


**Volume 47 Issue 10
(October 2008)**
[< Previous](#)
[Next >](#)


[Current Issue](#)
[Available Issues](#)
[Early Online Releases](#)
[Author Index](#)

Share this Article

[Share](#) |

Journal Information

Online ISSN: 1558-8432
 Print ISSN: 1558-8424
 Frequency: Monthly

[Special Collections](#)

[Staff and Editors](#)

[Instructions to Authors](#)

[Manuscript Submission](#)

[How to Subscribe](#)

[RSS \(What is this?\)](#)

Previously titled:

Journal of Applied Meteorology, 1988-2005

Journal of Climate and Applied Meteorology, 1983-1987

Journal of Applied Meteorology, 1962-1982

Table of Contents

- [Abstract](#)
- [1. Introduction](#)
- [2. Methodology](#)
- [3. Algeciras accidental release](#)
- [4. Results](#)
- [5. Conclusions](#)
- [Acknowledgments](#)
- [REFERENCES](#)

Related Articles

Articles Citing this Article

[Google Scholar](#)

Search for Other Articles

[< Previous Article](#)

Volume 47, Issue 10 (October 2008)

[Next Article >](#)
[Add to Favorites](#)
[Email](#)
[Download to Citation Manager](#)
[Track Citations](#)
[Glossary](#)
[Permissions](#)
[Abstract](#)
[PDF](#)
[Figures](#)

Delle Monache, Luca, and Coauthors, 2008: Bayesian Inference and Markov Chain Monte Carlo Sampling to Reconstruct a Contaminant Source on a Continental Scale. *J. Appl. Meteor. Climatol.*, 47, 2600–2613.
doi: <http://dx.doi.org/10.1175/2008JAMC1766.1>

Bayesian Inference and Markov Chain Monte Carlo Sampling to Reconstruct a Contaminant Source on a Continental Scale

Luca Delle Monache , Julie K. Lundquist , Branko Kosović , Gardar Johannesson , Kathleen M. Dyer , and Roger D. Aines

Lawrence Livermore National Laboratory, Livermore, California

Fotini K. Chow

University of California, Berkeley, Berkeley, California

Rich D. Belles , William G. Hanley , Shawn C. Larsen , Gwen A. Loosmore , John J. Nitao , Gayle A. Sugiyama , and Philip J. Vogt

Lawrence Livermore National Laboratory, Livermore, California

Abstract

A methodology combining Bayesian inference with Markov chain Monte Carlo (MCMC) sampling is applied to a real accidental radioactive release that occurred on a continental scale at the end of May 1998 near Algeciras, Spain. The source parameters (i.e., source location and strength) are reconstructed from a limited set of measurements of the release. Annealing and adaptive procedures are implemented to ensure a robust and effective parameter-space exploration. The simulation setup is similar to an emergency response scenario, with the simplifying assumptions that the source geometry and release time are known. The Bayesian stochastic algorithm provides likely source locations within 100 km from the true source, after exploring a domain covering an area of approximately 1800 km × 3600 km. The source strength is reconstructed with a distribution of values of the same order of magnitude as the upper end of the range reported by the Spanish Nuclear Security Agency. By running the Bayesian MCMC algorithm on a large parallel cluster the inversion results could be obtained in few hours as required for emergency response to continental-scale releases. With additional testing and refinement of the methodology (e.g., tests that also include the source geometry and release time among the unknown source parameters), as well as with the continuous and rapid growth of computational power, the approach can potentially be used for real-world emergency response in the near future.

Keywords: [Bayesian methods](#), [Pollution](#), [Dispersion](#)

Received: April 20, 2007; Final Form: January 24, 2008

Corresponding author address: Luca Delle Monache, Lawrence Livermore National Laboratory, 7000 East Avenue, L-103, Livermore, CA 94550. Email: ldm@llnl.gov

1. Introduction

Knowledge of the temporal and spatial evolution of a contaminant released into the atmosphere either accidentally or deliberately is fundamental to adopting efficient strategies to protect the public health and to mitigate the harmful effects of the dispersed material. In emergency response situations the source parameters may not be known. Typically a source is assumed, and assessment of the trajectory, spread, and ultimate fate of a contaminant plume is based on predictions from atmospheric dispersion models. The accuracy of these predictions is affected by uncertainties in several components of the plume prediction, including the atmospheric dispersion models themselves, the meteorological models used to drive the dispersion models, the atmospheric data assimilated by the meteorological models, and uncertainties in the parameters describing the contaminant source.

By keyword

- Bayesian methods
 Pollution
 Dispersion

By author

- Luca Delle Monache
 Julie K. Lundquist
 Branko Kosović
 Gardar Johannesson
 Kathleen M. Dyer
 Roger D. Aines
 Fotini K. Chow
 Rich D. Belles
 William G. Hanley
 Shawn C. Larsen
 Gwen A. Loosmore
 John J. Nitao
 Gayle A. Sugiyama
 Philip J. Vogt

Search in:

- AMS
 Google Scholar

Among these sources of uncertainty, those composing the initial state of the contaminant are often the most significant, and as such provide the central focus of this study. A methodology to solve the “inverse problem” is proposed to reconstruct unknown source parameters given a set of downwind measurements at a time after the release.

A comprehensive literature review of past works on solutions of the inverse problem for atmospheric contaminant releases can be found in [Keats et al. \(2007a\)](#). [Pudykiewicz \(1998\)](#) implemented an algorithm based on integrating the adjoint of a linear dispersion model backward in time to solve a reconstruction problem. [Johannesson et al. \(2004, 2005\)](#) introduced dynamic Bayesian modeling, and the Markov chain Monte Carlo (MCMC; [Gilks et al. 1996](#); [Gelman et al. 2003](#)) and the sequential Monte Carlo (SMC) ([Doucet et al. 2001](#)) sampling approaches to reconstruct a contaminant source for synthetic data. The same approach has also been applied to design sensor networks ([Lundquist et al. 2005](#)) and for the reconstruction of atmospheric releases employing the Urban Dispersion Model (UDM; [Neuman et al. 2006](#)). [Bocquet \(2005a, b\)](#) presented an approach based on the principle of maximum entropy on the mean. The approach was tested with synthetic and field data illustrating applicability to real-world source reconstruction problems. Similar skill has been demonstrated by the method described in [Allen et al. \(2007a, b\)](#), which is based on genetic algorithms and the optimization of variables using principles from genetics and evolution. [Deguillaume et al. \(2007\)](#) applied a forward version of the regional chemistry-transport model “CHIMERE” with a Bayesian Monte Carlo analysis for the purpose of inverse emissions modeling. Source reconstruction at urban scales using building-resolving models has been reported by [Chow et al. \(2008\)](#) and [Keats et al. \(2007a\)](#). [Keats et al. \(2007a\)](#) used an adjoint representation of the source-receptor relationship [with a linear dispersion model as in [Pudykiewicz \(1998\)](#)], but they used a Bayesian inference methodology in conjunction with MCMC sampling procedures for the source reconstruction in an urban environment, and they validated their approach using data from water channel simulations and a field experiment (Joint Urban 2003) in Oklahoma City, Oklahoma [[Yee \(2007\)](#) used a similar approach in the context of the European Tracer Experiment (ETEX) campaign, i.e., at the continental scale]. [Chow et al. \(2008\)](#) used an approach based on a set of a computational fluid dynamics (CFD) code forward runs along with Bayesian inference and MCMC sampling to reconstruct the source parameters of release experiments held during the Joint Urban 2003 field campaign. The Project Prairie Grass data were used for source inversion at local scales by [Keats et al. \(2007b\)](#) in the context of a nonconservative scalar.

Herein we present an inversion algorithm based on Bayesian inference combined with a MCMC procedure. The present algorithm is an improved version of the code previously applied to other dispersion cases and to design sensor networks, including canonical test cases, moving sources, and multiple sources, using Lagrangian particle models, Gaussian puff models, urban puff models, and CFD modes ([Johannesson et al. 2004, 2005](#); [Lundquist et al. 2005](#); [Neuman et al. 2006](#); [Chow et al. 2008](#)). To take advantage of the generality and flexibility of Bayesian inference, an unmodified forward model is used in our algorithm. Such an approach results in an algorithm that is adaptable and allows for simultaneous application of different dispersion models, including those for which it may not be possible to develop an adjoint model [e.g., those including significant empiricism; see [Errico \(1997\)](#)]. Integrating different forward models into the framework is a straightforward task, which makes the computational tool highly flexible, because predictions from different models may be needed for different dispersion scenarios.

The algorithm is applied here to a continental-scale accidental release of radioactive material from near Algeciras, Spain, during May 1998. This event affected much of continental Europe for several days, including locations a few thousands of kilometers downwind. The forward dispersion simulation is conducted in this study using the Lagrangian Operational Dispersion Integrator (LODI) model ([Ermak and Nasstrom 2000](#); [Nasstrom et al. 2000](#)), which was developed for operational emergency response dispersion predictions at the National Atmospheric Release Advisory Center (NARAC) at Lawrence Livermore National Laboratory (LLNL). The methodology provides a skillful, robust statistical characterization of the reconstructed source parameters (source location and strength) in the presence of a complex atmospheric flow field using only crude measurements. Future tests will also invert for the source release time and duration to replicate more closely the real-time emergency response scenarios where this information is also unknown.

The following section provides a detailed description of the reconstruction algorithm. [Section 3](#) discusses the Algeciras release incident, the reconstruction results of which are detailed in [section 4](#). [Section 5](#) presents the conclusions and follows with a discussion of computational issues and additional testing needed before using the methodology as an emergency response tool.

2. Methodology

The stochastic event reconstruction algorithm is based on Bayes’s theorem and an MCMC procedure to sample the unknown parameter space ([Gilks et al. 1996](#); [Gelman et al. 2003](#)). The following sections briefly summarize the theory on which the methodology is constructed and outline the main steps of the procedure. Further details can be found in [Johannesson et al. \(2004\)](#) and [Chow et al. \(2008\)](#).

a. Theoretical framework

Bayes’s theorem, as applied to a source reconstruction problem, can be stated as follows:

$$p(\mathbf{S}|\mathbf{M}) = \frac{p(\mathbf{M}|\mathbf{S})p(\mathbf{S})}{p(\mathbf{M})}. \quad (1)$$

Here, $p()$ is a probability distribution, $\mathbf{S} = (x, y, q)$ is the state vector formed by the point source parameters (x and y are the source horizontal coordinates and q is the emission rate), and \mathbf{M} is a vector formed by the measurements. Bayes’s theorem relates the posterior distribution $p(\mathbf{S}|\mathbf{M})$ to the product of the probability of the measurements given the source parameters $p(\mathbf{M}|\mathbf{S})$, also called the likelihood function, and the probability of the source parameters prior to any knowledge of the measurements $p(\mathbf{S})$, also called the prior. Here x , y , and q are assumed to be the unknown parameters, but in general Bayes’s theorem can be applied with \mathbf{S} including other parameters also, for example, z (the vertical coordinate), or the release time and duration.

Bayes’s theorem is often expressed alternatively as

$$p(\mathbf{S}|\mathbf{M}) \propto p(\mathbf{M}|\mathbf{S})p(\mathbf{S}). \quad (2)$$

This form avoids evaluation of the marginal distribution of \mathbf{M} , $p(\mathbf{M}) = \int p(\mathbf{S})p(\mathbf{M}|\mathbf{S}) d\mathbf{S}$, for which analytical solutions are rare and numerical computation is expensive.

In this study Bayes's theorem is applied to describe the conditional probability $p(\mathbf{S}|\mathbf{M})$ of a source described by x , y , and q , given the observed sensor measurements. To estimate the unknown source parameters, that is, to reconstruct the source using (2), the posterior distribution must be sampled. Sampling the posterior distribution [left-hand side of (2)] involves computing the probability distribution $p(\mathbf{M}|\mathbf{S})$ for any proposed \mathbf{S} realization. Here, $p(\mathbf{M}|\mathbf{S})$ quantifies the likelihood of a set of measurements \mathbf{M} given the source parameters \mathbf{S} . This likelihood is computed by running a forward dispersion model with the given source parameters \mathbf{S} and comparing the model-predicted concentrations with the sensor measurements [section 2b(1)]. The closer the prediction is to the measurements, the higher the likelihood of the source parameters.

b. The algorithm

The Lagrangian particle dispersion model LODI (Ermak and Nasstrom 2000; Nasstrom et al. 2000) is driven by meteorological data that can be obtained from a variety of sources, including real-time observations, atmospheric forecast models, atmospheric analysis fields, or any combination of the above. The posterior distribution is sampled with an MCMC procedure via a Metropolis–Hasting algorithm (Gilks et al. 1996; Gelman et al. 2003).

The main steps of the algorithm are shown in Fig. 1 and can be described as follows. Before executing the reconstruction procedure, the source term parameters x , y , and q are assigned prior distributions [$p(\mathbf{S})$] based on the limited information available about the release. Each parameter's prior distribution is bounded by a range specified from prior knowledge of the circumstances relevant to the problem being solved. The width of the prior distribution reflects the confidence in the initial estimate of the source parameters. Following setup, an initial value for each parameter is sampled from its prior and a number of iterations involving source-term sampling and dispersion simulation are executed until convergence to the posterior distribution is reached [see sections 2c(2) and 4c(2) for definition and a discussion on convergence criteria]. The accepted parameter values collected during all the iterations constitute a realization of a "Markov chain," which can be defined as a set of consecutive values of a parameter in which the probability of a value depends only upon the previous sampled value (Gilks et al. 1996). Once initiated, each iteration of the solution procedure consists of the following four steps:

- 1) A new value for the source term parameter x is proposed (x_{prop}) according to $x_{\text{prop}} = x + dx$. Here x is the current value and dx is the vector displacement from that value. The displacement is modeled as a random walk sampled from a Gaussian distribution with zero mean and a standard deviation specified from the current step size (discussed below). Hence, the source term parameter's (i.e., x) prior distribution $p(\mathbf{S})$ is used as a target distribution to estimate a "prior likelihood" of x_{prop} (see Tarantola 2005, appendix 6.11). If the prior likelihood of x_{prop} exceeds that of x , the proposed value x_{prop} replaces x . If not, a random (Bernoulli) "coin flip" [section 2b(1)] determines whether the new proposal, even with its lower prior likelihood, will be accepted.
- 2) Step 1 is repeated independently for y and q . Note that independently sampling each parameter helps ensure that they reflect their prior likelihoods better than sampling all of them at once and testing their likelihood together.
- 3) The forward dispersion simulation is conducted using the current values of the source term parameters x , y , and q .
- 4) The likelihood of the current values of x , y , and q , which forms the proposed state, is evaluated by comparing the agreement between the predicted concentration by running LODI using those source parameters, and the observed concentration at the sensor locations. This new likelihood is compared with that resulting from the previous forward LODI simulation. The proposed state likelihood $p(\mathbf{M}|\mathbf{S})$ should not be confused with the prior likelihood, as defined in step 1, that is based only on the source term parameters prior distributions, that is, $p(\mathbf{S})$, without considering the measurements (\mathbf{M}). If the proposed state likelihood is higher than the likelihood of the previous state, it is accepted. If not, then a random (Bernoulli) coin flip [section 2b(1)] is used to ensure that the search explores the entire posterior state space. Occasional acceptance of new proposals with lower likelihoods ensures that the reconstruction procedure continues to search the complete space of proposed source term parameter states, preventing the procedure from remaining indefinitely within the neighborhood of a local extreme.

1) The likelihood function and acceptance condition. The quality of agreement between the predicted and observed concentrations at the sensor locations is expressed in terms of a likelihood function (L). The present study utilizes a likelihood function of the form

$$\begin{aligned} \ln[p(\mathbf{M}|\mathbf{S})] &= \ln[L(\sigma, \mathbf{P}, \mathbf{M})] \\ &= -\frac{\sum_i^N [\log_{10}(P_i) - \log_{10}(M_i)]^2}{2\sigma^2}, \quad (3) \end{aligned}$$

where L is the likelihood function, P_i are the elements of the array \mathbf{P} of the predicted concentration values at the sensor locations, M_i are the elements of the array \mathbf{M} of the sensor measurements, σ is an error parameter chosen accordingly to expected errors in the observations and predictions, and N is the number of sensors. In this study the measured concentration values spread over several order of magnitudes. Taking the logarithm of both predicted and measured values prevents large concentrations (as for the Algeiras accident) from dominating in the likelihood computation. Both predictions and observations are set to a detection limit value if they fall below it, to prevent (3) from failing with 0 concentration values. The likelihood in (3) is relatively simple, where the N prediction errors are assumed independent and expected to be proportional in size to the predicted concentration (as a result of log transforming both the observed and predicted concentration).

After LODI is run with the new proposed state (i.e., a new set of the source term parameters x , y , and q), the proposed state is retained if

$$\ln(L_{\text{prop}}) \geq \ln(L) \quad \text{or} \quad \ln(L_{\text{prop}}) - \ln(L) \geq \ln\{\text{rnd}(0, 1)\},$$

(a)
(b)

(4)

where L_{prop} is the likelihood value of the proposed state, L is the previous likelihood value, and $\text{rnd}(0, 1]$ is a random number generated from a uniform distribution in the interval $(0, 1]$.

Condition (4a) is always satisfied when condition (4b) is satisfied, and therefore in a code only condition (4b) is implemented. Here condition (4a) is included to illustrate the difference between cases when the proposed state has a higher likelihood of the previous state [(4a)] and those when it does not [(4b)]. It is important to note that condition (4b) is more likely to be satisfied if the likelihood of the proposal is only slightly lower than the previous likelihood value. This aspect has important implications for improving the Bayesian event reconstruction algorithm efficiency, as explained in the next subsection. If large errors are expected in the prediction and/or measurements, large values of σ should be assumed, resulting in a broader–flatter posterior distribution [as evident from (3)]. In this study, an initial analysis was carried out for different values of σ . The results presented here are derived for σ fixed at 0.8. This value represents the lowest value that adequately represents the uncertainties in model predictions and measurements as reflected in the posterior distribution of source term parameters.

We note that the σ could be treated as an unknown parameter (in the same way as the source parameters), assigned a prior distribution, and sampled along with the source parameters. The extension of the current framework to include an adaptive LODI prediction error model (a likelihood) with unknown (stochastic) parameters that are sampled along with the source parameters is straightforward, but outside of the scope of this paper.

2) Cooling-off and adaptive-step procedures. A cooling-off procedure (i.e., annealing) is used to ensure a broader exploratory sampling of the posterior distribution values at the initial stages of the procedure (Neal 2001). The procedure consists of normalizing the log likelihood of the proposed values with a “temperature” constant that decreases (“cools off”) linearly with the number of iterations to unity. Initially, the normalization of the likelihood values increases the number of times condition (4b) is met, because the normalization reduces the relative differences between different likelihood values, allowing the acceptance of less likely proposals. As the temperature is decreased progressively, the acceptance rate resulting from condition (4b) decreases accordingly. This procedure allows the Markov chain to initially include a broader subset of the parameter range as accepted values. This study employed a cooling-off period consisting of the first 100 iterations of each Markov chain. For this case the 100 cooling-off iterations were sufficient to allow the chains to sample the regions of interest, as shown in section 4b. A different number may be needed for a different event.

An adaptive step-size algorithm has also been implemented to ensure that the magnitude of the displacement between current and proposed parameters is appropriate for the current stage of the search. Following Haario et al. (1999), after an initial transient during which the step size is held constant, the step size is computed as being directly proportional to the variance of the values sampled up to the current iteration. This approach guarantees a large step size in the initial stages of the search because of the large variance of the sampled values, encouraging a broad exploration of parameter space and identifying high likelihood regions more efficiently than a smaller step size would. During the latter stages of a search, after the Markov chains have converged to within small neighborhoods of the posterior extremes, the sampled variance is smaller, and the resulting smaller step size encourages a chain to explore the contours of the posterior within the neighborhood of its current extreme. It should be noted that the step-size adjustments affect only the rates of convergence and do not affect the posterior distribution (i.e., the shape of the returned sample). For the application presented here, the adaptive step-size algorithm resulted in a MCMC acceptance rate of approximately 20%.

It should be emphasized that the purpose of the algorithm is not to identify posterior extremes, but to reveal the probabilistic landscape of the posterior throughout the entire domain. The combination of the cooling-off procedure and the use of an adaptive step size results in chains that explore the entire parameter space yet repeatedly converge to the same subset, yielding probabilistic estimates of source parameter values.

c. Burn-in and convergence definitions

In this section two important aspects of the Bayesian MCMC stochastic algorithm are defined, the burn-in and the convergence criteria.

1) Burn in. The burn in is an important phase represented by an initial subset of the total iterations that can be defined as follows. It can be seen as the number of iterations needed for the current parameter distribution to relax from the initial state (Gilks et al. 1996). Burn-in samples are usually discarded to construct the parameters’ posterior distributions. Burn in is further discussed in section 4c(1).

2) Convergence. Convergence is reached in practice when more samples would not modify the resulting posterior distribution. Statistically, convergence to the posterior distribution can be estimated by computing between-chain and within-chain variance (Gelman et al. 2003). If there are m Markov chains of length n for a source parameter S , whose values are denoted by s , then the between-chain variance B can be computed as

$$B = \frac{n}{m-1} \sum_{j=1}^m (\bar{S}_j - \bar{S})^2, \quad (5)$$

where

$$\bar{S}_j = \frac{1}{n} \sum_{i=1}^n s_{ij} \quad (6)$$

and

$$\bar{S} = \frac{1}{m} \sum_{j=1}^m \bar{S}_j; \quad (7)$$

within-chain variance W is computed as

$$W = \frac{1}{m} \sum_{j=1}^m w_j^2, \quad (8)$$

where

$$w_j^2 = \frac{1}{n-1} \sum_{i=1}^n (s_{ij} - \bar{S}_i)^2. \quad (9)$$

The convergence value (\hat{R}) is given by

$$\hat{R} = \frac{\text{var}(S)}{W}, \quad (10)$$

where the variance of S is defined to be

$$\text{var}(S) = \frac{n-1}{n} W + \frac{1}{n} B. \quad (11)$$

The necessary (but not sufficient) condition for statistical convergence to the posterior distribution is that \hat{R} approaches unity (Gelman et al. 2003). In practice, simulation should be run until \hat{R} is smaller than 1.1 or 1.2 (Gilks et al. 1996). Convergence results are discussed in section 4c(2).

3. Algeciras accidental release

The Algeciras accident remained undetected for 3 days after the release and the source was unknown for 10 more days afterward. On 30 May 1998, a piece of medical equipment containing Cesium-137 (Cs-137) was accidentally melted in one of the furnaces at the Acerinox stainless steel production plant near Algeciras. This accident has been the subject of several prior modeling studies (e.g., Yamazawa 1998; Vogt et al. 1999; Buckley 2000; Graziani et al. 2000; Baklanov and Sorensen 2001).

According to the Spanish Nuclear Security Agency [Consejo de Seguridad Nuclear (CSN); see Vogt et al. 1999], the radioactivity of the material released was between 8 and 80 Ci (1 Ci = 3.7×10^{10} Bq) and the release occurred between 0100 and 0300 UTC 30 May. None of the detected radioactivity levels were sufficient to adversely affect human health or the environment. This event provides a real case to evaluate the source reconstruction methodology described in section 2 on a continental scale.

A detailed analysis of the synoptic weather patterns at the time of the release can be found in Vogt et al. (1999). Figure 2 qualitatively shows the radioactive plume transport as reproduced by LODI given realistic source parameters (36°10'N, 5°26'W location; 0130 UTC 30 May release time; 30-min release duration; and 50-Ci release amount). In each panel the source location is indicated by the white square; white circles represent locations of sensors used in this study that are located in northwest Italy and southeast France, approximately 1600 km northeast of the source location. Initially (top-left panel) the plume was caught in a westerly flow that advected the radioactive cloud over the Mediterranean Sea, leaving the accident undetected. Between 31 May and 1 June (top-right panel) the plume shifted to the north and was detected by the set of sensors used here. After 3 June, radioactivity was detected at a number of European stations; concentrations were reported from 24-h to 1-week averages (Vogt et al. 1999).

4. Results

This section describes the details of the problem setup and assumptions made to perform the simulations, followed by analysis of the results. The goal is to demonstrate the ability of the proposed method to reconstruct the source parameters for a contaminant release from available measurements in conditions similar (but with few assumptions, as described below) to a real emergency response scenario occurring at the continental scale.

a. Simulation setup and assumptions

The first measurements of contamination from the Algeciras release (24-h average concentrations of radioactive material) were obtained on 2 June, 3 days after the accident. From 3 to 14 June a variety of measurements (from 24-h averages to 1-week averages) were reported in several locations farther downwind, mostly in central and eastern Europe. For the present study, only the first set of available 24-h average measurements (17 values reported from 11 stations 2–3 June) is used to reconstruct the unknown source. These data are chosen to evaluate the ability of the proposed methodology to function in a situation similar to an emergency response, using only the first available data.

The sensors used for the source reconstruction simulation cover a very small portion of the European continent (Fig. 3). The compact area covered by the measurements presents a challenge to the reconstruction algorithm, because source parameters that are displaced relatively small distances from the true source may result in a predicted plume that does not encounter the sensors at all. Adding to this difficulty is the absence of “0” concentration readings. Concentrations near the threshold detection levels of the instruments were simply not reported, rather than being reported as 0 readings. Zero concentration readings provide useful information to the reconstruction algorithm; this information is lost when the locations fail to report those values.

The source geometry and duration were assumed known and were specified as a surface point source, which was active between 0130 and 0200 UTC 30 May 1998, as in Vogt et al. (1999). CSN reported the release to have happened between 0100 and 0300 UTC 30 May.

The meteorological data used to drive the dispersion model LODI was provided by the National Centers for Environmental Protection (NCEP) Aviation Model (AVN; Kanamitsu et al. 1991). The 6-h AVN analysis fields at 1°

horizontal spatial resolution were used as input for the Atmospheric Data Assimilation Parameterization Techniques (ADAPT) model (Sugiyama and Chan 1998), which generated 6-hourly spaced meteorological fields at 22-km horizontal resolution, which were input to LODI. Qualitatively this meteorological field was in fair agreement with the observed flow, particularly over the Mediterranean Sea, where the plume traveled before hitting the sensors. The dispersion simulation was conducted using 100 000 marker particles. To run LODI for this case it takes less than 2 min with about thirty 2.4-GHz CPU processors.

The prior distribution of the source location is indicated by the dashed box in Fig. 3. This domain (about 1800 km in the east–west direction and 3600 km in the north–south direction) was selected based on the predominantly westerly flow pattern over most of the sensor locations from 30 May (the day of the accident) through 2 June (when the first measurements were available), which implied that the source must have been located to the west of the sensors (western Europe or northwestern Africa). Further, given the wind speeds during the period, locations within 1800 km to the west would likely have passed beyond the sensors by the beginning of 2 June. The meridional extent of the prior distribution reflects similar confidence in potential values of the source latitude. For the first iteration, each location within this box was given approximately the same probability to be sampled, with the x and y values generated from two clipped Gaussian distributions with mean values equal to the box center coordinates and standard deviations equal to the box dimension in the east–west direction (i.e., approximately flat distributions substantially equivalent to a uniform distribution). The prior distribution for q was set similarly, that is, by giving at each value within the interval $[10^{13}, 10^{17}]$ approximately the same probability to be sampled.

b. Source reconstruction

The stochastic engine was run with three independent Markov chains. For a given number of sampled values, a higher number of chains would produce the same number of samples more rapidly, with the availability of a large computing cluster to run the independent chains all at the same time. Additional comments on the computational costs of the procedure can be found at the end of section 5.

Figure 3 depicts the locations explored by the three Markov chains. Given the assumption of a surface point source, the chains were allowed to explore only values at the surface. The gray dots represent accepted states during the burn-in phase [first 500 iterations, as defined in section 2c(1) and discussed in section 4c(1)], whereas the black dots represent those states accepted afterward. The chains effectively explore the given sampling box, constructing a posterior distribution for the sampled parameters that is most likely to include all of the dominant modes.

Figure 4 shows an expanded view of the region, including the locations of the states accepted after burn in. Different symbols (circle, five-point star, and triangle) represent the three chains. The posterior distribution reveals two distinct modes of high relative probability indicating likely source locations: one over land within 60 km north of the real source location (white square), and one over the Mediterranean Sea about 80 km east and 20 km south of Algeciras. The Markov chains sample both within and between the two modes providing evidence of chain convergence [as defined in section 2c(2) and discussed in section 4c(2)].

The probabilistic aspect of the answer provides a useful tool for a real emergency response scenario. The algorithm finds among all the possible solutions the few ones that are most consistent with the data available and its uncertainties. The information in Fig. 4 provides guidance for decision makers formulating an appropriate strategy. It would have suggested looking for potential source locations (e.g., steel mill) just north of Algeciras, and this would have led to rapid identification of the real source. Moreover, the mode downwind of Algeciras would have prompted a search for possible ships or any other floating or submerged bodies releasing radioactive material in that localized stretch of waters. The posterior distribution in Fig. 4 could also be used to construct an updated prior for a more detailed search. Repeated application of such an approach could yield successive refinements of the source location posterior distribution.

Figure 5 shows contours of location probabilities based on the accepted state's spatial distribution showed in Fig. 4. Also shown in Fig. 5 are the surface wind fields driving LODI at the time of the release. The probability distributions clearly show the location of the two modes of high relative probability and likewise indicate how the distributions are influenced by changes in the wind field. For instance, the distributions are stretched in approximately the along-wind direction, and the amount of the stretching is roughly proportional to the wind speed.

One explanation for this feature is the following. Measurements that are averaged over a given time interval are relatively insensitive to the plume's exact arrival and transit time, provided that the majority of the plume passes over the sensors within the given averaging interval. Because the flow pattern formed a nearly stationary arc between Algeciras and the sensors, potential sources spread over as much as 100 km along the streamline represented by the arc would have produced very similar 24-h-averaged predictions at the sensor locations. In fact, the along-wind length of the distribution shown in Figs. 4–5 is roughly the distance a parcel would have traversed given the model wind speed in 24 h. Similar features in the posterior distributions for source location have been observed in studies using comparable inversion algorithms (although applied at much smaller spatial scales) where analogous explanations hold (Lundquist et al. 2005; Chow et al. 2008).

More insights about the methodology performance can be inferred from Figs. 6–8. Figure 6 shows the surface location (x and y) and the release rate (q) values versus the iteration number. The range adopted in each panel for the vertical axes reflects the actual range spanned by the sampler. In the top and central panel the horizontal gray line represents the true value of the source location coordinates, whereas the two horizontal gray lines in the bottom panel show the range of the true source emission, as reported by CSN. In each panel, there are three black lines representing the values that the Markov chains assume during the iterations.

The first 100 iterations show the effect of the cooling-off procedure. Initially the chains span a wide portion of the given range for each parameter, but the acceptance rate decreases as the temperature linearly decreases with the iteration number. After 100 iterations, the chain values span only a limited subset of the possible range constrained by the likelihood of the proposed values as explained in section 2b(2).

The effects of the adaptive step procedure are less apparent than the cooling off, but they can be noticed in the top and bottom panel of Fig. 6. After the cooling-off period (i.e., after iteration 100), the values of x or q have a tendency to change slowly. This causes the variance of the iteration series to rapidly diminish, which in turn results in a decrease of

the step size for subsequent sampling [section 2b(2)]. With a smaller step size, the Bayesian event reconstruction algorithm proposes states with higher likelihood that are more frequently accepted (approximately between iterations 200 and 250) than for the previous iterations (between 100 and 200). Without the implementation of the adaptive step procedure, the proposed states' rejection rate would have been inefficiently high, requiring a much higher number of iterations to collect the same number of samples than when the adaptive step procedure is adopted.

As the Markov chains in Fig. 6 converge, the sampled parameter assumes the same values throughout consecutive iterations (horizontal black lines) with new proposed states or the individual proposed values (for x , y , and q) repeatedly rejected. Most of the values for x (Fig. 6 top panel) tend to be greater than the x coordinate of the real source, toward downwind locations. After about 500 iterations the chains start to explore a small subset of the initial range, corresponding to the values with the highest likelihood for this parameter. Using accepted values of x , a histogram can be constructed as shown in the top panel of Fig. 7 (vertical gray line represents the true source x -coordinate value). The majority of the values are within 100 km from the real value, with the dominant mode about 80 km downwind of it.

The central panel of Fig. 7 shows the quality of the prediction of the y values' posterior distribution (vertical gray line represents the true source y coordinate). Two modes can be recognized, just a few tens of kilometers north and south of the true source location, the latter being the dominant one. The y values converge faster to a much smaller subset of the initial range (Fig. 6, central panel). A higher rejection rate seems to apply to y as shown by more frequent intervals where y candidates are repeated, as compared to x . This behavior is due to the clustering and close spacing of the sensor locations; even a small (a few tens of kilometers) deviation from the real y value toward the south or toward the north leads to a plume that misses the sensor locations, which in turn results in a low likelihood value for the proposed state. Finally, this behavior indicates the quality of the meteorological field used to drive LODI. A less accurate meteorological field could lead to a plume hitting the sensor locations even if released from a source not close to Algeciras, resulting in an inaccurate source location reconstruction.

The bottom panel of Fig. 6 shows the q values of the accepted states. The ordinate is represented with a logarithmic scale spanning the wide range of values sampled for q (from 10^{13} to 10^{17} $\mu\text{Bq s}^{-1}$). There is a large uncertainty as to the real amount of the radioactive material released, with values ranging from 8 to 80 Ci, as shown by the two gray lines in the panel. From Fig. 6 it can be noted that q is the slowest parameter to reach a phase where the accepted states span only a subset of the provided range. The q values also have the highest variability reflecting higher uncertainty than the determination of the location.

There is a tendency to overpredict the 8–80-Ci range. The majority of the sampled values after the burn in [defined in section 2c(1) and discussed in section 4c(1)] are of the same order of magnitude as the upper limit of 80 Ci. Because the observed concentrations are 24-h averages, it is extremely challenging to obtain a tight posterior distribution and a limited number of modes for q , because these observed values could result from a variety of different choices of x , y , and q . These uncertainties also affect the rapidity with which convergence is reached, with q being the slowest converging parameter, as discussed in the next section.

c. Burn in and convergence

In this section two important aspects of the stochastic procedures are briefly discussed, the burn-in phase and the convergence criteria.

1) Burn in. Burn in [as defined in section 2c(1)] has been applied to the plots shown in Figs. 4 and 7. The burn-in phase was chosen as the first 500 iterations from a visual inspection of the parameter-sampled values variation with the number of iterations (Fig. 6) and the convergence criteria ($\hat{R} < 1.5$ in Fig. 8 for all parameters). Tests with burn in ranging from the first 300 to the first 1500 iterations lead to posterior distributions close to the one depicted in Fig. 7 (not shown).

2) Convergence. Figure 8 shows the convergence values \hat{R} [section 2c(2); see (10)] for the source parameter x , y , and q versus the number of iterations. After 2000 iterations the necessary statistical convergence condition of $\hat{R} < 1.2$ is met for all the parameters, with y being the fastest parameter reaching this condition and q being the slowest.

Note that after the first 100 iterations (i.e., during the cooling-off procedure), all three parameters appear to meet the convergence condition. Nevertheless, by visual examination of the parameters values at different stages of the iterative procedure (Figs. 3 and 6) it is clear that convergence has not been reached after 100 iterations. Indeed, expert judgment is also needed to assess the convergence of the simulation (Gilks et al. 1996). Practically, convergence is reached when it is clear that more samples would not modify the resulting posterior distribution.

5. Conclusions

A methodology to reconstruct a source given a set of measurements has been presented. The method combines Bayesian inference with Markov chain Monte Carlo sampling and produces posterior probability distributions of the parameters describing the unknown source. The methodology has been applied for the first time to a real accidental radioactive release at the continental scale that occurred in May 1998, near Algeciras, Spain.

The sampled parameters are the source location and strength. The source duration has been assumed to be 0130–0200 UTC 28 May, which falls within the time interval 0100–0300 UTC reported by the Spanish Nuclear Security Agency (Vogt et al. 1999). The release was also assumed to be a surface point source.

The source location is reconstructed as a roughly bimodal distribution, with modes located a few tens of kilometers north of Algeciras and within 100 km downwind of the real source location. The source strength is represented by a wider posterior distribution (reflecting the higher uncertainty of this parameter with respect to the source location), with a tendency of the Bayesian MCMC algorithm to overpredict the reported source strength. The majority of the sampled values of this parameter have the same order of magnitude of the estimated release. The poor time resolution of concentration measurements (24-h averages) prevents the stochastic procedure from obtaining a tight posterior distribution for the source strength.

The probabilistic aspect of the solution optimally combines a likely answer with the uncertainties of the available data. From several possible solutions, the Bayesian event reconstruction algorithm is able to find only the few that are more

consistent with the data available and its uncertainties. The source reconstruction performed in this study would have provided a decision maker with accurate information about the accident a few hours after the first sensor measurements were reported, assuming the availability of a powerful computing platform, as explained below. This would have led to timely and efficient actions to preserve public health, in case radioactive material had been released at harmful concentrations. Moreover, the results presented show that the methodology has skill despite the limited number of observations available.

To demonstrate the effectiveness of the methodology presented the stochastic engine was run with three Markov chains. Each Markov chain can be run independently at the same time. Moreover, each chain can be run in parallel on multiple processors. By using 30 chains the same results could be obtained in less than four wall-clock hours. This could be accomplished by independently running the 30 chains on about 900 processors, each having a 2.4-GHz CPU, providing decision makers with crucial information within a time frame appropriate for emergency response applications at the continental scale spanning several days. With additional testing and refinement of the methodology (e.g., tests that also include the source geometry and release time among the unknown source parameters), as well as with the continuous and rapid growth of computational power, in the foreseeable future the approach presented in this paper will soon be practical for real-world emergency response applications.

The uncertainties in the meteorology directly affect the quality of the source reconstruction. Both the mean and turbulent components of the wind strongly affect the performance of the dispersion model and the quality of the predicted concentrations. These uncertainties could be taken into account by combining Bayesian inference and MCMC sampling with ensemble techniques, an approach widely used in the weather forecast community (e.g., [Palmer and Hagedorn 2006](#)). This will be considered in future applications. Finally, future work will also focus on algorithm optimizations to improve efficiency and reduce the overall computational cost.

Acknowledgments

Thanks are given to Jeffrey Mirocha (LLNL) for carefully reviewing the manuscript and to Philip Cameron-Smith, Keith Grant, and Steve Homann (LLNL) for useful discussions that improved the quality of the presented material. The authors are also thankful to Art Mirin (LLNL) for providing software engineering improvements to the stochastic code, to Ernest Arnold, Robert Shectman, and Hoyt Walker (LLNL) for helping with the simulations setup, and to Matthew Simpson (LLNL) for his help with graphics. Thanks are also given to two anonymous reviewers for their valuable comments and suggestions.

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

REFERENCES

- Allen, C. T., G. S. Young, and S. E. Haupt, 2007a: Improving pollutant source characterization by better estimating wind direction with a genetic algorithm. *Atmos. Environ.*, **41**, 2283–2289. [[CrossRef](#)] 
- Allen, C. T., S. E. Haupt, and G. S. Young, 2007b: Source characterization with a genetic algorithm–coupled dispersion–backward model incorporating SCIPUFF. *J. Appl. Meteor. Climatol.*, **46**, 273–287. [[Abstract](#)] 
- Baklanov, A., and J. H. Sorensen, 2001: Parameterisation of radionuclide deposition in atmospheric long-range transport modelling. *Phys. Chem. Earth B Hydrol. Oceans Atmos.*, **26**, 787–799. [[CrossRef](#)] 
- Bocquet, M., 2005a: Reconstruction of an atmospheric tracer source using the principles of maximum entropy. I. Theory. *Quart. J. Roy. Meteor. Soc.*, **131**, 2191–2208. [[CrossRef](#)] 
- Bocquet, M., 2005b: Reconstruction of an atmospheric tracer source using the principles of maximum entropy. II. Applications. *Quart. J. Roy. Meteor. Soc.*, **131**, 2209–2223. [[CrossRef](#)] 
- Buckley, R. L., 2000: Modeling atmospheric deposition from a Cesium release in Spain using a stochastic transport model. Preprints, *11th Joint Conf. on the Applications of Air Pollution Meteorology with the Air and Waste Management Association*, Long Beach, CA, Amer. Meteor. Soc., 8B.3. 
- Chow, T., B. Kosovic, and S. Chan, 2008: Source inversion for contaminant plume dispersion in urban environments using building-resolving simulations. *J. Appl. Meteor. Climatol.*, **47**, 1553–1572. [[Abstract](#)] 
- Deguillaume, L., M. Beekmann, and L. Menut, 2007: Bayesian Monte Carlo analysis applied to regional-scale inverse emission modeling for reactive trace gases. *J. Geophys. Res.*, **112**.D02307, doi:10.1029/2006JD007518. 
- Doucet, A., N. de Freitas, and N. Gordon, 2001: *Sequential Monte Carlo Methods in Practice*. Springer, 581 pp. 
- Ermak, D., and J. Nasstrom, 2000: A Lagrangian stochastic diffusion method for inhomogeneous turbulence. *Atmos. Environ.*, **34**, 1059–1068. [[CrossRef](#)] 
- Errico, R. M., 1997: What is an adjoint model? *Bull. Amer. Meteor. Soc.*, **78**, 2577–2591. [[Abstract](#)] 
- Gelman, A., J. Carlin, H. Stern, and D. Rubin, 2003: *Bayesian Data Analysis*. Chapman & Hall/CRC, 668 pp. 
- Gilks, W., S. Richardson, and D. Spiegelhalter, 1996: *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC, 486 pp. 
- Graziani, G., S. Galmarini, and T. Mikkelsen, 2000: RTMOD: Real-Time Model evaluation. Risø National Laboratory/Ispra Joint Research Centre Joint Tech. Rep. Risø-R-1174(EN)/JRC-Ispra Rep. TN.1.00.11, 47 pp. [Available online at <http://www.risoe.dk/rispubl/VEA/ris-r-1174.htm>]. 

- Haario, H., E. Saksman, and J. Tamminen, 1999: Adaptive proposal distribution for random walk Metropolis algorithm. *Comput. Stat.*, **14**, 375–395. [[CrossRef](#)] 
- Johannesson, G., W. Hanley, and J. Nitao, 2004: Dynamic Bayesian models via Monte Carlo—An introduction with examples. Lawrence Livermore National Laboratory Tech. Rep. UCRL-TR-207173, 53 pp. [Available online at <http://www.llnl.gov/tid/lof/documents/pdf/312314.pdf>]. 
- Johannesson, G., and Coauthors, 2005: Sequential Monte-Carlo based framework for dynamic data-driven event reconstruction for atmospheric release. *Proc. of the Joint Statistical Meeting*, Minneapolis, MN, American Statistical Association and Cosponsors, 73–80. 
- Kanamitsu, M., and Coauthors, 1991: Recent changes implemented into the global forecast system at NMC. *Wea. Forecasting*, **6**, 425–435. [[Abstract](#)] 
- Keats, A., E. Yee, and F-S. Lien, 2007a: Bayesian inference for source determination with applications to a complex urban environment. *Atmos. Environ.*, **41**, 465–479. [[CrossRef](#)] 
- Keats, A., E. Yee, and F-S. Lien, 2007b: Efficiently characterizing the origin and decay rate of a non-conservative scalar using probability theory. *Ecol. Model.*, **205**, 437–452. [[CrossRef](#)] 
- Lundquist, J., B. Kosovic, and R. Belles, 2005: Synthetic event reconstruction experiments for defining sensors network characteristics. Lawrence Livermore National Laboratory Tech. Rep. UCRL-TR-217762, 34 pp. [Available online at <http://www.llnl.gov/tid/lof/documents/pdf/328798.pdf>]. 
- Nasstrom, J. S., G. Sugiyama, J. Leone, and D. L. Ermak, 2000: A real-time atmospheric dispersion modeling system. Preprints, *11th Joint Conf. on the Applications of Air Pollution Meteorology with the Air and Waste Management Association*, Long Beach, CA, Amer. Meteor. Soc., 5.1. 
- Neal, R., 2001: Annealed importance sampling. *Stat. Comput.*, **11**, 125–139. [[CrossRef](#)] 
- Neuman, S., L. Glascoe, B. Kosović, K. Dyer, W. Hanley, and J. Nitao, 2006: Event reconstruction for atmospheric releases employing urban puff model UDM with stochastic inversion methodology. Preprints, *Sixth Symp. on the Urban Environment*, Atlanta, GA, Amer. Meteor. Soc., J4.6. [Available online at <http://ams.confex.com/ams/pdfpapers/105467.pdf>]. 
- Palmer, T., and R. Hagedorn, 2006: *Predictability of Weather and Climate*. Cambridge University Press, 718 pp. 
- Pudykiewicz, J. A., 1998: Application of adjoint tracer transport equations for evaluating source parameters. *Atmos. Environ.*, **32**, 3039–3050. [[CrossRef](#)] 
- Sugiyama, G., and S. Chan, 1998: A new meteorological data assimilation model for real-time emergency response. Preprints, *10th Joint Conf. on the Applications of Air Pollution Meteorology with the Air and Waste Management Association*, Phoenix, AZ, Amer. Meteor. Soc., 7A.1. 
- Tarantola, A., 2005: *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics, 342 pp. 
- Vogt, P., B. Pobanz, F. Aluzzi, R. Baskett, and T. Sullivan, 1999: ARAC simulation of the Algeciras, Spain steel mill CS-137 release. Lawrence Livermore National Laboratory Tech. Rep. UCRL-JC-131330, 51 pp. [Available online at <http://www.llnl.gov/tid/lof/documents/pdf/235247.pdf>]. 
- Yamazawa, H., 1998: Long-range dispersion analysis on accidental atmospheric release of Cesium-137 at Algeciras (in Japanese). *J. Nucl. Sci. Technol.*, **41**, 114–116. 
- Yee, E., 2007: Bayesian probabilistic approach for inverse source determination from limited and noisy chemical or biological sensor concentration measurements. *Chemical and Biological Sensing VIII*, A. W. Fountain III, Ed., International Society for Optical Engineering (SPIE Proceedings, Vol. 6554), 65540W, doi:10.1117/12.721630. 



FIG. 1.
Flow diagram of the algorithm (described in [section 2b](#)).

[View larger version \(41K\)](#)

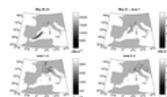


FIG. 2.
Qualitative illustration of the plume fate the 4 days following the release (30 May–3 Jun 1998), as simulated with the LODI model. The white square represents the real source location ($36^{\circ}10'N$, $5^{\circ}26'W$), whereas the white circles are the sensor locations. Contoured concentrations ($\mu\text{Bq m}^{-3}$) are 24-h averages. (Note: shading bars have different ranges)

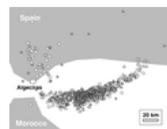
[View larger version \(51K\)](#)



[View larger version \(46K\)](#)

FIG. 3.

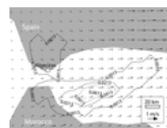
Proposed accepted state locations. The dashed box is the starting sampling area, provided as input (i.e., the location prior distribution), covering western Europe and northwestern Africa. The gray dots are the pre-burn-in states (iterations 1–500), and the black dots represent the post-burn-in states (iterations >500). White square and circles are as in [Fig. 2](#).



[View larger version \(38K\)](#)

FIG. 4.

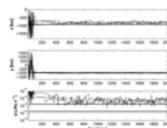
Zoom in of [Fig. 3](#) nearby the real-source location. Each point represents an accepted state post-burn in (iterations >500), where different symbols (circle, five-point star, and triangle) correspond to different chains.



[View larger version \(46K\)](#)

FIG. 5.

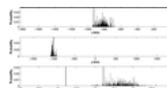
Location probability spatial distribution built with post-burn-in (iterations >500) accepted states. Arrows represent surface winds driving LODI at the time of the release.



[View larger version \(53K\)](#)

FIG. 6.

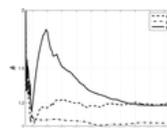
The three chain values vs the number of iterations, for (top) x (km), (middle) y (km), and (bottom) q ($\mu\text{Bq s}^{-1}$) (black lines). Values of q are reported on a logarithmic scale. Horizontal gray lines represent the true values of x and y . The two horizontal gray lines in the bottom panel are the estimated likely range of the source strength (according to the Spanish Nuclear Security Agency; see [Vogt et al. 1999](#)).



[View larger version \(26K\)](#)

FIG. 7.

Posterior distribution as inferred by the Bayesian event reconstruction algorithm for (top) x (km), (middle) y (km), and (bottom) $\log(q)$ [$\log(\mu\text{Bq s}^{-1})$]. Vertical gray lines represent the true values for x and y and the reported (according to the Spanish Nuclear Security Agency; see [Vogt et al. 1999](#)) range for q .



[View larger version \(39K\)](#)

FIG. 8.

Convergence values vs the number of iterations for x (dashed line), y (dot-dashed line), and q (solid line).

Cited by

Victor Winiarek, Marc Bocquet, Nora Duhanyan, Yelva Roustan, Olivier Saunier, Anne Mathieu. (2014) Estimation of the caesium-137 source term from the Fukushima Daiichi nuclear power plant using a consistent joint assimilation of air concentration and deposition observations. *Atmospheric Environment* **82**, 268-279.
Online publication date: 1-Jan-2014.

[CrossRef](#)

Václav Šmidl, Radek Hofman. (2013) Efficient Sequential Monte Carlo Sampling for Continuous Monitoring of a Radiation Situation. *Technometrics*, 131121101310000.
Online publication date: 21-Nov-2013.

[CrossRef](#)

Derek Wade, Inanc Senocak. (2013) Stochastic reconstruction of multiple source atmospheric contaminant dispersion events. *Atmospheric Environment* **74**, 45-51.
Online publication date: 1-Aug-2013.

[CrossRef](#)

Nan Wu, Rui Yang, Hui Zhang, Lifeng Qiao. (2013) Decentralized Inverse Model for Estimating Building Fire Source Location and Intensity. *Journal of Thermophysics and Heat Transfer* **27**:3, 563-575.
Online publication date: 1-Jul-2013.

[CrossRef](#)

Paul E. Bieringer, Steven Hanna, George Young, Branko Kosovic, John Hannan, Ryohji Ohba. (2013) Methods For Estimating The Atmospheric Radiation Release From The Fukushima Dai-Ichi Nuclear Power Plant. *Bulletin of the American Meteorological Society* **94**:1, ES1-ES4.

Online publication date: 1-Jan-2013.
[Citation](#) . [Full Text](#) . [PDF \(1678 KB\)](#)

Mohammad Reza Koohkan, Marc Bocquet, Lin Wu, Monika Krysta. (2012) Potential of the International Monitoring System radionuclide network for inverse modelling. *Atmospheric Environment* **54**, 557-567.
Online publication date: 1-Jul-2012.

[CrossRef](#)

Victor Winiarek, Marc Bocquet, Olivier Saunier, Anne Mathieu. (2012) Estimation of errors in the inverse modeling of accidental release of atmospheric pollutant: Application to the reconstruction of the cesium-137 and iodine-131 source terms from the Fukushima Daiichi power plant. *Journal of Geophysical Research* **117**:D5.
Online publication date: 1-Jan-2012.

[CrossRef](#)

Andrew J. Annunzio, George S. Young, Sue Ellen Haupt. (2012) Utilizing state estimation to determine the source location for a contaminant. *Atmospheric Environment* **46**, 580-589.
Online publication date: 1-Jan-2012.

[CrossRef](#)

José Ruy Porto de Carvalho, Eduardo Delgado Assad, Hilton Silveira Pinto. (2011) Kalman filter and correction of the temperatures estimated by PRECIS model. *Atmospheric Research* **102**:1-2, 218-226.
Online publication date: 1-Oct-2011.

[CrossRef](#)

Guido Cervone, Pasquale Franzese. (2011) Non-Darwinian evolution for the source detection of atmospheric releases. *Atmospheric Environment* **45**:26, 4497-4506.
Online publication date: 1-Aug-2011.

[CrossRef](#)

Victor Winiarek, Julius Vira, Marc Bocquet, Mikhail Sofiev, Olivier Saunier. (2011) Towards the operational estimation of a radiological plume using data assimilation after a radiological accidental atmospheric release. *Atmospheric Environment* **45**:17, 2944-2955.

Online publication date: 1-Jun-2011.

[CrossRef](#)

Guido Cervone, Pasquale Franzese. (2010) Monte Carlo source detection of atmospheric emissions and error functions analysis. *Computers & Geosciences* **36**:7, 902-909.
Online publication date: 1-Jul-2010.

[CrossRef](#)

Guido Cervone, Pasquale Franzese, Allen P. K. Keesee. (2010) Algorithm quasi-optimal (AQ) learning. *Wiley Interdisciplinary Reviews: Computational Statistics* **2**:2, 218-236.
Online publication date: 1-Mar-2010.

[CrossRef](#)

P. Robins, V. E. Rapley, N. Green. (2009) Realtime sequential inference of static parameters with expensive likelihood calculations. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **58**:5, 641-662.
Online publication date: 1-Dec-2009.

[CrossRef](#)



© 2014 American Meteorological Society [Privacy Policy and Disclaimer](#)
Headquarters: 45 Beacon Street Boston, MA 02108-3693
DC Office: 1120 G Street, NW, Suite 800 Washington DC, 20005-3826
amsinfo@ametsoc.org Phone: 617-227-2425 Fax: 617-742-8718
[Allen Press, Inc.](#) assists in the online publication of AMS journals

Technology Partner -  Atypon Systems, Inc.