

Bayesian estimation of airborne fugitive emissions using a Gaussian plume model

Bamdad Hosseini^{a,*}, John M. Stockie^a

^a*Department of Mathematics, Simon Fraser University, 8888 University Drive, Burnaby, BC, V5A 1S6, Canada*

Abstract

A new method is proposed for estimating the rate of fugitive emissions of particulate matter from multiple time-dependent sources via measurements of deposition and concentration. We cast this source inversion problem within the Bayesian framework, and use a forward model based on a Gaussian plume solution. We present three alternate models for constructing the prior distribution on the emission rates as functions of time. Next, we present an industrial case study in which our framework is applied to estimate the rate of fugitive emissions of lead particulates from a smelter in Trail, British Columbia, Canada. The Bayesian framework not only provides an approximate solution to the inverse problem, but also quantifies the uncertainty in the solution. Using this information we perform an uncertainty propagation study in order to assess the impact of the estimated sources on the area surrounding the industrial site.

Keywords: Pollutant dispersion, Gaussian plume, Particle deposition, Inverse problem, Bayesian estimation

PACS: 92.60.Sz, 93.85.Bc

2010 MSC: 65F20, 65M06, 65M32, 76Rxx, 86A10

1. Introduction

Dispersion of airborne pollutants emitted from anthropogenic sources and their effect on the surrounding environment have been a growing source of concern over the past several decades. Both primary polluters and government monitoring agencies dedicate significant resources to tracking and controlling the release of particulate emissions from industrial operations. Atmospheric dispersion modelling, which is the study of mathematical models and numerical algorithms for simulating processes involved in dispersion of pollutants from a known source, is a vital tool for monitoring of existing industrial operations as well as assessing the potential risk and impact of future operations. Many dispersion models address the situation where a pollutant source has already been identified and the source emission rate is known approximately, and many industry standard software packages such as AERMOD [9] and CALPUFF [38] are already available to solve this problem.

In many practical settings the main problem of interest is not to determine the impact of known sources but rather to solve the source identification problem, which refers to determining the emission properties and possibly also locations for a collection of unknown sources. Inverse source identification is particularly prominent in the study of fugitive sources, which are particulate or gaseous emissions that derive from leaks or other unintended releases from building windows or vents, entrainment from debris piles, or operations of trucks and loading equipment. In these situations, it is usually not possible to obtain direct measurements of the fugitive source emissions by installing sensors; this is in contrast with emissions from stacks and exhaust vents where such measurements are routine. Nevertheless, it is often still possible to take *indirect*

*Corresponding author

Email addresses: bhossein@sfu.ca (Bamdad Hosseini), stockie@math.sfu.ca (John M. Stockie)

URL: <http://www.math.sfu.ca/~stockie> (John M. Stockie)

measurements of fugitive emissions, for example by measuring concentration of a pollutant at a remote location some distance from the source. In this case, the challenge is to estimate the rate of emissions based on indirect measurements, which can be posed as an inverse problem [19, 20].

Atmospheric dispersion modelling is a well-developed area of research, and for a comprehensive overview we refer the reader to the work of Zlatev [53] or Dimov et al. [12]. The use of partial differential equation (PDE) models based on the advection-diffusion equation for modelling short-range transport of pollutants dates back to the work of Taylor [45], Roberts [36] and Sutton [44]. In simple scenarios involving constant emissions, steady-state transport, point or line sources, or flat topography, it is possible to derive approximate analytical solutions to the governing equations [3, 31, 33, 39]. These analytical solutions, referred to collectively as Gaussian plume solutions, have the distinct advantage of being simple and relatively cheap to compute, and consequently form the basis of many standard monitoring tools (including AERMOD and CALPUFF). For more realistic situations where one is interested in incorporating effects such as topographical variations or more complex time-varying wind patterns, the only recourse is to solve the governing PDEs directly, typically using approaches based on finite volume [16, 17], finite difference [23, 30] or finite element schemes [1].

In comparison with the atmospheric dispersion models just described for solving the forward problem, the source inversion problem has attracted less attention in the literature. The monograph by Vogel [50] is a notable reference that lays out the general mathematical theory for inverse problems as well as common solution approaches. More specific to the context of atmospheric transport, the work of Rao [34] and Enting [13] provides an overview of methods for solving the source inversion problem. For example, inversion of fugitive sources was considered recently in [37] where the authors coupled a finite volume solver within an optimization algorithm in order to estimate source emission rates. Another powerful and promising approach to the solution of inverse problems is based on Bayesian methods, whose mathematical formulation is well-described by Kaipio and Somersalo [21] and Stuart [43]. Although some attempts have been made to apply Bayesian methods in the context of atmospheric dispersion problems, their application is much less extensive than other approaches. Some examples include Huang et al. [18] who employ a Bayesian method for identifying the location and emission rate for a single point source by incorporating a Gaussian puff solution, while Keats et al. [22] use Bayesian inference to identify emissions in a more complex urban environment. In both cases, the authors use an adjoint approach to efficiently solve the advection-diffusion PDE and evaluate the likelihood function using a Markov Chain Monte Carlo algorithm. Senocak et al. [40] and Wade and Senocak [51] use Bayesian inference along with a Gaussian plume model in order to reconstruct multiple sources in an atmospheric contamination scenario.

In this paper, we aim to develop an accurate and efficient Bayesian approach for solving the inverse source identification problem. We aim to estimate fugitive particulate emissions from various areas of an industrial site based on measurements of contaminant concentration and particulate mass deposited at a distance from the suspected sources. We consider a scenario wherein material from fugitive sources is dispersed by the wind and then deposited on the ground due to a combination of diffusive transport and gravitational settling. Some sources may reasonably be approximated as constant in time, but we are particularly interested in the study of time-dependent sources arising for example from dust entrained during loading operations that are performed on a rotating shift schedule. In order to monitor emissions in such a scenario, various types of measurements are typically performed at fixed locations in the vicinity of the known or suspected sources. We are particularly interested in two classes of measurements, deriving from either total accumulated deposition of particles over a long time period (on the order of one month) or short-time averaged concentrations (taken over a period of one hour, which can be considered essentially instantaneous in comparison with long-term measurements). We utilize a Gaussian plume model for short-range dispersion of pollutants and incorporate this model within the Bayesian framework for solution of inverse problems. We split the industrial site into a number of areas that are suspected to contain the most important fugitive sources. The Bayesian framework provides a natural setting for estimating emission rates and also quantifying the uncertainty associated with the estimates. This study was motivated by a collaboration with Teck Resources Ltd., in which we studied particulate emissions from a lead-zinc smelter located in Trail, British Columbia, Canada [16, 27].

In contrast with some other studies, we do not consider the problem of determining either the number

or location of sources. Instead, we consider a given number of areas corresponding to suspected fugitive emission sources, and approximate each area source by a single point source located at the area centroid; however, we do allow the emission rate for each source to vary as a continuous function of time. We also incorporate multiple measurement types and develop a unified framework in which the forward model relates the entire measurement data set to the emission rates. The main challenge we encounter is in terms of the low quality of the measured data that derives from two main features: first, the most abundant measurements are from dust-fall jars, which measure only monthly accumulated deposition and so unfortunately provide no information about short-time variations in particulate emission; and second, although we do have access to a few real-time measurement devices, these sensors provide useful data only when the wind blows in the direction linking them to the sources of interest. The main advantage of our Bayesian framework is its ability to obtain a solution even when the measured data is of relatively low quality and the problem is severely under-determined. Another benefit of our approach is that we obtain an estimate of the accompanying uncertainty in the solution. Finally, we stress that our framework can easily be extended to deal with a general class of atmospheric dispersion problems including applications such as seed or odour dispersal [7, 26], natural disasters such as volcanic eruptions [48], and nuclear or chemical accidents [29].

The remainder of this article is organized as follows. In Section 2 we present the forward model which is based on a Gaussian plume solution. In Section 3 we develop the Bayesian framework for solving the inverse problem by considering three different instances of the inverse problem corresponding to different models for the prior distribution. Section 4 is dedicated to an industrial case study in which we apply our framework using a physical dataset in order to estimate sources of fugitive emissions from a lead-zinc smelter. After solving the inverse problem we also assess the impact of the estimated sources by propagating the solution of the Bayesian inverse problem through the forward model.

2. The forward problem

Dispersion and deposition of airborne particulates can be described mathematically using the advection–diffusion equation

$$c_t + \mathbf{u} \cdot \nabla c - \nabla \cdot (\mathbf{K} \nabla c) = Q, \quad (1)$$

which is a partial differential equation for the unknown particle concentration, $c(\mathbf{x}, t)$ [kg m^{-3}]. The other quantities appearing in this equation are the wind advection velocity $\mathbf{u}(\mathbf{x}, t)$ [m s^{-1}], diffusion tensor $\mathbf{K}(\mathbf{x}, t) := \text{diag}(k_x, k_y, k_z)(\mathbf{x}, t)$ [$\text{m}^2 \text{s}^{-1}$], and emission source term $Q(\mathbf{x}, t)$ [$\text{kg m}^{-3} \text{s}^{-1}$], which are all assumed to be given functions. We denote the spatial coordinates by $\mathbf{x} = (x, y, z)^T$ [m] and look for solutions of (1) defined on $\Omega \times [0, T]$, where the spatial variable $\mathbf{x} \in \Omega = \{-\infty \leq x, y \leq \infty, 0 \leq z \leq \infty\}$ is in the half-space (assuming the ground is flat and located along $z = 0$), and the time interval of interest is $t \in [0, T]$. We impose $c(\mathbf{x}, 0) = 0$ so that initially there is no contaminant within the domain, and specify the consistent far-field boundary conditions $c(\mathbf{x}, t) \rightarrow 0$ as $|\mathbf{x}| \rightarrow \infty$. At ground level we apply a mixed (Robin) boundary condition for the deposition flux

$$\left(K_z \frac{\partial c}{\partial z} + W_{\text{set}} c \right) \Big|_{z=0} = W_{\text{dep}} c|_{z=0}, \quad (2)$$

which is commonly applied for capturing deposition of airborne particulates onto the ground surface. Here, W_{dep} [m s^{-1}] is the deposition velocity [39], which is assumed to be constant and dependent only on the individual particles under consideration. There is a corresponding settling velocity for each particle type, W_{set} [m s^{-1}], which is also constant and is approximated using the Stokes law for spherical particles

$$W_{\text{set}} = \frac{\rho g d_p^2}{18\mu}, \quad (3)$$

where ρ [kg m^{-3}] is the mass density of the particle, d_p [m] is particle diameter, $\mu = 1.8 \times 10^{-5} \text{ kg m}^{-1} \text{s}^{-1}$ is viscosity of air, and $g = 9.8 \text{ m s}^{-2}$ is gravitational acceleration. Since wind measurements are typically only available at or near ground level, any study of atmospheric particulate transport requires choosing

some appropriate model for the wind velocity \mathbf{u} , such as a power-law or logarithmic profile for the vertical variation [3, 39]. Likewise, many models are available to represent the diffusion coefficients \mathbf{K} as functions of height and other solution variables, such as the Monin-Obukhov model [39].

Once all parameters are specified, a common approach to solving (1) is by discretizing Ω on a finite computational domain, imposing some suitable artificial boundary conditions, and then applying a numerical method that is capable of handling variable-coefficient parabolic partial differential equations. Examples of commonly-used algorithms for atmospheric dispersion models include finite volume, finite element and spectral methods [1, 8, 17, 25, 28]. However, in this study our interest is in solving inverse problems, which requires many evaluations of the solution of (1) and hence a direct application of any such algorithm is often too expensive. Instead, we will make use of an approximate analytical solution of (1) known as the *Gaussian plume solution*. Methods based on the Gaussian plume approximation are widespread in the literature, as well as in both commercial and public domain software used for regulatory and monitoring purposes [2, 15, 39, 49].

2.1. Gaussian plume solution

We begin by imposing several simplifying assumptions on the problem that allow us to use a large class of Gaussian plume solutions:

Assumptions

- A1. *Wind velocity takes the form $\mathbf{u}(\mathbf{x}, t) = (u_x(\mathbf{x}, t), u_y(\mathbf{x}, t), -W_{set})$ so that aside from a small vertical settling component, the wind blows only in the horizontal direction.*
- A2. *Variations in topography are negligible, so that the ground may be treated as a horizontal plane.*
- A3. *Only dry deposition occurs, and we neglect any enhancement in deposition due to rainfall or other events.*
- A4. *The transient behavior of contaminant plumes can be neglected, and the solution is assumed to be in a quasi-steady state during each time interval in which wind measurements are taken.*
- A5. *The site of interest can be divided into a number of disjoint area sources, and the emissions within each area are well-approximated by a single point source located at the area centroid.*

The first three Assumptions A1–A3 are imposed due to limits on the availability of data. For example, wind data are often measured at only a few meteorological stations, and typically only the horizontal wind components are available. Variations in topography are often difficult to measure accurately and, even more importantly, it is difficult (or impossible) to reconstruct a wind field over variable topography when wind measurements are sparse. Finally, the theory of wet deposition introduces significant complications and requires meteorological data that are simply not available in many applications. Assumption A4 is imposed so that a Gaussian plume approximation for (1) can be employed within each time step of a simulation. In effect, the Gaussian plume solution is the Green’s function of the steady state advection-diffusion equation for a specific choice of wind profile and diffusion tensor [39, 42]. Assumption A5 is imposed to allow treatment of a general class of sources, including well-defined areas of the site over which emissions are distributed (such as large debris piles), and also sources whose location cannot be accurately determined (such as large buildings having many open vents, windows, etc.). Assumption A5 also permits us to write the source term Q as a linear superposition of singular point sources

$$Q(\mathbf{x}, t) = \sum_{i=1}^{N_s} q_i(t) \delta(\mathbf{x} - \mathbf{x}_i), \quad (4)$$

where $q_i(t)$ is the (possibly time-varying) total emission rate from the i th area, \mathbf{x}_i is the location of the area centroid, N_s is the total number of sites, and $\delta(\mathbf{x})$ represents the three-dimensional delta distribution. Owing to the linearity of both Equation (1) and the source term, we may consider separately the contribution to the concentration field $c_i(\mathbf{x}, t)$ arising from each individual point source i , and then write $c = \sum_{i=1}^{N_s} c_i$.

We also consider a scenario where wind measurements are available at a sequence of times t_j , and each interval $[t_j, t_{j+1}]$ is chosen so that Assumption A4 is satisfied; in practice, taking a constant time interval of 10 minutes is reasonable [15, 27]. We then introduce transformed spatial coordinates $\tilde{\mathbf{x}}_{i,j} = (\tilde{x}, \tilde{y}, \tilde{z})$ with

$$\tilde{\mathbf{x}}_{i,j} = \mathbf{R}_j(\mathbf{x} - \mathbf{x}_i), \quad \mathbf{R}_j = \begin{bmatrix} \cos(\theta_j) & -\sin(\theta_j) & 0 \\ \sin(\theta_j) & \cos(\theta_j) & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (5)$$

and $\theta_j = \tan^{-1}(u_y(t_j)/u_x(t_j))$. Note that $\tilde{\mathbf{x}}_{i,j}$ represents a rotation of \mathbf{x} so that the \tilde{x} -axis points in direction of the horizontal wind vector at time t_j , composed with a translation that shifts the i th source to the origin. After this transformation, we may apply the Gaussian plume solution of Ermak [14, 42] for a point source located at the origin with wind speed $U = (u_x(t_j)^2 + u_y(t_j)^2)^{1/2}$ directed along the \tilde{x} -axis to obtain

$$c_i(\tilde{\mathbf{x}}_{i,j}, t) = \frac{q_i(t_j)}{2\pi U \sigma_{\tilde{y}} \sigma_{\tilde{z}}} \exp\left(-\frac{\tilde{y}^2}{2\sigma_{\tilde{y}}^2} - \frac{W_{\text{set}} \tilde{z}}{2K_{\tilde{z}}} - \frac{W_{\text{set}}^2 \sigma_{\tilde{z}}^2}{8K_{\tilde{z}}^2}\right) \times \left[\exp\left(-\frac{\tilde{z}^2}{2\sigma_{\tilde{z}}^2}\right) + \exp\left(-\frac{(\tilde{z} + 2z_i)^2}{2\sigma_{\tilde{z}}^2}\right) - \exp\left(\frac{W_o(\tilde{z} + 2z_i)}{K_{\tilde{z}}} + \frac{W_o^2 \sigma_{\tilde{z}}^2}{2K_{\tilde{z}}^2}\right) \operatorname{erf}\left(\frac{W_o \sigma_{\tilde{z}}}{\sqrt{2}K_{\tilde{z}}} + \frac{\tilde{z} + 2z_i}{\sqrt{2}\sigma_{\tilde{z}}}\right) \right]. \quad (6)$$

Here, $W_o := W_{\text{dep}} - \frac{1}{2}W_{\text{set}}$, the time interval is $t \in [t_j, t_{j+1}]$, and the standard deviations of concentration are given by

$$\sigma_{\tilde{y},\tilde{z}}^2(\tilde{x}) = \frac{2}{U} \int_0^{\tilde{x}} K_{\tilde{y},\tilde{z}}(s) ds, \quad (7)$$

in terms of the diffusion coefficients $K_{\tilde{y},\tilde{z}}$ in the $\tilde{\mathbf{x}}$ coordinates. Note that the standard deviations $\sigma_{\tilde{y},\tilde{z}}$ are assumed to be functions of downwind distance \tilde{x} , and in practice are much easier to measure than $K_{\tilde{y},\tilde{z}}$. Indeed, many formulas have been proposed in the literature for $\sigma_{\tilde{y},\tilde{z}}$ based on theory and experiment, and we will use one of the more common parameterizations due to Briggs [6] with

$$\sigma(\tilde{x}) = a\tilde{x}(1 + b\tilde{x})^{-c}. \quad (8)$$

Values of the parameters a , b and c are listed in Table 1 for different Pasquill atmospheric stability classes.

Table 1: Value of parameters in (8) for different Pasquill atmospheric stability classes [39].

Stability class	$\sigma_{\tilde{y}}$			$\sigma_{\tilde{z}}$		
	a	b	c	a	b	c
A	0.22	1.0e-4	0.50	0.20	0.0	0.0
B	1.60	1.0e-4	0.50	1.2	0.0	0.0
C	0.11	1.0e-4	0.50	0.08	2.0e-4	0.5
D	0.08	1.0e-4	0.50	0.06	1.5e-3	0.5
E	0.06	1.0e-4	0.50	0.03	3.0e-4	1.0
F	0.04	1.0e-4	0.50	0.016	3.0e-4	1.0

Note that equation (6) remains linear in the emission rates $q_i(t)$ so that we can simply rotate the Ermak solution back into the original \mathbf{x} -coordinate system to obtain

$$c(\mathbf{x}, t) = \sum_{i=1}^{N_s} q_i(t_j) \mathcal{G}_i(\mathbf{x}; \mathbf{u}(\mathbf{x}, t_j), \mathbf{x}_i, W_{\text{set}}, W_{\text{dep}}, \sigma_{\tilde{x}}, \sigma_{\tilde{y}}), \quad t \in [t_j, t_{j+1}), \quad (9)$$

for suitably defined functions \mathcal{G}_i . From this point on, we will suppress the dependence of \mathcal{G}_i on all parameters other than \mathbf{x} in order to obtain a cleaner notation. This is also convenient because for a given set of wind data and other parameters, the \mathcal{G}_i are independent of $q_i(t_j)$.

Equation (9) concludes our approximation of the solution to the forward problem, which we use to compute the concentration at any given time based on knowledge of emission rates and other parameters. We have so far assumed that t_j are the times at which the wind velocity is measured. However, we can easily generalize to the situation where wind, concentration and emission rates are at different times by making use of either interpolation or averaging. We will return to this issue in Section 4 and construct an arbitrary time mesh by first fitting a Gaussian process [52] to the wind data and then evaluating the mean of this process on a discrete time grid.

2.2. Modelling observations

We next describe how to incorporate various types of measurements of deposition or concentration that are commonly encountered in practical applications, aiming for a general framework that is capable of handling both short-time (instantaneous) and long-time (accumulated or integrated) sensor measurements. For this purpose, we construct a linear operator that represents the transformation between unknown emission sources $q_i(t_j)$ and known deposition measurements.

Suppose that a total of N_r sensors are located throughout the domain in the vicinity of the sources of interest, and let $\bar{\mathbf{x}}_k$ for $k = 1, \dots, N_r$ represent the sensor locations. We let $d_{\ell,k}$ represent the corresponding measurement at $\bar{\mathbf{x}}_k$, where the subscript $\ell = 1, \dots, m_k$ refers to data from the k -th sensor corresponding to the ℓ -th time interval. We assume that each measurement can be represented as a time integral

$$d_{\ell,k} = \int_0^T f_{\ell,k}(t) c(\bar{\mathbf{x}}_k, t) dt, \quad (10)$$

where the $f_{\ell,k}(t)$ are “window functions” that pick out the corresponding time interval during which each sensor actively measures the concentration or deposition. A different kernel $f_{\ell,k}$ must be chosen for each measurement device, and these functions are discussed in the next two subsections for the dust-fall jars and real-time measurement devices used in this study. In practice, one needs to approximate the above integrals numerically, for which we propose using a time discretization based on a uniform grid $\{t_i\}_{i=1}^{N_T}$ that coincides with the time intervals used for computing concentration in (9). We then employ a simple one-sided quadrature rule to approximate (10).

Collecting all sensor measurements into a vector

$$\mathbf{d}_k := (d_{1,k}, d_{2,k}, \dots, d_{m_k,k}) \in \mathbb{R}^{m_k}, \quad (11)$$

we then write a matrix representation that connects concentrations to sensor measurements:

$$\mathbf{d}_k = \mathbf{M}_k \mathbf{c}_k \quad \text{where} \quad \mathbf{c}_k := (c(\bar{\mathbf{x}}_k, t_1), \dots, c(\bar{\mathbf{x}}_k, t_{N_T}))^T. \quad (12)$$

Exploiting the linearity of (9), we may then define a block diagonal matrix form of the functions \mathcal{G}_i

$$\mathbf{G}_k := \mathbf{diag}(\mathbf{g}_{k,1}, \mathbf{g}_{k,2}, \dots, \mathbf{g}_{k,N_T}) \quad \text{where} \quad \mathbf{g}_{k,j} := (\mathcal{G}_1(\bar{\mathbf{x}}_k; t_j), \dots, \mathcal{G}_{N_s}(\bar{\mathbf{x}}_k; t_j)), \quad (13)$$

which in turn allows us to write

$$\mathbf{d}_k = \mathbf{M}_k \mathbf{G}_k \mathbf{q} \quad \text{where} \quad \mathbf{q} := (q_1(t_1), \dots, q_{N_s}(t_1), q_1(t_2), \dots, q_{N_s}(t_2), \dots, q_1(t_{N_T}), \dots, q_{N_s}(t_{N_T}))^T. \quad (14)$$

We may then concatenate the \mathbf{d}_k vectors for all measurement devices, and similarly for the \mathbf{M}_k and \mathbf{G}_k matrices, which allows us to rewrite (14) in the compact form

$$\mathbf{d} = \mathbf{M} \mathbf{G} \mathbf{q} =: \mathbf{F} \mathbf{q}, \quad (15)$$

where \mathbf{d} is a long vector containing all available data from measurement devices, \mathbf{q} is the discretization of $q_i(t)$ functions in time, and we refer to $\mathbf{F} = \mathbf{M} \mathbf{G}$ as the *observation map* or *matrix*. In the next two sections, we define the window functions $f_{\ell,k}$ appearing in (10) for the two specific types of measurement devices used in this study.

2.2.1. Dust-fall jars

Dust-fall jars provide a cheap and convenient means of measuring long-term deposition of solid particulate matter such as the lead and zinc oxides that are of particular concern here. These jars are simply cylindrical plastic containers that are filled with water and left out in the open, usually on an elevated platform, for an extended time period (in our case, for one month). Sometimes a mesh is placed on top of the jar to prevent large pieces of unwanted material from contaminating the sample (e.g., leaves, garbage or insects). At the end of a measurement period, the jars are sent to a laboratory for chemical analysis, yielding a list of particles of interest and their total mass. Each dust-fall jar provides a single data point for every sampling period, which corresponds to taking $\ell = \{1\}$ in (10). Suppose that each jar has cross-sectional area A [m²] and recall from (2) that the ground-level deposition flux is given by $W_{\text{dep}}c(\mathbf{x}, t)$. Then a dust-fall jar can be treated within the general framework developed above by taking $f_{\ell,k} = AW_{\text{dep}}\mathbf{1}_{(0,T]}$ where $\mathbf{1}_{(0,T]}$ is the indicator function for the entire time interval of interest.

2.2.2. Real-time measurement devices

The other major class of measurement devices are real-time sensors such as the Xact ambient metals monitor [10] or Andersen high-volume air sampler [46] that measure average concentration of particulates over a shorter time period ranging from minutes to hours. These devices are often operated automatically according to a preset schedule, coming on-line and taking measurements at a set of specified times, and otherwise remaining inactive. Suppose that such a sensor is scheduled to initiate measurements at a sequence of times $\{\tau_1, \tau_2, \dots, \tau_p\}$ and that each measurement of average concentration is taken over a time period of length $\Delta\tau$. Then we can write $f_{\ell,k} = \frac{1}{\Delta\tau}\mathbf{1}_{(\tau_\ell, \tau_\ell + \Delta\tau]}$ with $\ell = \{1, 2, \dots, p\}$.

3. The inverse problem

In the previous section, we considered the forward problem and derived an explicit form for the linear mapping that takes as input the emission rates from all sources and returns as output the ground-level deposition at a number of specified sensor locations. Our actual aim is to solve the inverse problem, which is equivalent to inverting this linear mapping and corresponds to taking a set of sensor observations at given locations and estimating the corresponding emission rates. However, before getting into the details we first need to address the issue of noise due to measurement errors, which can have a major impact on the solution to the inverse problem.

Assuming that real measurements $\tilde{\mathbf{d}}$ are actually noisy realizations of the predictions of our model, we need to replace Equation (15) with

$$\mathbf{d} = \mathbf{F}\mathbf{q} + \boldsymbol{\epsilon}, \quad (16)$$

where $\boldsymbol{\epsilon}$ is a vector of independent, additive noise. Throughout this paper we will assume that the noise is a zero-mean multivariate Gaussian random variable, $\boldsymbol{\epsilon} \sim N(0, \boldsymbol{\Sigma})$. Note that the covariance operator $\boldsymbol{\Sigma}$ is taken to be diagonal owing to the independence assumption; however, the diagonal elements may still vary between data points owing to differences in the accuracy of each sensor.

Our approach to solving the source inversion problem with additive noise will proceed in three steps, corresponding to three problems that impose progressively more realistic constraints on the solution and hence become increasingly more challenging. The main reason for using this approach is that the data available for the problem of interest is actually quite sparse (which we describe in Section 4). Therefore, the prior knowledge of the solution (modelled via the prior distribution defined below) has a strong influence on the final solution. In order to come up with a good prior distribution, we first employ the maximum likelihood (MLE) or posterior mean estimate for simpler problems in order to construct a prior distribution for the next, more complex, problem.

3.1. Constant emissions

Our first step in solving the source inversion problem assumes that emission rates are constant in time. In order to apply the framework outlined above, it is convenient to define a vector $\tilde{\mathbf{p}}_c$ of length N_s containing the list of constant emission rates and then introduce a matrix \mathbf{A} such that $\tilde{\mathbf{q}}_c := \mathbf{A}\tilde{\mathbf{p}}_c$, where \mathbf{A} contains

N_T copies of each entry in the emissions vector $\tilde{\mathbf{p}}_c$ so that $\tilde{\mathbf{q}}_c$ corresponds to a straightforward discretization of the constant-in-time emissions problem (hence the subscript c).

We assumed that the noise $\boldsymbol{\epsilon}$ is a multivariate zero-mean Gaussian random variable that is distributed according to the density

$$\pi_{\boldsymbol{\epsilon}}(\mathbf{v}) = ((2\pi)^M |\boldsymbol{\Sigma}|)^{-1/2} \exp\left(-\frac{1}{2} \|\boldsymbol{\Sigma}^{-1/2} \mathbf{x}\|_2^2\right), \quad \mathbf{v} \in \mathbb{R}^M. \quad (17)$$

Combining this with (16) gives

$$\pi(\mathbf{d} | \tilde{\mathbf{q}}_c) = \pi(\mathbf{d} | \tilde{\mathbf{p}}_c) = \pi_{\boldsymbol{\epsilon}}(\mathbf{d} - \mathbf{F}\mathbf{A}\tilde{\mathbf{p}}_c) = ((2\pi)^M |\boldsymbol{\Sigma}|)^{-1/2} \exp\left(-\frac{1}{2} \|\boldsymbol{\Sigma}^{-1/2}(\mathbf{F}\mathbf{A}\tilde{\mathbf{p}}_c - \mathbf{d})\|_2^2\right), \quad (18)$$

where M denotes the length of the vector \mathbf{d} , $\boldsymbol{\Sigma}^{1/2}$ is defined in the sense of the square root of non-negative definite matrices, and $\|\cdot\|_2$ is the usual Euclidean norm. Given that the parameter space in this setting is small (with size equal to the number of point sources) it is sufficient to use a maximum likelihood estimator. We therefore consider the maximizer of the distribution on the right hand side of (18) which corresponds to the solution of the following constrained least squares problem:

$$\mathbf{q}_c := \mathbf{A}\mathbf{p}_c, \quad \mathbf{p}_c := \underset{\tilde{\mathbf{p}}_c}{\operatorname{argmin}} \left\| \boldsymbol{\Sigma}^{-1/2}(\mathbf{F}\mathbf{A}\tilde{\mathbf{p}}_c - \mathbf{d}) \right\|_2^2 \quad \text{subject to } \tilde{\mathbf{p}}_c \geq 0. \quad (19)$$

This problem can be solved efficiently using standard software such as MATLAB's `lsqmin` function (which we use here).

3.2. Unconstrained emission rates

In the second step, we employ the solution just derived for the constant emissions case to construct a prior distribution for the less restrictive problem in which the emission rates are not constant but rather vary smoothly in time (we leave the non-negativity constraint for the third and last step). This results in a linear problem for which we can still write an explicit analytical solution. We first rewrite the likelihood distribution appearing in (18) as

$$\pi(\mathbf{d} | \tilde{\mathbf{q}}_s) = ((2\pi)^M |\boldsymbol{\Sigma}|)^{-1/2} \exp\left(-\frac{1}{2} \|\boldsymbol{\Sigma}^{-1/2}(\mathbf{F}\tilde{\mathbf{q}}_s - \mathbf{d})\|_2^2\right). \quad (20)$$

Our intention is to use Bayes' rule [21] to estimate the variable emission rates \mathbf{q}_s given the sensor data \mathbf{d} , where the subscript "s" stands for "smoothness". We then use the MLE solution from (19) to construct a Gaussian prior distribution on \mathbf{q}_s as

$$\pi_{\text{prior}}(\tilde{\mathbf{q}}_s) = ((2\pi)^N |\mathbf{C}|)^{-1/2} \exp\left(-\frac{1}{2} \|\mathbf{C}^{-1/2}(\tilde{\mathbf{q}}_s - \mathbf{q}_c)\|_2^2\right), \quad (21)$$

where the choice of the covariance matrix \mathbf{C} is the key to constructing an appropriate prior. Here, we assume that \mathbf{C} has a block diagonal structure

$$\mathbf{C} = \mathbf{I}_{N_s} \otimes \mathbf{L}^{-2}, \quad (22)$$

where \mathbf{I}_{N_s} is the identity matrix of size $N_s \times N_s$ (recall N_s is the number of sources), \otimes denotes the Kronecker product, and \mathbf{L} is a finite difference discretization of the differential operator

$$\mathcal{L} := \alpha(\mathcal{I} - \gamma\partial_{tt}) \quad \text{on } [0, T]. \quad (23)$$

Here, $\alpha, \gamma > 0$ are fixed constants, \mathcal{I} is the identity map, and ∂_{tt} is the Laplacian with homogeneous Neumann boundary conditions. Intuitively, this choice implies that the second derivative of the emission rates is distributed as a Gaussian random variable with bounded variance at each point and so it yields sufficiently smooth sample functions.

It remains to choose a suitable discretization of the operator \mathcal{L} in (23), for which we use finite differences to get

$$\mathbf{L} = \alpha \sqrt{\frac{\Delta t}{T}} (\mathbf{I} - \gamma \mathbf{\Delta}) \quad \text{and} \quad \mathbf{\Delta} := \left(\frac{T}{\Delta t} \right)^2 \begin{bmatrix} -1 & 1 & 0 & \dots \\ 1 & -2 & 1 & 0 & \dots \\ & \ddots & \ddots & \ddots & \\ \dots & 0 & 1 & -2 & 1 \\ \dots & 0 & 0 & 1 & -1 \end{bmatrix}. \quad (24)$$

The parameter α controls the prior variance whereas γ controls the bandwidth of samples in the Fourier domain. Our choice of scaling $\sqrt{\Delta t/T}$ ensures that we obtain a proper discretization of the biharmonic operator, in which sample variances are independent of the discretization parameter Δt (for more details on discretization of random functions, see [21, Sec. 5.7] and [24]). Throughout the rest of this paper, we choose $\alpha = 1$ and $\gamma = 5 \times 10^{-3}$ for the two control parameters in (24). Figure 1 depicts two samples from the zero-mean prior distribution for this parameter choice, and we draw the reader's attention to both the continuity and smoothness of the samples.




Figure 1: Two samples from the zero-mean distribution $N(\mathbf{0}, \mathbf{I}_{N_s} \otimes \mathbf{L}^{-2})$ for the synthetic data problem, with $\alpha = 1$ and $\gamma = 5 \times 10^{-3}$.

With the prior distribution identified, we can now use Bayes' rule to write the posterior distribution of $\tilde{\mathbf{q}}_s$ given the data as

$$\pi_{\text{posterior}}(\tilde{\mathbf{q}}_s | \mathbf{d}) \propto \exp\left(-\frac{1}{2} \left\| \mathbf{\Sigma}^{-1/2} (\mathbf{F} \tilde{\mathbf{q}}_s - \mathbf{d}) \right\|_2^2\right) \exp\left(-\frac{1}{2} \left\| \mathbf{C}^{-1/2} (\tilde{\mathbf{q}}_s - \mathbf{q}_c) \right\|_2^2\right). \quad (25)$$

Because (9) is linear in $\tilde{\mathbf{q}}_s$, both the likelihood and the prior are Gaussian and so the posterior must also be Gaussian. In this setting, we may write the posterior analytically [52, 21] as

$$\pi_{\text{posterior}}(\tilde{\mathbf{q}}_s | \mathbf{d}) = N(\mathbf{q}_s, \mathbf{C}_s), \quad (26)$$

where

$$\mathbf{q}_s := \mathbf{q}_c + \mathbf{C} \mathbf{F}^T (\mathbf{\Sigma} + \mathbf{F} \mathbf{C} \mathbf{F}^T)^{-1} (\mathbf{d} - \mathbf{F} \mathbf{q}_c), \quad (27a)$$

$$\mathbf{C}_s := \mathbf{C} - \mathbf{C} \mathbf{F}^T (\mathbf{\Sigma} + \mathbf{F} \mathbf{C} \mathbf{F}^T)^{-1} \mathbf{F} \mathbf{C}, \quad (27b)$$

which completely determines the posterior distribution for the emission rates. Here, we take \mathbf{q}_s to be our next best guess of the true emission rates, which is intuitive because the mean of a Gaussian distribution coincides with the maximizer of its distribution so that \mathbf{q}_s is also the point of maximum probability of the posterior distribution (this time with the smoothness prior). We can also think of \mathbf{q}_s as an improvement on the constant-emissions estimate \mathbf{q}_c , which permits the emission rates to vary in time but still aims to keep the average emission rates close to \mathbf{q}_c .

3.3. Non-negative emission rates

In the third and final step, we impose a positivity (or more accurately non-negativity) constraint on the emission rates, for which the forward problem now becomes nonlinear. We define the function

$$h : \mathbb{R}^N \rightarrow \mathbb{R}^N, \quad \mathbf{q} = h(\mathbf{v}) \quad \text{then} \quad \mathbf{q}_i = \max\{0, \mathbf{v}_i\}, \quad (28)$$

which permits us to pose the inverse problem for an auxiliary vector \mathbf{v} and rewrite (25) as

$$\pi_{\text{posterior}}(\mathbf{v} | \mathbf{d}) \propto \exp\left(-\frac{1}{2} \|\boldsymbol{\Sigma}^{-1/2}(\mathbf{F}h(\mathbf{v}) - \mathbf{d})\|_2^2\right) \exp\left(-\frac{1}{2} \|\mathbf{C}^{-1/2}(\mathbf{v} - h(\mathbf{q}_s))\|_2^2\right). \quad (29)$$

Note that because h is a nonlinear function, the posterior distribution is no longer Gaussian and so (27) no longer applies. In this case, we turn to Markov Chain Monte Carlo (MCMC) algorithms to generate samples from the posterior distribution and compute the expectation of certain functions. A detailed discussion of MCMC algorithms is outside the scope of this article and we refer the reader to the monograph [35] for an introduction to MCMC and also to [21] for applications of these algorithms to inverse problems.

In this paper, we choose the preconditioned Crank-Nicolson (pCN) algorithm of Cotter et al. [11], which is a variation of the usual random walk Metropolis-Hastings algorithm. Consider a general setting where we would like to compute the expectation of some function $f(\mathbf{v})$ under a Gaussian prior distribution of the form $\pi_{\text{prior}}(\mathbf{v}) = N(\mathbf{0}, \mathbf{C})$, and define

$$\phi(\mathbf{v}, \mathbf{d}) = \frac{1}{2} \|\boldsymbol{\Sigma}^{-1}(\mathbf{F}(h(\mathbf{v}) - \mathbf{d}))\|_2^2.$$

Then we have that

$$\int f(\mathbf{v}) \pi_{\text{posterior}}(\mathbf{v} | \mathbf{d}) d\mathbf{v} \approx \frac{1}{K} \sum_{k=1}^K f(\mathbf{v}^{(k)}), \quad (30)$$

where the samples $\mathbf{v}^{(k)}$ are generated by the following algorithm:

1. Take $\mathbf{v}^{(0)} \sim N(\mathbf{0}, \mathbf{C})$.
2. While $k \leq K$,
 - (i) Generate $\mathbf{v} \sim N(\mathbf{0}, \mathbf{C})$.
 - (ii) Set $\tilde{\mathbf{v}} = \sqrt{1 - \beta^2} \mathbf{v}^{(k)} + \beta \mathbf{v}$.
 - (iii) With probability $a(\mathbf{v}^{(k)}, \tilde{\mathbf{v}}) = \min\{1, \exp[\phi(\mathbf{v}^{(k)}, \mathbf{d}) - \phi(\tilde{\mathbf{v}}, \mathbf{d})]\}$, accept the sample and set

$$\mathbf{v}^{(k+1)} = \tilde{\mathbf{v}}.$$

Otherwise, reject the sample and set

$$\mathbf{v}^{(k+1)} = \mathbf{v}^{(k)}.$$

- (iv) $k \rightarrow k + 1$.

The parameter K must be chosen large enough to ensure that the Markov chain converges and that the error in the approximation of the integral (30) is sufficiently small. The parameter $0 < \beta < 1$ is also user-specified and controls the rate of convergence of the algorithm. In general, smaller values of β result in a larger acceptance probability in step 2(iii) of the above algorithm. This means that the algorithm accepts more samples but the resulting Markov chain will make correspondingly smaller jumps; therefore, when β is small, the iteration will be slower to traverse the posterior distribution. If β is large, then the Markov chain employs longer jumps which is more desirable for exploring the parameter space; however, the acceptance probability is also reduced which means that we will reject most of the samples and occasionally get stuck

at certain points. In practice it is desirable to choose β so that the acceptance probability from step 2(iii) lies in the range $[0.25, 0.35]$, which is close to the optimal acceptance rate of the random walk algorithm [35]. Finally, we note that the above algorithm can be used to sample using non-centered priors by applying a simple linear shift of \mathbf{v} in the definition of ϕ .

Once the samples are generated, we can store them in memory and use them to estimate the posterior mean and standard deviation via

$$\mathbf{v}_{\text{PM}} := \int \mathbf{v} \pi_{\text{posterior}}(\mathbf{v} | \mathbf{d}) d\mathbf{v} \quad \text{and} \quad \mathbf{C}_{\mathbf{v}} := \int (\mathbf{v} - \mathbf{v}_{\text{PM}})(\mathbf{v} - \mathbf{v}_{\text{PM}})^T \pi_{\text{posterior}}(\mathbf{v} | \mathbf{d}) d\mathbf{v}. \quad (31)$$

We then define

$$\mathbf{q}_{\text{sp}} := h(\mathbf{v}_{\text{PM}}) \quad \text{and} \quad \mathbf{C}_{\text{sp}} := \int (h(\mathbf{v}) - \mathbf{q}_{\text{sp}})(h(\mathbf{v}) - \mathbf{q}_{\text{sp}})^T \pi_{\text{posterior}}(\mathbf{v} | \mathbf{d}) d\mathbf{v}, \quad (32)$$

and take \mathbf{q}_{sp} as our estimate of the non-negative emission rates, where the subscript “sp” stands for “smoothness and positivity”.

To summarize, (19) defines the posterior mean \mathbf{q}_{c} under the prior assumption of constant emissions which is used in (27) to derive the posterior mean \mathbf{q}_{s} under the smoothness assumption. This estimate is then used in (31) and (32) to derive the posterior mean \mathbf{q}_{sp} when both smoothness and positivity assumptions are imposed.

4. An industrial case study

We now apply the methodology developed in the previous two sections to an industrial case study in which airborne particles, most notably lead, are emitted from a lead-zinc smelter operated by Teck Resources in Trail, British Columbia, Canada. These particulate emissions are of a type called “fugitive emissions” that are not easily identified because they are caused by accidental releases or come from buildings or other areas of the industrial operation that are not amenable to direct measurement approaches. It is therefore of particular interest to identify the location and amount of emitted material from such fugitive sources in order to prioritize capital intensive projects to control sources.

This study was performed over the period August 20 to September 19, 2013. The company had installed sensors on the main stacks and other non-fugitive sources that were already known to be historically major emitters of lead particulates. Prior to the date of this study, fugitive dust emissions reported by the company to Environment Canada’s National Pollutant Release Inventory (NPRI) were based on engineering calculations listed in the NPRI reporting Toolbox (EPA method 42, which does not account for local conditions). Figure 2 shows an overhead view of the Trail smelter site along with certain areas that were selected by the company’s environmental engineering team as the most likely sources of fugitive emissions (and the centroid of each area is also identified).

We were provided with meteorological data consisting of measurements of wind velocity and direction at 10-minute intervals throughout the entire one-month period. There were also several types of data related to particulate material obtained from a number of measurement devices, each with a different sampling schedule:

- “Dust-fall”: a set of 30 dust-fall jars deployed throughout the site, which yield measurements of the total deposited mass of lead (and various other particulates) accumulated over the entire monthly sampling period.
- “Xact”: real-time PM10 measurements (particulate material up to 10 μm in diameter), which were taken using an Xact 620 ambient metals monitor [10] and reported as hourly-averaged concentrations;
- “TSP”: total suspended particulate measurements (up to 100 μm in diameter), taken with an Andersen high-volume air sampler [46] as hourly-averaged measurements once every second day at midnight;



Figure 2: The smelter site in Trail, BC, highlighting the mostly likely areas to be sources of fugitive emissions. The centroid of each area source is indicated by a point and labelled q1–q7.

- “PM10”: the subset of the Andersen hi-vol TSP measurements that correspond to PM10 particulates (up to $10\ \mu\text{m}$), taken as hourly-averaged measurements but only recorded once per week.

The locations of the dust-fall jars and real-time sensors are indicated in Figure 3, and Table 2 summarizes the type of measurements each sensor provides. The sensor accuracy is characterized using signal to noise ratio (SNR), which is the ratio of the variance of the signal to that of the noise. The given values in Table 2 are chosen based on discussions with the environmental engineering team at Teck Trail Operations.

Table 2: Measurement characteristics for each sensor type.

Type	Dust-fall	TSP	PM10	Xact
Schedule	monthly	12:00am, every 2 days	12:00am, weekly	hourly
SNR	10	100	100	100
# measurements/mo.	1	16	6	761

4.1. Parameters and wind data

In this section, we summarize the input data specific to the Teck case study that pertains to wind measurements and physical parameters for lead particulates. For simplicity, we restrict our attention to a single particulate type corresponding to the most abundant form of lead found by the sensors: lead monoxide (PbO). A chemical analysis performed by the company suggested that PbO particulate material has an average diameter of $5 \times 10^{-6}\ \text{m}$ [5] which we use as the size of all particulates in this study. This value lies within the range of $0.8 - 20 \times 10^{-6}\ \text{m}$ that is reported in [32] for lead dust. Based on this value of particle size, we may then estimate the deposition velocity W_{dep} using the data provided in [41, Figure 2]. Our assumed values of physical parameters for PbO are summarized in Table 3.

The other essential inputs to the model are the wind velocity and direction as functions of time. As mentioned earlier, wind data is available at 10-minute intervals from a single meteorological station identified

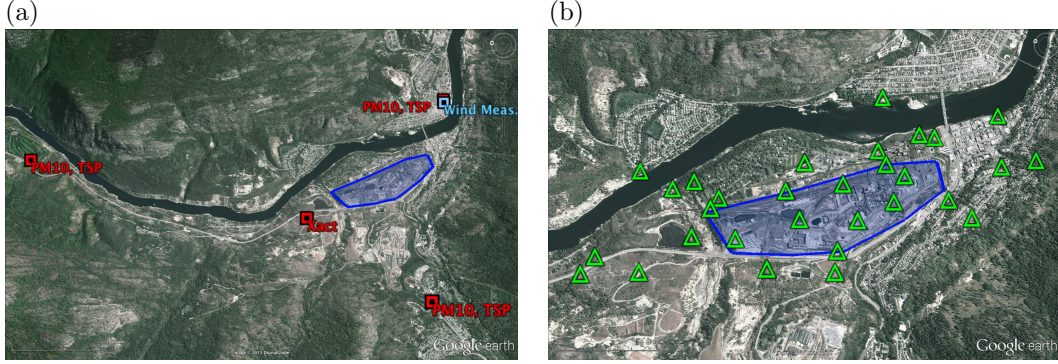


Figure 3: (a) Locations of the real-time measurement samples (Xact, PM10 and TSP systems) and the meteorological post where wind measurements were obtained, relative to the smelter site (outlined in blue). The TSP and PM10 measurements are taken with the same Andersen hi-vol device. (b) A zoomed-in view of the outlined smelter site depicts the dust-fall jar locations with green triangles.

Table 3: Physical parameters for lead monoxide (PbO), the dominant lead particulate encountered in this study.

Property	Density	Diameter	Deposition velocity	Settling velocity
Symbol	ρ	d_p	W_{dep}	W_{set}
Units	$kg\ m^{-3}$	m	$m\ s^{-1}$	$m\ s^{-1}$
Value	9530	5×10^{-6}	0.005	0.0026

in Figure 3a, while Figure 4 presents the speed histogram and wind-rose diagram corresponding to the month of interest. The wind-rose data indicate that the wind during this period blows consistently toward the south-east, although there are significant variations in wind speed. The raw wind data contains substantial measurement errors and as a result they cannot be input directly into our forward model. Instead, we must first apply some form of regularization. For this purpose, we take the y -axis to point towards north and split the wind data into velocity components in the x - and y -directions. We then treat the wind coordinates as independent variables and fit a Gaussian process separately to each coordinate data set. The fit is done by taking a Gaussian covariance operator for each process with the variance of the operator determined using ten-fold cross-validation. The mean of the Gaussian process obtained for each data set is then used as the regularized wind input data for the forward model, and the resulting wind components are pictured in Figure 5. Since this regularization procedure is based on a standard approach in machine learning, we refer the interested reader to [52, Ch. 2, Ch. 3 and Sec. 5.3] and [4, Sec. 6.4] for details on the use of Gaussian processes in regression.

4.2. A simple test with synthetic data

Before presenting results with the measured data, we start by testing our framework on a synthetic data set in order to validate the method and to also gain some insight into what type of reconstructions we can expect from our algorithm. The main challenge to be dealt with here, and particularly in a real industrial setting, is that the data is sparse and so one cannot expect to reconstruct emission rates with high fidelity.

4.2.1. Generating the test data

We consider the artificial emission rates pictured in Figure 6a,b, which are imposed at seven sources numbered q1–q7. Each emission rate is a simple sinusoidal function of time with a different amplitude and frequency, and with a positivity constraint imposed if necessary. Our motivation in making this choice is to have time-varying emission rates that incorporate both high and low frequencies, and which also have




Figure 4: Wind speed histogram (left) and wind-rose diagram (right) for the period August 20–September 19, 2013. The compass direction in the wind rose diagram denotes the direction that the wind is blowing from.




Figure 5: Comparison of regularized wind direction and velocity (solid lines) and raw measurements (points).

significant variations in magnitude. To generate the synthetic deposition measurements, we apply the actual regularized wind data shown in Figure 5 and discretize the forward problem with uniform time increments of size $\Delta t = 1800$ s. Note that this step size enters all aspects of the forward problem including the Gaussian plume solution, wind regularization, and measurement operators. After solving the forward problem (2), (3), (6) and (9) to generate the artificial deposition measurements, we introduce noise artificially by perturbing the data by a vector of independent Gaussian random numbers having standard deviations equal to the measurement errors listed in Table 2.

4.2.2. Solving the inverse problem

In order to avoid the “inverse crime” of choosing the same mesh to generate the artificial data and solve the inverse problem [21], we instead solve the inverse problem using a coarser time increment $\Delta t = 3600$ s, which ensures that our wind data, forward map and measurement operators are different from those used to generate the data. We will also consider a setting where the measurement noise is mis-specified and construct the Σ matrix using measurement errors equal to one-half the value listed in Table 2.

We begin by solving the inverse problem under the assumption that emission rates are constant, following the methodology outlined in Section 3.1. The resulting solution \mathbf{q}_c is pictured in Figure 6b, where the grey bars (labeled “Artificial”) depict the average of the emission rates in Figure 6a. Our estimate indicates that the constant emissions assumption is relevant if one is interested in the time-averaged source emission rates, although no information is provided about the time dependence.

We next turn our attention to the case of smooth (non-constant) emission rates without a positivity constraint. We use the same noise covariance matrix Σ as in the constant case and substitute the vector \mathbf{q}_c of Figure 6b into (27) to compute the posterior mean and covariance. Results of this computation are presented in Figures 6c,d, which depict the posterior mean and standard deviation. Note that the posterior mean is not strictly a physically reasonable estimate of the emission rates because one source (q6) exhibits negative

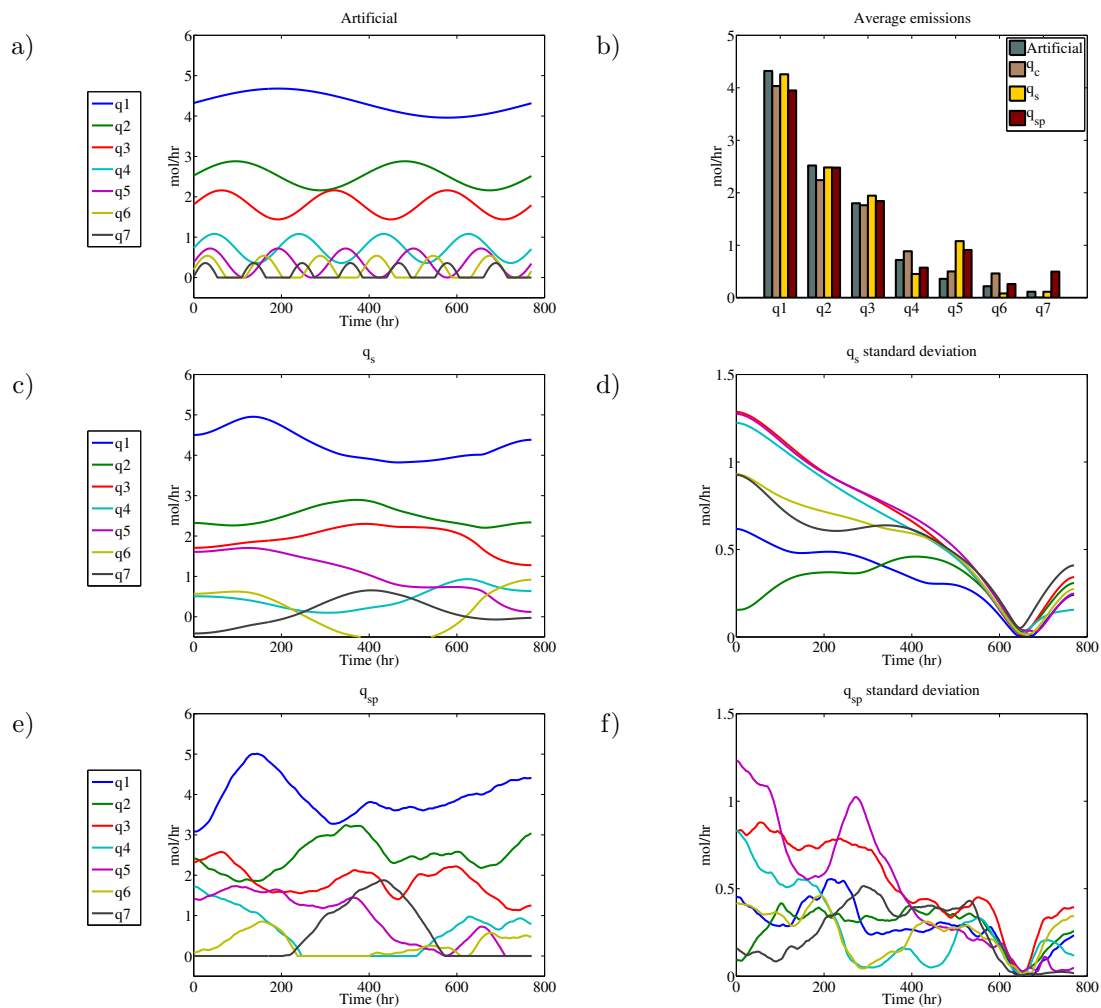


Figure 6: Summary of results using the artificial input data set. (a) Artificial emission rates used to generate the data. (b) Average emission rate of each source obtained using the three different priors for constant emissions, smooth emissions and smooth and positive emissions. (c) Posterior mean of sources as functions of time using the smoothness assumption only (no positivity constraint) along with (d) the standard deviations. (e) Posterior mean of estimated emissions with the both smoothness and positivity assumptions along with (f) corresponding standard deviations.

emissions. Nonetheless, our algorithm still correctly identifies the largest sources, for which the magnitude and overall shape of the emission curves for the larger sources is captured quite well. However, we note that the smaller sources are not resolved by our reconstruction, which is apparent in the standard deviation plots of Figure 6d. By considering the standard deviation as a measure of uncertainty in the reconstructions, we see that the uncertainty is larger for the smaller sources (q3–q7) and significantly smaller for the two main sources (q1, q2), meaning that we should have less confidence in our reconstruction for the smaller sources. This is due to the scarcity of the data as well as our choice of prior which prefers smooth estimates.

As a final test, we impose the positivity constraint using the approach described in Section 3.3. As mentioned before, there is no longer an analytic expression for the posterior distribution and so we employ an MCMC algorithm to sample the posterior and compute the mean and standard deviations. Figures 6e,f depict the estimate of \mathbf{q}_{sp} and its standard deviation using $K = 5 \times 10^5$ and $\beta = 0.6$, which yields an average acceptance probability of 0.32. These estimates with the positivity constraint appear slightly better than the ones without, which is reflected in the smaller values of standard deviation.

4.3. Results for actual measured deposition data

We now turn our attention to solving the inverse problem using the actual measured deposition values at the Trail smelter site. We use the same parameter values as for the artificial problem above and results are summarized in Figure 7. The constant emissions estimate in Figure 7 suggests that q2 and q5 are by far the main sources of lead particulates. This is to be expected for q5, since it is a loading area where piles of material are mixed and stored and hence is expected to generate significant quantities of airborne dust, whereas q2 is the area immediately to the north of the smelter building. On the other hand, q1, q6 and q7 are all predicted to emit minimal amounts of lead. Figures 7a,c show the posterior mean of the estimates using the smoothness prior with and without a positivity constraint. Averaged emission rates pertaining to these cases are presented in Figure 7e. The relative importance of the various sources does not differ dramatically between the three different algorithms, perhaps with the exception of the estimates for q6 and q7.

It is interesting that Figures 7a,c both exhibit a surge in certain emission rates (q2, q5, q6) near the end of the month. Looking at the standard deviation plots in Figures 7b,d, we note that these surging month-end values are estimated with a higher level of confidence, compared to the beginning of the month when uncertainties are larger. Note also that the uncertainties are smaller in Figure 7d as compared to 7b which indicates that imposing the positivity constraint yields improved estimates. In the next section, we will discuss in more detail the uncertainties in solution estimates and how they can be used to assess the impact of emissions on the surrounding area.

Another insightful result of this case study is the sensitivity of the solution to the different types of measurements. The Xact device is the most expensive instrument, but it also provides nearly six times the measurement data than all other instruments combined. There are thus clear advantages to having more Xact data, although deploying new Xact devices may be difficult to justify due to their high cost and maintenance requirements. Figure 7f depicts the \mathbf{q}_c estimates obtained with and without the Xact data, clearly demonstrating that Xact measurements have a major impact on the outcome. Without the Xact device, our approach greatly over-estimates emission rate from source q5 and also fails to identify q2 as the primary smelter building contributing to lead emissions. This clearly demonstrates the importance of having real-time measurements available in such an emissions study. Indeed, in the absence of time-varying data from such a real-time sensor, one simply obtains a posterior that is very similar to the prior and hence it is not possible to reconstruct general time-varying sources with any degree of accuracy.

4.4. Impact assessment and uncertainty propagation

Now that we have obtained the solution to the inverse problem, it is natural to go one step further and study the broader implications of the emissions estimates in terms of how lead particulates are distributed over the area surrounding the smelter site. These effects can be quantified using various measures such as total annual lead emissions or monthly ground-level deposition. The total annual emission of a given pollutant is typically of great interest to a company such as Teck because it forms part of their annual reporting

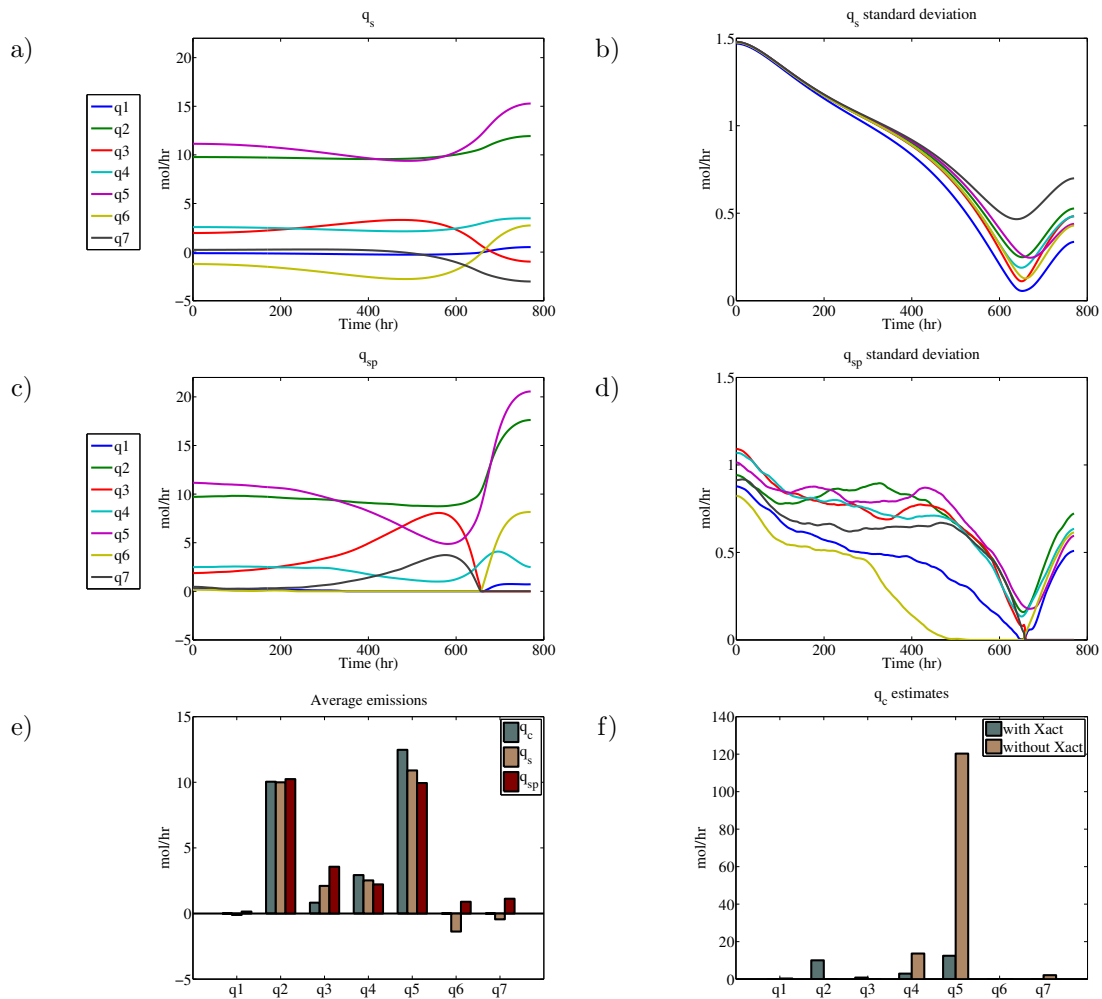


Figure 7: Results obtained using actual measured depositions at the Trail smelter. (a,b) Posterior mean and standard deviation of the emissions using the smoothness prior with no positivity constraints. (c,d) Posterior mean and standard deviation obtained using the smoothness and positivity prior. (e) Average emission rate of each source for the duration of the study, using three choices of prior. (f) Comparison of the q_c estimate obtained using the entire data set and the estimate when the Xact data is excluded.

requirements to environmental monitoring bodies, not to mention that companies often set ambitious goals for reducing the total emission figure. It is straightforward to extrapolate our average emission rates from Figure 7e to estimates of the annual emission rate for the entire smelter. Table 4 shows the results of our computation along with annual emissions that were given in an independent study that was performed by the company [47]. The fact that our estimates are close to the values in the independent study increases our confidence in the obtained solutions. The reported values of standard deviations were obtained by first approximating the posterior with a Gaussian (this is only needed in the case of \mathbf{q}_{sp}). We also present an estimate of the 90% probability interval, which is the interval around the mean that contains 90% of the probability mass. This interval is also computed based on the Gaussian approximation to the posterior.

Table 4: Extrapolated total annual emission rates (in tonne/yr) for lead particulates. The standard deviation and 90% probability intervals are computed by first estimating the posterior with a Gaussian.

Approximation	[tonne/yr]	Standard deviation	90% probability interval
\mathbf{q}_c	51.0		
\mathbf{q}_s	46.1	5	± 8.2
\mathbf{q}_{sp}	54.95	5	± 8.2
Independent study [47]	45		

A more interesting problem is that of computing the total lead deposited at ground level using our estimated emission rates for the period of August 20 to September 19, 2013. This is an important problem because it allows us to assess in more detail the impact of emissions on the area surrounding the industrial site. In order to perform this computation we take a domain consisting of a rectangular ground patch, $[L_x, U_x] \times [L_y, U_y] \subset \mathbb{R}^2$, which is discretized on an $n_x \times n_y$ grid of equally-spaced points. We use the Gaussian plume solution (6) to compute lead concentration at each grid point, which is then integrated in time and multiplied by the deposition velocity to obtain monthly deposition values. This procedure is similar to our construction of the \mathbf{F} operator in (15). We then let \mathbf{b} denote the vector containing grid point values of the monthly depositions and define an operator \mathbf{H} such that

$$\mathbf{b} = \mathbf{H}\mathbf{q} \tag{33}$$

for a given vector of emission rates \mathbf{q} . We can now solve the forward problem using our estimate \mathbf{q}_{sp} to obtain $\mathbf{b}_{\text{sp}} = \mathbf{H}\mathbf{q}_{\text{sp}}$. The resulting solution is depicted as a contour plot of lead mass per unit area in Figure 8a, which indicates that most deposition occurs close to the sources and within the boundaries of the smelter site itself.

Next, we propagate the posterior uncertainty through the forward model to obtain error bounds on the monthly deposition values, with our main aim being to obtain a framework that provides a reasonable estimate of solution uncertainty on a fine grid. To make the computations as efficient as possible, we would like to avoid sampling algorithms and so we approximate the posterior distribution for the non-negative solution using a Gaussian distribution $N(\mathbf{q}_{\text{sp}}, \mathbf{C}_{\text{sp}})$. Our aim is to exploit the fact that a linear transformation of a Gaussian distribution is also Gaussian. Recall that \mathbf{C}_{sp} can be estimated using the MCMC algorithm of Section 3.3. Because (33) is linear, we have that $\mathbf{b}_{\text{sp}} \sim N(\mathbf{H}\mathbf{q}_{\text{sp}}, \mathbf{H}\mathbf{C}_{\text{sp}}\mathbf{H}^T)$. Computing this covariance matrix on a fine grid might still be intractable because \mathbf{H} is a large matrix and \mathbf{C}_{sp} is dense. However, Figure 9 shows that the eigenvalues of \mathbf{C}_{sp} decay rapidly in time, which suggests replacing it by a low-rank approximation by simply truncating the spectrum. To this end, let $\mathbf{C}_{\text{sp}} = \mathbf{L}\mathbf{D}\mathbf{L}^T$ be the usual eigenvalue decomposition of the covariance, let n_e be the number of eigenvalues we want to retain, and take $\tilde{\mathbf{D}}$ as the first $n_e \times n_e$ sub-matrix of \mathbf{D} . Then let $\tilde{\mathbf{L}}$ denote the tall matrix containing the first n_e columns of \mathbf{L} . We can now define $\tilde{\mathbf{C}}_{\text{sp}} := \tilde{\mathbf{L}}\tilde{\mathbf{D}}\tilde{\mathbf{L}}^T$ and approximate the distribution of the depositions as $\mathbf{b}_{\text{sp}} \sim N(\mathbf{H}\mathbf{q}_{\text{sp}}, \mathbf{H}\tilde{\mathbf{C}}_{\text{sp}}\mathbf{H}^T)$, and the covariance can now be estimated using n_e solves of the forward problem.

In Figure 10 we depict a few selected eigenvectors of the posterior covariance, where larger eigenvalues indicate directions of higher uncertainty. Our low-rank approximation of the covariance matrix preserves the eigen-directions and dismisses the more oscillatory directions arising from smaller eigenvalues. Figure 8

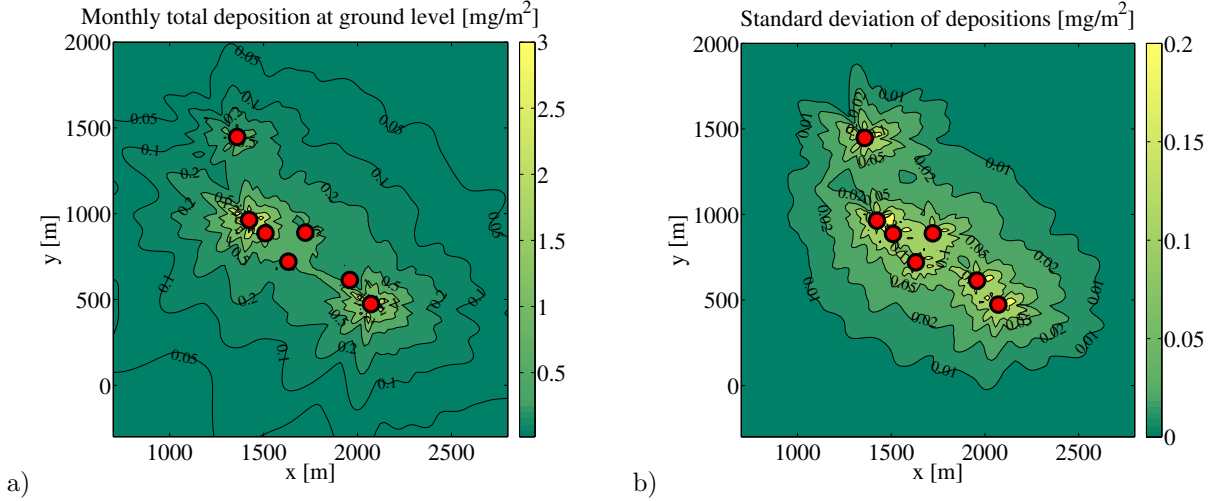


Figure 8: (Left) Deposition in mg/m^2 obtained by solving the forward problem with the posterior mean \mathbf{q}_{sp} . (Right) Standard deviation of deposition using a low-rank approximation of the posterior covariance. The solution is computed on a 100×100 spatial grid using only the first 100 eigen-pairs.

shows the final result of our uncertainty propagation study, with the standard deviation of the depositions computed by retaining $n_e = 100$ eigenvectors of \mathbf{C}_{sp} , and presented on the same 100×100 spatial grid that was used in Figure 8a. The contours of standard deviation indicate the spatial variation in the uncertainty of the deposition estimates around the industrial site. It is clear from Figure 8b that the deposition estimates become more uncertain closer to the sources. However, uncertainties are very small outside the boundaries of the smelter site, which means that we can be confident in our estimate of the impact of emissions on the surrounding area, even though we observed relatively large uncertainties in estimates of the emission rates in Section 4.3.




Figure 9: Eigenvalues of the posterior covariance for the solution of the inverse problem with positivity constraint.

5. Conclusions

In this paper we studied the inverse problem corresponding to estimating the rate of fugitive emissions for airborne particulate matter from an industrial site. We restrict ourselves to short-range transport of particles (on the order of a few kilometres from the source) and make simplifying assumptions that allow

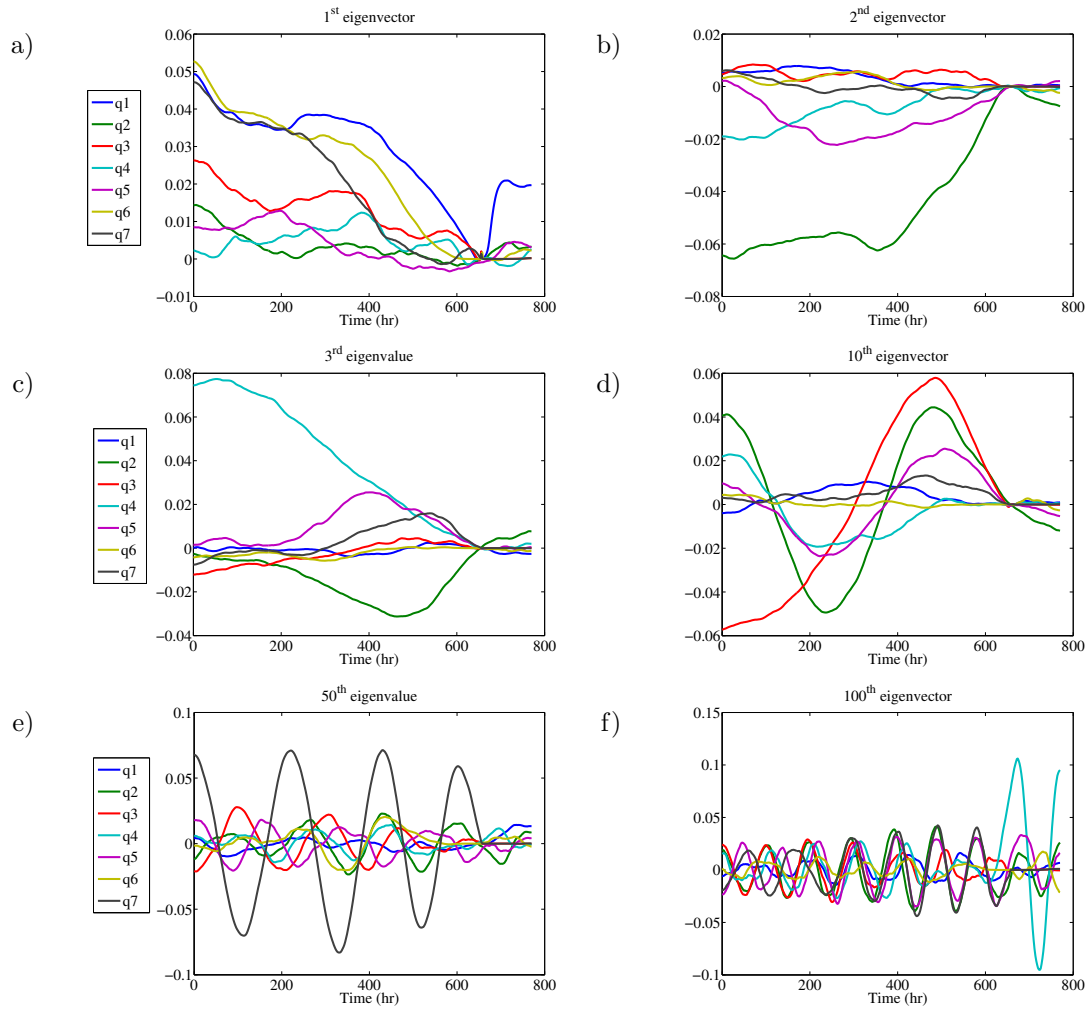


Figure 10: Selected normalized eigenvectors of the posterior covariance for the solution with positivity constraint. The first few eigenvalues correspond to directions of maximum uncertainty in the solution.

us to use an efficient Gaussian plume type solution as our forward solver. We then model our measurement devices and construct a linear forward solver.

We solved the inverse problem using a Bayesian framework, developing a solution approach that employs three possible prior assumptions, based on the emission rates being: constant and positive; smoothly-varying but not necessarily positive; and smoothly-varying and positive. The solution for these three cases becomes successively more flexible, but also more expensive to compute. We discuss methods for efficient solution of the inverse problem in each case. Finally, we apply our solution framework to a concrete study of fugitive emissions of lead particulates from a lead-zinc smelter in Trail, British Columbia, Canada. We obtain estimates of emissions from seven suspect area sources contained within the boundaries of the industrial site and then extrapolate our solution to obtain a total annual emission rate. Our estimates are consistent with the results of an independent study that was performed by the company in 2013 [47], which gives us a high degree of confidence in our choice of priors and the solution methodology.

There are several obvious areas for future research that aim to extend the framework developed in this study. For example, one can consider estimating emission rates for more than just one particulate. In many applications certain tracer particulates can be connected with specific operations or sources, so that incorporating these tracers within the framework has the potential to improve estimates of the emission rates by distinguishing depositions that originate from specific sources. Another interesting problem is that of identifying a singular emission event, or in other words estimating emissions from a source that emits a large amount of particulate material during a relatively short time period. In this case, one would have to consider a different prior that can incorporate singular sources in contrast with the smoothness priors used in this study. Finally, the optimal experimental design for source inversion problems is of great interest in practical applications. For example, determining optimal locations for dust-fall jars and real-time measurement devices could significantly improve both the data quality and the emission estimates. It is crucial to have a good strategy for deploying sensors, and this strategy will be subject to practical constraints and may depend on factors such as wind patterns that are not known at the time of the measurements. One promising approach is to apply statistical learning techniques to predict unknown parameters based on existing data, which is related to our method in Section 4.1 for regularizing and extending wind data using Gaussian processes.

Acknowledgements

We would like to thank Peter Golden, Cheryl Darrah and Mark Tinholt from Teck’s Trail Operations for many valuable discussions. This project was supported by an NSERC Discovery Grant (to JMS) and an Accelerate Internship Grant from Mitacs and Teck Resources (to BH).

References

- [1] R. A. S. Albani, F. P. Duda, and L. C. G. Pimentel. On the modeling of atmospheric pollutant dispersion during a diurnal cycle: A finite element study. *Atmospheric Environment*, 118:19–27, 2015.
- [2] S. P. Arya. Modeling and parameterization of near-source diffusion in weak winds. *Journal of Applied Meteorology*, 34:1112–1122, 1995.
- [3] S. P. Arya. *Air Pollution Meteorology and Dispersion*. Oxford University Press, New York, 1999.
- [4] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
- [5] R. J. Blaskovich. Environment dustfall particulate characterization. Technical Report 1224-915537, Teck Resources Ltd. Operations at Trail, 18 October 2013.
- [6] G. A. Briggs. Diffusion estimation for small emissions. Atmospheric Turbulence and Diffusion Laboratory Contribution, ATDL File No. 79, National Oceanic and Atmospheric Administration, Oak Ridge, TN, May 1973.

- [7] B. Bunton, P. O’Shaughnessy, S. Fitzsimmons, J. Gering, S. Hoff, M. Lyngbye, P. S. Thorne, J. Wasson, and M. Werner. Monitoring and modeling of emissions from concentrated animal feeding operations: Overview of methods. *Environmental Health Perspectives*, 115(2):303–307, 2007.
- [8] O. Christensen and L. P. Prahm. A pseudospectral model for dispersion of atmospheric pollutants. *Journal of Applied Meteorology*, 15(12):1284–1294, 1976.
- [9] A. J. Cimorelli, S. G. Perry, A. Venkatram, J. C. Weil, R. J. Paine, R. B. Wilson, R. F. Lee, W. D. Peters, and R. W. Brode. AERMOD: A dispersion model for industrial source applications. Part I: General model formulation and boundary layer characterization. *Journal of Applied Meteorology*, 44(5):682–693, 2005.
- [10] Cooper Environmental, Beaverton, OR. *Ambient monitoring: Xact 625 (product description)*, 2015. <http://cooperenvironmental.com/ambient-monitoring>.
- [11] S. L. Cotter, G. O. Roberts, A. M. Stuart, D. White, et al. MCMC methods for functions: Modifying old algorithms to make them faster. *Statistical Science*, 28(3):424–446, 2013.
- [12] I. Dimov, K. Georgiev, T. Ostromsky, and Z. Zlatev. Computational challenges in the numerical treatment of large air pollution models. *Ecological Modelling*, 179(2):187–203, 2004.
- [13] I. G. Enting. *Inverse Problems in Atmospheric Constituent Transport*. Cambridge Atmospheric and Space Science Series. Cambridge University Press, 2002.
- [14] D. L. Ermak. An analytical model for air pollutant transport and deposition from a point source. *Atmospheric Environment*, 11(3):231–237, 1977.
- [15] S. R. Hanna, G. A. Briggs, and R. P. Hosker Jr. Handbook on atmospheric diffusion. Technical Report DOE/TIC-11223, Atmospheric Turbulence and Diffusion Lab, National Oceanic and Atmospheric Administration, Oak Ridge, TN, 1982.
- [16] B. Hosseini. Dispersion of pollutants in the atmosphere: A numerical study. Master’s thesis, Department of Mathematics, Simon Fraser University, Burnaby, BC, Canada, Aug. 2013. Available at <http://summit.sfu.ca/item/13646>.
- [17] B. Hosseini and J. M. Stockie. A finite volume scheme for short-range dispersion and deposition of airborne particulates, 2016. Preprint.
- [18] Z. Huang, Y. Wang, Q. Yu, W. Ma, Y. Zhang, and L. Chen. Source area identification with observation from limited monitor sites for air pollution episodes in industrial parks. *Atmospheric Environment*, 122:1–9, 2015.
- [19] V. Isakov. *Inverse Source Problems*. American Mathematical Society, Providence, RI, 1990.
- [20] S. I. Kabanikhin. *Inverse and Ill-posed Problems: Theory and Applications*, volume 55 of *Inverse and Ill-Posed Problems Series*. Walter De Gruyter, Berlin, 2011.
- [21] J. P. Kaipio and E. Somersalo. *Statistical and Computational Inverse Problems*. Springer, New York, 2005.
- [22] A. Keats, E. Yee, and F.-S. Lien. Bayesian inference for source determination with applications to a complex urban environment. *Atmospheric Environment*, 41(3):465–479, 2007.
- [23] R. Lange. ADPIC – A three-dimensional particle-in-cell model for the dispersal of atmospheric pollutants and its comparison to regional tracer studies. *Journal of Applied Meteorology*, 17(3):320–329, 1978.
- [24] M. Lassas, E. Saksman, and S. Siltanen. Discretization-invariant Bayesian inversion and Besov space priors. *Inverse Problems and Imaging*, 3:87–122, 2009.

- [25] R. L. Lee, J. R. Albritton, D. L. Ermak, and J. Kim. Computational fluid dynamics modeling for emergency preparedness and response. *Environmental Modelling & Software*, 12(1):43–50, 1997.
- [26] S. A. Levin, H. C. Muller-Landau, R. Nathan, and J. Chave. The ecology and evolution of seed dispersal: A theoretical perspective. *Annual Review of Ecology, Evolution and Systematics*, 34:575–604, 2003.
- [27] E. Lushi and J. M. Stockie. An inverse Gaussian plume approach for estimating atmospheric pollutant emissions from multiple point sources. *Atmospheric Environment*, 44:1097–1107, 2010.
- [28] G. J. McRae, W. R. Goodin, and J. H. Seinfeld. Numerical solution of the atmospheric diffusion equation for chemically reacting flows. *Journal of Computational Physics*, 45(1):1–42, 1982.
- [29] C. W. Miller and L. M. Hively. A review of validation studies for the Gaussian plume atmospheric dispersion model. *Nuclear Safety*, 28(4):522–531, 1987.
- [30] J. Nikmo, J.-P. Tuovinen, J. Kukkonen, and I. Valkama. A hybrid plume model for local-scale atmospheric dispersion. *Atmospheric Environment*, 33:4389–4399, 1999.
- [31] S. Okamoto, H. Ohnishi, T. Yamada, T. Mikami, S. Momose, H. Shinji, and T. Itohiya. A model for simulating atmospheric dispersion in low wind conditions. *International Journal of Environment and Pollution*, 16(1-6):69–79, 2001.
- [32] M. K. Owen, D. S. Ensor, and L. E. Sparks. Airborne particle sizes and sources found in indoor air. *Atmospheric Environment*, 26A(12):2149–2162, 1992.
- [33] Y.-S. Park and J.-J. Baik. Analytical solution of the advection-diffusion equation for a ground-level finite area source. *Atmospheric Environment*, 42:9063–9069, 2008.
- [34] K. S. Rao. Source estimation methods for atmospheric dispersion. *Atmospheric Environment*, 41:6964–6973, 2007.
- [35] C. P. Robert and G. Casella. *Monte Carlo Statistical Methods*. Springer Texts in Statistics. Springer, New York, second edition, 2004.
- [36] O. F. T. Roberts. The theoretical scattering of smoke in a turbulent atmosphere. *Philosophical Transactions of the Royal Society of London, Series A*, 104(728):640–654, 1924.
- [37] V. Sanf elix, A. Escrig, A. L opez-Lilao, I. Celades, and E. Monfort. On the source inversion of fugitive surface layer releases. Part I. Model formulation and application to simple sources. *Atmospheric Environment*, 109:171–177, 2015.
- [38] J. S. Scire, D. G. Strimaitis, and R. J. Yamartino. *A User’s Guide for the CALPUFF Dispersion Model (Version 5)*. Earth Tech Inc., Concord, MA, Jan. 2000.
- [39] J. H. Seinfeld and S. N. Pandis. *Atmospheric Chemistry and Physics: From Air Pollution to Climate Change*. John Wiley & Sons, New York, 1998.
- [40] I. Senocak, N. W. Hengartner, M. B. Short, and W. B. Daniel. Stochastic event reconstruction of atmospheric contaminant dispersion using Bayesian inference. *Atmospheric Environment*, 42(33):7718–7727, 2008.
- [41] S. A. Slinn and W. G. N. Slinn. Predictions for particle deposition on natural waters. *Atmospheric Environment*, 14(9):1013–1016, 1980.
- [42] J. M. Stockie. The mathematics of atmospheric dispersion modeling. *SIAM Review*, 53(2):349–372, 2011.
- [43] A. M. Stuart. Inverse problems: A Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.

- [44] O. G. Sutton. A theory of eddy diffusion in the atmosphere. *Proceedings of the Royal Society of London, Series A*, 135(826):143–165, 1932.
- [45] G. I. Taylor. Eddy motion in the atmosphere. *Philosophical Transactions of the Royal Society of London, Series A*, 215:1–26, 1915.
- [46] Thermo Scientific, Waltham, MA. *High-volume air samplers (product description)*, 2015. <http://www.thermoscientific.com/en/product/high-volume-air-samplers.html>.
- [47] Trail Area Health and Environment Committee. Trail Area Health & Environment Program Document, September 9, 2014. <http://www.thep.ca/pages/reports/>.
- [48] R. Turner and T. Hurst. Factors influencing volcanic ash dispersal from the 1995 and 1996 eruptions of Mount Ruapehu, New Zealand. *Journal of Applied Meteorology*, 40(1):56–69, 2001.
- [49] U.S. Environmental Protection Agency, Washington, DC. *Guideline on Air Quality Models*, 2010. Appendix W to Part 51, Title 40: Protection of the Environment, Code of Federal Regulations. Source: <http://www.gpo.gov>.
- [50] C. R. Vogel. *Computational Methods for Inverse Problems*, volume 23 of *Frontiers in Applied Mathematics*. SIAM, Philadelphia, PA, 2002.
- [51] D. Wade and I. Senocak. Stochastic reconstruction of multiple source atmospheric contaminant dispersion events. *Atmospheric Environment*, 74:45–51, 2013.
- [52] C. K. I. Williams and C. E. Rasmussen. *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA, 2006.
- [53] Z. Zlatev and I. Dimov. *Computational and Numerical Challenges in Environmental Modelling*. Elsevier, Amsterdam, 2006.