Nonlinear Reduced Order Source Identification under Uncertainty

Reza Khodayi-mehr and Michael M. Zavlanos

***Abstract*—We propose a tractable stochastic model-based approach for identification of chemical sources that relies on the Advection-Diffusion (AD) PDE to model the transport phenomenon and utilizes Markov Chain Monte Carlo sampling to obtain the posterior distribution of the source parameters considering uncertainty in the parameters of the PDE and the sensor data. To make the algorithm tractable, we model the sources using nonlinear basis functions and utilize a model reduction method to obtain closed-form approximate solutions for the AD-PDE. The former idea drastically reduces the dimension of the sampling space while the latter facilitates the evaluation of the likelihood function. We present extensive numerical experiments that demonstrate that our algorithm can estimate the desired source parameters and provide uncertainty bounds for them.**

I. Introduction

The problem of Source Identification (SI) refers to the estimation of the properties of a source using a set of measurements of a quantity that is generated under the action of that source. The SI problem has various applications ranging from environmental protection to human safety. In many of these applications, there is uncertainty in the parameters of the model and the measurements used to identify the sources. Thus, it is beneficial to obtain a posterior distribution for the source parameters instead of a point estimate.

The SI problem in AD transport systems is known as chemical plume tracing and odor localization in the robotics literature and has been investigated since the early 80s [1]. The algorithms differ depending on the dispersal mechanism, i.e., diffusion- or turbulence-dominated, and are specialized for the particular types of sensors used to take the measurements. They are often bio-inspired and try to mimic the behavior of different bacteria [2], insects [3], or crabs [4]. Generally, the main idea is to stay in the plume and move upwind, in the concentration ascent direction, or a combination of the two. In the literature, the former approach is called anemotaxis while the latter is called chemotaxis [1].

Such heuristic approaches to SI are often successful in practice but do not offer a systematic approach that can handle the localization task under a wide range of conditions and only provide information about the source locations. These limitations can be addressed if the underlying physics is properly incorporated in the formulation, which leads to model-based SI methods. In [5], the localization of a single point source in steady-state AD transport for a semi-infinite domain is considered. More general problems that involve sources of arbitrary shapes in arbitrary domains are typically solved numerically using, e.g., the Finite Element (FE) method. The authors in [6] use the FE method along with total variation regularization to solve the SI problem. Similarly, in [7] we proposed an iterative sparse recovery approach to the SI problem. Despite generality, numerical methods such as the FE method become computationally demanding as the size of the domain grows. In [8] and [9] respectively, we proposed and experimentally validated a nonlinear formulation for the SI problem that is computationally tractable for a mobile robot.

The SI methods discussed so far do not provide uncertainty bounds on the estimates that they return since obtaining such bounds is in general computationally expensive. In [10], the authors use Bayesian hypothesis testing combined with closed-form solutions of the AD model to localize a landmine. There also exists a more dedicated set of works that obtain posterior distributions for the source parameters. For instance, in [11] we utilized Stochastic Reduced Order Models (SROMs) to compute approximate posterior distributions for the source parameters. SROMs, although use a small number of samples, require solving a nonlinear optimization problem for the corresponding probability weights which can be challenging.

In this work, we propose a stochastic, computationally light model-based method to solve the SI problem under parameter and measurement uncertainty. To the best of our knowledge, this is the first approach that can be used to solve realistic SI problems under uncertainty in near real-time onboard mobile robots. Moreover, the proposed SI framework is most general in that it can solve problems in non-convex domains containing multiple sources of arbitrary shapes and, more importantly, it can accommodate arbitrary probability distributions and secondary parameters like boundary conditions that are often uncertain in practice. Specifically, given a prior distribution for the source parameters and secondary parameters, we utilize the MCMC method to sample the posterior distribution of these parameters conditioned on the available measurements. To obtain a tractable solution, we significantly reduce the dimension of the sampling space by modeling the source term using nonlinear basis functions. Furthermore, to efficiently calculate the sample likelihood for the MCMC algorithm, we employ a model reduction method that allows us to obtain closed-form solutions for the AD model. Finally, to construct the proposal distribution of the MCMC algorithm, we utilize deterministic point estimates obtained using the techniques proposed in [8]. This allows us to capture the desired posterior distributions with a considerably smaller number of samples. Our algorithm provides posterior distributions that are used to obtain point estimates for the parameters as well as uncertainty bounds for these estimates.

The rest of this paper is organized as follows. In Section II, we define the stochastic SI problem and in Section III we discuss our approach to efficiently solve it. Section IV contains the numerical experiments and Section V concludes the paper.

Reza Khodayi-mehr and Michael M. Zavlanos are with the Department of Mechanical Engineering and Materials Science, Duke University, Durham, NC 27708, USA, {reza.khodayi.mehr, michael.zavlanos}@duke.edu. This work is supported in part by NSF under grant CPS #1837499.

II. Problem Definition

Let $\Omega \subset \mathbb{R}^{d_x}$ denote the domain of interest and consider a velocity vector field $q : \Omega \to \mathbb{R}^{d_x}$ and its corresponding diffusivity field $\kappa : \Omega \to \mathbb{R}_+$. Then, the transport of a quantity of interest $c : \Omega \to \mathbb{R}$, e.g., a chemical concentration, in this domain is captured by the Advection-Diffusion (AD) PDE

$$-\nabla \cdot (-\kappa \nabla c + q\, c) + s = 0, \tag{1}$$

where $s : \Omega \to \mathbb{R}$ is the source field. Given an appropriate set of Boundary Conditions (BCs), the solution of the AD-PDE (1) is well-posed and unique [12]. We refer to the parameters that appear in the AD-PDE (1), i.e., the velocity and diffusivity fields as well as the BCs, as the input data. In real world problems, these input data are uncertain due to modeling and measurement errors.

Let $x \in \Omega$ denote the spatial coordinates and $z_k \in \mathbb{R}$ denote the k-th measurement of the concentration field $c(x)$ at location x_k for $1 \leq k \leq m$. We consider a measurement model with additive Gaussian noise given as

$$z_k = c(x_k) + \bar{\epsilon}, \tag{2}$$

where $\bar{\epsilon} \sim \mathcal{N}(0, \bar{\sigma}^2)$ and $\bar{\sigma} > 0$ is the standard deviation of the noise. We assume that the noise components for different measurements are independent and identically distributed.

Let $p \in \mathbb{R}^{d_p}$ parameterize the source function s and $\pi(P)$ encode our prior belief about the source parameters, where we use upper case letters to denote Random Variables (RVs) and lower case letters to denote their realizations. Furthermore, let $\theta \in \mathbb{R}^{d_\theta}$ denote a RV that collects the uncertain secondary parameters including the parameters that the PDE input-data depend on and $\pi(\Theta)$ denote our prior belief about these secondary parameters. Note that θ is independent of the source parameters p. Then, we define the desired stochastic SI problem as follows.

Problem 2.1 (Stochastic Source Identification Problem): Given the prior distributions $\pi(P)$ and $\pi(\Theta)$ and a set of m measurements $z_{1:m}$, find the joint posterior distribution $\pi(P, \Theta | z_{1:m}) = \pi(P|z_{1:m})\pi(\Theta|z_{1:m})$ of the source and secondary parameters.

Once we have this distribution, we can acquire point estimates as well as uncertainty bounds for the desired parameters; see Section III-C.

Although the AD-PDE (1) has a unique solution $c(x)$, it is often impossible to obtain this solution in closed-form and only approximations can be obtained using discretization-based methods, e.g. the FE method [12]. Given a realization of the RVs p and θ, let the function $\mu : \Omega \to \mathbb{R}$ denote an approximate solution of the AD model (1). Then, we rewrite the measurement model (2) as

$$z_k = \mu(x_k; p, \theta) + \epsilon, \tag{3}$$

where $\epsilon \sim \mathcal{N}(0, \sigma^2)$ encapsulates the model error in addition to the sensor noise and thus $\sigma \geq \bar{\sigma}$; note that $\bar{\epsilon}$ in (2) models only the sensor noise. We include σ in the vector of secondary parameters θ. Using the Bayes' rule, we have $\pi(P, \Theta | z_{1:m}) \propto \pi(z_{1:m}|P, \Theta)\pi(P)\pi(\Theta)$ where

$$z_{1:m}|p, \theta \sim \mathcal{N}(\mu(x_{1:m}), \sigma^2 \mathcal{I}_m) \tag{4}$$

Algorithm 1 Metropolis-Hastings Algorithm

Require: The measurements $z_{1:m}$;
Require: The model $\mu(\cdot)$ and priors $\pi(P)$ and $\pi(\Theta)$;
Require: Current samples p^i and θ^i;
Require: The proposal distribution $\hat{\pi}(P, \Theta | p^i, \theta^i, z_{1:m})$;
1: Draw a sample $(\bar{p}^i, \bar{\theta}^i)$ from $\hat{\pi}(P, \Theta | p^i, \theta^i, z_{1:m})$;
2: Compute the acceptance probability:

$$\rho^{i+1} = \min\left\{1,\ \frac{\pi(\bar{p}^i, \bar{\theta}^i | z_{1:m})}{\pi(p^i, \theta^i | z_{1:m})} \frac{\hat{\pi}(p^i, \theta^i | \bar{p}^i, \bar{\theta}^i, z_{1:m})}{\hat{\pi}(\bar{p}^i, \bar{\theta}^i | p^i, \theta^i, z_{1:m})}\right\};$$

3: Set $(p^{i+1}, \theta^{i+1}) = (\bar{p}^i, \bar{\theta}^i)$ with probability ρ^{i+1} and $(p^{i+1}, \theta^{i+1}) = (p^i, \theta^i)$ with probability $1 - \rho^{i+1}$.

is the likelihood, $x_{1:m}$ denotes the sequence of m measurement locations, $\mu(x_{1:m}) = (\mu(x_1), \ldots, \mu(x_m))$, and $\mathcal{I}_m \in \mathbb{R}^{m \times m}$ denotes the identity matrix.

Often the target posterior $\pi(P, \Theta | z_{1:m})$ has an unknown distribution that depends on the priors $\pi(P)$ and $\pi(\Theta)$. Furthermore, the set A_p of admissible values of the source parameters is often compact and non-convex. This also can be true for the set A_θ of admissible secondary parameters. Consequently, it is in general impossible to obtain the posterior as a known distribution in closed-form. In the next section we utilize a Markov Chain Monte Carlo (MCMC) sampling approach to approximate $\pi(P, \Theta | z_{1:m})$.

III. Stochastic Source Identification

A. Sampling the Posterior Distribution

Given the measurements $z_{1:m}$ and a sample (p^i, θ^i), let $\hat{\pi}(P, \Theta | p^i, \theta^i, z_{1:m})$ denote a proposal distribution that we can easily sample from. Then, we use the Metropolis-Hastings (MH) algorithm to sample the target posterior $\pi(P, \Theta | z_{1:m})$. Algorithm 1 shows iteration $i + 1$ of this algorithm during which a new sample $(\bar{p}^i, \bar{\theta}^i)$ is drawn from the proposal distribution, its acceptance probability ρ^{i+1} is computed, and this sample is randomly accepted with probability of ρ^{i+1}. Running Algorithm 1 for n iterations generates the set of samples $\{p^i, \theta^i\}_{i=1}^n$ that approximate the desired posterior distribution. More specifically, for the MH algorithm the Markov chain corresponding to these samples is ergodic and stationary with respect to the posterior distribution $\pi(P, \Theta | z_{1:m})$ and converges to it in distribution; see [13], [14] for details.

Efficient implementation of Algorithm 1, requires addressing three challenges. First, the dimensions of the admissible sets A_p and A_θ should be as small as possible since the number of samples required to explore them grows as these dimensions increase. Second, each evaluation of the likelihood (4) in line 2 to compute the acceptance probability ρ^{i+1}, requires one solution of the AD model. This can be computationally prohibitive since the MCMC methods require a large number of samples for convergence. Finally, we need a proposal distribution that effectively samples the high probability regions of the admissible sets.

To address the first challenge, we propose a parameterization for the source term that drastically reduces the number of

required unknowns compared to a discretization-based formulation; cf. [7]. Particularly, we approximately represent the source function $s(x)$ in (1) as a linear combination of n_s tower functions

$$s(x) = \sum\nolimits_{j=1}^{n_s} \beta_j \phi(x; \underline{x}_j, \bar{x}_j), \tag{5}$$

where $\beta_j > 0$ denotes the intensity of the j-th tower function and $\underline{x}_j$ and $\bar{x}_j$ denote the bounds on its support where $\underline{x}_j < \bar{x}_j$. Each tower function in (5) is defined as $\phi(x; \underline{x}, \bar{x}) = 1$ if $\underline{x} \leq x \leq \bar{x}$ and zero otherwise. Given this parameterization of the source function, we have $p = (\beta_1, \underline{x}_1, \bar{x}_1, \ldots, \beta_{n_s}, \underline{x}_{n_s}, \bar{x}_{n_s})$ and $d_p = (2d_x + 1)n_s$. This particular parameterization considerably decreases the dimension of the admissible set $A_p \subset \mathbb{R}^{d_p}$.

Next, to speed-up the evaluation of the AD model, required for the computation of the likelihood (4), we construct a closed-form approximation $\mu(x)$ of the field $c(x)$ for a given realization of the secondary parameters θ by utilizing the Proper Orthogonal Decomposition method [15]. The idea is to perform the expensive computations offline by solving the AD model (1) many times. These solutions, called snapshots, are then used to construct a small set of optimal basis functions that span the solution of the AD model. Given the source parameters p and secondary parameters θ, we can approximate the concentration field $c(x; p, \theta)$ by $\mu(x; p, \theta)$ efficiently and with acceptable accuracy; see [9] for more details.

Finally, using a proposal distribution that samples from high probability regions of the admissible sets A_p and A_θ is crucial for the efficient implementation of the MH Algorithm 1. Since the low dimensional representation of the source discussed above renders the relationship between the approximate concentration field $\mu(x; p, \theta)$ and parameters p nonlinear, finding high probability regions of A_p is difficult; see Section IV-A. In this paper, we build on our previous work [8] to construct a proposal distribution centered around a deterministic solution of the SI problem which allows us to directly sample important regions of A_p. Particularly, let $\hat{p}$ denote the mean of the proposal distribution. Then, given that the RVs P and Θ are independent, we define the desired proposal distribution as

$$\hat{\pi}(P, \Theta | p^i, \theta^i, z_{1:m}) \propto \begin{cases} \mathcal{N}(\hat{p}, \hat{\sigma}^2 \mathcal{I}_{d_p})\, \pi(\Theta) & \text{if } P \in A_p,\ \Theta \in A_\theta \\ 0 & \text{o.w.} \end{cases} \tag{6}$$

where $\hat{\sigma} > 0$ is the standard deviation of the proposal distribution. Various MH algorithms differ in the selection of the mean $\hat{p}$ of the proposal distribution. Here we utilize the random-walk MH algorithm where we set $\hat{p} = p^i$ in Algorithm 1 to draw sample $i+1$. We initialize the algorithm by setting p^0 equal to a deterministic solution of the SI problem obtained from [8] for the measurements $z_{1:m}$, where we solve the deterministic problem for the most probable value of the secondary parameters θ obtained from $\pi(\Theta)$. Note that in (6), we confine the support of the proposal distribution to the admissible set A_p and we directly use the prior distribution $\pi(\Theta)$ for the secondary parameters. This is based on the assumption that Θ can be broken down into small, preferably one-dimensional independent RVs such that we can directly sample $\pi(\Theta)$ without needing a proposal distribution for it.

B. Modeling Uncertainty in the Secondary Parameters

Solving real-world SI problems requires the consideration of various sources of uncertainty pertaining to the mathematical model and the PDE input-data, which here are captured in the RV Θ. The uncertainty in the model includes the simplifying assumptions required to derive it as well as the approximations made afterwards, e.g., the discretization of the domain and the model reduction. The standard deviation σ of the measurement model in (3) captures these uncertainties. Uncertainty in the PDE input-data is another practical concern; see [16] for a detailed study of uncertainty quantification for flow fields. A major challenge in incorporating uncertainty in the PDE input-data is that for every new input sample, we need to build a new AD model since the reduced model $\mu(x; p, \theta)$ is obtained for a given value of these input data. In the following, we present an efficient approach based on Stochastic Reduced Order Models (SROMs) [11] to sample the distribution of the secondary parameters θ that appear in the PDE input-data.

For simplicity assume that the first $\bar{d}_\theta$ elements of $\theta \in \mathbb{R}^{d_\theta}$ correspond to parameters that appear in the input data whereas the rest of the entries correspond to other hyper-parameters, e.g., the standard deviation σ in (3). With slight abuse of notation, let Θ_l denote the l-th entry of Θ that is independent of all other entries where $l \in \{1, \ldots, \bar{d}_\theta\}$. For instance, θ_l can denote the uncertain inlet velocity BC that is required to obtain the velocity field $q(x)$ in (1). Given the marginal prior distribution $\pi(\Theta_l)$, we can use SROMs to obtain a discrete approximation $\{\alpha_{l,j}, \theta_{l,j}\}_{j=1}^{n_{\theta,l}}$ for this distribution using $n_{\theta,l}$ samples of the RV Θ_l with corresponding probability weights $\alpha_{l,j}$. Specifically, let $\Xi_{l,j} \subset \mathbb{R}$ for $j \in \{1, \ldots, n_{\theta,l}\}$ denote a partition of the feasible set A_{θ_l} corresponding to the sample $\theta_{l,j}$. Then, the distribution $\pi(\Theta_l)$ can be approximated by the piecewise constant probability density function

$$\pi(\theta_l) \approx \sum\nolimits_{j=1}^{n_{\theta,l}} \alpha_{l,j} \mathbb{1}(\theta_l \in \Xi_{l,j}), \tag{7}$$

where $\mathbb{1}(\theta_l \in \Xi_{l,j})$ is an indicator function. If Θ_l is dependent on any other secondary parameters, then we consider the joint distribution. The SROM approximation of a multivariate distribution is handled exactly the same way; see [11] and the references therein for theoretical details.

Given a set of SROM approximations $\{\alpha_{l,j}, \theta_{l,j}\}_{j=1}^{n_{\theta,l}}$, we construct the joint SROM representation of all secondary parameters $\theta_{1:\bar{d}_\theta}$ that appear in the PDE input-data as $\left\{\prod_{l=1}^{n_\theta} \alpha_{l,j}, \theta_{1:\bar{d}_\theta,j}\right\}_{j=1}^{n_\theta}$ where $n_\theta = \prod_{l=1}^{n_\theta} n_{\theta,l}$ is the total number of SROM samples $\theta_{1:\bar{d}_\theta,j}$ obtained from all combinations of individual samples $\theta_{l,j}$ for $l \in \{1, \ldots, \bar{d}_\theta\}$. Given these samples in the joint space, we approximate the desired distribution $\pi(\Theta_{1:\bar{d}_\theta})$ of the secondary parameters with an equation similar to (7) in which the partition $\Xi_{1:\bar{d}_\theta,j}$ is the intersection of the individual partitions $\Xi_{l,j}$.

The use of unequal weights in SROMs allows for estimation of $\pi(\Theta_{1:\bar{d}_\theta})$ with a small number of samples. This is the key to quantifying uncertainty in the presence of secondary parameters. Specifically, given a sample $\theta_{1:\bar{d}_\theta,j}$, we construct a reduced AD model $\mu(x; p, \theta_j)$ corresponding to the PDE input-data specified by that sample. We use direct sampling to obtain the posterior distribution $\pi(\Theta_{1:\bar{d}_\theta} | z_{1:m})$ of the secondary

parameters. Specifically, we update the SROM probability weights $\alpha_{l,j}$ while keeping the samples $\theta_{l,j}$ fixed.

C. Point Estimators for Source Parameters

In this section we briefly discuss the procedure to obtain point estimates and uncertainty bounds for the source parameters P given a set of n MCMC samples generated by the MH Algorithm 1 for the target distribution $\pi(P|z_{1:m})$. Specifically, we can calculate the *mean* as a point estimate as $\text{avg}[P] = 1/n \sum_{i=1}^{n} p^i$. On the other hand, the computation of the variance is not as straight-forward since the MCMC samples are correlated. The correlation decreases the efficiency of sampling so that n samples provide less accurate estimates than if the samples were independent. In order to remove the correlation among the samples, we utilize the so-called batch-means method [17]. Specifically, we divide the samples into n_b batches of length n_l so that $n \approx n_b\, n_l$. [1] Let $\text{avg}[P]_b$ be the mean of the samples in batch b. Then, the sample variance of these mean values asymptotically estimates σ^2/n_l. Thus, we can estimate the *variance* of the source parameters as

$$\text{var}[P] = \frac{n_l}{n_b} \sum_{b=1}^{n_b} (\text{avg}[P]_b - \text{avg}[P])(\text{avg}[P]_b - \text{avg}[P])^\top.$$

Note that as $n \to \infty$, the uncertainty in the point estimator $\text{avg}[P]$ will converge to zero while $\text{var}[P]$ will converge to the true value. Nevertheless, due to correlation, the convergence is slower than if the samples were independent. It is possible to define an effective sample size that measures the efficiency of the sampling process. Specifically, let $\text{corr}(\tau)$ denote the autocorrelation with lag $\tau \in \mathbb{N}$ among the samples. Then, the effective sample size is defined as

$$n_{\text{eff}} = \frac{n}{1 + 2\sum_{\tau=1}^{\infty} \text{corr}(\tau)}. \tag{8}$$

Remark 3.1 (Alternative Estimators): Given the posterior distribution $\pi(P|z_{1:m})$, we can define different point estimators for the source parameters, e.g., the mean estimator. Another estimator of interest is the *maximum a posteriori* (*map*) estimator that is defined as the source parameter values with maximum posterior probability [18].

IV. Results

In this section we demonstrate the performance of our proposed stochastic Source Identification Algorithm 1 through numerical experiments implemented using MATLAB on a computer with 64GB of memory and 4.20GHz clock speed. Specifically, we consider a non-convex domain depicted in Figure 1. We assume that the air flows into the domain through the sides with a constant velocity. Then, an in-house fluid dynamics code is utilized to simulate the steady-state velocity filed $q(x)$ in the domain with a Finite Element (FE) mesh with 15034 grid points. In the following results, we set the inlet velocity BC to $q_{\text{in}} = 0.025$ and use a constant diffusivity value of $\kappa = 10^{-4}$ as the ground truth.

To select the standard deviation $\hat{\sigma}$ of the proposal distribution (6), we scale the source parameters so that they all have similar ranges and a universal value of $\hat{\sigma}$ can be used for all of them. Furthermore, we assume that the sensor noise

[1] We set $n_l = \lfloor n^{0.5} \rfloor$ as per the recommendation in [13, ch. 8].

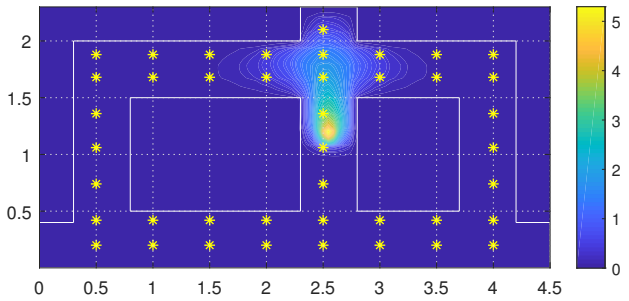

Fig. 1: True concentration field generated by the rectangular source and the corresponding lattice of $m = 42$ measurement locations denoted by the stars.

standard deviation in (2) is $\bar{\sigma} = 0.03$. We use a uniform prior for the measurement error standard deviation σ in (3), i.e., we set $\pi(\sigma) = \text{Uniform}[0, \sigma_{\max}]$ where $\sigma_{\max} \in \mathbb{R}_+$ is an upper-bound. For simplicity we assume that the m measurements required for SI, are collected over lattices that uniformly sample the domain. Unless specified otherwise, we use $m = 42$ measurements and utilize the Sensitivity Analysis (SA), presented in [8], for the selection of p^0 in the proposal distribution (6). Note that the SA method is used as an efficient approach to initialize the deterministic solver in [8]. To assess the identification accuracy, we define the error term $\text{err}_{\text{avg}} = \|p^* - \text{avg}[P]\| \,/\, \|p^*\|$ where p^* denotes the true source parameters. We define similarly the error err_{map} of the *map* estimator. Finally, we discard 10% of the samples generated by Algorithm 1 as burned-in samples; see the next section for justification.

To diagnose the convergence of the algorithm, in addition to the effective sample size (8), we use the *acceptance rate* which we define as the mean acceptance probability given as $\bar{\rho} = \text{avg}[\rho]$. Since independent sampling of multiple Markov chains in parallel is extremely efficient, we simulate n_c chains simultaneously using MATLAB parallel toolbox. These chains provide additional diagnosis tools, e.g., the shrink factor; see [13, ch. 8]. Nevertheless, caution should be taken when interpreting multi-chain samples as they might result in misleading conclusions [17].

A. Parameter Study

Our objective in this section is to examine the ability of our proposed stochastic SI algorithm to capture the desired posterior distribution $\pi(P, \Theta|z_{1:m})$ as a function of the number of measurements m, different initialization approaches, and model inaccuracies. To do so, we consider an experiment with one rectangular source with parameters $p^* = (0.2, 2.50, 2.60, 1.10, 1.20)$ creating the concentration filed given in Figure 1. We obtain a ground truth model for the concentration field by solving the AD-PDE (1) for the true values of the diffusivity and inlet velocity BC using another in-house FE code with 1230 grid points. Then, given this ground truth FE model, we build an approximate closed-form model $\mu_{\text{true}}(x)$ using the Proper Orthogonal Decomposition method; see Section III-A. To investigate the effect of model inaccuracy, we also construct a second reduced model $\mu_1(x)$ for PDE input-data different than the true ones. Finally, we compare two different initialization approaches for the selection of p^0

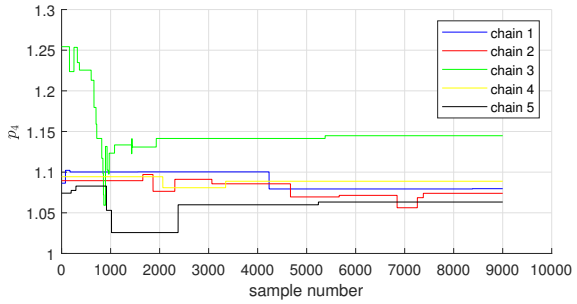

Fig. 2: Trace plot of MCMC samples for the lower-bound $p_4 = \underline{x}_2$ of the source support for simulation 3 where the true value is 1.10.

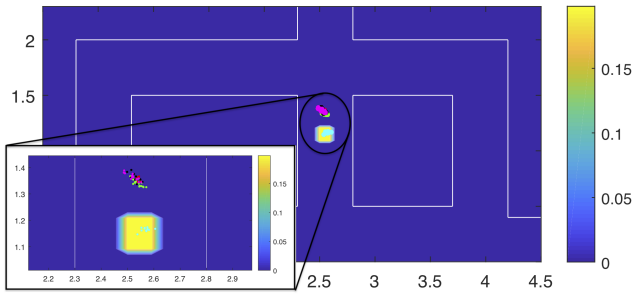

Fig. 3: Scatter plot of the source centers for MCMC samples in simulation 5. Note that only one of the chains has converged to true source support.

in the proposal distribution (6): (i) the SA method presented in [8] and, (ii) the deterministic solution obtained after solving the SI problem as per the discussions in [8]. In these case studies, the only uncertain secondary parameter is the error standard deviation in the measurement model (3), i.e., $\theta = \sigma$.

Table I shows the simulation results for different combinations of the parameters where we average the values of n_{eff} and $\bar{\rho}$ over all n_c chains. Comparing the first three rows of Table I, we see that the computation time increases as the number of samples and chains increase. Nevertheless, all runtimes are in the order of seconds (or a few minutes), which illustrates the tractability of our method. Note that the acceptance rates are generally small due to the complicated non-convex form of the admissible set A_p. Nevertheless, for simulation 1, $\bar{\rho}$ is an order of magnitude larger since for $n = 10^3$, the samples are still converging to the high probability region of A_p resulting in higher acceptance probabilities. Once they reach this region, most new samples are rejected since they have smaller likelihood than the current ones; see Algorithm 1. The fact that the samples in simulation 1 have not converged is also evident from the larger error values err_{avg} and err_{map} and standard deviation value $\text{avg}[\sigma]$. Figure 2, shows the trace plot of the lower-bound $p_4 = \underline{x}_2$ of the source support for simulation 3 which is identical to simulation 1 except that $n = 10^4$. Observe that the early samples have larger variations. Note that 10% of the initial samples are discarded (burned-in) since they correspond to low probability regions of A_p.

Next in simulation 4, we compare the effect of the initialization method used to select p^0 in the proposal distribution (6). In this simulation, SI refers to the deterministic solution obtained using the formulation presented in [8]. Note that the deterministic solution adds to the computation time but does not seem to improve the error values. Next, simulation 5 uses a smaller number of measurements $m = 28$. This results in a less accurate initialization which means that more samples n are required for the algorithm to discover high-probability regions of A_p. Figure 3 depicts the scatter plot of source centers corresponding to samples overlaid on the exact source field. Note that only one of the chains converges to the true source support. Thus, when the initialization is inaccurate, the algorithm will require larger number of samples for convergence. This shows the importance of a properly constructed proposal distribution (6). Finally, the last row in Table I corresponds to an inaccurate model that is built with wrong input data. Note that the algorithm successfully detects this and assigns a much larger measurement error value $\text{avg}[\sigma]$ to this case. Referring to Table I, notice that $\text{avg}[\sigma]$ as a point estimate of the RV σ is larger than the sensor noise $\bar{\sigma} = 0.03$ in all cases. This is expected since σ captures model errors in addition to sensor noise; see Section II.

B. Uncertainty in PDE Input Data

Next, we investigate the performance of our algorithm for uncertain PDE input-data. Specifically, we consider the case of one source, as in the previous section, and assume uncertainty in diffusivity κ and inlet velocity BC q_{in}, i.e., $\theta = (\kappa, q_{\text{in}}, \sigma)$. Since κ and q_{in} appear in the PDE input-data, we use SROMs to capture their distribution as per the discussion in Section III-B. Specifically, we use 5 samples $\pi(\kappa) \approx \{1/5; 10^{-5}, 5\times10^{-5}, 10^{-4}, 5\times10^{-4}, 10^{-3}\}$ for the constant diffusivity field and 5 samples for the inlet velocity BC as $\pi(q_{\text{in}}) \approx \{1/5; 0.001, 0.002, 0.003, 0.004, 0.005\}$, resulting in 25 models $\mu(x; p, \theta_j)$ with the prior distribution $\pi(\mu) \approx \{1/25, \mu(x; p, \theta_j)\}_{j=1}^{25}$. Note that in the absence of any prior knowledge, we have used uniform distribution over the models. Furthermore, note that the true value of the inlet velocity BC is $q_{\text{in}} = 0.025$, which is not included in the SROM samples meaning that none of the models is constructed based on the true input data.

We use $n = 5\times10^4$ samples and $n_c = 2$ chains here. Figure 4 depicts the posterior probability weights of the SROM samples for both the diffusivity field and inlet velocity BC. It can be seen that the true diffusivity value $\kappa = 10^{-4}$ has a high posterior probability. Since the true inlet velocity BC $q_{\text{in}} = 2.5\times10^{-3}$ is not included in the SROM samples, the closest value has a high probability. Also, the average standard deviation value is $\text{avg}[\sigma] = 0.231$, which is considerably high reflecting the large model error due to inaccurate input data and model reduction. It is worth mentioning that the diffusivity range used here is very large (two magnitude orders). This would seldom happen in practice.

C. Multiple Sources

Finally, in this section we consider two sources, a circular source centered at $(2.5, 0.3)$ with a radius of 0.08 and intensity of 0.2 and a rectangular source with parameters $p^*_{6:10} = (0.25, 3.85, 3.95, 1.0, 1.15)$. We use $n = 2\times10^4$ samples and $n_c = 2$ chains to explore the feasible set $A_p \subset \mathbb{R}^{10}$. Given the samples of the posterior distribution

TABLE I: Parameter study for the proposed stochastic SI Algorithm 1.

No.	parameters							results				diagnostics	
	n	n_c	m	p^0	model	$\hat{\sigma}$	$\sigma_{\max}$	err_{avg}	err_{map}	$\text{avg}[\sigma]$	time(sec)	n_{eff}	$\bar{\rho}$
1	10^3	2	42	SA	μ_{true}	0.01	0.3	0.046	0.059	0.103	8.88	23	0.021
2	10^4	2	42	SA	μ_{true}	0.01	0.3	0.019	0.025	0.035	84.7	212	0.003
3	10^4	5	42	SA	μ_{true}	0.01	0.3	0.034	0.035	0.036	226	212	0.003
4	10^4	5	42	SI	μ_{true}	0.01	0.3	0.024	0.044	0.041	236	213	0.001
5	10^4	5	28	SA	μ_{true}	0.01	0.3	0.071	0.078	0.042	216	211	0.003
6	10^4	5	42	SA	μ_1	0.01	0.3	0.23	0.22	0.29	225	211	0.002

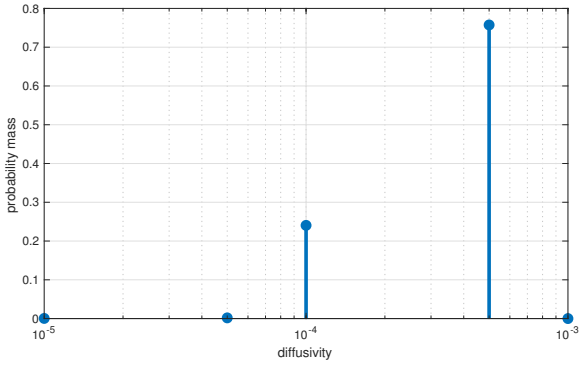

(a) constant diffusivity filed

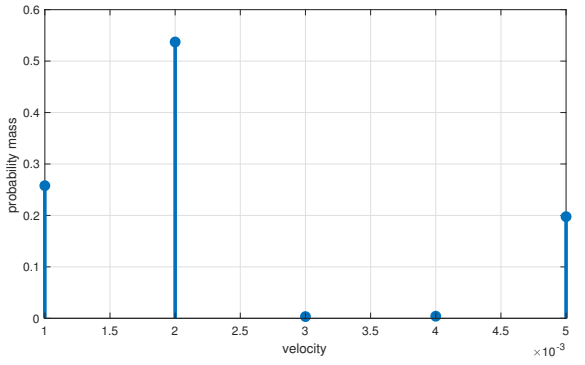

(b) inlet velocity BC

Fig. 4: The posterior probability weights of the SROM samples.

$\pi(P, \Theta|z_{1:m})$, we have $\text{avg}[P] = (0.10, 2.23, 2.56, 0.20, 0.36; 0.24, 3.82, 4.10, 1.06, 1.14)$ and $\text{map}[P] = (0.095, 2.24, 2.62, 0.22, 0.35; 0.19, 3.74, 4.08, 1.06, 1.13)$ with error values $\text{err}_{\text{avg}} = 0.040$ and $\text{err}_{\text{map}} = 0.043$. These small error values for the point estimators demonstrate that Algorithm 1 successfully approximates the posterior distribution. The standard deviation of the mean is $\text{sd}[\text{avg}[P]] = (0.01, 0.02, 0.03, 0.02, 0.01; 0.01, 0.03, 0.03, 0.02, 0.01)$ indicating high confidence in this point estimator. Moreover, the standard deviation of source parameters is $\text{sd}[P] = (0.24, 0.37, 0.70, 0.43, 0.10; 0.27, 0.54, 0.54, 0.42, 0.17)$ which is a measure of the variation of P around the mean.

V. Conclusion

In this paper we proposed the first model-based stochastic method for the identification of chemical sources that can be implemented in near real-time to solve realistic problems. Our approach relied on the Advection-Diffusion (AD) PDE to model the transport phenomenon and we utilized Markov Chain Monte Carlo sampling to obtain the posterior distribution of the source parameters considering uncertainty in the parameters of the AD-PDE and the sensor data. We presented extensive numerical experiments that demonstrate the ability of our proposed algorithm to correctly identify the desired source parameters and provide uncertainty bounds for them.

References

[1] G. Kowadlo and R. A. Russell, "Robot odor localization: a taxonomy and survey," *International Journal of Robotics Research*, vol. 27, no. 8, pp. 869–894, 2008.

[2] L. Marques, U. Nunes, and A. T. de Almeida, "Olfaction-based mobile robot navigation," *Thin solid films*, vol. 418, no. 1, pp. 51–58, 2002.

[3] R. A. Russell, A. Bab-Hadiashar, R. L. Shepherd, and G. G. Wallace, "A comparison of reactive robot chemotaxis algorithms," *Robotics and Autonomous Systems*, vol. 45, no. 2, pp. 83–97, 2003.

[4] D. Webster, K. Volyanskyy, and M. Weissburg, "Bioinspired algorithm for autonomous sensor-driven guidance in turbulent chemical plumes," *Bioinspiration & biomimetics*, vol. 7, no. 3, p. 036023, 2012.

[5] J. Matthes, L. Groll, and H. B. Keller, "Source localization by spatially distributed electronic noses for Advection and Diffusion," *IEEE Transactions on Signal Processing*, vol. 53, pp. 1711–1719, May 2005.

[6] V. Akçelik, G. Biros, O. Ghattas, K. R. Long, and B. van Bloemen Waanders, "A variational Finite Element method for source inversion for Convective-diffusive transport," *Finite Elements in Analysis and Design*, vol. 39, pp. 683–705, May 2003.

[7] R. Khodayi-mehr, W. Aquino, and M. M. Zavlanos, "Model-based sparse source identification," in *Proceedings of American Control Conference*, pp. 1818–1823, July 2015.

[8] R. Khodayi-mehr, W. Aquino, and M. M. Zavlanos, "Nonlinear reduced order source identification," in *Proceedings of American Control Conference*, pp. 6302–6307, July 2016.

[9] R. Khodayi-mehr, W. Aquino, and M. M. Zavlanos, "Model-based active source identification in complex environments," *IEEE Transactions on Robotics*, 2019.

[10] A. Jeremic and A. Nehorai, "Landmine detection and localization using chemical sensor array processing," *IEEE Transactions on Signal Processing*, vol. 48, pp. 1295–1305, May 2000.

[11] L. Calkins, R. Khodayi-mehr, W. Aquino, and M. Zavlanos, "Stochastic model-based source identification," in *Proceedings of IEEE Conference on Decision and Control*, pp. 1272–1277, 2017.

[12] B. D. Reddy, *Introductory functional analysis: with applications to boundary value problems and finite elements*, vol. 27. Springer, 2013.

[13] C. P. Robert, G. Casella, and G. Casella, *Introducing Monte Carlo methods with R*, vol. 18. Springer, 2010.

[14] C. P. Robert and G. Casella, "The Metropolis-Hastings algorithm," in *Monte Carlo Statistical Methods*, pp. 231–283, Springer, 1999.

[15] J. A. Atwell and B. B. King, "Proper Orthogonal Decomposition for reduced basis feedback controllers for parabolic equations," *Mathematical and Computer Modeling*, vol. 33, no. 1, pp. 1–19, 2001.

[16] R. Khodayi-mehr and M. M. Zavlanos, "Model-based learning of turbulent flows using mobile robots," *IEEE Transactions on Robotics*, 2018. (submitted). [Online]. Available: https://arxiv.org/pdf/1812.03894.pdf.

[17] C. Geyer, *Introduction to Markov chain Monte Carlo*. 2011.

[18] J. Martin, L. C. Wilcox, C. Burstedde, and O. Ghattas, "A stochastic newton MCMC method for large-scale statistical inverse problems with application to seismic inversion," *SIAM Scientific Computing*, vol. 34, no. 3, pp. 1460–1487, 2012.