensor (observation) model considered in this paper assumes the following: (i) there is no background energy; (ii) sensors are assumed to be independent of each other; and (iii) given the true status, signals sent by the same sensor are independent from one time period to another. Under the earlier assumption, the likelihood of a single sensor i at time t is given by

$$f(z_t^{(i)} | x_t, r) = \eta^{d_t^{(i)} s_t^{(i)}} (1-\eta)^{(1-d_t^{(i)}) s_t^{(i)}} \gamma^{(1-d_t^{(i)})(1-s_t^{(i)})} (1-\gamma)^{d_t^{(i)}(1-s_t^{(i)})}, \tag{4}$$

where x_t is the unknown source location and $s_t^{(i)}$ is defined in Equation 3 for each sensor i and at time point t, where the distance is the distance between x_t and $y_t^{(i)}$. We denote $TP_t = \sum_{i=1}^{n_t} d_t^{(i)} s_t^{(i)}$ the number of true-positive signals at time t, $FN_t = \sum_{i=1}^{n_t} (1 - d_t^{(i)}) s_t^{(i)}$ the number of false-negative signals at time t, $TN_t = \sum_{i=1}^{n_t} (1 - d_t^{(i)})(1 - s_t^{(i)})$ the number of

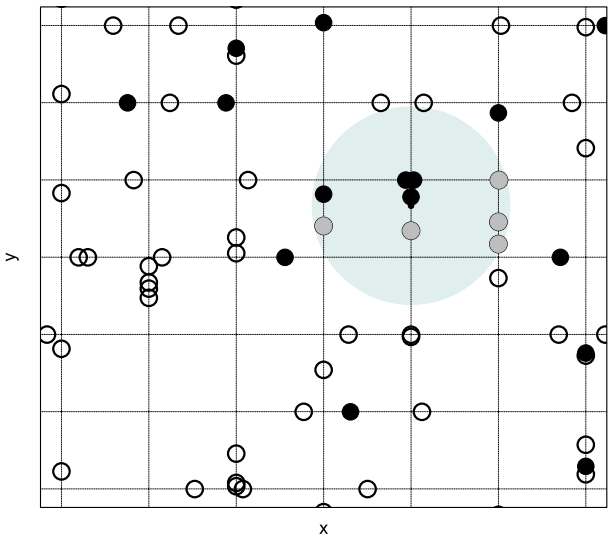

Figure 1. Mapping of the received signals at a particular time. A sensor is represented by a dot. A solid black dot represents a positive signal, and the red circles and red dots are negative signals. The shaded area represents detectable region. The red circles are false negatives. [Colour figure can be viewed at wileyonlinelibrary.com]

true-negative signals at time t, and $FP_t = \sum_{i=1}^{n_t} d_t^{(i)}(1 - s_t^{(i)})$ the number of false-positive signals at time t. Because the signals emitted by the sensors are assumed to be independent, the likelihood for all signals at time t is given by

$$f(\mathbf{z}_t|x_t, r) = \eta^{TP_t}(1-\eta)^{FN_t}\gamma^{TN_t}(1-\gamma)^{FP_t}. \tag{5}$$

Let $A(x_t, r)$ be the circle centered at x_t with radius r, then TP_t and FN_t are the numbers of positive and negative signals in $A(x_t, r)$, respectively, and FP_t and TN_t are the numbers of positive and negative signals outside $A(x_t, r)$, respectively.

The likelihood of the fused data from time 1 to T is

$$f(\mathbf{z}_{1:T}|x_{1:T}, r) = \prod_{t=1}^{T} f(\mathbf{z}_t|x_t, r),$$

as the signals are also assumed to be independent over time.

Remark 1

Although it is reasonable to assume that the sensors operate independently of each other, conditioned on whether a source is in the range or not, it is questionable to assume that the signals from the same sensor are independent over time. For example, a malfunctioned sensor will consistently produce false signals. Unfortunately, to take account of such dependency in the analysis would require tracking the motion of each individual sensor and estimating the sensitivity and specificity of each sensor at each time. The computational cost would be too high for online processing and monitoring. On the other hand, the impact of ignoring the dependency is minimum, as the sensors are moving and the sensors within the detection range of the source (or in any given local area) are always changing over time, hence producing almost independent signals within a small local area (whitening by localization). Such an independence approximation makes statistical analysis much simpler and faster. This is in fact one of the advantages of using mobile sensors. Analysis of a fixed sensor network would need to take the dependency into account.

2.3. Source motion model

The source is assumed to be placed in a fixed location or be carried by somebody walking in the street or in a moving car. We assume the source stays on a grid system that is similar to the road system in New York City. The motion model associated with the source, corresponding to the state distribution in (1), will need to address the different types of movements.

Our proposed motion model relies on weak assumptions, but it is flexible enough to accommodate all types of motions on a grid road system. It is well suited for a walking object or a vehicle moving in a traffic congestion, which is often the case in large cities such as New York City. Specifically, we assume the location of the source is always at the street level on a street map. Given x_{t-1}, the source location at time $t-1$, the distribution of the next location x_t is constructed sequentially as follows. First, the Manhattan distance between x_t and x_{t-1} is determined, following a uniform distribution,

that is, $p(v) = \text{Unif}[0, B]$, where $v = ||x_t - x_{t-1}||$ and B is the maximum distance a source may travel between the time period $[t-1, t]$. Given the distance travelled, the exact location of x_t is chosen from the set $C(x_{t-1}, v)$, which contains all locations on the grid that is v distance away from x_{t-1}, including all possible turns, but within the monitoring area. In our setting, we assume the source does not leave the area. Let $g_{C(x_{t-1},v)}(x_t|x_{t-1}, v)$ denote the selection probability of x_t from $C(x_{t-1}, v)$. We use the following assignment. First, the source has a probability δ to move forward on the street it was on at time $t-1$ and probability $1-\delta$ to make a U-turn. In the direction of travel, if the distance v allows the source to reach and pass an intersection, it will have equal probability to make all possible turns (e.g., straight, left, and right if not on the boundary). Figure 2 depicts such a scenario, and Figure 3 shows a simulated trajectory of a source moving for 30 min (60 steps). The combination of distance v and location x_{t-1} specifies the conditional distribution in (1):

$$h(x_t|x_{t-1}) = p(v)g_{C(x_{t-1},v)}(x_t|x_{t-1}, v).$$

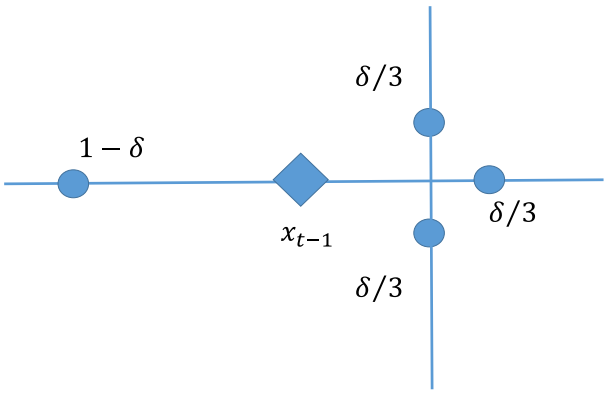

Figure 2. An example of the distribution of $g_{C(x_{t-1},v)}(x_t|x_{t-1}, v)$. As depicted, the source can go east with probability δ. If it goes east, it crosses an intersection; hence, three different locations can be reached on this side, each with probability $\delta/3$. [Colour figure can be viewed at wileyonlinelibrary.com]

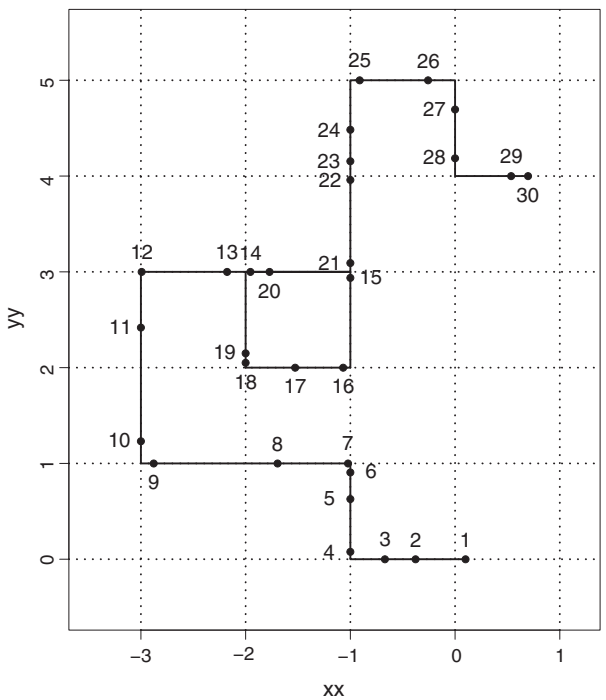

Figure 3. An example of the trajectory of a moving source, with forward probability 0.95, and distance follows a uniform distribution between 0 and 1.2 city blocks per 30 s.

It is quite complicated to model the traffic pattern in a large city because this depends on the location, the street, the time of the day, and many other issues. For a city with heavy congestion, with stop-and-go traffic and many intersections with traffic lights, our model provides a simple description. Our study shows that the result is not sensitive when the maximum distance B allowed in each time period changes in a reasonable range. More sophisticated models that involve speed, traffic lights, and other possible scenarios are possible and may provide more accurate results, but they may also demand significantly more computational resources for the control center and may not be necessary. The simple motion model covers all motion behaviors. If the road system is not on a regular grid, a detailed map can be used under our framework without much complication.

3. Inference algorithm

The motion model (Section 2.3) and the observation equation (Section 2.2) form a state–space model where the source location is the unobserved state.

3.1. Estimation of the location with a known range parameter r

Under the state–space model framework, the estimation of the source location x_t at time t is a typical filtering problem, that is, we are interested in $p(x_t|\mathbf{z}_1, \cdots, \mathbf{z}_t)$. As the system is nonlinear and non-Gaussian, we use SMC algorithm [30, 31]. SMC methods consist of simulating from a sequence of distributions, all defined on the same space [32]. A sample of the states, called 'particles', is moved from one distribution to the next using an iterated importance sampling scheme. Periodic resampling is used to ensure filtering quality. Estimation of the distributions or quantities related to the distributions is obtained by a Monte Carlo approximation.

Briefly, at time 0, particles $\{x_0^{(j)}\}_{j=1,\cdots,N}$ are drawn from an initial distribution. In our case, they are generated from a uniform distribution on the map. The particles are then 'moved' from time $t-1$ to time t, using a proposal distribution $q(x_t^{(j)}|x_{t-1}^{(j)})$. Then importance weights $W_t^{(j)}$ are calculated to ensure $\{x_t^{(j)}\}_{j=1,\cdots,N}$ are properly weighted with respect to $p(x_t|\mathbf{z}_t)$ [30]. That is, the importance weight $W_t^{(j)}$ has the property that, for any integrable function $h(\cdot)$, we have

$$\frac{\sum_{j=1}^{N} h(x_t^{(j)}) W_t^{(j)}}{\sum_{j=1}^{N} W_t^{(j)}} \to E[h(x_t) \mid \mathbf{z}_t] \quad \text{as} \quad N \to \infty.$$

The performance of the method depends on the design of the proposal distribution $q(\cdot)$. The optimal $q(\cdot)$ combines both the source motion model and the information from the observations. In our implementation, we use $q(x_t|x_{t-1}) \propto p(v) g_{C(x_{t-1},v)}(x_t|x_{t-1}, v) f(\mathbf{z}_t \mid x_t, r)$, where $p(v)$ and $g_{C(x_{t-1},v)}$ are defined in Section 2.3, and $f(\mathbf{z}_t \mid x_t, r)$ is defined in (5) in Section 2.2. That is, we first generate v from the uniform distribution $U[0, B]$ then determine the possible locations $C(x_{t-1}, v)$. We then sample x_t from the set $C(x_{t-1}, v)$ with probability

$$p(x) = \frac{g_{C(x_{t-1},v)}(x|x_{t-1}, v) f(\mathbf{z}_t \mid x, r)}{\sum_{C(x \in x_{t-1},v)} g_{C(x_{t-1},v)}(x|x_{t-1}, v) f(\mathbf{z}_t \mid x, r)}.$$

This proposal distribution effectively combines the motion model with the observed signals. The corresponding weight associated with the sample is $w_t = \sum_{C(x \in x_{t-1},v)} g_{C(x_{t-1},v)}(x|x_{t-1}, v) f(\mathbf{z}_t \mid x, r)$.

Resampling is an important component of SMC [30, 48–50]. Resampling consists of sampling with replacement N particles with probability proportional to the weights $W_t^{(j)}$. The newly resampled particles are assigned with weight 1.

Based on the output, at time t, an estimate of the posterior distribution of the location can be obtained using the empirical distribution of the particles $x_t^{(j)}$. The maximum a posteriori estimator of the location $\widehat{x}_t$ corresponds to

$$\widehat{x}_t = x_t^{(j^*)}, \text{ where } j^* = \arg\max_j W_t^{(j)}. \tag{6}$$

To be more robust and to avoid isolated samples with extremely large weights, one can obtain a smoother estimator of $p(x)$ using a kernel smoother

$$\hat{p}(x) = \frac{\sum_{j=1}^{N} W_t^{(j)} K_h(x_t^{(j)} - x)}{\sum_{j=1}^{N} K_h(x_t^{(j)} - x)}$$

and finding the mode of $\hat{p}(x)$ as the maximum a posteriori for x_t.

3.2. Range parameter and source detection

The effective range r of the sensors deployed is linked to the unknown total energy E of the source as well as the sensor property. There are several approaches of fixed parameter estimation in state–space model with SMC, including particle Markov chain Monte Carlo [51] and maximum likelihood estimation [52]. However, for fast online processing and without much loss of accuracy in estimating the source location, we adopt a simpler approach. Specifically, we allow the range parameter r to take only a small set of possible values, such as $\{0.5, 0.75, 1, 1.5, 2\}$ city blocks. This is a rough approximation. The empirical study presented later shows that such an approximation does not have a significant impact on detection power. One can increase the number of possible values to reduce the approximation error at the expense of more computational time.

For detection, we first establish a baseline model, hereafter referred to as model M_0, which assumes that no nuclear source is in the region of interest. It can be seen as a special case of the model defined in Section 2.2 with the range r equals to 0. Because of the absence of source, signals can only be either false positive $FP_t = \sum_{i=1}^{n_t} d_t^{(i)}$ or true negative $TN_t = \sum_{i=1}^{n_t}(1 - d_t^{(i)})$. The likelihood is then

$$f(\mathbf{z}_t|x_t, r = 0) = (1-\eta)^{FP_t}\gamma^{TN_t}. \tag{7}$$

The alternative model, M_1, is the model with one source with unknown range taking value in the set, say, $\{r_1, r_2, \cdots, r_k\}$. With a prior distribution on r and model averaging over the different ranges in the set, a Bayesian model selection procedure can then be used for detection. In a Bayesian framework, model comparison is often based on the Bayes factor [53] defined as

$$BF(\mathbf{z}) = \frac{\pi(M=a)}{\pi(M=b)}\frac{\pi(\mathbf{z}|M=a)}{\pi(\mathbf{z}|M=b)} = \frac{m_a(\mathbf{z})}{m_b(\mathbf{z})}, \tag{8}$$

where $k \in \{a, b\}$ is the model index, $\pi(M = k)$ is the prior probability of model k, and $m_k(\mathbf{z})$ is the marginal likelihood of model k. A Bayes factor, $BF(\mathbf{z})$, higher than 1 indicates a preference for model a. According to Jeffrey's scale, strong evidence for a model corresponds to $\log BF(\mathbf{z}) > 3$ [53–55]. In our dynamic setting, the Bayes factor can be updated easily from $t-1$ to t. Specifically, we have

$$\begin{aligned} BF_t(\mathbf{z}_1, \cdots, \mathbf{z}_t) &= \frac{m_a(\mathbf{z}_1, \cdots, \mathbf{z}_t)}{m_b(\mathbf{z}_1, \cdots, \mathbf{z}_t)} \\ &= \frac{m_a(\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}{m_b(\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}\frac{m_a(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}{m_b(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})} \\ &= BF_{t-1}(\mathbf{z}_1, \ldots, \mathbf{z}_t) IBF_t(\mathbf{z}_1, \ldots, \mathbf{z}_t), \end{aligned}$$

where IBF_t is the incremental Bayes factor

$$IBF_t = \frac{m_a(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}{m_b(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}. \tag{9}$$

Bayes factor and marginal likelihoods of complex systems are often estimated using MCMC techniques [56–59]. SMC has the special property that the marginal likelihood $m_k(\mathbf{z}_t)$ can be obtained and updated without extra computational cost [32]. Specifically, we have

$$\hat{m}_k(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1}) = \frac{1}{N}\sum_{j=1}^{N} W_t^{(j)}. \tag{10}$$

When resampling was performed at time $t_1, \ldots, t_k < t$, we have

$$\hat{m}_k(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1}) = \left[\prod_{i=1}^{k}\frac{1}{N}\sum_{j=1}^{N} W_{t_i}^{(j)}\right]\frac{1}{N}\sum_{j=1}^{N} W_t^{(j)}.$$

An estimator of the incremental Bayes factor can be obtained by $\widehat{IBF}_t = \frac{\hat{m}_a(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}{\hat{m}_b(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1})}$.

For model M_1, obtained by model averaging over the values of r, the marginal likelihood is defined as

$$m_1(\mathbf{z}_t|\mathbf{z}_1, \cdots, \mathbf{z}_{t-1}) = \sum_{i=1}^{k} m_1(\mathbf{z}_k|r = r_i, \mathbf{z}_1, \cdots, \mathbf{z}_{t-1})\mathbb{P}(r = r_i),$$

where $m_1(\mathbf{z}_k|r = r_i, \mathbf{z}_1, \cdots, \mathbf{z}_{t-1})$ can be estimated using Equation 10. The Bayes factor comparing model M_0 and model M_1 can then be updated at each time t.

4. Detection system with multiple panels

The aim of this paper is to propose a dynamic, real-time detection procedure. Therefore, computational time is one of the key system considerations, and observations at time t should be processed before time $t + 1$ observations arrive. Another control parameter is the threshold of triggering an alarm. In practice, long delay shortens the time for security personnel to locate and process the potential source. On the other hand, false alarms should also be minimized. Here, we introduce a fast detection system that continuously monitors the entire region and generates an alarm if the risk level, measured by the Bayes factor, exceeds a threshold. The detection criterion is built from the outputs of several parallel subsystems, called panels hereafter, to increase the reliability of the system.

Inevitably, there will be a slight delay between the introduction of the source to the surveillance region and the time of detection. Our experiences show that the delay is shorter when the source is already present in the region when our SMC detection algorithm starts. When running SMC under model M_1, assuming that there is one source in the region, but in fact there is none, particles may be attracted by a cluster of false-positive signals. Thus, it may take many time periods before some particles (samples of the source location) are generated around the newly arrived source because the motion model only allows the source to travel with a realistic speed and does not allow for jumps. To overcome this difficulty, we propose to start multiple panels at different times on a rolling basis. In this way, there is always a panel started close to the introduction of the source, resulting in shorter delays of detection. Each panel only runs for a finite period of time in order to maintain a fixed number of panels.

Information from these panels is fused together to make effective detections. Specifically, a detection decision is made based on the Bayes factors of all concurrent panels. We study the evolution of the Bayes factor on the rolling basis. Note that Bayes factor $BF_t(\mathbf{z}_1, \cdots, \mathbf{z}_t)$ in Section 3.2 compares models 0 and 1 based on $(\mathbf{z}_1, \cdots, \mathbf{z}_t)$ at time t. It measures the likelihood of the presence of a source during the entire period 1 to t. When a source is indeed present, BF_t increases with t. Although the value of BF_t is informative, its evolution is more revealing if the goal is to detect a sudden appearance of a

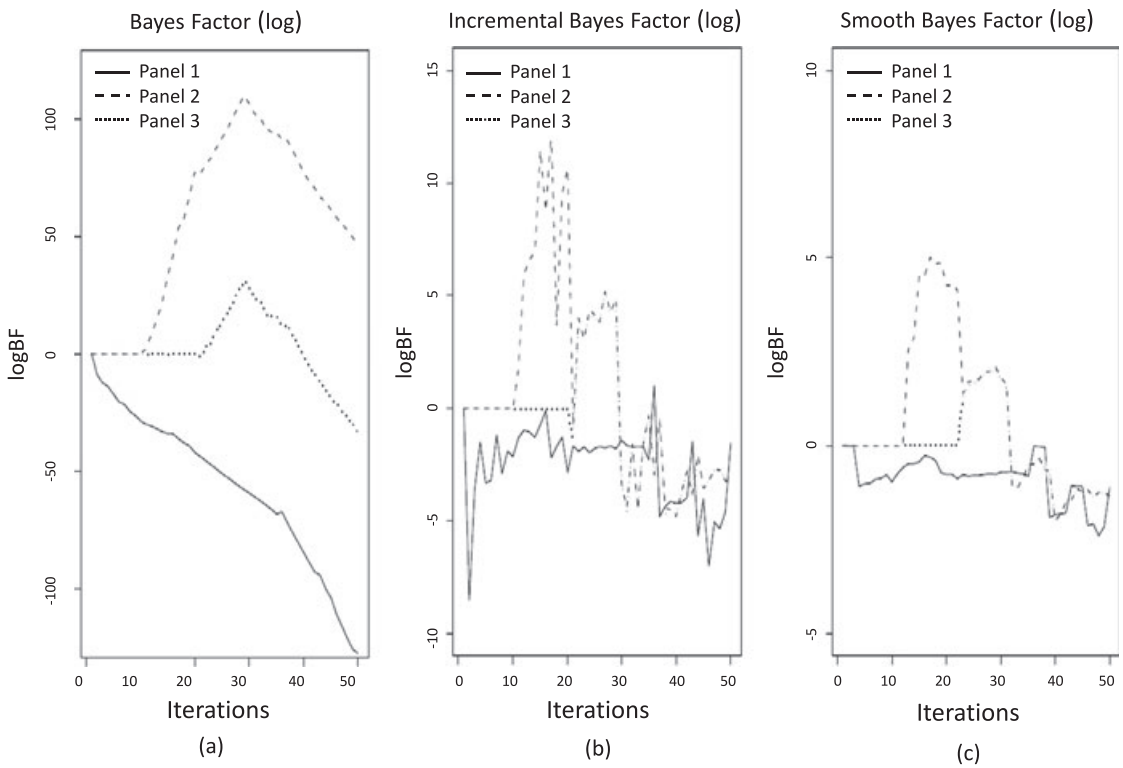

Figure 4. Evolution of log Bayes factor (a), incremental log Bayes factor (b), and $RIBF_t(5)$ (c), for three panels starting at time $t = 1, 11$, and 21. Sensitivity and specificity are equal to 0.95 and number of particles and taxicabs to 1500. Source was introduced at $t = 8$ and removed at $t = 25$.

source. Figure 4(a) shows the log BF_t of the panels starting at time $t = 1, 11$, and 21, respectively. The source is introduced at $t = 8$. Figure 4(b) shows the corresponding log IBF_t. Note that when a source is present, log IBF_t is positive. We define a *reversed-interval BF* as

$$RIBF_t(\ell) = \sum_{i=0}^{\ell} \log IBF_{t-i}.$$

It measures the cumulative strength of information of the presence of the source at time t. Detection can be made at time t if $\max_{L_1 \le \ell \le L_2} RIBF_t(l) \ge c$. This is a CUSUM type of measure, in reverse of time. L_1 is the minimum smoothing window to ensure a single outlier will not trigger an alarm, and L_2 is the maximum smoothing window to save computing and storage. In our simulation, we will fix $L_1 = L_2 = 5$ for simplicity.

Note that in Figure 4(a), log BF_t for panel 1 continues to decrease after the source is introduced. This is because, when there is no source in the region for a long time, the Monte Carlo samples trying to track the nonexisting source tend to concentrate in certain areas. Because of the constraint of source motion model, these samples have difficulty moving towards the source when the source is introduced, thus failing to produce accurate localization of the source. In a way, when the source is introduced, our 'prior' distribution of the source could be far away from the true location of the source. This is in fact exactly the reason we need multiple panels to properly initialize the filter algorithm at different times.

The use of RIBF also solves the problem of multiple panels starting at different times whose BFs are based on different sets of data. RIBF measures the 'instantaneous' information, although it is extracted from a different 'history' utilized in the panel. For multiple panels, we use the majority vote to fuse information and make detection decisions.

5. Simulation

This section demonstrates the effectiveness of the proposed approach using simulated data sets. The performance of the procedure is assessed under a number of settings with combinations of different values of several key elements including range, number of sensors, sensitivity and specificity, and the number of particles used in SMC.

5.1. Simulation design

We use a monitoring region of 25×25 city blocks similar in size to Manhattan in New York City. As there are approximately 13,000 taxicabs registered in New York City, representing over 50,000 drivers, here, we consider a network of 1500 active sensors on average, or about $1500/25^*25 = 2.4$ sensors per block. Of course, detection performance increases as the number of active sensors increases; this impact will be shown in the succeeding texts. We assume the locations are provided by GPS devices, and the binary signals by sensors similar to Thermo Fisher Scientific and ICX technologies products are sent to a control center every 30 s. From our discussion with field experts, we anticipate the actual sensitivity and specificity values close to 0.85 for an inexpensive sensor in practice.

The taxicab motion model used to generate the simulated data mimics the traffic pattern observed in big cities. Specifically, the speed varies among a set of values according to a Markov chain with possible transitions from the current speed to the two closest levels or remaining the same. Vehicles can only turn (left or right) if speed is low enough while U-turns are not allowed anywhere. To simulate/mimic the car movement in Manhattan, speed is on average greater on avenues (north/south direction) than on streets (east/west direction). This taxicab motion model is different from our source motion model as defined before. In our simulation study and to detect the source, we only need to know the taxicab positions and the signals they send but not the actual motion model that is used to generate the location data.

We assume the sensors are turned on only when the car is moving and assume the sensors are moving independently. Again, this is only used in simulating the data and is not used in the detection algorithm, because we do observe the location of sensors at every point. We consider driving, walking, or stationary sources that are always located on the map. If the source is being driven by a car, we use the taxicab motion model earlier to simulate the movement of the source as well. A walking source also uses the same model but with a different set of slower speeds.

For detection and localization, we run multiple panels in parallel. A new panel is started every 10 time periods (5 min), and an old panel is retired after it runs for 30 time periods (15 min). The panels overlap, and there are always three active panels at any moment of time. To initialize a new panel, the proposal distribution for $p(x_0)$ consists of a mixture of the current posterior distributions of the existing panels and the initial source distribution (Section 2.3). When at least two panels' $RIBF_t(5)$ exceeds the threshold, c, a detection is made. The source was placed randomly between $t = 20$ and 30 (10 and 15 min from the start) when three panels were already in place.

The performance of our detection algorithm is assessed through detection power and detection size. Detection power is defined as the proportion of experiments in which the proposed algorithm detected the presence of a nuclear source when the source was actually present in the study region. Detection size is defined as the proportion of experiments for which a source is detected when there is none (type I error). If the detection criterion is met, the location of the source is estimated.

Data were simulated under different settings with different values of sensor range, sensitivity, specificity, and the number of sensors (taxicabs). When not specified, the following default values are used: range $r = 1$ block, sensitivity and specificity $(\eta, \gamma) = (0.85,\ 0.85)$, and number of sensors $n = 1500$. For each experiment with different settings, we analyzed 100 simulated data sets, among which 50 has a static source and 50 had a moving source. In the detection algorithm, the number of particles $N = 1500$ are used throughout the simulation study. To simplify the presentation, in all experiments, we use equal sensitivity and specificity. In the early stage of our research, we experienced some settings of unequal sensitivity and specificity values. The results did not seem to alter the conclusion of the simulations presented here.

5.2. Results

We present here the results from four sets of experiments. The first is designed to determine the detection criterion. The rest are designed to assess the performance in terms of percentage of detection for various given network parameters, accuracy of the location estimation, and delay of detection.

Detection threshold: In order to determine the impact of the detection threshold, we ran simulations under different sensitivity and specificity $\eta = \gamma \in \{0.75, 0.80, 0.85, 0.90\}$, while all other parameters are set to the default values specified earlier. Figure 5 shows the detection power and size for various Bayes factor thresholds. When the threshold is set to 3, the power is very high while the size is less than 5% for all the cases, which is also equivalent to strong evidence on Jeffrey's scale as we discussed in Section 3.2. Hence, $c = 3$ is used as the threshold in further analysis.

The impact of the detection range, *r*, and the number of sensors: Table I shows the detection power for various combinations of the sensor range and the number of active sensors, with all other parameter set to their default values. The last row of the table shows the size when there is no source in the region.

The number of sensors is important because one would like to know how many cars should be equipped. As expected, detection percentage improves as the number of sensors increases as shown in Table I. For 1500 or more sensors, the detection percentage ranges from 76% to 100%. But for a lower number of sensors and a smaller range r, such as 800 and 0.5 block, the detection percentage is only 42%. As the detection range increases, detection percentage improves even with a low number of sensors in the area. To ensure that the detection system works well, we recommend using at least 1500 sensors for an area of that size or about $1500/25^2 = 2.4$ sensors per block.

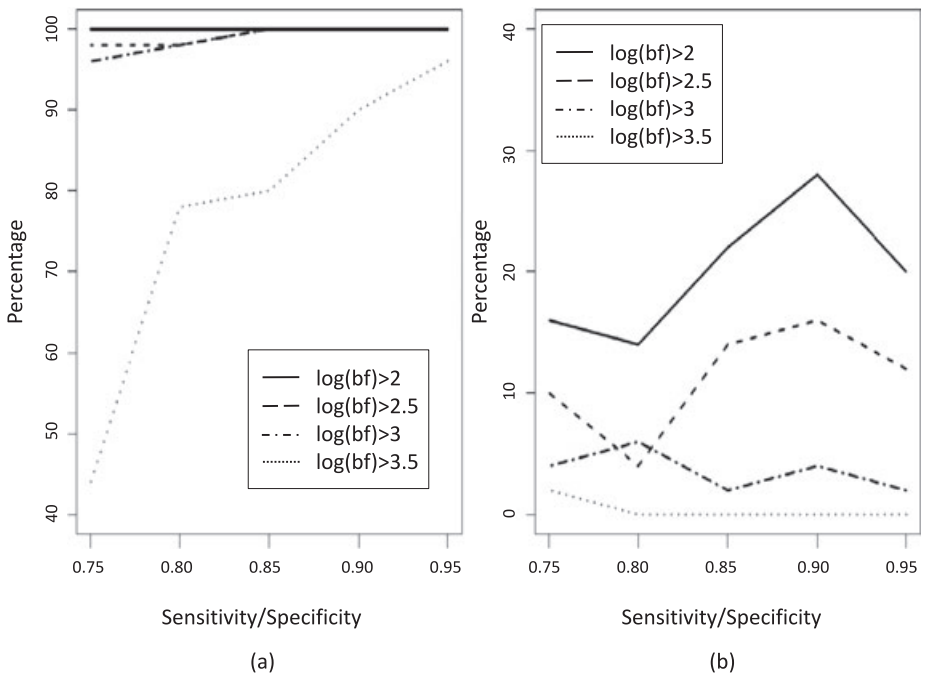

Figure 5. Detection power (a) and size (b) given equal sensitivity and specificity using a threshold equal to 2, 2.5, 3, and 3.5.

Table I. Detection power and size given the true range (row) and the number of sensors (column).

Range r (block)	Number of sensors 800	1000	1500	2000
0.50	0.42	0.56	0.76	0.88
0.75	0.58	0.66	0.96	1.00
1	0.78	0.88	1.00	1.00
1.25	1.00	1.00	1.00	1.00
1.50	1.00	1.00	1.00	1.00
0	0.04	0.02	0.04	0.02

A true range of 0 corresponds to data simulated under model M_0 and provides the size of the test. The number of particles is 1500, and the sensitivity/specificity is 0.85. Results are based on 50 simulations.

Table II. Detection power and size given true range (row) and sensitivity/specificity (column) when source is moving.

Range r (block)	Sensitivity and specificity (0.95, 0.95)	(0.90, 0.90)	(0.85, 0.85)	(0.80, 0.80)	(0.75, 0.75)
0.50	0.98	0.98	0.72	0.44	0.36
0.75	1.00	1.00	0.98	0.82	0.60
1	1.00	1.00	1.00	1.00	0.96
1.25	1.00	1.00	1.00	1.00	1.00
1.50	1.00	1.00	1.00	1.00	1.00
0	0.02	0.00	0.04	0.08	0.04

A true range of 0 corresponds to data simulated under model M_0 and provides the size of the test. The number of particles is 1500, and the number of sensors is 1500. Results are based on 50 simulations.

The impact of range and sensitivity and specificity: Table II shows the detection power for a moving source under different combinations of detection range and sensitivity and specificity. It is seen that when r is 1 block or higher or for high sensitivity and specificity (greater than 0.85), the detection power is very close to 100%. For small ranges, the power decreases as sensitivity and specificity decrease, respectively, from 100% to 60% and from 98% to 36% for ranges equal to 0.75 block and 0.5 block. The result for a stationary source is similar. Here, to simplify the simulation setting and to follow Cheng *et al*. [28], the values of sensitivity and specificity are set to be the same in each of the settings. A preliminary simulation study in our early exploratory stage (result not included in this paper) also suggested that the use of unequal sensitivity and specificity values did not seem to have much impact on the conclusions.

Estimation of the location: To measure the accuracy of location estimation, Euclidean distance between the estimated location and the true location is obtained. We refer this distance as an *estimation error*. If the estimation error is smaller than the range, the source is located within the area of positive detection, and the estimation is classified as accurate. Figure 6 shows the empirical distribution function of the estimation error with true sensor range r being 0.75 and 1.25, respectively.

Time till detection: The delay between the introduction of the source to the area and its detection is critical for the emergency response team. Figure 7 shows the cumulative proportion of detection over time for different ranges and specificity/sensitivity. For the 50 sets of simulations, the time of target introduction is randomly selected between 21 and 30, when there are three running panels, two of them started much earlier than the introduction of the target. It seems that the third panel, which started at $t = 20$, was not able to 'out-vote' the two other panels most of the time. Once a new panel was started at time $t = 30$, and with certain time delay, a firm detection was made, mostly around time $t = 37$. Our detection criterion is based on $RIBF_t(5)$; hence, in theory, the detection is more likely to be made between times 26 and 35. Because we assume data are collected every 30 s, the delay is, on average, 6 min. The delay can be shortened with more panels and hence a higher frequency of starting times. The results agree with the theory in that we can observe the first alarms around time 30. There is then a big jump at time 37. At this moment, the fourth panel, initiated after the source, has been present for six time periods.

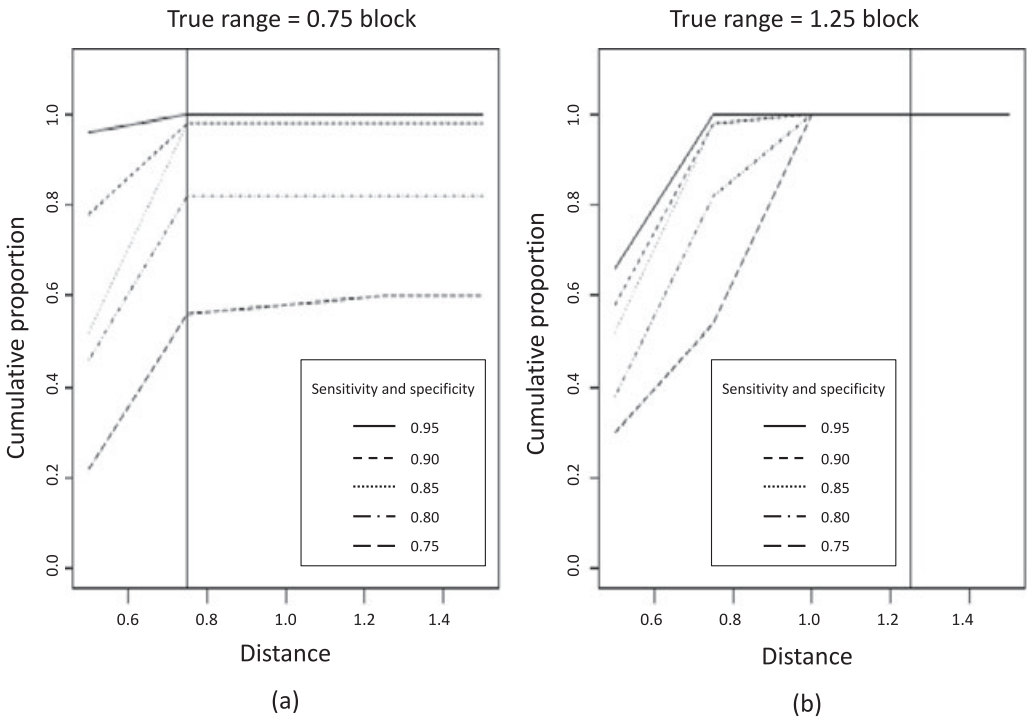

Figure 6. CDF of estimation error for two range r values.

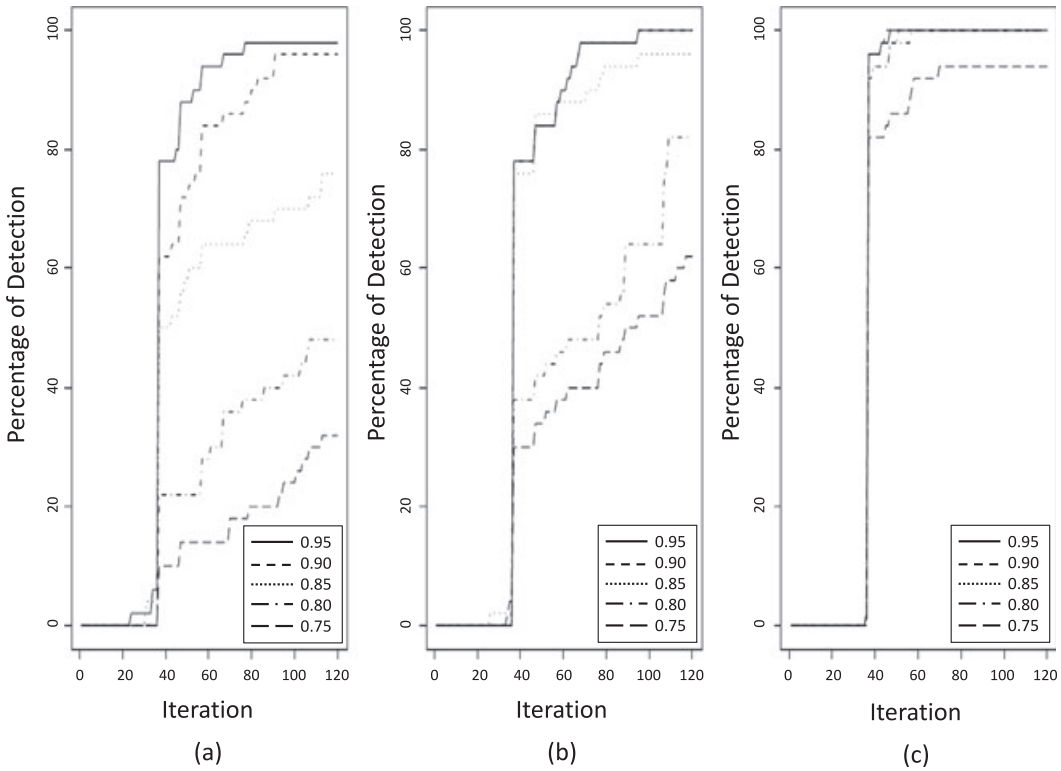

Figure 7. Cumulative proportion of detection over time. Each graph corresponds to different range value (0.5, 0.75, and 1 block, respectively) and each curve to a different sensitivity/specificity value. Source was introduced between $t = 21$ and 30.

Table III. Detection power given true range (row) and fixed source range used in the model (column).

True range (block)	Range (block) 0.50	0.75	1	1.25	1.50
0.50	1.00	0.72	0.32	0.16	0.08
0.75	0.82	1.00	1.00	0.96	0.88
1	0.78	0.94	1.00	1.00	1.00
1.25	0.66	0.94	1.00	1.00	1.00
1.50	0.72	0.76	0.96	1.00	1.00

Sensitivity and specificity are set to (0.85, 0.85), and the number of sensors n and particles N are both 1500. Results are based on 50 simulations.

Impact of misspecification of the range: We assess here the detection power for different true source ranges under a fixed range model as discussed in Section 3.2. The true range is chosen from a set $\{0.5, 0.75, 1, 1.25, 1.5\}$ blocks, and the detection power is evaluated for each of these values. It allows us to assess the effect of using a range value that is different from the true one. Other parameters are set at the default values. Table III shows its impact on the detection power and size.

It is seen that when the true range is used (diagonal), perfect detection is achieved under this particular setting. The proportions stay high (on average around 80%) when the assumed range r is different but close to the true value. The first row of the table, corresponding to a true range equal to 0.5 block, shows that detection power is sensitive to the range assumption when the true range is small. Performance is usually better when the assumed ranged is larger than the true one.

Impact of the number of particles and computational time: The number of particles N used in SMC controls the estimation accuracy and computational time. Similar to most of Monte Carlo methods, the standard error of SMC estimators decreases at the rate of $N^{-1/2}$, where N is the Monte Carlo sample size, and the computational time increases linearly with N, which is mostly confirmed by our simulation study. The results shown earlier are mostly based on $N = 1500$, which takes about 10 min of computational time to analyze and monitor signals in a period of 30 min in real time, on a standard laptop with a relatively straightforward implementation in MATLAB. The process averages about 10 s CPU time to update each time step (30 s in real time). Hence, real-time monitoring is achievable on a standard computer with reasonable performance. A larger number of panels and a larger N would result in higher performance but may require a more optimized implementation on a fast computer for online monitoring.

6. Discussion

This paper proposes a new fusion learning strategy aimed at detecting the presence of nuclear material in big cities relying on data provided by a network of mobile sensors in real time. We work within a Bayesian framework, and our estimation procedure relies on the SMC algorithm.

Our approach differs from all the previous ones in that it is based on the data from different time periods. It is consequently difficult to compare the results. Cheng *et al.* [28] proposed the closest approach to ours in that the collected data are the same but the analysis does not use a dynamic setup. The detection analysis (called testing) has been tested for high values of sensitivity and specificity such as $(0.95, 0.95)$ and $(0.98, 0.98)$. From the results presented in Section 5, our detection percentage is close to 100% in all these cases except the one with the smallest range (half a block) where the source can be detected in 98% of the cases with the same sensitivity and specificity. Thus, our method achieves better level of detection power for high values of sensitivity and specificity. Lower values of sensitivity and specificity (as low as $(0.75, 0.75)$) are more realistic as inexpensive sensors are less accurate. As mentioned before, in practice, sensitivity and specificity values are close to 0.85 for a reasonably priced sensor. Although the simulations provide several insights to assess the validity and the performance of the new method, it is desirable to test the practical viability of the proposed methods using experimental data. In addition, our limited simulation study did not fully explore the impact of interactions among all the control parameters on the detection power and size. Given a specific budget, and a cost function of a detection error (cost of a false alarm and cost of time delay of detection), it is possible to utilize computer experimental design tools to carefully study the interactions and find an optimal setting that minimized the expected loss. This is an interesting problem, and further research is needed.

It is possible that terrorists may use multiple bombs within a limited area. To extend this methodology to a multiple target detection problem, one can easily build a state–space model with a fixed number of nonoverlapping sources and determine the number of sources as a model selection problem in the same fashion as the detection in Section 3.2. Although the state–space model can be easily extended, filtering and tracking multiple sources presents a special challenge in that the trackers often coalesce when the sources get too close to each other in their path. More sophisticated sampling methods, along with many more panels and adaptive restarting scheme, are needed. It is expected that the computational burden would be greatly increased. Finally, we also anticipate future advances in sensor development eventually leading to sensors with accurate quantitative values that provide more information about both the level of the signal and its associated uncertainty. This additional information should be able to be incorporated in the Bayesian framework and SMC estimation procedure with some modifications in future model development.

Acknowledgements

The authors acknowledge grant support from the NSF (NSF-DMS1513483, DMS1209085, and DMS1107012) and DHS (2008-DN-077-ARI012). They wish to thank Drs. Fred Roberts and Jerry Cheng and several anonymous officials for their discussions and feedback on the research. They also wish to thank two anonymous referees and an AE for their detailed and insightful comments, which led to significant improvement of the paper, and also Suzanne Thornton for her assistance and careful proofreading of the entire manuscript.

References

1. Wrobel M. Sigma, 2014. https://www.darpa.mil/program/sigma. [Accessed July 2017].
2. NUSTL. National Urban Security Technology Laboratory, 2015 annual report, 2015. https://www.dhs.gov/sites/default/files/publications/NUSTL_2015_Annual_Report_FINAL_508.PDF. Department of Homeland Defense. [Accessed July 2017].
3. Vilim R, Klann R. RadTrac: a system for detecting, localizing, and tracking radioactive sources in real time. *Nuclear Technology* 2009; **168**: 61–73.
4. Brennan SM, Mielke AM, Torney DC. Radioactive source detection by sensor networks. *IEEE Transactions on Nuclear Science* 2005; **52**: 813–819.
5. Qian L, Fuller J, Chang I. Quickest detection of nuclear radiation using a sensor network. *2012 IEEE Conference on Technologies for Homeland Security (HST)*, Waltham, MA, USA, 2012;648–653.
6. Nemzek RJ, Dreicer JS, Torney DC, Warnock TT. Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Nuclear Science* 2004; **51**:1693–1700.
7. Coulon R, Kondrasovs V, Boudergui K, Normand S. Moving sources detection system. *2013 3rd International Conference on Advancements in Nuclear Instrumentation, Measurement Methods and their Applications (ANIMMA)*, Marseille, France, 2013;1–4.
8. Coulon R, Kondrasovs V, Boudergui K, Normand S. Moving sources detection algorithm for radiation portal monitors used in a linear network. *IEEE Transactions on Nuclear Science* 2014; **61**:2189–2194.
9. Boudergui K, Kondrasovs V, Coulon R, Corre G, Normand S. New monitoring system to detect a radioactive material in motion. *Proc. ANIMMA, 2013*, Marsaille, France, 2013,1–5.
10. Carpenter T, Cheng J, Roberts F, Xie M. Sensor management problems of nuclear detection. In *Safety and Risk Modeling and its Applications*, Pham H (ed.), Springer Series in Reliability Engineering. Springer London, 2011; 299–323.
11. Hochbaum D, Fishbain B. Nuclear threat detection with mobile distributed sensor networks. *Annals of Operations Research* 2011; **187**:45–63.
12. Borozdin KN, Klimenko AV, Priedhorsky WC, Hengartner N, Alexander CC, Cortez RA, Tanner HG, Papageorgiou X. Optimized strategies for smart nuclear search. *2006 IEEE Nuclear Science Symposium Conference Record*, vol. 2, San Diego, CA, USA, 2006;926–928.
13. Klimenko AV, Priedhorsky WC, Tanner H, Borozdin KN, Hengartner N. Intelligent sensor management in nuclear searches and radiological surveys. *Transactions of the American Nuclear Society* 2006; **95**:21–22.
14. DARPA. Sigmaradioactive threat detection system completes emergency vehicle test deployment in nation's capital, 2017. https://www.darpa.mil/news-events/2017-03-01. [Accessed July 2017].
15. Akyildiz I, Su W, Sankarasubramaniam YE, Cayiric E. A survey on sensor networks. *IEEE Communications Magazine* 2002; **40**:102–114.
16. Kulldorff M, Nagarwalla N. Spatial disease clusters: detection and infection. *Statistics in Medicine* 1995; **14**:799–810.
17. Diggle P, Rowlingson B, Sun TL. Point processes methodology for on-line spatiotemporal disease surveillance. *Environmetrics* 2005; **16**: 423–434.
18. Glaz J, Balakrishnan N. *Scan Statistics and Applications*. Birkhauser: Boston, 1999.
19. Glaz J, Naus J, Wallenstein S. *Scan Statistics and Applications*. Springer: New York, 2001.
20. Fu JC, Lou WY. *Distribution Theory of Runs and Patterns and its Applications: A Finite Markov Chain Embedding Approach*. World Scientific: Singapore, 2003.
21. Naus J, Wallenstein S. Multiple window and cluster size scan procedures. *Methodology and Computing in Applied Probability* 2004; **6**:389–400.
22. Lawson A. Markov chain Monte Carlo methods for putative pollution source problems. *Environmental Epidemiology* 1995; **14**:2473–2486.
23. Ghosh M, Natarajan K, Waller LA, Kim D. Hierarchical Bayes GLMs for the analysis of spatial data: an application to disease mapping. *Journal of Statistical Planning and Inference* 1999; **75**:305–318.
24. Gangnon R, Clayton M. Bayesian detection and modeling of spatial disease maps. *Biometrics* 2000; **56**:922–935.
25. Denison D, Holmes C. Bayesian partitioning for estimating disease risk. *Biometrics* 2001; **57**:143–149.
26. Gangnon R, Clayton M. A hierarchical model for spatially clustered disease rates. *Biometrics* 2003; **22**:3213–3228.

27. Xie M, Sun Q, Naus J. A latent model to detect multiple temporal clusters. *Biometrics* 2009; **65**:1011–1020.
28. Cheng J, Xie M, Chen R, Roberts F. A latent source model to detect multiple spatial clusters with application in a mobile sensor network for surveillance of nuclear materials. *Journal of the American Statistical Association* 2013; **108**:902–913.
29. Varshney PK. *Distributed Detection and Data Fusion*. Springer Science & Business Media: New York, 2012.
30. Liu J, Chen R. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association* 1998; **93**:1032–1044.
31. Doucet A, Godsill S, Andrieu C. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing* 2000; **10**: 197–208.
32. Del Moral P, Doucet A, Jasra A. Sequential Monte Carlo samplers. *Journal of Royal Statistical Society, Series B* 2006; **68**:411–436.
33. Wang X, Chen R, Liu J. Monte Carlo Bayesian signal processing for wireless communications. *Journal of Signal processing* 2002; **30**:89–105.
34. Rao NS, Sen S, Prins NJ, Cooper DA, Ledoux RJ, Costales JB, Kamieniecki K, Korbly SE, Thompson JK, Batcheler J, Brooks RR, Wu CQ. Network algorithms for detection of radiation sources. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 2015; **784**:326–331. Symposium on Radiation Measurements and Applications 2014 (SORMA XV).
35. Sun J, Tanner HG, Poulakakis I. Active sensor networks for nuclear detection. *2015 IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, 2015;3549–3554.
36. Gordon N, Salmond D, Smith A. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEEE Proceedings Radar and Signal Processing* 1993; **140**:107–113.
37. Avitzour D. Stochastic simulation Bayesian approach to multitarget tracking. *IEE Proceedings - Radar, Sonar and Navigation*, vol. 142, 1995, 41–44.
38. Mc Ginnity S, Irwin G. Multiple model bootstrap filter for maneuvering target tracking. *IEEE Transactions on Aerospace and Electronic Systems* 2000; **36**:1006–1012.
39. Salmond D, Gordon N. Particles and mixtures for tracking and guidance. In *Sequential Monte Carlo Methods in Practice*: Springer-Verlag, New York, 2001; 517–532.
40. Hue C, Le Cadre J, Pérez P. Sequential Monte Carlo methods for multiple target tracking and data fusion. *IEEE Transactions on Signal Processing* 2002; **50**:309–325.
41. Arulampalam M, Maskell S, Gordon N, Clapp T. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* 2002; **50**:174–188.
42. Orton M, Fitzgerald W. A Bayesian approach to tracking multiple targets using sensor arrays and particle filters. *IEEE Transactions on Signal Processing* 2002; **50**:216–223.
43. Gustafsson F, Gunnarsson F, Bergman N, Forssell U, Jansson J, Karlsson R, Nordlund P. Particle filters for positioning, navigation, and tracking. *IEEE Transactions on Signal Processing* 2002; **50**:425–437.
44. Wein L, Wilkins A, Baveja M, Flynn S. Preventing the importation of illicit nuclear materials in shipping containers. *Risk Analysis* 2006; **26**:1377–1393.
45. Venere E, Gardner E. Domestic nuclear threat security initiative, 2009.
46. Boros E, Fedzhora L, Kantor PB, Saeger K, Stroud P. A large-scale linear programming model for finding optimal container inspection strategies. *Naval Research Logistics* 2009; **56**:404–420.
47. Shentu Y, Xie M. A note on dichotomization of continuous response variable in the presence of contamination and model misspecification. *Statistics in Medicine* 2010; **49**:2200–2214.
48. Doucet A, de Freitas JFG, Gordon NJ. *Sequential Monte Carlo Methods in Practice*. Springer-Verlag: New York, 2001.
49. Chopin N. A sequential particle filter method for static models. *Biometrica* 2002; **89**:539–552.
50. Chopin N. Central limit theorem for sequential Monte Carlo methods and its applications to Bayesian inference. *The Annals of Statistics* 2004; **32**:2385–2411.
51. Andrieu C, Doucet A, Holenstein R. Particle Markov chain Monte Carlo methods. *Journal of Royal Statistic Society, Series B* 2010; **72**:269–342.
52. Poyiadjis G, Doucet A, Singh SS. Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika* 2011; **98**(1):65–80.
53. Kass RE, Raftery AE. Bayes factor and model uncertainty. *Journal of American Statistic Association* 1995; **90**:773–795.
54. Jeffreys H. *Theory of Probability*. Oxford Classic Texts in the Physical Sciences: New York, 1939.
55. Efron B, Gous A. Scales of evidence for model selection: Fisher versus Jeffreys. In *Model selection*, Lahiri P. (ed.), Lecture Notes–Monograph Series, vol. 38. Institute of Mathematical Statistics: Beachwood, OH, 2001; 208–246.
56. Meng XL, Wong WH. Simulating ratios of normalizing constants via a simple identity: a theorical exploration. *Statistica Sinica* 1996; **6**: 831–860.
57. Carlin BP, Chib S. Bayesian model choice via Markov chain Monte Carlo methods. *Journal of Royal Statistic Society, Series B* 1995; **57**: 473–484.
58. Han C, Carlin B. MCMC methods for computing Bayes factors: a comparative review. *Journal of Royal Statistic Society, Series B* 2001; **96**: 1122–1132.
59. Green PJ. Reversible jump MCMC algorithm computation and Bayesian model determination. *Biometrika* 1995; **82**(4):711–732.