18th International Conference on
Harmonisation within Atmospheric Dispersion Modelling for Regulatory Purposes
9-12 October 2017, Bologna, Italy

COUPLING ADJOINT AND BAYESIAN APPROACHES FOR SOURCE TERM ESTIMATION

Paul B. Westoby[1], Luca Delle Monache[2] and Daniel Silk[1]

[1]Defence Science & Technology Laboratory, Porton Down, UK
[2]National Center for Atmospheric Research, Boulder, Colorado, USA

Abstract: A range of situations exist in both the regulatory and emergency response fields in which it is desirable to be able to determine the location and nature of an emission into the atmosphere. The development of methods to address this problem is known as inverse dispersion modelling or source term estimation (STE). Given the potential value of being able to identify and characterize sources so that the hazard they pose to populations can be determined, considerable research effort has been dedicated to the development of STE algorithms.

In this work we have sought to exploit the strengths of two separate algorithms to provide a more effective overall STE system. Bayesian numerical methods can provide high resolution probabilistic source term estimates that account for sources of uncertainty in the data, model and environment. However, their real-time application is potentially limited by demanding computational requirements. Reducing the computational time for Bayesian STE is therefore critical particularly for urban scenarios, where more complex dispersion modelling is required to improve accuracy.

In contrast to Bayesian methods, adjoint-based methods are relatively quick-to-compute source term estimation algorithms. For example, the Variational Iterative Refinement Source term estimation Algorithm (VIRSA) method provides prompt approximations for release location, time and mass. As with Bayesian algorithms, VIRSA is data-driven; the parameters of interest are optimised to minimise the discrepancy between observed and simulated observations, as quantified by a (least-squares) cost function. The output of VIRSA is thus a point estimate representing the source term that provides the best fit to the data.

A novel model combination algorithm has been implemented that allows the Bayesian algorithm to take and refine VIRSA solutions in order to produce fully probabilistic estimates for the source term and hazard and benefits. Execution time benefits of the Bayesian inference are linked to accuracy of VIRSA solutions.

Key words: *Source term estimation; Bayesian inference; VIRSA; model combination; hazard prediction*.

INTRODUCTION

Hazardous substances released into the atmosphere pose both an immediate and delayed risk to human health. In the event of such a release it is necessary to predict where the material will disperse and deposit, as this will enable the first responder to undertake appropriate mitigation strategies. Atmospheric dispersion models are run forward in time to generate hazard predictions, but these require complete and precise knowledge of the release, or source term, parameters that characterise the type of release and the conditions under which it is made. In many situations, however, the source term is not known well and must be inferred from, for example, sensor measurements in the field. Source term estimation is therefore a classic inverse problem, but one in which there are often strong time and data constraints.

Bayesian approaches have been widely used to solve the source term estimation problem (e.g., Delle Monache *et al*., 2008; Robins *et al*., 2009; Yee, 2012), since they account for sources of uncertainty in the data, model and environment (Kennedy & O'Hagan, 2001; Chkrebtii *et al*. 2013). The real-time application of Bayesian methods to STE is currently limited, however, by demanding computational requirements due to expensive likelihood calculations. Conversly, adjoint-based methods (e.g., Sykes *et al*. 2008, Annunzio *et al*. 2012, Bieringer *et al*. 2015), which return point estimates for the release parameters, can be computed in quick time, but do not enable uncertainties to be reflected in resulting hazard predictions.

This paper describes an approach for combining an adjoint-based STE method with a Bayesian STE algorithm to reduce the computation time of the Bayesian inference.

Bayesian STE

A Bayesian approach to STE produces a posterior probability distribution for the source term parameters, θ, based on readings received from sensors, $\mathbf{D}$. Using Bayes' rule, the posterior distribution, indicating how likely the source term parameters are given the data, is calculated as:

$$P(\theta|\mathbf{D}) = \frac{P(\theta)P(\mathbf{D}|\theta)}{\int P(\theta)P(\mathbf{D}|\theta)\mathrm{d}\theta} \tag{1}$$

A more common representation of Bayes' theorem is:

$$p(\theta|\mathbf{D}) \propto p(\theta)p(\mathbf{D}|\theta) \tag{2}$$

where $p(\theta)$ represents the prior distribution – the assumed probability of the source term parameters before any data has been received, and $p(\mathbf{D}|\theta)$ is the likelihood distribution – a measure of how likely a particular dataset is, given a proposed source term.

Due to the numerical intractability of the normalising constant (denominator) in equation 1, Bayesian approaches to STE seek a statistical solution, rather than an analytic one. Monte Carlo methods (Gelman *et al.* 2014) are therefore employed to generate samples from the posterior distribution of source terms.

The Bayesian approach used in this paper is that of Robins, Rapley & Green (2009). In this approach, at every time instant, t, define a large collection of N weighted random samples $\left\{\theta_t^{(i)}, w_t^{(i)}\right\}$ for $i = 1, \ldots, N$, such that $w_t^{(i)} > 0$ for all i and all t, and $\sum_{i=1}^{N} w_t^{(i)} = 1$. A hypothesis at time t is denoted $\theta_t^{(i)}$ and $w_t^{(i)}$ is the associated weight, which reflects the importance of that hypothesis relative to the complete set of hypotheses.

A Markov Chain Monte Carlo (MCMC) algorithm is then adopted to generate hypotheses so that as the number of hypotheses increases and the algorithm is run for sufficient time, the empirical distribution converges asymptotically to the target posterior distribution of the parameters, θ.

Adopting a weighted sample approach for the source term estimation problem enables the data to be processed on-line, as measurements are made, under the assumption that the likelihood of θ, will factorise as:

$$p(\mathbf{D}|\theta) = \prod_i p(\mathrm{D}_i|\theta) \tag{3}$$

The weighted sample approach therefore allows for continuous updating of the posterior probability density. This ensures that whenever the user requests output, e.g. a hazard prediction, the current source term posterior probability density distribution from which samples are drawn, is the most accurate one possible.

Variational Iterative Refinement STE Algorithm

The Variational Iterative Refinement Source Term Estimation Algorithm (VIRSA; Bieringer *et al.* 2015) is a relatively quick-to-compute STE algorithm that provides a prompt approximation for release location, time and mass, as well as surface wind speed. VIRSA consists of a combination of modelling and STE systems including: a back trajectory-based source inversion method; a forward Gaussian puff dispersion model – the Second Order Closure Integrated PUFF (SCIPUFF) model (Sykes *et al.* 2008); and a variational minimization approach, which uses both a simple forward dispersion model that is a surrogate for a more complex Gaussian puff model, and a formal adjoint of this surrogate model.

In this approach, the backward trajectory method is used to provide a 'first guess' source estimate given the observations, while a variational data assimilation approach is used to refine the first guess via minimization of a cost function.

INTEGRATION OF VIRSA SOLUTIONS INTO BAYESIAN STE

If available, pre-computed point estimates could be used within Bayesian STE to guide high-dimensional MCMC sampling towards areas of high likelihood and avoid expensive dispersion model runs from unlikely regions of parameter space. Previous attempts at coupling adjoint and Bayesian methods for this purpose (e.g. Brown, Robins & Green, 2010), have encountered two main challenges. Firstly, model discrepancies exist between the STE methodologies. This means that the adjoint solution is unlikely to coincide with the maximal likelihood/*posteriori* estimate of the Bayesian methodology. Secondly, only a subset of the source term parameters inferred by the Bayesian approach described here are estimated by adjoint-based methods. Methodologies that seek to fix or initialise the source term parameters to their adjoint estimates can fail if they ignore multivariate effects between non-estimated and estimated source term parameters (Silk et al. 2015).

Approach

The approach described here overcomes these challenges by viewing the VIRSA solution as an additional observation. An associated likelihood can then be defined that accounts for the expected disagreement between the STE procedures, and naturally allows the partial 'observation' of the source term to be fused into the Bayesian inference. Implemented naively, this approach would violate a fundamental rule of Bayesian statistics; that data should only be used once. The uncertainty in the source term posterior distribution would be artificially reduced by incorrectly considering the VIRSA solution and the sensor data observations as independent sources of information. The approach described here, therefore, also considers how the information contained within the VIRSA solution can be removed from the posterior distribution. The VIRSA solution effectively acts as a temporary 'crutch' to the Bayesian inference, guiding the sampling until adequate information has been assimilated via the actual sensor data.

Conditioning on the VIRSA solution

A VIRSA solution consists of point estimates for: release mass, *m;* release latitude, *x;* release longitude, *y;* and release time, *t*. Within the Bayesian STE approach, this solution is treated as a direct, partial observation of the source term, *θ*. The partial observation is denoted $v = (x_v, y_v, m_v, t_v)$, and can be processed similarly to other data points by defining an associated likelihood model,

$$p(v|\theta) = \phi(x; x_v, \sigma_x^2)\phi(y; y_v, \sigma_y^2)\phi_c(m; m_v, \sigma_m^2)\phi(t; t_v, \sigma_t^2) \quad (4)$$

where $\phi(\cdot\,; \mu, \sigma^2)$ is a Gaussian density with mean μ and variance σ^2, and $\phi_c(m; m_v, \sigma_m^2)$ is a Gaussian density clipped at $m = 0$, with m_v and σ_m^2 the mean and variance of the underlying Gaussian distribution.

Assuming a multi-step Metropolis-Hastings algorithm, as described in Robins *et al.* 2009, a VIRSA event is then used to influence the acceptance or rejection of new hypotheses, θ^*, through the likelihood ratio in the Metropolis-Hastings. The new hypothesis, θ^*, is accepted if:

$$U(0,1) \le \left(\frac{\prod_{j=1}^{N} P(D_j|\theta^*)}{\prod_{j=1}^{N} P(D_j|\theta_t^i)}\right)\left(\frac{p(v|\theta^*)}{p(v|\theta_t^i)}\right)^{\frac{1}{T_v}} \quad (5)$$

where T_v is a tempering parameter that controls the influence of the VIRSA event. High values of T_v equate to low influence, with $\left(\frac{p(v|\theta^*)}{p(v|\theta_t^i)}\right)^{\frac{1}{T_v}} \to 1$ as $T_v \to \infty$ for all θ^* and θ_t^i.

Adaptive control of VIRSA influence

The dependence of Bayesian STE sampling on the VIRSA solution is controlled by T_v. It is desirable that T_v is low at the start of an inference, when the Bayesian STE inference engine has processed few data points, and high near the end of an inference, when the majority of the sensor information has been incorporated into the posterior distribution. Given that sensor measurements may vary in how informative they are about the source term, control over T_v is based upon an information measure on the posterior distribution, rather than the number of processed data points.

The control algorithm is defined as follows:

1. Initialise T_v=1;

2. After the MCMC sampler converges given the VIRSA solution, define V_0 as the volume[1] of the current marginal posterior distribution $p(x, y, m, t|\mathbf{D}, v)$ where $\mathbf{D}$ represents the data points that have been processed so far;
3. When convergence is achieved for each further data point, if the new volume of the marginal distribution $V < V_0$, set $T_v = 2T_v$ and update the hypothesis weights as described below.

Updating importance weights

In between stages of MCMC sampling, the current (VIRSA-assisted) posterior distribution $p(\theta|\mathbf{D}, v, T_v)$ is represented by a set of weighted hypotheses, (θ_i, w_i). When T_v is updated to T_v' (reducing the influence of VIRSA), the importance weights are updated as:

$$w_i' = w_i \frac{p(v|\theta)^{\frac{1}{T_v'}}}{p(v|\theta)^{\frac{1}{T_v}}} \tag{6}$$

and the hypothesis likelihood updated as:

$$p(\mathbf{D}, v, T_v'|\theta) = p(\mathbf{D}, v, T_v|\theta) \frac{p(v|\theta)^{\frac{1}{T_v'}}}{p(v|\theta)^{\frac{1}{T_v}}} \tag{7}$$

Removal of VIRSA information

Removal of the VIRSA event is achieved by deleting the VIRSA event from the list of data points in the validity window, and updating the importance weights as:

$$w_i' = \frac{w_i}{p(v|\theta)^{\frac{1}{T_v}}} \tag{8}$$

and updating the likelihoods as:

$$p(\mathbf{D}|\theta) = \frac{p(\mathbf{D}, v, T_v|\theta)}{p(v|\theta)^{\frac{1}{T_v}}} \tag{9}$$

This recovers an importance sample from the posterior distribution, $p(\theta|\mathbf{D})$. However, as with processing or removing sensor data, the Bayesian STE should be allowed to re-converge before outputs are obtained. Complete removal of the VIRSA is triggered if one of the following occurs:

1. The user requests removal;
2. There are no further data points to process and an output has been requested.

DISCUSSION

This paper has presented a novel approach for utilising adjoint-based STE solutions within a more rigorous Bayesian STE inference engine. One of the original motivations for developing this approach was to facilitate more efficient data processing. In the Bayesian STE method described above, data is processed one at a time, with convergence of the posterior distribution achieved prior to processing the next data point. Attempting to process all the data in one go, however, could potentially lead to the inference getting trapped in local maxima. Coupling the Bayesian approach with an adjoint-based method, provides the potential to move more quickly to the appropriate parameter space, or provide the ability to overcome situations where the Bayesian inference has become stuck in regions of narrow density.

A prototype implementation of this method has been produced and testing is ongoing. Early indications are that the success of combining the adjoint and Bayesian approaches for STE – measured in terms of runtime and accuracy – is inevitably highly dependent on the accuracy of the adjoint solution and the point at which it is injected into the Bayesian process. If the adjoint solution is far from the truth, then it will pull the Bayesian inference away from the truth. This will inhibit the Bayesian process, as it will need to wait for the adjoint 'event' to become exhausted (i.e. $T_v \rightarrow \infty$) or removed by the user, before it

[1] Used to assess convergence.

can move to the correct regions of parameter space. This does, however, illustrate the importance of using the adjoint solution as only a temporary input into the Bayesian inference.

REFERENCES

Annunzio, A. J., Young, G. S., Haupt, S. E., 2012: Utilizing state estimation to determine the source location for a contaminant. *Atmos. Environ.* 46, 580-589.

Bieringer, P.E., Rodriguez, L.M., Vandenberghe, F., Hurst, J.G., Bieberbach, G., Sykes, I., Hannan, J.R., Zaragoza, J., Fry, R.N., 2015: Automated source term and wind parameter estimation for atmospheric transport and dispersion applications. *Atmos. Environ.*, 122, 206–219.

Chkrebtii, O., Campbell, D. A., Girolami M. A., and Calderhead, B., 2013: Bayesian uncertainty quantification for differential equations. arXiv:1306.2365v1 [stat.ME] 10 Jun 2013.

Delle Monache, L., Lundquist, J. K., Kosovic, B., Johannesson, G., Dyer, K. M., Aines, R. D., Chow, F. K., Belles, R. D., Hanley, W. G., Larsen, S. C., Loosmore, G. A., Nitao, J. J., Sugiyama, G. A., and Vogt, P. J., 2008: Bayesian inference and Markov Chain Monte Carlo sampling to reconstruct a contaminant source at continental scale. *J. Appl. Meteorol. Clim*, 47, 2600-2613.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B., 2014: Bayesian Data Analysis, Volume 2, Taylor & Francis.

Kennedy, M.C. and O'Hagan, A., 2001: Bayesian calibration of computer models. *J. Royal Statistical Society*. Series B, 63(3), 425-464.

Robins, P., Rapley, V. E., and Green, N., 2009: Realtime sequential inference of static parameters with expensive likelihood calculations. *J. Appl. Stats.,* 58(5), 641-662.

Yee, E. 2012: Inverse Dispersion for an Unknown Number of Sources: Model Selection and Uncertainty Analysis. *ISRN Appl. Math.* 2012, Article ID 465320.

Brown, G., Robins, P., & Green, N., 2010: Feasibility study into the amalgamation of the SCIPUFF Adjoint and Monte Carlo Bayesian Data Fusion source term estimation algorithms, *Dstl Technical Report* DSTL/CR40382.

Silk, D., Stenning, M., Westoby, P., Herring, S. J., Ahmad, S., Webb, A., 2015: Development of an Urban Source Term Estimation Capability: Final Report OY2, *Dstl Technical Report* DSTL/CR90902.

Sykes, R. I., Parker, S. F., Henn, D. S., Chowdhury, B., 2008: SCIPUFF Version 2.4 Technical Documentation.