

Discussion

Bayesian source detection and parameter estimation of a plume model based on sensor network measurements[‡]

Chunfeng Huang^{1,*†}, Tailen Hsing², Noel Cressie³, Auroop R. Ganguly⁴,
Vladimir A. Protopopescu⁴ and Nageswara S. Rao⁴

¹*Department of Statistics, Indiana University, Bloomington, IN, U.S.A.*

²*Department of Statistics, University of Michigan, Ann Arbor, MI, U.S.A.*

³*Department of Statistics, The Ohio State University, Columbus, OH, U.S.A.*

⁴*Oak Ridge National Laboratory, Oak Ridge, TN, U.S.A.*

SUMMARY

We consider a network of sensors that measure the intensities of a complex plume composed of multiple absorption–diffusion source components. We address the problem of estimating the plume parameters, including the spatial and temporal source origins and the parameters of the diffusion model for each source, based on a sequence of sensor measurements. The approach not only leads to multiple-source detection, but also the characterization and prediction of the combined plume in space and time. The parameter estimation is formulated as a Bayesian inference problem, and the solution is obtained using a Markov chain Monte Carlo algorithm. The approach is applied to a simulation study, which shows that an accurate parameter estimation is achievable. Copyright © 2010 John Wiley & Sons, Ltd.

Received 15 February 2010; Revised 3 May 2010; Accepted 22 June 2010

KEY WORDS: Bayesian statistics; Markov chain Monte Carlo; partial differential equation; plume model; sensor networks

1. INTRODUCTION

Advances in monitoring and communication technologies have enabled environmental monitoring and security assessment using sensor networks, as evidenced by their deployments in a wide variety of contexts [1, 2]. While the specific configuration of any particular sensor network depends on the context of the problem, sensor networks typically generate potentially complex and unstructured

*Correspondence to: Chunfeng Huang, Department of Statistics, Indiana University, Bloomington, IN, U.S.A.

†E-mail: huang48@indiana.edu

‡A preliminary version of this work was presented at 9th ONR/GTRI Workshop on Target Tracking in Sensor Fusion, 2006.

data sets. These data sets are a crucial part of our understanding of the phenomena that the networks monitor. Hence, our ability to analyze and extract useful information determines the effectiveness of the sensor networks.

After the initial explosion/release at a certain time and location, the results of a dirty bomb, a chemical leak, or the release of biological agents take the form of an atmospheric plume that will undergo transport, advection, diffusion, adsorption, radioactive decay, and delayed re-emission. To detect, study, identify, and/or track the plume, we must study its spatio-temporal evolution. In this paper, we propose a Bayesian statistical approach that relies on both physical laws and statistical models in space and time. That is, they are physical statistical models in the sense of [3].

Spatially disperse networks of sensors hold an enormous potential for the detection, identification, tracking, and prediction of these point-source phenomena. In particular, the problem of detecting and tracking sources was addressed using sensor networks for chemical dispersions in [4–6], and nuclear radiation in [7]. However, there has been limited effort in incorporating the analytical and statistical nature of the plume and sensor models to analyze the measurements collected by such networks. In [8], a linear array of detectors was employed to detect a moving radiation source against the background radiation. A sensor-network solution for detecting a moving radiological dispersion device was presented using a Bayesian formulation in [7]. Cost–benefit analysis of using sensor networks for detecting the moving radiation sources was carried out in [9]. These contributions mainly focus on detecting low-level radiation emitted by sources prior to the explosion, and consequently they do not explicitly consider the dispersion dynamics. A recursive algorithm was presented in [10] to locate a single source and track the plume intensity. It is not clear to us how this method can be extended to a multiple-source situation.

This paper is structured as follows. Section 2 gives a technical formulation of the problem, introduces a PDE plume model, and discusses the Bayesian inferential approach for this model. Section 3 contains a simulation study, and the discussion and conclusions are given in Section 4.

2. PDE PLUME MODEL AND ESTIMATION METHOD

2.1. Formulation of the problem

The objective of this paper is to introduce a statistical approach to systematically analyze and estimate the parameters of a plume resulting from a chemical leak (say), based on a sequence of sensor-network measurements. The establishment of such networks is expected to accelerate, and the roles that they play in addressing significant issues will continue to expand in the near future. We consider the following scenario of K pollution sources released into the environment at different times t_{01}, \dots, t_{0K} and spatial locations s_{01}, \dots, s_{0K} , respectively, for some $K > 1$. For simplicity, we assume that the sensor locations are on the two-dimensional plane. Let

$$u_k(x, y, t) \equiv \text{true plume intensity at time } t \text{ and location } (x, y),$$

due to the k th source (1)

clearly $u_k(x, y, t) = 0$ if $t < t_{0k}$. In most situations, such as dirty bombs and chemical leaks, the plume intensities are usually low and sources are not close to each other. We assume that plume

intensities are additive, so that the true total plume intensity at location (x, y) and time t is,

$$\sum_{k=1}^K u_k(x, y, t).$$

Suppose that sensors are placed at locations $(x_i, y_i); i = 1, 2, \dots, n_s$, and the (total) plume intensity is observed at times $t_{i,j}; j = 1, 2, \dots, n_i$, for sensor i . The plume intensity is measured by the sensors; we assume that there is an additive random error $\varepsilon_{i,j}$ associated with each measurement, so that we observe

$$z_{i,j} = \sum_{k=1}^K u_k(x_i, y_i, t_{i,j}) + \varepsilon_{i,j}, \tag{2}$$

where the errors $\{\varepsilon_{i,j}\}$ may have a spatio-temporal dependence structure.

An example of u_k is the simple product-form model in [11, 12]

$$u_k(x, y, t) = \begin{cases} \beta e^{-[(x-x_{0k})^2+(y-y_{0k})^2]^{1/2}}, & t \in [t_{0k}, t_{0k} + T], \\ \beta e^{\alpha(t-t_{0k}-T)-[(x-x_{0k})^2+(y-y_{0k})^2]^{1/2}}, & t > t_{0k} + T, \end{cases}$$

where α, β, T are unknown parameters. Other plume models considered in the literature can be found in [4, 5, 13, 14]. The product-form model provides a convenient analytic formulation of the spatio-temporal evolution of the plume, but it does not account for propagation or drift due to exogenous factors such as wind. Indeed, within this model the consequence of the explosion is instantaneous in the whole space, followed by absorption after a delay T . On the other hand, one could attempt to simulate each particle and track them through space and time, with the sensor measurements providing parameter estimates on the simulation. An intermediary step between the extremely simple product model given above and the realistic, but often cumbersome numerical simulations of particles, is given by analytic PDE solutions, adapted to a statistical-modeling framework. These provide an excellent basis for evaluating statistical methods, carrying out parameter identification, and refining sensor-fusion techniques.

The main purpose of this paper is to consider the statistical inference of such a PDE plume model based on the spatio-temporal data $\{z_{i,j}\}$ described in (2). Owing to the large number of unknown parameters, a Bayesian perspective will be adopted and the posterior distribution will be obtained from Markov chain Monte Carlo simulations. As mentioned in Section 1, the Bayesian approach has been used previously in detecting moving radiation sources (characterized by simple diffusions) using sensor networks; see [7, 8, 15].

2.2. A PDE plume model

The most important features of the physical phenomenon described in Section 1 are well captured by an absorption-drift-diffusion model. When the space is two dimensional, a homogeneous version of the model is given by

$$\frac{\partial u}{\partial t} = -bu + v_1 \frac{\partial u}{\partial x} + v_2 \frac{\partial u}{\partial y} + c^2 \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} \right), \tag{3}$$

where $u(x, y, t)$ models the plume intensity at the point (x, y) at time t , b is the absorption coefficient, and v_1 and v_2 are the components of the advection velocity. Considering that the

propagation usually takes place within a relatively small field in the atmosphere, an isotropic medium is assumed with a diffusion coefficient c . This model also assumes that the propagation takes place essentially at the surface; accounting for the third dimension would be an important next step.

We assume that the medium is infinite and the initial condition is:

$$u(x, y, t = t_0) = C\delta(x - x_0)\delta(y - y_0), \quad (4)$$

where $\delta(\cdot)$ is the Dirac delta function, which corresponds to a very sharply localized release of intensity C at time $t = t_0$, at the point (x_0, y_0) . Then the solution to (3) and (4) can be calculated exactly as

$$u(x, y, t) = \frac{a}{t - t_0} \exp \left\{ -b(t - t_0) - \frac{(x - v_1 t - x_0)^2 + (y - v_2 t - y_0)^2}{4c^2(t - t_0)} \right\}, \quad (5)$$

where a is a normalization constant depending on the intensity of the initial source, C , and on the diffusion constant, c ; for example, see [16]. Formula (5), which represents the Green's function of Equation (3) in an infinite medium, can be generalized to higher dimensions and various anisotropic situations. Notice that it is not of product (time \times space) form, and hence the model is non-separable in space and time.

Propagation in inhomogeneous and/or non-stationary media (i.e. accounting for the presence of buildings, mountains, forests, changing meteorological conditions, etc.) can be described by replacing the constant coefficients with appropriate functions of time and position and/or supplementing the equation with appropriate boundary conditions. While in general these situations do not lend themselves to simple analytical solutions (see [17]), numerical approaches are usually possible and have been pursued to a great extent (e.g. [18]). Note that (stochastic) PDE models for environmental and ecological processes have been used in the hierarchical Bayesian framework (e.g. [19, 20]), although their goal was not inference on the location of sources.

In the remainder of our paper, we rely on knowing the explicit solution (5). There are several things to note. Should the PDE be generalized to account for spatial or temporal inhomogeneity, our approach would need an analytical solution. A similar consideration holds if we generalize from two dimensions to three dimensions. Finally, the PDE plume model (3) has no uncertainty associated with it (often expressed in the form of a *stochastic* PDE).

The advantages of the class of models illustrated by Equation (3) include (i) their general mathematical properties are very well understood [16, 17]; (ii) the evolution described by the continuous versions preserves the required physical properties, such as positivity and conservation laws [16, 17]; (iii) due to the parabolic character of the equation, the discretized versions are stable.

2.3. Methodology

From now on we shall focus on the PDE-based plume model (5). We assume here that the total number of plume sources K is known; an approach to determine K from data will be discussed in Section 4. To simplify notation, write

$$u(x, y, t; \tau_k) = I(t > t_{0k}) \frac{a_k}{t - t_{0k}} \exp \left\{ -b_k(t - t_{0k}) - \frac{(x - v_{1k}t - x_{0k})^2 + (y - v_{2k}t - y_{0k})^2}{c_k(t - t_{0k})} \right\},$$

where $\tau_k = (a_k, b_k, c_k, v_{1k}, v_{2k}, t_{0k}, x_{0k}, y_{0k})'$ is the vector of parameter values for the k th plume source. Let $\tilde{\tau}_k \equiv (\tilde{a}_k, \tilde{b}_k, \tilde{c}_k, \tilde{v}_{1k}, \tilde{v}_{2k}, \tilde{t}_{0k}, \tilde{x}_{0k}, \tilde{y}_{0k})'$ denote the corresponding true parameter values.

Observe that the true intensity function $u_k(x, y, t)$ defined in (1) can be written in terms of the model as $u(x, y, t; \tilde{\tau}_k)$.

Assume that the errors $\{\varepsilon_{i,j}\}$ in (2) are independent and normally distributed with mean 0 and variance σ^2 ; let the true variance be $\tilde{\sigma}^2$. We consider the robustness of this assumption below in Section 3. Thus, the combined parameter of the K plume sources is

$$\theta \equiv (\tau_1, \tau_2, \dots, \tau_K, \sigma^2)',$$

which is a vector of length $8K + 1$. In addition, denote by \mathbf{z} the data $\{z_{i,j}, 1 \leq j \leq n_i, 1 \leq i \leq n_s\}$ defined by (2). Then the likelihood of θ is

$$p(\mathbf{z}|\theta) \propto (\sigma^2)^{-N/2} \exp \left[-\frac{1}{2\sigma^2} \sum_{i=1}^{n_s} \sum_{j=1}^{n_i} \left\{ z_{i,j} - \sum_{k=1}^K u(x_i, y_i, t_{i,j}; \tau_k) \right\}^2 \right], \tag{6}$$

where $N = \sum_{i=1}^{n_s} n_i$. The high dimensionality of θ makes it attractive to regularize the likelihood by putting a prior on θ . Our approach in this paper is to put all the uncertainty about the PDE onto its parameters τ , and express that uncertainty through a prior distribution. Hence, the posterior distribution can (in principle) be obtained, and inference carried out on such parameters as the plume's source.

For prior distribution consideration, if the plume sources are dirty bombs, then highly populated areas may be more likely to contain the bombs' source locations than lowly populated areas, in which case it makes sense for the prior distribution to reflect this. In addition, certain values of a_k, b_k, c_k may be more likely than others, as determined by the physical conditions of the environment. However, if no prior information is available or if one chooses to ignore it, then a uniform distribution on the $6K$ -dimensional vector, $(\tau'_1, \dots, \tau'_K)$, often called a non-informative prior, could be used. One could also use a normal prior with large variances for the elements of $\tau_k, k = 1, \dots, K$.

Bayesian inference focuses on the properties of the posterior distribution, $p(\theta|\mathbf{z})$, the conditional distribution of θ given the data \mathbf{z} . Once $p(\theta|\mathbf{z})$ is available, inference based on it can be made with regard to the unknown model parameters. Furthermore, the overall plume intensity at any location (x, y) and any time t can be predicted using its posterior mean (say):

$$\sum_{k=1}^K \int_{\theta} u(x, y, t; \tau_k) p(\theta|\mathbf{z}) d\theta. \tag{7}$$

The posterior distribution can sometimes be derived in a closed form but often has to be computed numerically. In this paper, the posterior distribution will be computed by Markov chain Monte Carlo (MCMC) simulations. MCMC methods are widely used in statistics, and there are many different approaches. See [21] for an overview of MCMC and discussions of important issues.

3. NUMERICAL RESULTS

In this section, we present some numerical results for the statistical inference on θ from the PDE-based plume model. We are not aware of any real-world data available, although there may be some that are classified. Therefore, we illustrate our approach on the simulated data. For convenience, we shall make the simplification that, for the advection velocities, $v_{1k} = v$ and $v_{2k} = 0$, is the same

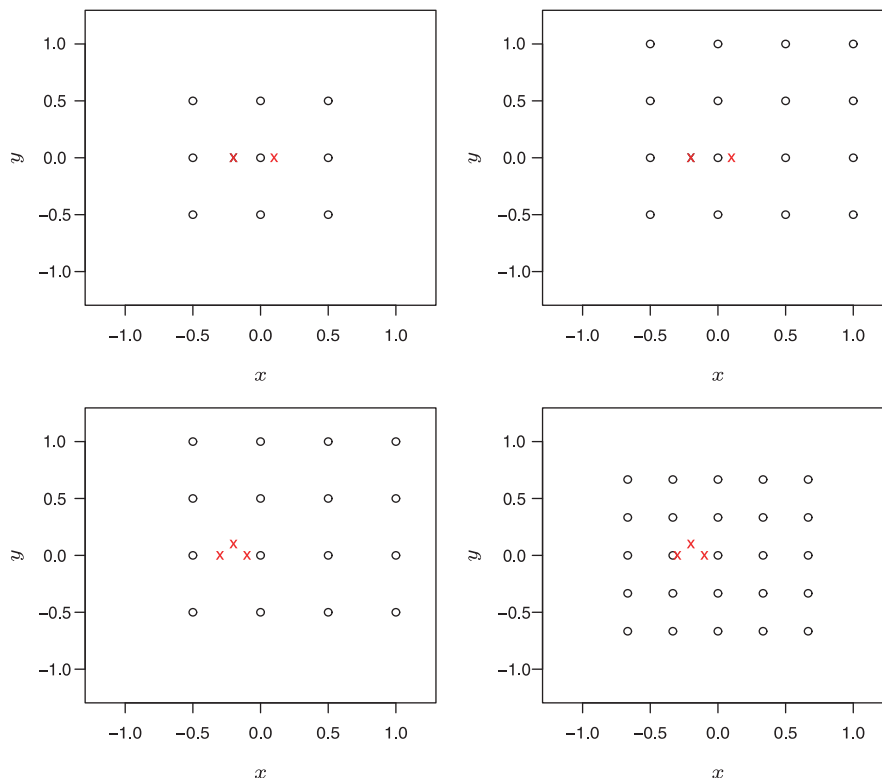


Figure 1. The upper-left/right panels give the source locations (\times) and the 9/16 sensor locations (\circ) for the examples with two sources. The lower-left/right panels give the source locations (\times) and the 16/25 sensor locations (\circ) for the examples with three sources.

for all sensors, which reduces the number of parameters from $8K + 1$ to $6K + 2$. The assumption means that the wind is blowing at more or less constant velocity in the direction of the x -axis in the region containing the sensors. Note that this direction can be replaced by any other known direction, and we do not assume that the wind speed is known. The locations of the plume sources (unknown) and sensors (known) in our examples are given by the configurations in the upper-left panel of Figure 1, where sensor locations are denoted by ' \circ ' and source locations are denoted by ' \times '.

First, we consider an example for two sources (i.e. $K = 2$) and $n_s = 9$ sensors. Assume that the true PDE plume model is specified by the parameters in the top row of Table I, and we simulate data from the model (5) followed by (2). The data $\mathbf{z} = \{z_{i,j}\}$ are observed at 20 equally spaced time points from 0 to 2 at each of the 9 sensor locations shown in the top-left panel of Figure 1. In this simulation, we assumed that the distribution of the measurement errors $\{\varepsilon_{i,j}\}$ is independent and identically distributed normal with mean 0 and standard deviation $\tilde{\sigma} = 0.05$. We first conducted a single MCMC simulation run; see Figure 2, which shows the data $z_{i,j}$ observed at the 9 sensors at 20 time points.

Based on the data, we conducted the statistical inference using the Bayesian approach described in Section 2.3. We assumed that the prior distributions for each of the parameters were mutually

Table I. Estimated posterior means of the parameters for $K = 2$. The top row shows the parameter setting for the simulations. 'Independent' indicates that the measurement errors are independent, while 'Spatio-temporal' indicates that the measurement errors are correlated in space and time.

Parameter Setting	$\tilde{a}_1 = 3$ $\tilde{a}_2 = 3$	$\tilde{b}_1 = 1$ $\tilde{b}_2 = 1$	$\tilde{c}_1 = 1$ $\tilde{c}_2 = 1$	$(\tilde{x}_{01}, \tilde{y}_{01}) = (-0.2, 0)$ $(\tilde{x}_{02}, \tilde{y}_{02}) = (0.1, 0)$	$\tilde{f}_{01} = 0.4$ $\tilde{f}_{02} = 0.5$	$\tilde{v} = 0.6$ $\tilde{\sigma}$
$n_s = 9$ Independent	2.972 2.984	1.105 0.907	1.015 1.009	(-0.199, -0.003) (0.101, 0.004)	0.402 0.500	0.588 0.056
$n_s = 9$ Spatio-temporal	2.979 3.011	1.132 0.902	1.004 1.002	(-0.198, 0.001) (0.100, -0.003)	0.402 0.499	0.587 0.040
$n_s = 9$ Independent	2.893 2.942	1.099 0.901	1.036 1.034	(-0.202, -0.012) (0.103, 0.015)	0.407 0.503	0.591 0.222
$n_s = 9$ Spatio-temporal	3.097 2.941	1.327 0.789	0.953 1.036	(-0.202, 0.002) (0.099, -0.012)	0.403 0.500	0.590 0.159
$n_s = 16$ Independent	3.114 2.896	1.071 0.938	0.974 1.032	(-0.203, 0.001) (0.101, -0.002)	0.397 0.503	0.600 0.053
$n_s = 16$ Spatio-temporal	3.067 2.927	0.886 1.123	0.965 1.022	(-0.205, -0.004) (0.102, 0.001)	0.397 0.502	0.616 0.043
$n_s = 16$ Independent	3.165 2.857	1.174 0.884	0.960 1.067	(-0.208, 0.004) (0.101, -0.009)	0.396 0.504	0.602 0.213
$n_s = 16$ Spatio-temporal	3.095 2.856	0.761 1.329	0.926 1.061	(-0.212, -0.019) (0.102, 0.010)	0.398 0.504	0.636 0.172

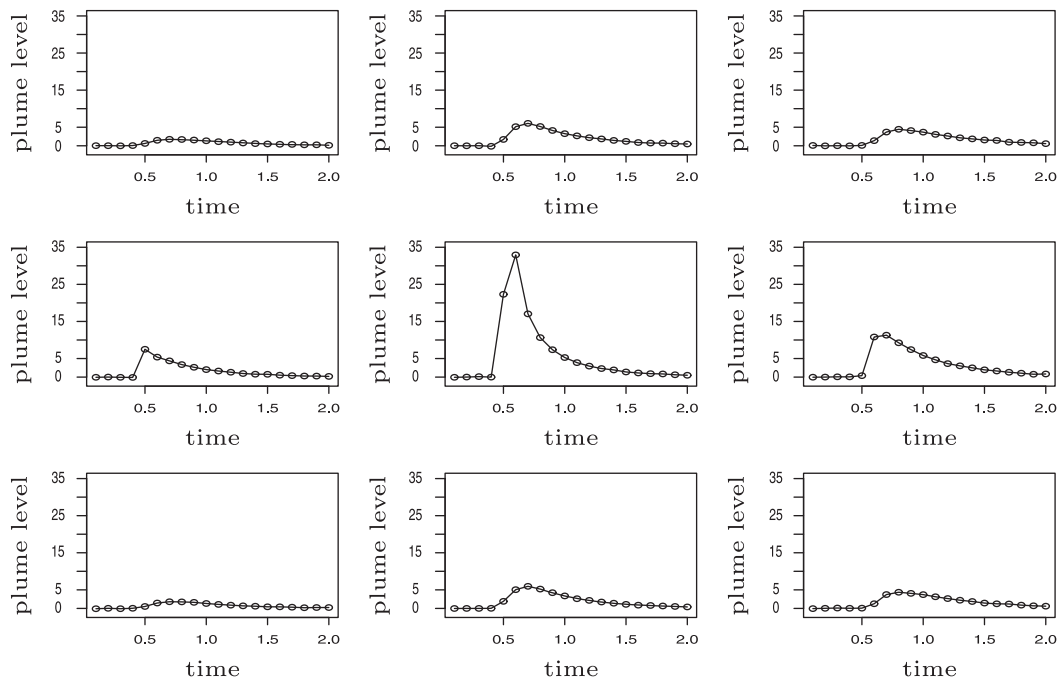


Figure 2. The plots show the data collected by the nine sensors for one simulation.

independent. A scaled inverse- χ^2 distribution with scale parameter 0.01 and degrees of freedom 1 was used for σ^2 , while for x_{0i} , y_{0i} , v , a_i , b_i , c_i , t_i we used uniform priors with truncated ranges $(-1, 1)$, $(-1, 1)$, $(-1, 1)$, $(1, 3)$, $(1, 3)$, $(1, 3)$, and $(0.2, 0.8)$, respectively. Note that in practice, if more information is available, one may choose some more informative priors. While σ^2 was updated directly through a conjugate posterior distribution, all other parameters were updated using the Metropolis algorithm (see [21]). MCMC simulations were implemented with 100 000 iterations; after a burn-in of 50 000, the last 50 000 simulation results were used to compute the empirical posterior distributions. The distributions obtained from a single MCMC run are summarized in the plots in Figure 3, where the bar charts are relative-frequency histograms. The superimposed smooth curves are the corresponding density estimates obtained by the kernel smoothing, where Silverman's bandwidth-selection rule-of-thumb (see [22]) is implemented. It can be observed that the estimated posterior distributions are fairly tight and include the true parameters in their ranges. The estimated posterior means are reported in Table I on the row that begins with ' $n_s = 9$, $\tilde{\sigma} = 0.05$, Independent'. All the estimates are quite close to their corresponding true values.

We then used (7) to predict the true plume intensity at time $t = 2.5$ and over the spatial region $[-1, 1] \times [-1, 1]$. Note that this time point and a portion of the spatial region are beyond the coverage of the sensors, as reflected by the data \mathbf{z} . The predicted intensities are displayed in the left panel of Figure 4, and the difference between the predicted and the true plume intensities is displayed in the right panel of Figure 4. Clearly, the plume levels are predicted extremely well over the entire region at this time point.

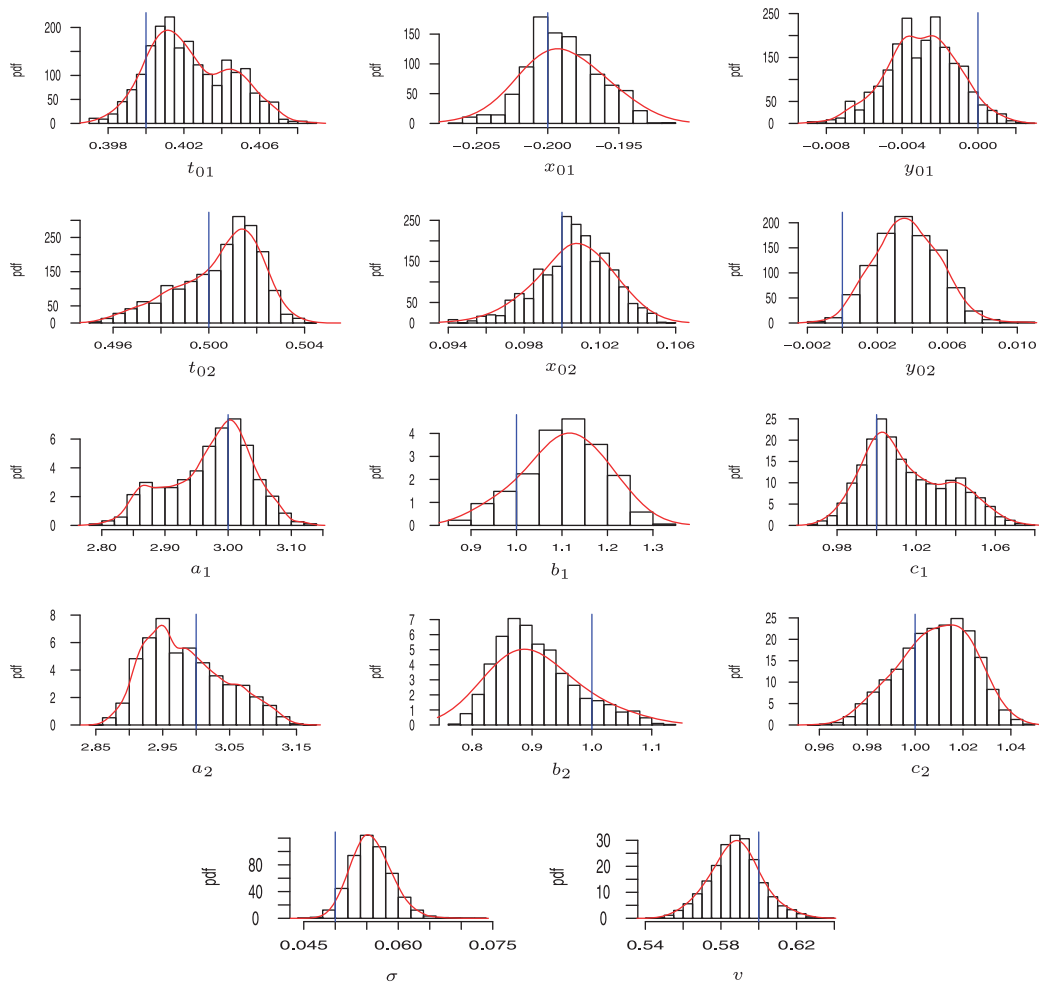


Figure 3. Estimated posterior distributions of the parameters based on a single MCMC run for $K=2$. The curves denote kernel-smoothed densities. The true values are indicated by the vertical lines.

In practice, the sensor-network measurements may not be independent. For instance, one could assume that the measurement errors have spatio-temporal covariance:

$$\text{cov}(\varepsilon_{i_1, j_1}, \varepsilon_{i_2, j_2}) = \sigma^2 C_1(s, \phi_1) C_2(t, \phi_2), \tag{8}$$

where s is the spatial distance between the sensor locations (x_{i_1}, y_{i_1}) and (x_{i_2}, y_{i_2}) , and t is the time difference between the observation times t_{i_1, j_1} and t_{i_2, j_2} . In (8), σ^2 is the overall variance; $C_1(s, \phi_1)$ is the spatial correlation function, where ϕ_1 indicates the spatial dependency; and $C_2(t, \phi_2)$ is the temporal correlation function with ϕ_2 representing the temporal dependency. In (6), let \mathbf{w} denote the vector whose elements are:

$$z_{i,j} = \sum_{k=1}^K u(x_i, y_i, t_{i,j}; \boldsymbol{\tau}_k), \quad 1 \leq i \leq n_s, \quad 1 \leq j \leq n_i.$$




Figure 4. The left panel shows the predicted plume levels. The right panel shows the difference between the predicted plume levels and the true plume levels.

Then, under the independent-measurement-error assumptions made earlier,

$$p(\mathbf{z}|\boldsymbol{\theta}) \propto (\sigma^2)^{-N/2} \exp\left\{-\frac{1}{2\sigma^2} \mathbf{w}'\mathbf{w}\right\}. \quad (9)$$

Now, if $\{\varepsilon_{i,j}\}$ are dependent, having the covariance matrix Σ (e.g. obtained from (8)), then the likelihood function becomes:

$$p(\mathbf{z}|\boldsymbol{\theta}) \propto |\Sigma|^{-1/2} \exp\left\{-\frac{1}{2} \mathbf{w}'\Sigma^{-1}\mathbf{w}\right\} \quad (10)$$

note that $\boldsymbol{\theta}$ is then augmented to include the parameters of Σ . In each iteration of the MCMC procedure, an inversion of the matrix Σ needs to be computed. This can complicate the procedure and can be time consuming.

It is of great interest to investigate the robustness of our procedure in this situation. With the same set up as above, we generated the random error terms according to the spatio-temporal structure (8) with the spatial correlation function $C_1(s, \phi_1) = e^{-s/\phi_1}$ and the temporal correlation function $C_2(t, \phi_2) = e^{-t/\phi_2}$, where $\phi_1 = \phi_2 = 0.4$. The same MCMC procedure was conducted to obtain the (estimated) posterior distributions. We first present the estimated posterior means on the row that begins with ' $n_s = 9$, $\tilde{\sigma} = 0.05$, spatio-temporal' in Table I. It can be observed that these values are very similar to their 'independent' counterparts and are close to the true values. To compare the estimated posterior distributions with those in Figure 3, Q-Q plots were drawn for each parameter and are included in Figure 5. These plots show that the estimated posterior distributions of each parameter for the two models are similar to each other. In conclusion, this limited study demonstrates that an assumption of independent measurement errors in (2), even when the measurement errors are correlated as in (8), can still result in valid inferences. Note that this approach can be viewed as a Bayesian procedure based on the quasi-likelihood [23]; we use (9) instead of the full likelihood (10), and the parameters (ϕ_1, ϕ_2) in the spatio-temporal correlation functions are not considered. Conducting inference based on quasi-likelihoods is common in non-Bayesian statistics; see [24] for a general discussion.

To investigate further the performance of the MCMC procedure for the case of two plume sources, we expanded the simulation study by considering all combinations of $\tilde{\sigma} = 0.05$ and 0.2;

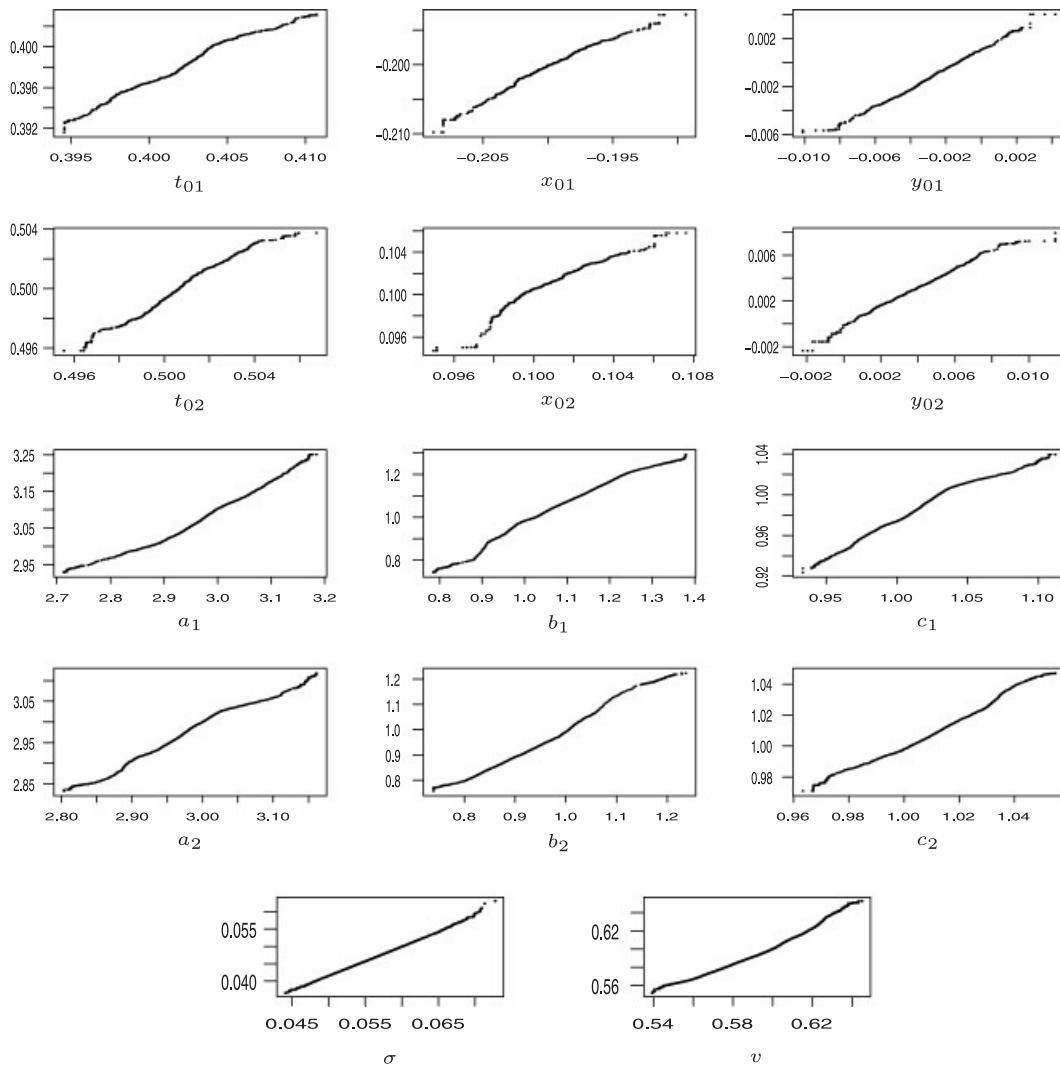


Figure 5. Q–Q plots of the estimated posterior distributions of parameters when measurement errors are correlated in space and time versus those when measurement errors are independent; $K = 2$.

$n_s = 9$ and 16; and independent and spatio-temporally correlated measurement errors. For example, the plume-source/sensor configuration for $n_s = 16$ is shown in the upper-right panel of Figure 1. The estimated posterior means are reported in Table I. For all the settings, the estimated posterior means estimate the true parameter values quite well, especially for estimating the spatial and temporal origins (x_0, y_0) and t_0 , respectively, of the sources. The robustness of the procedure can be observed by comparing the entries between the independent-measurement-error setting and the correlated-measurement-error setting. Note that the procedure performs slightly better for $\tilde{\sigma} = 0.05$ than for $\tilde{\sigma} = 0.2$, as expected, while the results are similar when the number of sensors changes

from $n_s = 9$ to 16. A more complete simulation study would determine the value of $\tilde{\sigma}$ (expressed in terms of a signal-to-noise ratio) at which inferences deteriorate.

Next, we consider the situation when there are three sources (i.e. $K = 3$). First assume that the true PDE plume model is specified by the parameters in the top row of Table II; the measurement errors $\{\varepsilon_{i,j}\}$ are assumed to be normally distributed with mean 0 and standard deviation $\tilde{\sigma} = 0.05$, and the data $\{z_{i,j}\}$ are observed at 20 equally spaced time points, as before. Consider the case of 25 sensor locations shown in the lower-right panel of Figure 1, for which we obtained data from a single MCMC run. The estimated posterior means are presented in Table II, and the estimated posterior distributions for the parameters are summarized in the plots in Figure 6. Some plots seem to suggest that the procedure gives biased estimates; for example, the true value of t_{02} is not in the range of the estimated posterior distribution. However, the true value of t_{02} is $\tilde{t}_{02} = 0.6$; considering the empirical posterior distribution ranges from 0.6002 to 0.6008 with an estimated posterior mean of 0.601, the bias is actually quite small.

Similar to the case of two plume sources, we also studied the performance of the procedure under various settings. Specifically, we considered all combinations of $\tilde{\sigma} = 0.05$ and 0.2; $n_s = 16$ and 25; and independent and spatio-temporally correlated measurement errors. The plume-source/sensor configuration for $n_s = 16$ is shown in the lower-left panel of Figure 1, and recall that for $n_s = 25$ is shown in the lower-right panel. The estimated posterior means using the MCMC procedure are presented in Table II. Again, the MCMC procedure worked very well in all the settings.

All the studies up to this point are based on a single MCMC run. One may wonder whether this approach would consistently perform well when multiple MCMC runs are conducted. In this part of simulation, we conducted 50 MCMC runs for two examples: one with two sources ($K = 2$) and one with three sources ($K = 3$). Both examples have the same set up: $n_s = 16$, $\tilde{\sigma} = 0.05$ and measurement errors are independent. For each example, we simulated the data $z_{i,j}$, $1 \leq i \leq 16$, $1 \leq j \leq 20$, 50 independent times, and for each run we went through the MCMC simulation to obtain the (estimated) posterior distributions. For each parameter θ_ℓ , we computed the (estimated) posterior mean $\int \theta_\ell p(\theta|\mathbf{z}) d\theta$. Consequently, for each example, we obtained 50 posterior means, one from each of the 50 simulations; these are presented in Figure 7 ($K = 2$) and Figure 8 ($K = 3$). It can be seen that they are all in a tight range and that they cover the corresponding true values. In particular, the spatial locations (x_0, y_0) and the temporal origins t_0 of the sources are estimated very well for all 50 MCMC runs in both figures.

4. DISCUSSION

In this paper, we present a Bayesian statistical approach for identifying the parameters, particularly the source locations, of a PDE-based plume model from sensor-network data. We showed that even with a moderate amount of data, the model parameters can be estimated using an MCMC approach. The approach can also be used for spatial and temporal prediction of plume levels beyond the sensor range.

In our analysis, we assumed that the true number of sources is known. In practice, the number may have to be determined from data. In the literature, the deviance information criterion (DIC) has been used (e.g. [21, 25]):

$$\text{DIC} = \hat{D}_{\text{ave}} + p_D,$$

Table II. Estimated posterior means of the parameters for $K=3$. The top row shows the parameter setting for the simulation. 'Independent' denotes that the measurement errors are independent, while 'Spatio-temporal' indicates that the measurement errors are correlated in space and time.

Parameter Setting	$\hat{a}_1=3$ $\hat{a}_2=3$ $\hat{a}_3=3$	$\hat{b}_1=1$ $\hat{b}_2=1$ $\hat{b}_3=1$	$\hat{c}_1=1$ $\hat{c}_2=1$ $\hat{c}_3=1$	$(\hat{x}_{01}, \hat{y}_{01}) = (-0.2, 0.1)$ $(\hat{x}_{02}, \hat{y}_{02}) = (-0.3, 0.0)$ $(\hat{x}_{03}, \hat{y}_{03}) = (-0.1, 0.0)$	$\hat{\tau}_{01}=0.4$ $\hat{\tau}_{02}=0.6$ $\hat{\tau}_{03}=0.8$	$\hat{\nu}=0.6$ $\hat{\sigma}$
$n_s = 16$ Independent $\tilde{\sigma}=0.05$	2.925 3.022 3.019	0.958 0.991 1.053	1.016 0.988 0.998	$(-0.198, 0.101)$ $(-0.301, 0.003)$ $(-0.100, 0.001)$	0.401 0.600 0.800	0.601 0.054
$n_s = 16$ Spatio-temporal $\tilde{\sigma}=0.05$	2.936 3.022 3.006	0.940 1.046 0.976	1.014 0.994 1.001	$(-0.198, 0.103)$ $(-0.300, 0.004)$ $(-0.100, 0.003)$	0.402 0.600 0.800	0.602 0.066
$n_s = 16$ Independent $\tilde{\sigma}=0.2$	2.822 2.786 3.152	1.232 0.683 1.125	1.093 1.119 0.917	$(-0.203, 0.112)$ $(-0.297, -0.002)$ $(-0.129, 0.000)$	0.413 0.602 0.802	0.648 0.199
$n_s = 16$ Spatio-temporal $\tilde{\sigma}=0.2$	2.827 2.950 3.104	0.822 1.045 1.018	1.041 1.004 0.983	$(-0.195, 0.112)$ $(-0.297, 0.019)$ $(-0.102, 0.012)$	0.408 0.604 0.797	0.611 0.138
$n_s = 25$ Independent $\tilde{\sigma}=0.05$	3.182 2.822 3.113	1.084 0.792 1.157	0.960 1.033 0.972	$(-0.204, 0.100)$ $(-0.296, -0.000)$ $(-0.103, 0.000)$	0.395 0.601 0.798	0.604 0.060
$n_s = 25$ Spatio-temporal $\tilde{\sigma}=0.05$	2.989 2.980 3.079	0.955 0.974 1.078	1.001 1.004 0.988	$(-0.199, 0.098)$ $(-0.298, -0.000)$ $(-0.100, 0.000)$	0.400 0.600 0.798	0.588 0.048
$n_s = 25$ Independent $\tilde{\sigma}=0.2$	3.201 2.929 2.957	0.943 0.998 1.062	0.950 1.002 1.003	$(-0.204, 0.098)$ $(-0.301, 0.000)$ $(-0.103, 0.001)$	0.394 0.600 0.802	0.598 0.211
$n_s = 25$ Spatio-temporal $\tilde{\sigma}=0.2$	3.246 3.069 2.924	1.153 0.977 0.980	0.955 0.987 1.007	$(-0.204, 0.100)$ $(-0.304, -0.005)$ $(-0.103, -0.002)$	0.394 0.598 0.801	0.612 0.206

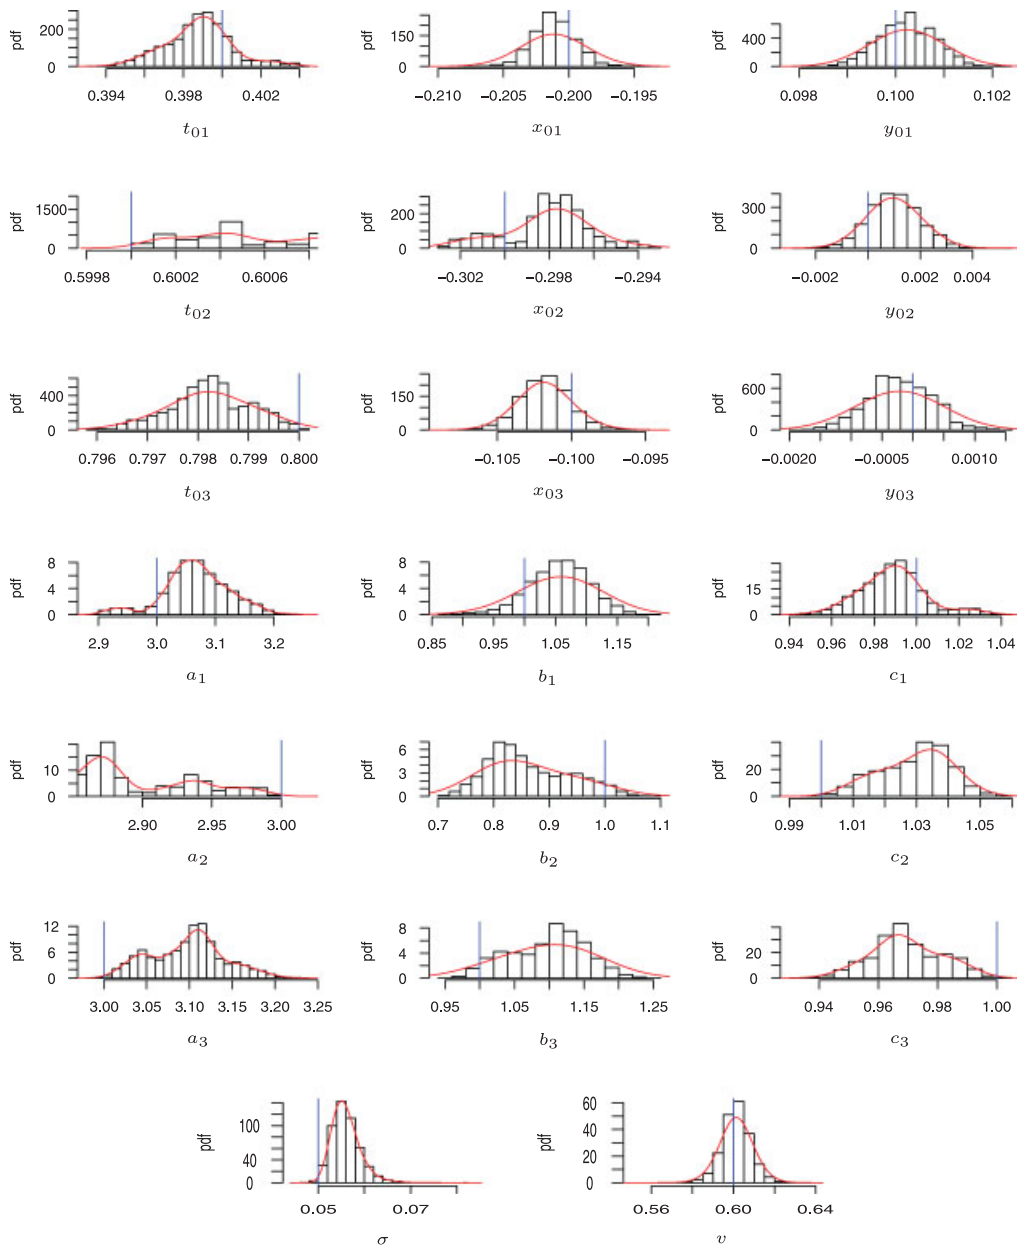


Figure 6. Estimated posterior distributions of the parameters based on one MCMC run for $K = 3$. The curves denote kernel-smoothed densities. The true values are indicated by the vertical lines.

where D is the deviance defined as -2 times the log-likelihood, and $p_D = \hat{D}_{\text{ave}} - D_{\hat{\theta}}$ is a measure of the effective number of parameters. The \hat{D}_{ave} is computed as the average of the deviance functions for all MCMC iterations, and $D_{\hat{\theta}}$ is the deviance at the parameters' posterior means.

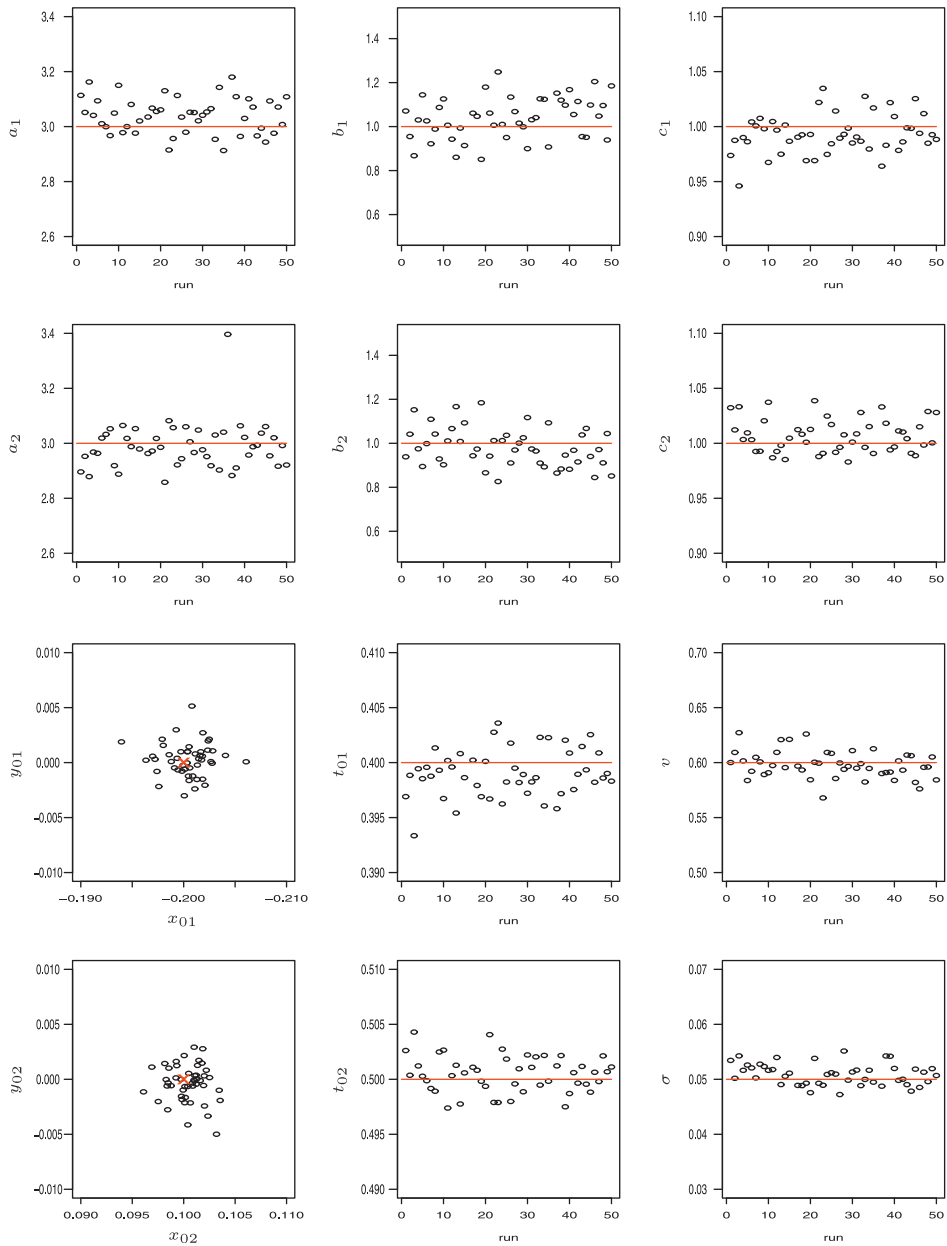


Figure 7. Estimated posterior means of the parameters based on 50 MCMC runs for $K = 2$. The crosses indicate the true source locations and the lines show the true values.

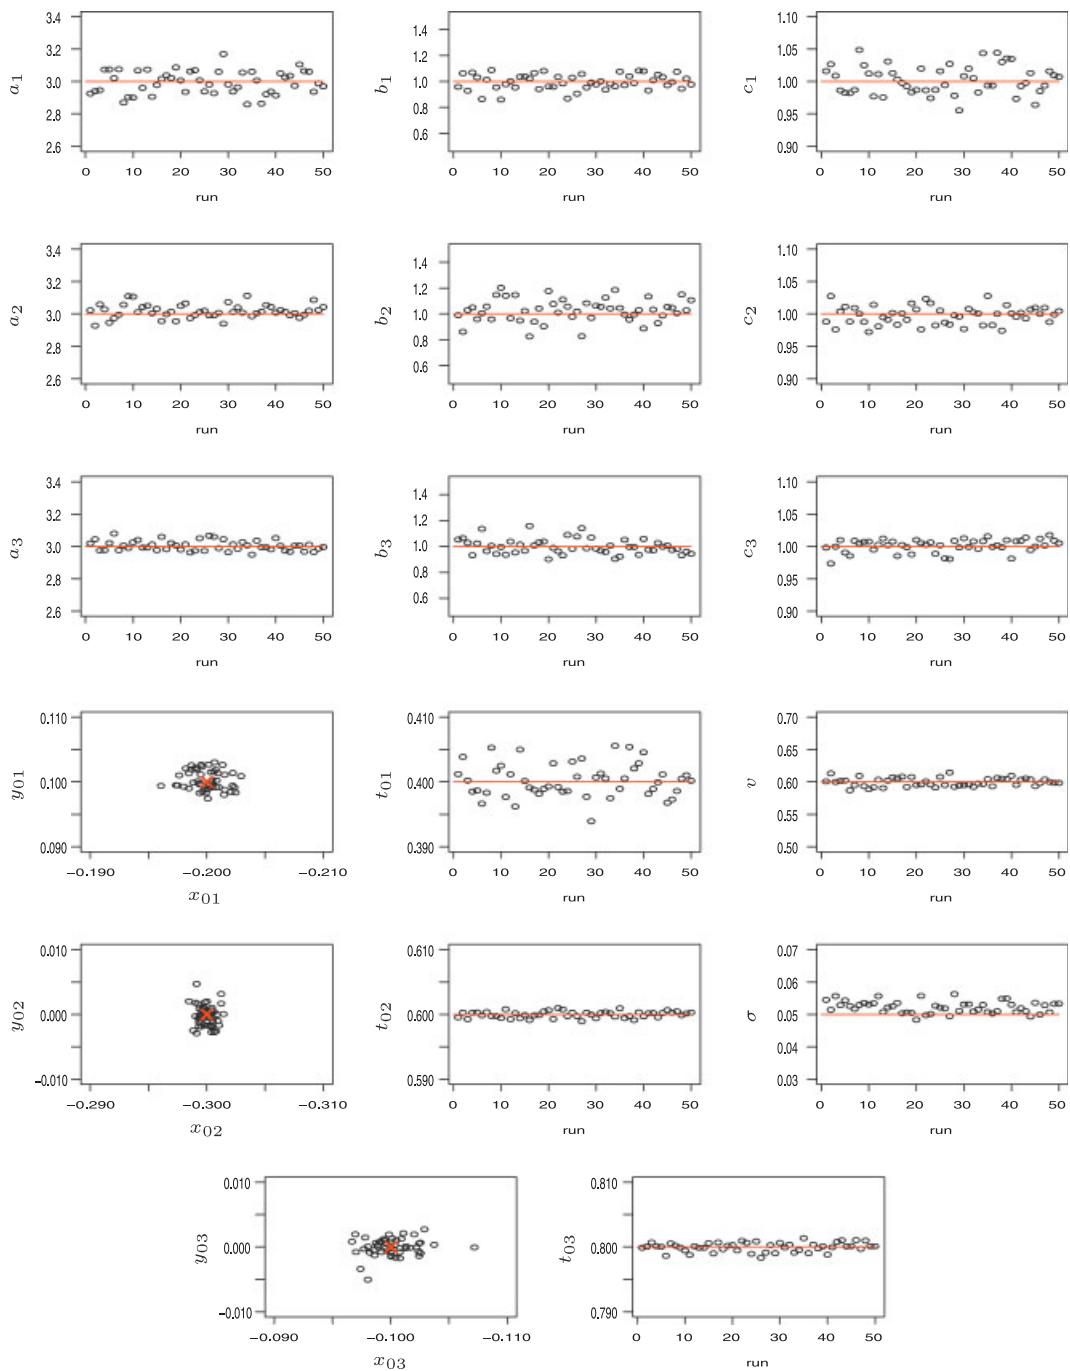


Figure 8. Estimated posterior means of the parameters based on 50 MCMC runs for $K = 3$. The crosses indicate the true source locations and the lines show the true values.

Table III. The \hat{D}_{ave} , $D_{\hat{\theta}}$, p_D and DIC values for three models with $K = 1, 2, 3$. Lower values of DIC indicate a better fit. The true model is $K = 2$ with 14 parameters.

Model	\hat{D}_{ave}	$D_{\hat{\theta}}$	p_D	DIC
$K = 1$	155.0	156.7	1.7	158.3
$K = 2$	-873.2	-857.4	15.8	-841.6
$K = 3$	-875.3	-858.1	17.2	-840.9

Based on the data of the two-source example in Figure 2, the \hat{D}_{ave} , $D_{\hat{\theta}}$, p_D and DIC values are computed for $K = 1, 2, 3$ and are shown in Table III. In this case, the true model has two sources, and it is clear from Table III that the model with $K = 2$ is preferable than others, since it has the smallest DIC. Note that the chosen model ($K = 2$) has a p_D that is closer to the true number of parameters.

A number of important extensions would require further research:

- (i) Various computational issues should be considered. The MCMC simulations that we presented for the two-source example took about 90 seconds to implement on a Linux workstation (Dual Pentium 4 Xeon at 3 GHz with 4 GB RAM) using Fortran code. The computational speed could be substantially improved using a more efficient computing platform. The more interesting question is, how can the computations be performed in real time? In other words, as new plume-evolution information becomes available, it is desirable to update the estimation/prediction based on the previous results without restarting them from scratch. One possible way to achieve this goal is to use sequential Monte Carlo algorithms as in [26].
- (ii) In our simulations, we assumed that the sensors are placed on a grid. This is because our trial runs revealed that this scheme of sensor placement leads to the most satisfactory statistical inference on average. An important question is if we have a fixed number of sensors with limited query capabilities, how should we optimize the design of the network? Spatial designs of this type are discussed in [27], and spatio-temporal designs (assuming mobile sensors) are discussed in [28, 29].

ACKNOWLEDGEMENTS

This work was supported by the SensorNet Program at the Oak Ridge National Laboratory under UT-Battelle LLC contract 4000042084, and the Office of Naval Research under contract number N00014-08-01-0464. The authors thank the editors and the referees for their helpful comments.

REFERENCES

1. Brooks RR, Iyengar SS. *Frontiers of Distributed Sensor Networks*. CRC Press: Boca Raton, FL, 2004.
2. Culler D, Estrin D, Srivastava M. Overview of sensor networks. *IEEE Computer* 2004; **37**:41–49.
3. Berliner LM. Physical-statistical modeling in geophysics. *Journal of Geophysical Research* 2003; **108**:D24, STS 3-1-ST5 3-10.
4. Chen Y, Moore K, Song Z. Diffusion boundary determination and zone control via mobile actuator-sensor networks (mas-net): challenges and opportunities. *Proceedings of SPIE: Intelligent Computing: Theory and Applications* 2004; **5421**:102–113.

5. Ishida H, Nakamoto T, Moriizumi T, Kikas T, Janata J. Plume-tracking robots: a new application of chemical sensors. *Biological Bulletin* 2001; **200**:222–226.
6. Christopoulos VN, Roumeliotis S. Adaptive sensing for instantaneous gas release parameter estimation. *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005)*, Barcelona, Spain, 18–22 April 2005.
7. Brennan SM, Mielke AM, Torney DC, Maccabe AB. Radiation detection with distributed sensor networks. *IEEE Computer* 2004; **37**:57–59.
8. 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.
9. Stephens DL, Peurrung AJ. Detection of moving radioactive sources using sensors networks. *IEEE Transactions on Nuclear Science* 2004; **51**:2273–2278.
10. Ram SS, Veeravalli VV. Localization and intensity tracking of diffusing point sources using sensor networks. *Proceedings of the IEEE Global Telecommunications Conference (IEEE GLOBECOM 2007)*, Washington, DC, 26–30 November 2007.
11. Rao NSV. Identification of a class of simple product-form plumes using sensor networks. *Innovations and Commercial Applications of Distributed Sensor Networks Symposium*, Bethesda, MD, 2005.
12. Rao NSV. Identification of simple product-form plumes using networks of sensors with random errors. *Proceedings of the 9th International Conference on Information Fusion (Fusion 2006)*, Florence, Italy, 10–13 July 2006.
13. Li W, Farrell JA, Carde R. Tracking of fluid-advected odor plumes: strategies inspired by insect orientation to pheromone. *Adaptive Behavior* 2001; **9**:143–170.
14. Sykes RI, Lewellen WS, Parker SF. A Gaussian plume model of atmospheric dispersion based on second-order closure. *Journal of Climate and Applied Meteorology* 1986; **25**:322–331.
15. Brennan SM, Mielke AM, Torney DC. Radiation source detection by sensor networks. *IEEE Transactions on Nuclear Science* 2005; **52**:813–819.
16. Greenberg W, van der Mee CVM, Protopopescu V. *Boundary Value Problems in Abstract Kinetic Theory*. Birkhauser: Boston, 1987.
17. Morse PM, Feshbach H. *Methods of Mathematical Physics*. McGraw-Hill: New York, 1953.
18. Sykes RI. *PC-Scipuff, Version 1.3 Technical Documentation*. Titan Corporation, A.R.A.P. Report no. 725, 2000.
19. Malmberg A, Arelleno A, Edwards DP, Flyer N, Nychka D, Wikle CK. Interpolating fields of carbon monoxide data using a hybrid statistical-physical model. *Annals of Applied Statistics* 2008; **2**:1231–1248.
20. Wikle CK. Hierarchical Bayesian models for predicting the spread of ecological processes. *Ecology* 2003; **84**:1382–1394.
21. Gelman A, Karlin JB, Stern HS, Rubin DB. *Bayesian Data Analysis* (2nd edn). Chapman & Hall: Boca Raton, FL, 2003.
22. Silverman BW. *Density Estimation*. Chapman & Hall: Boca Raton, FL, 1986.
23. White H. Maximum likelihood estimation of misspecified models. *Econometrica* 1982; **50**:1–25.
24. Cox DR, Reid N. A note on pseudolikelihood constructed from marginal densities. *Biometrika* 2004; **91**:729–737.
25. Spiegelhalter D, Best NG, Carlin BP, van der Linde A. Bayesian measures of model complexity and fit. *Journal of Royal Statistical Society B* 2002; **64**:583–639.
26. Wikle CK, Berliner LM. A Bayesian tutorial for data assimilation. *Physica D* 2007; **230**:1–16.
27. Cressie N. *Statistics for Spatial Data* (Revised Edn). Wiley: New York, 1993.
28. Wikle CK, Royle JA. Space-time dynamic design of environmental monitoring networks. *Journal of Agriculture, Biological, and Environmental Statistics* 1999; **4**:489–507.
29. Royle JA, Wikle CK. Efficient statistical mapping of avian count data. *Environmental and Ecological Statistics* 2005; **12**:225–243.