

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA



***National Institute for Applied Statistics Research
Australia***

University of Wollongong, Australia

Working Paper

09-23

**Emulation of Greenhouse-Gas Sensitivities using
Variational Autoencoders**

Laura Cartwright, Andrew Zammit-Mangion, and Nicholas Deutscher

*Copyright © 2023 by the National Institute for Applied Statistics Research Australia, UOW.
Work in progress, no part of this paper may be reproduced without permission from the Institute.*

National Institute for Applied Statistics Research Australia, University of Wollongong,
Wollongong NSW 2522, Australia T: +61 2 42215076. E: karink@uow.edu.au

Emulation of greenhouse-gas sensitivities using variational autoencoders

Laura Cartwright¹, Andrew Zammit-Mangion¹, and Nicholas M. Deutscher²

¹School of Mathematics and Applied Statistics, University of Wollongong, Wollongong, Australia

²Centre for Atmospheric Chemistry, School of Earth, Atmospheric and Life Sciences, University of Wollongong, Wollongong, Australia

Correspondence: Laura Cartwright (lcartwri@uow.edu.au)

Abstract

Flux inversion is the process by which sources and sinks of a gas are identified from observations of gas mole fraction. The inversion often involves running a Lagrangian particle dispersion model (LPDM) to generate sensitivities between observations and fluxes over a spatial domain of interest. The LPDM must be run backward in time for every gas measurement, and this can be computationally prohibitive. To address this problem, here we develop a novel spatio-temporal emulator for LPDM sensitivities that is built using a convolutional variational autoencoder (CVAE). With the encoder segment of the CVAE, we obtain approximate (variational) posterior distributions over latent variables in a low-dimensional space. We then use a spatio-temporal Gaussian process emulator on the low-dimensional space to emulate new variables at prediction locations and time points. Emulated variables are then passed through the decoder segment of the CVAE to yield emulated sensitivities. We show that our CVAE-based emulator outperforms the more traditional emulator built using empirical orthogonal functions and that it can be used with different LPDMs. We conclude that our emulation-based approach can be used to reliably reduce the computing time needed to generate LPDM outputs for use in high-resolution flux inversions.

1 Introduction

It is becoming increasingly urgent to not only identify, but also quantify, the contributors to climate change (see, e.g., Annan, 2002; Flannery, 2009; Zidanšek et al., 2009; Reddy et al., 2010). The major contributor to the enhanced greenhouse effect is the emission of gases that have the capacity for trapping heat in the atmosphere. Methane (CH_4) is the second most prevalent anthropogenic greenhouse gas (GHG) in the atmosphere after carbon dioxide (CO_2); however, by mass it is much more potent than CO_2 (1 metric tonne of CH_4 is approximately 28-34 times more effective at trapping energy than 1 metric tonne of CO_2 over a hundred year period (Myhre et al., 2013)). Therefore, identifying

and quantifying GHG emissions, especially CH₄ emissions, is of considerable interest to atmospheric scientists and policy makers alike. This is, however, a difficult problem: while measurements of CH₄ mole fraction can be done reasonably accurately using sophisticated instruments, estimating the spatial distribution of sources of CH₄ from those measurements is an ill-posed problem.

Flux inversion is at the core of many methods for quantifying GHG emissions (recent examples include Humphries et al., 2012; Luhar et al., 2014; Alexe et al., 2015; Ganesan et al., 2015; Zammit-Mangion et al., 2015; Houweling et al., 2017; Lucas et al., 2017; Bergamaschi et al., 2018; Cressie, 2018; Feitz et al., 2018; Cartwright et al., 2019). It can be summarised as the process of comparing observed GHG mole fractions to those simulated using an atmospheric transport model (ATM), in order to estimate the fluxes that are contributing to the observed mole fractions. Often, this process involves “backward simulation”, where one simulates particle trajectories at an observation location backward in time through an ATM. The derived relationship between each observation and the flux from this simulation is a *sensitivity* that evolves in both space and time. These sensitivities can be time-integrated over the duration of the simulation, to yield a *sensitivity plume* that identifies the sensitivity of a specific observation to the (here, assumed time-invariant) flux field. These sensitivity plumes can then be used in a Bayesian inference framework to locate and quantify the sources or sinks (e.g., Zammit-Mangion et al., 2016).

Many highly sophisticated ATMs exist, and in this work we consider a family of models called Lagrangian Particle Dispersion Models (LPDMs). Some examples of LPDMs are the Numerical Atmospheric Modelling Environment (NAME, Jones et al., 2007), the FLEXible PARTicle Dispersion Model (FLEXPART, Pisso et al., 2019), and the Stochastic Time-Inverted Lagrangian Transport Model (STILT, Lin et al., 2003). LPDMs are used to simulate atmospheric dispersion of particles moving through the air. Trajectories are computed based on the Brownian motion of the particles (see De Visscher, 2013, Chapter 9, for a thorough explanation). LPDMs involve a large number of parameters (often well over 50), and though highly representative of the physical environment, their large computational cost can be prohibitive in flux inversion applications (Francom et al., 2019). This can lead, for example, to flux inversions being performed over coarser spatio-temporal resolutions than desired, restricting one from resolving fine-scale components of interest.

The computational burden associated with running LPDMs has led to the investigation of ways in which they can be used to inform the construction of surrogate models. For example, Lucas et al. (2017) use decision/regression trees to predict outputs from a hybrid of the FLEXPART and Weather Research and Forecasting (WRF) models, FLEXPART-WRF. Over small (local-scale) regions, simpler models such as the Gaussian plume dispersion model are often used (e.g., Borysiewicz et al., 2012; Feitz et al., 2018; Cartwright et al., 2019). While much less sophisticated than an LPDM, the Gaussian plume dispersion model is quick to evaluate, allowing it to be easily calibrated (a procedure in which numerical model parameters are adjusted so that the simulated concentrations are a good fit to measurements) during the inversion. However, such an approach cannot be reliably used over regions larger than a few square kilometres, since the assumptions of the Gaussian plume dispersion model (e.g., constant windspeed and constant vertical Eddy

diffusion) are not reasonable over large distances (Riddick et al., 2017).

The core idea of emulating LPDM output is not new in this article. For example, Harvey et al. (2018) present a polynomial-based model, while Francom et al. (2019) use Bayesian multivariate adaptive regression splines as the emulator basis. In the former case, the emulator is used directly on a subset of inputs to the LPDM, while in the latter, the emulator also takes information from the principal components of simulated outputs, using empirical orthogonal functions (EOFs). However, in both cases, the emulated values represent atmospheric pollutant mole fractions downwind from a *single* point source, at a small number of locations and time points during the simulation. Further, measurements of the atmospheric pollutant (volcanic ash, for example) used in these works are unlikely to be affected by other sources. Our more general case is different, as measurements of methane are likely to be affected by multiple sources, and we need to consider a flux field consisting of many (not necessarily known) sources and sinks. To accommodate these requirements, we propose “borrowing strength” from sensitivity plumes obtained via an LPDM at a small number of locations in order to construct additional plumes at new locations. That is, we extract information about the shape/structure of plumes, and then use this information to emulate the LPDM output at additional spatio-temporal locations. The motivation for this can be seen in Figure 1, where we note that sensitivity plumes for observations of methane that are nearby in space and time appear to be highly correlated, likely due to the smoothly varying nature of the governing meteorological influences (e.g., wind speed, wind direction, air pressure). Our approach allows us to construct time-integrated methane sensitivity plumes over the entire spatio-temporal domain of interest with relative ease.

While it is possible to build an emulator that predicts the LPDM sensitivity at each individual cell in the model grid, it would be difficult to construct in practice, since popular emulators such as the Gaussian process (GP) emulator (see Rasmussen and Williams, 2006, for example) do not scale well with the dimensionality of the parameter space. We instead propose a two-fold approach. First, we extract “important features” of the LPDM-built sensitivity plumes, and use these features to reduce the dimension of the output down to a small number, with each dimension representing a prominent feature of the sensitivity plumes. Second, we emulate variables associated with these features at new spatio-temporal locations. By applying a reverse transformation on the features we then produce predicted LPDM-sensitivity-plumes on the original space.

For the first, dimension-reduction, step, a number of methods are available, including principal component analysis (PCA) or EOF analysis (e.g., Higdon et al., 2008; Zhang et al., 2018; Francom et al., 2019), gradient-based kernels (e.g., Liu and Guillas, 2017), sliced inverse regression (e.g., Li, 1991), and variable selection algorithms (e.g., Wolf and Bileschi, 2005; Harvey et al., 2018). Machine learning methods are also popular for dimension reduction. When carrying out dimensionality reduction where reconstruction quality of inputs is important, in machine learning one often considers the autoencoder (see Goodfellow et al., 2016, chapter 14). An autoencoder first reduces the dimension of inputs by mapping them, via an encoder, onto a latent space. Via a decoder, it then takes values on the latent space and constructs an approximation to the original input via a reverse mapping. In this article we propose the use of a convolutional variational autoencoder (CVAE); see Kingma and Welling (2014) for a full introduction to variational

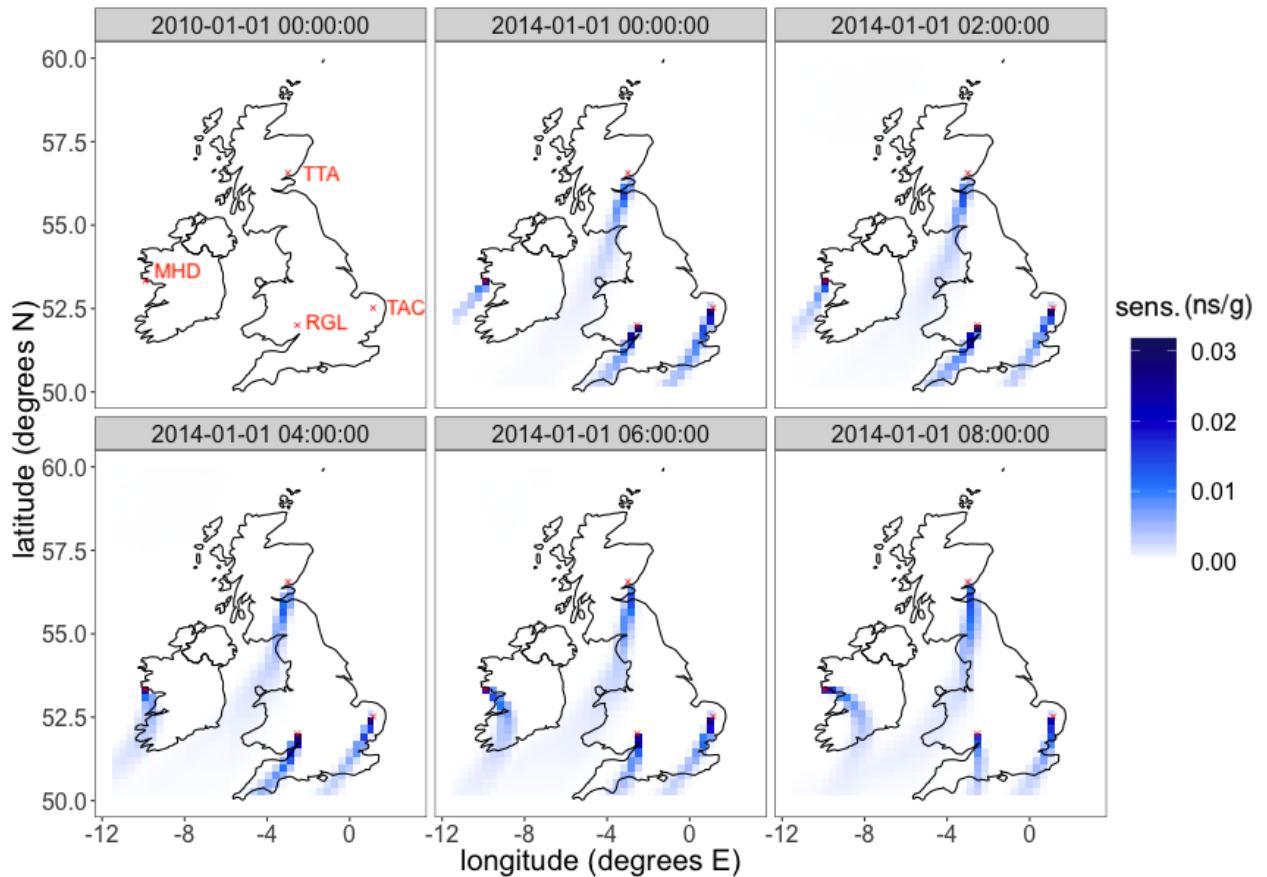


Figure 1: Top-left panel: Four methane measurement sites from the UK and Ireland, located at Mace Head (MHD), Ridge Hill (RGL), Tacolneston (TAC), and Angus (TTA). Other panels: The sensitivity of mole fraction to the flux over two-hourly time intervals on 1 January 2014, based on a release from each of the four sites, as output from the Numerical Atmospheric Dispersion Modeling Environment (NAME). Here “sens.” represents the maximum sensitivity at each grid cell, where the maximum is taken over each of the four NAME outputs from the corresponding time point (one per measurement site).

autoencoders. The “variational” portion of the name arises from the variational-Bayes-based probabilistic model used to build the autoencoder, which allows us to avoid overfitting and increase the generalisability of the fitted model.

For the second, emulation, step, we use spatio-temporal GP emulators on the latent space, which allow us to quantify the uncertainty arising from not running the LPDM at new spatio-temporal locations of interest. We see improvement in the emulated plumes, both visually and in terms of mean-squared error, when the CVAE is used for dimensionality reduction instead of EOFs.

The remainder of the article is organised as follows. In Section 2, we describe the methodology behind EOF analysis and the CVAE, and give a brief description of the GP emulators we employ. In Section 3, we present and discuss the results we obtain when emulating the sensitivity plumes from two LPDMs. Section 4 concludes.

2 Methodology

Consider a rectangular spatial domain $D \subset [-180, 180]^\circ \times [-90, 90]^\circ$. Let $\mathbf{s} \equiv (s_1, s_2)'$ be a spatial location in D with units of degrees, and $\mathbf{w} \equiv (\mathbf{s}'; t)'$, where t denotes time in UTC. Define $Y_m(\mathbf{w})$ to be the spatio-temporal methane mole fraction process at \mathbf{w} , and $Y_f(\mathbf{s})$ to be the (assumed time-invariant) spatial methane flux process at \mathbf{s} . The *sensitivity* of Y_m to Y_f depends on a bivariate function $b(\cdot, \cdot)$, and is given by (Zammit-Mangion et al., 2015),

$$Y_m(\mathbf{w}) = \int_D b(\mathbf{w}, \mathbf{r}) Y_f(\mathbf{r}) d\mathbf{r}. \quad (1)$$

In practice, the integral in (1) can rarely be done in closed form, and is instead approximated using an ATM (such as an LPDM in this case). Specifically, the sensitivity is evaluated numerically over a fine space-time grid.

Let D^G be a rectangular spatial gridding of D , consisting of K equally-sized grid cells, and let $\mathbf{s}_k \equiv (s_{k,1}, s_{k,2})'$, $k \in 1, 2, \dots, K$, be the lower-left corner of the k th grid cell in D^G , where $s_{k,1}$ corresponds to degrees longitude, and $s_{k,2}$ to degrees latitude. We are interested in how the spatio-temporal mole fraction process $Y_m(\cdot)$ at a collection of N points, $\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N$, relates to the spatial flux process, $Y_f(\cdot)$, where $\mathbf{w}_i \equiv (\tilde{\mathbf{s}}_i'; t_i)'$, $\tilde{\mathbf{s}}_i \equiv (\tilde{s}_{i,1}, \tilde{s}_{i,2})' \in D$, and t_i is a time point. Following the discretisation of (1), we have the linear relationship

$$\mathbf{Y}_m = \mathbf{B} \mathbf{Y}_f, \quad (2)$$

where $\mathbf{Y}_m \equiv (Y_m(\mathbf{w}_i): i = 1, 2, \dots, N)'$, and $\mathbf{Y}_f \equiv (Y_f(\mathbf{s}_k): k = 1, 2, \dots, K)'$. The i th row of \mathbf{B} , denoted \mathbf{b}_i , describes the sensitivity of $Y_m(\mathbf{w}_i)$ to the flux at all $\mathbf{s}_k, k = 1, 2, \dots, K$. We determine each \mathbf{b}_i using an LPDM.

Suppose now that we have an additional set of M points, $\mathbf{w}_1^*, \mathbf{w}_2^*, \dots, \mathbf{w}_M^*$, at which we wish to model the mole fraction in response to \mathbf{Y}_f , where for $j = 1, 2, \dots, M$, $\mathbf{w}_j^* \equiv (\tilde{\mathbf{s}}_j^*; t_j^*)$, $\tilde{\mathbf{s}}_j^* \equiv (\tilde{s}_{j,1}^*, \tilde{s}_{j,2}^*)'$ is a point in D , t_j^* is a point in time, $t_1 \leq t_1^*$, and $t_N \geq t_M^*$. The primary aim of this work is to construct an emulator that can be used to construct a new matrix $\hat{\mathbf{B}}^*$ from \mathbf{B} , where each row, $\hat{\mathbf{b}}_j^*$, is the *predicted* sensitivity of $Y_m(\mathbf{w}_j^*)$ to the flux at all $\mathbf{s}_k \in D^G$. The construction of such an emulator is the subject of the following sections.

2.1 Plume feature extraction

Before constructing the emulator, we first focus on reducing the dimension of the plumes. We do this since emulating a K -dimensional quantity for large K is difficult, and also wasteful due to sensitivity vectors containing a lot of structure. As can be seen from Figure 1, much of the sensitivity over D^G is often zero, and the ‘‘important information’’ about the sensitivity of $Y_m(\mathbf{w}_i)$ is contained in only a small subset of \mathbf{b}_i . This suggests that we could capture the important information regarding the shape (structure) and intensity of these plumes through a small set of features.

We now describe two methods for feature extraction/dimension reduction, one using EOFs, and one using a CVAE.

2.1.1 Empirical orthogonal functions

EOF analysis (Halldor and Venegas, 1997) refers to the process of decomposing the spatio-temporal data into a set of orthogonal basis vectors that can help describe the spatial structure of the data. EOF analysis is used to extract spatial features from data indexed over time, and hence can be thought of as a spatio-temporal version of principal component analysis (PCA, Wikle et al., 2019, Chapter 2). The basis vectors are constructed successively in such a way that each is orthogonal to those constructed previously, and in a way that they account for as much spatial variability in the data as possible. That is, the first basis vector explains the largest amount of spatial variability in the data, while the second basis vector explains the second largest amount of spatial variability, and so on.

We derive the EOFs via the singular value decomposition (SVD) of \mathbf{B} , which is of size $N \times K$. The SVD of a matrix has three components, one of which contains the singular values of the matrix. Specifically, when $N > K$,

$$\mathbf{B} = \mathbf{U}\mathbf{D}\mathbf{V}',$$

where \mathbf{U} is an $N \times K$ orthonormal matrix of left singular vectors, \mathbf{D} is a $K \times K$ matrix with the singular values along its main diagonal, and \mathbf{V} is a $K \times K$ orthonormal matrix of right singular vectors. It can be easily shown that the columns of \mathbf{V} are the K eigenvectors of $\mathbf{B}'\mathbf{B}$, and that the diagonal of \mathbf{D} contains the square-root of the eigenvalues. The product $\mathbf{D}\mathbf{V}'$ makes up the EOFs, and each column in \mathbf{U} represents the coefficients for one EOF.

We construct an approximate representation of \mathbf{B} using the first r EOFs, by computing

$$\mathbf{B} \approx \mathbf{U}_r\mathbf{D}_r\mathbf{V}_r', \quad (3)$$

where \mathbf{U}_r , and \mathbf{V}_r are the first r columns of the matrices \mathbf{U} and \mathbf{V} , respectively, and \mathbf{D}_r , consists of the first r rows and columns of \mathbf{D} . The i th row in \mathbf{U}_r , denoted \mathbf{u}_i , represents the coefficients for the first r EOFs for the i th plume. When using EOFs, we therefore emulate (Section 2.2) based on the information in the rows of \mathbf{U}_r , to predict sensitivities at new locations. Specifically, for spatio-temporal locations $\mathbf{w}_1^*, \mathbf{w}_2^*, \dots, \mathbf{w}_M^*$, we obtain an $M \times r$ matrix of coefficients through emulation, $\hat{\mathbf{U}}_r^*$. We then construct $\hat{\mathbf{B}}^*$ (of size $M \times K$) by computing

$$\hat{\mathbf{B}}^* = \hat{\mathbf{U}}_r^*\mathbf{D}_r\mathbf{V}_r'.$$

This approach to emulation is often seen in the computer experiment literature and has been shown to be successful in various application settings (e.g., Higdon et al., 2008; Holden et al., 2015; Zhang et al., 2018).

2.1.2 Convolutional variational autoencoder

While EOFs are widely used in the atmospheric sciences, they can also be restrictive. For example, Hannachi et al. (2007) notes that the spatial orthogonality of EOFs restricts their ability to describe localised patterns within the domain, and can lead to difficulty in pattern interpretation as well as domain-dependence issues. In addition, the SVD yields a linear decomposition; as a result, underlying non-linear relationships between the “inputs”

and their “features” could be overlooked (Bishop and Winn, 2000). In this section we explore the use of a deep learning architecture to extract the basis functions with which to do emulation. A deep learning model is a collection of simple, interconnected layers of “neurons”, often used to extract important features from a dataset. Neural networks feature a collection of weights that are optimised according to one or more criteria during the training of the network (Gurney, 2003, Chapter 1). In a fully-connected neural network (i.e., when each neuron in one layer is connected to all neurons in the next layer) for image data, a very large number of weights are required. As images have a known, gridded structure, it is possible to use convolutional layers instead, which are not fully connected (Soydaner, 2019). In a convolutional neural network (CNN), some or all of the layers use a convolution in place of the standard matrix multiplication used in fully-connected neural networks.

In practice, we often not only wish to find features of an image, but also to use these features to reconstruct the original image. The architecture used for this purpose is known as a *convolutional autoencoder*, which can be thought of as a two-piece CNN: an encoder, and a decoder. The encoder maps each image onto a latent space of variables \mathbf{u} , which represent the value of the extracted features. Then, these features can be mapped back onto the original space (reconstructing the original image) via the decoder. The CNN autoencoder has been used for prediction elsewhere (e.g., Xu and Duraisamy, 2020), however it does not require that the latent space be contiguous (Song et al., 2019). This can be problematic if one wishes to use the decoder for emulation, as in our application, since, as with most emulators, we rely on the output varying smoothly with the input when predicting at a new location.

The *variational autoencoder* (VAE, Kingma and Welling, 2014) differs from the standard autoencoder in that the latent variables are no longer deterministic, but are now random variables. In addition to the mean squared error, the CVAE objective function also includes a Kullback-Leibler divergence (KL-divergence), which penalises for (and, in practice, does not allow) discontinuities in the latent space.

Consider N vectorised training images $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N$, each of dimension $K \times 1$, and an r -dimensional latent variable $\mathbf{u}_i \equiv (u_{i,1}, u_{i,2}, \dots, u_{i,r})'$ for each $\mathbf{b}_i, i = 1, 2, \dots, N$. In a CVAE, we seek to maximise a lower bound for the marginal likelihood $p(\mathbf{b}_i; \boldsymbol{\theta})$, where $\boldsymbol{\theta} \equiv (\boldsymbol{\theta}'_e, \boldsymbol{\theta}'_d)'$, and $\boldsymbol{\theta}_e$ and $\boldsymbol{\theta}_d$ are the encoder and decoder parameters, respectively, that need to be estimated. In particular, we have that

$$\begin{aligned} p(\mathbf{b}_i; \boldsymbol{\theta}) &= \int p(\mathbf{b}_i, \mathbf{u}_i; \boldsymbol{\theta}) \, d\mathbf{u}_i \\ &= \int q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) \frac{p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d) p(\mathbf{u}_i)}{q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)} \, d\mathbf{u}_i, \end{aligned}$$

where $q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)$ is the approximation to the posterior distribution $p(\mathbf{u}_i | \mathbf{b}_i, \boldsymbol{\theta}_d)$ when making inference using variational Bayes. By Jensen’s inequality, we therefore have that

$$\begin{aligned} \log(p(\mathbf{b}_i; \boldsymbol{\theta}_e)) &\geq \int q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) \log \left(\frac{p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d) p(\mathbf{u}_i)}{q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)} \right) \, d\mathbf{u}_i \\ &= \mathbb{E}_{q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)} [\log(p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d))] - \text{KL}(q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) || p(\mathbf{u}_i)). \quad (4) \end{aligned}$$

Maximising (4) is equivalent to minimising

$$-\mathbb{E}_{q(\mathbf{u}_i|\mathbf{b}_i;\boldsymbol{\theta}_e)} [\log(p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d))] + \text{KL} (q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) || p(\mathbf{u}_i)). \quad (5)$$

We let $\mathbf{u}_i \sim \text{Gau}(\mathbf{0}, \mathbb{I}_r)$ and $q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) = \text{Gau}(\mathbf{e}(\mathbf{b}_i; \boldsymbol{\theta}_e), \boldsymbol{\Sigma}_e(\mathbf{b}_i; \boldsymbol{\theta}_e))$, $i = 1, 2, \dots, N$, where $\mathbf{e}(\cdot; \boldsymbol{\theta}_e) \equiv (e_1(\cdot; \boldsymbol{\theta}_e), e_2(\cdot; \boldsymbol{\theta}_e), \dots, e_r(\cdot; \boldsymbol{\theta}_e))'$ is the encoder mean, and

$$\boldsymbol{\Sigma}_e(\mathbf{b}_i; \boldsymbol{\theta}_e) \equiv \text{diag}(\sigma_{e,1}^2(\mathbf{b}_i; \boldsymbol{\theta}_e), \sigma_{e,2}^2(\mathbf{b}_i; \boldsymbol{\theta}_e), \dots, \sigma_{e,r}^2(\mathbf{b}_i; \boldsymbol{\theta}_e)),$$

is the (diagonal) encoder covariance matrix, with $\sigma_{e,\gamma}^2(\mathbf{b}_i; \boldsymbol{\theta}_e)$, $\gamma = 1, 2, \dots, r$, nonnegative functions of \mathbf{b}_i . Under these assumptions, the second term on the right-hand side of (5) becomes

$$\text{KL} (q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) || p(\mathbf{u}_i)) = -\frac{1}{2} \sum_{\gamma=1}^r (\log(\sigma_{e,\gamma}^2(\mathbf{b}_i; \boldsymbol{\theta}_e)) + 1 - \sigma_{e,\gamma}^2(\mathbf{b}_i; \boldsymbol{\theta}_e) - e_\gamma(\mathbf{b}_i; \boldsymbol{\theta}_e)^2).$$

Therefore, under the Gaussian prior and the chosen approximate form of the posterior distribution, the KL-divergence between the approximate posterior distribution and the prior on \mathbf{u}_i can be calculated using the encoder-determined mean and variance (in practice, log-variance) for each $u_{i,\gamma}$ (Kingma and Welling, 2014).

For the first term on the right-hand side of (5), we assume a Gaussian likelihood,

$$p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d) = \text{Gau}(\mathbf{d}(\mathbf{u}_i; \boldsymbol{\theta}_d), \boldsymbol{\Sigma}_d(\mathbf{u}_i)),$$

where $\mathbf{d}(\cdot; \boldsymbol{\theta}_d)$ is the decoder mean, and $\boldsymbol{\Sigma}_d(\mathbf{u}_i)$ is the decoder covariance matrix. Thus, we assume that the reconstructed images are corrupted with Gaussian noise which, for simplicity, we assume is homoscedastic and uncorrelated; that is, we let $\boldsymbol{\Sigma}_d(\mathbf{u}_i) = c\mathbb{I}_M$, $c > 0$. Now, minimisation of

$$-\mathbb{E}_{q(\mathbf{u}_i|\mathbf{b}_i;\boldsymbol{\theta}_e)} [\log(p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d))]$$

cannot be done analytically since $\mathbf{d}(\cdot; \boldsymbol{\theta}_d)$ is an arbitrary non-linear function of its arguments. We therefore approximate this expectation using Monte Carlo. In particular,

$$-\mathbb{E}_{q(\mathbf{u}_i|\mathbf{b}_i;\boldsymbol{\theta}_e)} [\log(p(\mathbf{b}_i | \mathbf{u}_i; \boldsymbol{\theta}_d))] \approx \frac{1}{L} \sum_{l=1}^L \frac{1}{2c} \left\| \mathbf{b}_i - \mathbf{d}(\mathbf{u}_i^{(l)}; \boldsymbol{\theta}_d) \right\|_2^2 + \text{const.}, \quad (6)$$

where $\mathbf{u}_i^{(l)}$ is the l th ($l = 1, 2, \dots, L$) random sample from $q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)$, and ‘‘const.’’ does not depend on the decoder parameters. In our implementation, we therefore train the CVAE by minimising

$$\sum_{i=1}^N \left[\frac{1}{L} \sum_{l=1}^L \left\| \mathbf{b}_i - \mathbf{d}(\mathbf{u}_i^{(l)}; \boldsymbol{\theta}_d) \right\|_2^2 + \lambda \text{KL} (q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e) || p(\mathbf{u}_i)) \right], \quad (7)$$

with respect to $\boldsymbol{\theta}$, where $\lambda = 2c$ can be thought of as a tuning parameter. Typically, L is set to be small to encourage exploration of the objective function when optimising. As in Kingma and Welling (2014), we set $L = 1$.

Once trained (i.e., estimates of $\boldsymbol{\theta}$ are available), known sensitivity plumes, $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_N$, can be put through the encoder to obtain their corresponding variational means and

variances on the latent space. The spatio-temporal emulator (Section 2.2) is then used on this latent space to obtain predictions and prediction variances for new variational means, corresponding to $\mathbf{w}_j^*, j = 1, 2, \dots, M$.

2.2 Emulation

Here, we describe the way we use the information gathered from feature extraction of \mathbf{B} (Section 2.1) to build an emulator that can predict plumes at new spatio-temporal locations, which can then be used to construct the matrix $\widehat{\mathbf{B}}^*$. With a slight abuse of notation, denote the r coefficients for the features extracted for the i th plume by one of the two methods in Section 2.1 as $\mathbf{u}_i \equiv (u_{i,1}, u_{i,2}, \dots, u_{i,r})'$, where $i = 1, 2, \dots, N$ (these would correspond to either rows of the matrix \mathbf{U} for the EOFs, or variational means for the CVAE). We describe below the emulators we use to predict the elements of each coefficient in \mathbf{u}_j^* , at a new spatio-temporal location \mathbf{w}_j^* , for $j = 1, 2, \dots, M$.

Let $\mathbf{W} \equiv (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_N)'$, and collect the coefficients of feature γ of the plumes in $[\mathbf{u}]_\gamma$; that is, define $[\mathbf{u}]_\gamma \equiv (u_{i,\gamma} : i = 1, \dots, N)'$. We use independent Gaussian process emulators (Rasmussen and Williams, 2006, Chapter 2), so that the prediction of the coefficient of feature γ for a plume at \mathbf{w}_j^* is normally distributed with mean

$$S_\gamma(\mathbf{w}_j^*, \mathbf{W})S_\gamma(\mathbf{W}, \mathbf{W})^{-1}[\mathbf{u}]_\gamma, \quad \gamma = 1, 2, \dots, r, \quad (8)$$

and variance

$$S_\gamma(\mathbf{w}_j^*, \mathbf{w}_j^*) - S_\gamma(\mathbf{w}_j^*, \mathbf{W})S_\gamma(\mathbf{W}, \mathbf{W})^{-1}S_\gamma(\mathbf{W}, \mathbf{w}_j^*), \quad \gamma = 1, 2, \dots, r, \quad (9)$$

where $S_\gamma(\cdot, \cdot)$ is the spatio-temporal squared-exponential kernel. Specifically, for two spatio-temporal points $\mathbf{w}_a = (\tilde{\mathbf{s}}'_a; t_a)'$ and $\mathbf{w}_b = (\tilde{\mathbf{s}}'_b; t_b)'$,

$$S_\gamma(\mathbf{w}_a, \mathbf{w}_b) = \sigma_{st,\gamma}^2 \exp\left(-\frac{\|\tilde{\mathbf{s}}_a - \tilde{\mathbf{s}}_b\|_2^2}{2l_{s,\gamma}} - \frac{|t_a - t_b|}{2l_{t,\gamma}}\right)^2, \quad \sigma_{st,\gamma}^2, l_{s,\gamma}, l_{t,\gamma} \in \mathbb{R}^+,$$

where $l_{s,\gamma}$, $l_{t,\gamma}$, and $\sigma_{st,\gamma}^2$ denote the spatial length scale, the temporal length scale, and the variance, of the Gaussian process for feature γ , respectively. We obtain estimates for \mathbf{u}_j^* , denoted $\widehat{\mathbf{u}}_j^*$, via (8), and the corresponding covariance matrix via (9). This gives us a Gaussian distribution over each emulated feature. We then draw samples from this distribution and pass them through the decoder in order to obtain a Monte Carlo approximation of the predictive distribution of the sensitivity plume at \mathbf{w}_j^* .

3 Results & Discussion

In this section, we present and discuss the results obtained when applying our methodology described in Section 2 to sensitivities computed from LPDMs.

3.1 Dimension reduction and reconstruction

To train the CVAE and extract the EOFs, we obtained 20000 sensitivity plumes from individual simulation runs using FLEXPART. In each simulation, 2000 particles were released from a spatio-temporal location, and their trajectories followed backward in time

Table 1: Details of the four regions over which FLEXPART simulations were run. Spatio-temporal locations for simulations were chosen on a subset of the simulation region in order to capture plume tails. Latitude is given in degrees North, longitude in degrees East, and time is of the format YYYY-MM-DD hh:mm:ss UTC.

Region	Min lat	Max lat	Min lon	Max lon	Min time	Max time
West Europe	46.000	75.952	-3.000	42.056	2011-02-02 00:00:00	2011-07-01 00:00:00
Australia	-40.000	-10.048	112.000	157.056	2012-09-02 00:00:00	2012-12-31 00:00:00
North-West America	30.000	59.952	-116.000	-70.944	2013-02-02 00:00:00	2013-07-01 00:00:00
United Kingdom & Ireland	36.000	65.792	-15.000	30.056	2014-09-02 00:00:00	2014-12-31 00:00:00

for 30 days. Five thousand release locations were randomly chosen from within each of four spatio-temporal regions. Releases occurred over West Europe, Australia, North-West America, and the United Kingdom/Ireland, though the simulation region in each case was larger to allow plume tails to be captured. The spatial domain for each simulation was a 128×128 grid of resolution $0.352^\circ \times 0.234^\circ$, with boundaries described in Table 1. At regular temporal intervals over the 30 days, the length of time which particles spent within the first 100 m above ground level was recorded for each grid cell, in units of $\text{m}^3 \text{s g}^{-1}$ (Seibert and Frank, 2004). We then converted these quantities to time-integrated sensitivities (sometimes referred to as footprints, e.g., Oney et al. (2015)) in units of ns g^{-1} (see Appendix A for details). Finally, any sensitivity plumes deemed to have weak signal (defined as plumes with at most 10 grid cells with a sensitivity above the lowest 99.5% of all sensitivities across all simulations) were removed from the dataset. To minimise spatial variability between all plumes as much as possible before training, we then subtracted the location of each plume’s origin from all $\mathbf{s}_k, k = 1, 2, \dots, K$, and rotated each plume such that its departure angle was 0° (the horizontal (East-West) axis in our images). This required us to estimate the original departure angle of each plume (details on how this was done are given in Appendix B). After rotation, inverse distance weighting was used to reduce the resolution of the images down to 64×64 , to help reduce the computational burden when training the CVAE.

The EOFs were obtained using 70% (11470) of the filtered set of simulated FLEXPART plumes. We retained the first $r = 20$ singular values (\mathbf{D}_{20}) and eigenvectors (\mathbf{V}_{20}), along with their corresponding coefficients (\mathbf{U}_{20}). To assess the error due to dimensionality reduction, we computed a mean squared error (MSE) for each plume by averaging the sum

Table 2: The sumMSEs (in units of $\text{ns}^2 \text{g}^{-2}$) between the plumes and their reconstructions. The FLEXPART sumMSEs are calculated after reducing the dimension of the plumes down to $r = 20$, and then reconstructing them again, either using EOFs or via the CVAE. The NAME sumMSEs are calculated after emulation, rotation, and translation.

	FLEXPART (Training)	NAME (Application)
EOF	0.00697	0.000312
CVAE	0.00274	0.000260

of the differences between the true values and the reconstructed values of the sensitivities at each grid cell. As shown in the top-left cell of Table 2, the sum of MSEs (sumMSE) for all plumes after EOF reconstruction was $0.00697 \text{ ns}^2 \text{g}^{-2}$. The middle panel in both rows of Figure 2 shows the reconstruction from the first 20 EOFs of two of the plumes used to obtain the EOFs. As can be seen, particularly in the second case, the tail end of the plume is not well captured by the EOF reconstruction.

The CVAE was trained using the same 70% of simulated FLEXPART plumes as used to obtain the EOFs. The remaining 30% were used for validation during training, and we set the latent dimension to $r = 20$. The encoder portion of the CVAE architecture consisted of six 2D convolutional layers with max pooling and scaled exponential linear unit (selu) activation functions, a flattening layer, and then two dense layers with leaky rectified linear unit (leaky relu) activation functions, one for the latent mean vector, and one for the latent log-variance vector. The decoder portion of the CVAE architecture consisted of a single dense layer, followed by a layer which reshaped the variables into 2D images, and then six 2D convolutional transpose layers with selu activation functions. A schematic of the full architecture is shown in Figure 3 for a latent space of dimension $r = 1$. Using cross-validation, we found that our model was largely insensitive to small values of λ , and that it underfit for large λ . This suggests that it is relatively easy to estimate the mean and variance encoding maps, as well as the mean decoder map, from the data, and that not much regularisation is needed in our case. We therefore fixed λ to 10^{-9} ; note that a small value for λ does not mean the resulting network is a standard autoencoder since each \mathbf{u}_i is still associated with some non-trivial variational mean and variance in $q(\mathbf{u}_i | \mathbf{b}_i; \boldsymbol{\theta}_e)$. Training was done using batches constructed from random samples without replacement of inputs of a fixed size ($< N$) from the full training data set. The CVAE objective shown in (7) was then computed by summing over the batch rather than the full data set. This procedure was repeated until the full training set was sampled several times (a single run through the entire data set is known as an *epoch*). Parameters and weights were randomly initialised, and variables were randomly sampled on the latent space from distributions defined via the encoded means and log-variances throughout training (see (6)). Due to the complexity of the objective function and the sensitivity of the parameter estimates to initial conditions, we elected to train the network 10 times over 500 epochs using an adaptive moment (Adam) optimiser. The trained network with the lowest training MSE




Figure 2: Left: Rotated and shifted sensitivity plumes as produced by FLEXPART, at the indicated time and date. Middle: EOF reconstruction of the same plumes, using 20 singular values, eigenvectors, and their corresponding coefficients. Right: CVAE reconstruction of the same plumes, using a 20-dimensional latent space. The MSE (in units of $\text{ns}^2 \text{g}^{-2}$) is shown above each reconstruction.

was then chosen for emulation.

Occasionally, our dimension-reduction procedures yielded negative sensitivities for some grid cells after reconstruction. For example, one randomly selected plume from the CVAE yielded 90 negative sensitivities out of a total of 4096, with the minimum and maximum sensitivities being $-0.0013 \text{ ns g}^{-1}$, and 0.052 ns g^{-1} , respectively. Another randomly chosen plume yielded 6 negative sensitivities, with the minimum and maximum sensitivities being $-0.00084 \text{ ns g}^{-1}$, and 0.018 ns g^{-1} , respectively. Negative sensitivities do not make physical sense, and we therefore set all negative sensitivities (from both dimension-reduction procedures), when these arose after reconstruction, to zero. The right panels in both rows of Figure 2 show the resulting CVAE reconstructions for two training plumes after truncation. When compared to the EOF reconstructions, we see that the CVAE is able to retain more of the important structural information about the plumes during the dimension reduction stage. This is further supported by the sumMSE value shown in the bottom-left cell of Table 2 ($0.00274 \text{ ns}^2 \text{g}^{-2}$), which is less than half the sumMSE for the EOFs.




Figure 3: Visual representation of the CVAE architecture used in Section 3.1. The vector \mathbf{b}_i represents the true (in practice, rotated and shifted) plumes, while $\hat{\mathbf{b}}_i$ represents the reconstructed plumes. In this figure, u_i is a one-dimensional latent variable. We also show the make-up of the encoder: convolutional layers followed by max-pooling layers and a flattening layer. The decoder is made up similarly with convolutional 2D transpose layers.

Having demonstrated the ability for the CVAE to reduce, and reconstruct, plumes with superior accuracy to the conventional EOF-based approach, we now turn our attention to emulation. We note that the methods we present for dimension reduction are not LPDM-dependent. That is, we do not need to use FLEXPART plumes, despite it having been the LPDM used to generate the data that the CVAE was trained on and that the EOFs were obtained from. As we show next, our models are general enough that we can apply them to plumes from other LPDMs.

3.2 Application

To demonstrate the utility of our surrogate plume models and subsequent emulation, we apply them to plumes simulated using the Numerical Atmospheric-Dispersion Modelling Environment (NAME), which were used for flux inversion in previous works (Ganesan et al., 2015, Zammit-Mangion et al., 2015, Zammit-Mangion et al., 2016). These plumes were simulated over the United Kingdom and Ireland every two hours between 1 January 2014 00:00:00, and 30 April 2014 22:00:00, using a 128×128 spatial grid of resolution $0.352^\circ \times 0.234^\circ$. The particle trajectories were followed backward in time for 30 days. Each sensitivity plume originated at one of the four measurement sites shown in the top-left panel of Figure 1.

Every second plume was removed from the dataset, and those plumes that were removed were then emulated using solely those plumes that were retained. With the CVAE, those which were retained were passed through the encoder in order to find the corresponding variational means. The mean vectors corresponding to the plumes were then used to emulate mean vectors for those plumes we removed from the dataset. We did this by fitting GP emulators to each plume “feature” (see Section 2.2); the parameters $l_{s,\gamma}$, $l_{t,\gamma}$, and $\sigma_{st,\gamma}$ in each case were fit via maximum likelihood estimation (MLE). We also used two additional GP emulators: one for the East-West component of the plume departure angle, and one for the North-South component. We emulated values for both these components for each of the plumes removed from the NAME dataset. In order to account for the uncertainty in the plume reconstruction and the rotation and translation back to the original space, we sampled 100 realisations from each Gaussian distribution on the latent space, and used these realisations to construct, and rotate, 100 emulated plumes to their predicted departure angles. The origins of the plumes were then shifted from $(0, 0)$ to the appropriate spatial location (i.e., the coordinates of one of the measurement sites), and inverse distance weighting was used to bring the spatial resolution back up to 128×128 . Finally, we took the mean and standard deviation over the grid cells to obtain a mean plume and associated prediction standard errors, respectively. After these transformations, the sumMSE, as shown in the bottom-right cell of Table 2, was $0.000260 \text{ ns}^2 \text{ g}^{-2}$.

For comparative purposes, the same analysis was performed using the EOFs. Using the existing matrices \mathbf{D}_r and \mathbf{V}_r found from the FLEXPART training dataset, we used linear regression to find the corresponding coefficients for the 50% of plumes left in the NAME dataset. GP emulators were then fit to each coefficient and both angle components, with parameters again found via MLE. As before, to build an emulated plume, 100 samples from the distribution of each emulated coefficient were drawn. These samples were then combined with the existing singular values and eigenvectors to construct emulated plumes. The mean and variance over the grid cells were computed to construct a mean plume and the associated prediction standard errors. As shown in the top-right cell of Table 2, the sumMSE obtained was $0.000312 \text{ ns}^2 \text{ g}^{-2}$, which is comparable, but slightly worse than what we obtained using the CVAE.

Figure 4 shows the emulated plumes for one prediction location and time point, using both the CVAE and EOFs for dimension reduction. More examples are given in Appendix C.

4 Conclusion

In this article we present a novel method for emulating LPDM sensitivity plumes using a CVAE. The method is not LPDM dependent, making it flexible across many applications in atmospheric science. We show that the CVAE is able to capture important structural information within a collection of FLEXPART plumes in a 20-dimensional latent space, and hence produce reasonable reconstructions of these plumes. We also show through GP emulation on a set of known NAME plumes that we are able to reasonably reconstruct additional new plumes given information from nearby, known plumes. The effectiveness of the CVAE in dimension reduction when compared to the use of EOFs is highlighted. We find that the CVAE outperforms the EOFs both visually and via sumMSE metrics.




Figure 4: Top: Sensitivity plume as produced by NAME on 1 April 2014 at 20:00:00, followed by the emulated reconstruction using the EOFs, and the CVAE, respectively, along with their MSEs (in units of $\text{ns}^2 \text{g}^{-2}$). Bottom: Prediction standard errors for both the EOF and CVAE methods. The emulated plumes and prediction standard errors were obtained from 100 Monte Carlo samples from the conditional distributions following GP emulation. In each case, samples were used to construct, rotate, and translate an emulated plume.

Our framework will allow researchers to save time when performing flux inversions, and allow for the generation of numerous, higher-resolution sensitivity plumes than what was previously possible due to computational constraints. It is difficult to precisely quantify the amount of computational time that could be saved; however, the computing resources needed for this study are revealing. Using all cores on a 64-core server, it took us approximately two days to obtain 3000 plumes using FLEXPART. This includes the time we needed to obtain meteorological data using the software `flex_extract`, and the time needed to post-process the simulation output. This is already substantial, despite the fact that our simulations only involved computing hourly calculations based on the position of the released particles, and ERA-Interim meteorology data which is only at a six-hourly temporal resolution. Using higher-resolution meteorological data would slow down the simulations considerably. Further, if a smaller simulation time step is desired, then the number of released particles must be increased in a commensurate manner (P. Seibert,

communication via FLEXPART mailing list, June 26 2020): Separate timing tests revealed that if sensitivities were needed at five-minute intervals instead of hourly intervals, the time taken for every individual simulation would increase eight-fold. On the other hand, setting up and predicting with our emulator framework described in Section 3.2, took approximately 10 hours. About one hour was needed to perform MLE to optimise each $l_{s,j}$, $l_{t,j}$, and σ_j^2 , and approximately nine hours were needed to perform the MC sampling, and plume rotations and translations with code that was not optimised for speed. The time taken to do the actual emulation was less than a minute. The key attraction of our framework is that the timings associated with our emulator are independent of those associated with the FLEXPART simulations: If the number of particles or the frequency of sampling were to be increased in FLEXPART (resulting in an increase in computational time), the time taken to obtain plumes via our methodology would remain the same. This framework thus presents a huge opportunity for large savings in computational time for more complex simulations.

In future work, it would be interesting to see whether this methodology could be used for emulating plumes associated with column-averaged greenhouse-gas measurements, which is needed when doing inversions from remote-sensing data. If so, this methodology could drastically reduce the computational time associated with flux inversions using data from satellite missions, which retrieve data at many spatial locations and time points, and hence would require many sensitivity plumes if used in a flux inversion setup. It should be noted, however, that our approach may not be useful if plumes become less correlated in space and time as a result of rapidly varying meteorological factors. The utility of our approach to remote-sensing data and to space-time regions with highly variable meteorology are the subject of future research.

Acknowledgements

LC acknowledges the support of the Australian Government Research Training Program Scholarship. AZ-M's research was supported by the Australian Research Council (ARC) Discovery Early Career Research Award (DECRA) DE180100203, and by Discovery Project DP190100180. NMD's research was supported by the ARC Future Fellowship FT180100327. For his help with the FLEXPART and `flex_extract` installations, the authors thank Yi Cao. For provision of the NAME simulations, the authors thank Alistair Manning. For their discussions related to the research, the authors thank Anita Ganesan, Rachel Tunnicliffe, Dan Pagendam, and Michael Bertolacci.

Code and data availability

The code required to do the emulation described in this article are available at <https://github.com/Lcartwright94/LPDM-Sensitivity-Emulation>. Instructions for accessing and downloading the accompanying datasets are also available via this link.

Conflict of interest statement

The authors declare they have no conflict of interest.

References

- Alexe, M., Bergamaschi, P., Segers, A., Detmers, R., Butz, A., Hasekamp, O., Guerlet, S., Parker, R., Boesch, H., Frankenberg, C., Scheepmaker, R. A., Dlugokencky, E., Sweeney, C., Wofsy, S. C., and Kort, E. A. (2015). Inverse modelling of CH₄ emissions for 2010–2011 using different satellite retrieval products from GOSAT and SCIAMACHY. *Atmospheric Chemistry and Physics*, 15(1):113–133.
- Annan, K. A. (2002). Toward a sustainable future. *Environment: Science and Policy for Sustainable Development*, 44:10–15.
- Bergamaschi, P., Karstens, U., Manning, A. J., Saunio, M., Tsuruta, A., Berchet, A., Vermeulen, A. T., Arnold, T., Janssens-Maenhout, G., Hammer, S., Levin, I., Schmidt, M., Ramonet, M., Lopez, M., Lavric, J., Aalto, T., Chen, H., Feist, D. G., Gerbig, C., Haszpra, L., Hermansen, O., Manca, G., Moncrieff, J., Meinhardt, F., Necki, J., Galkowski, M., O’Doherty, S., Paramonova, N., Scheeren, H. A., Steinbacher, M., and Dlugokencky, E. (2018). Inverse modelling of european ch4 emissions during 2006–2012 using different inverse models and reassessed atmospheric observations. *Atmospheric Chemistry and Physics*, 18:901–920.
- Bishop, C. M. and Winn, J. M. (2000). Non-linear Bayesian image modelling. In *Computer Vision – ECCV 2000*. Dublin, Ireland. Springer. https://doi.org/10.1007/3-540-45054-8_1.
- Borysiewicz, M., Wawrzynczak, A., and Kopka, P. (2012). Stochastic algorithm for estimation of the model’s unknown parameters via Bayesian inference. In Denisiuk, A., editor, *Proceedings of the 2012 Federated Conference on Computer Science and Information Systems*, pages 501–508. Wroclaw, Poland. <https://annals-csis.org/proceedings/2012/>.
- Cartwright, L., Zammit-Mangion, A., Bhatia, S., Schroder, I., Phillips, F., Coates, T., Negandhi, K., Naylor, T., Kennedy, M., Zegelin, S., Wokker, N., Deutscher, N. M., and Feitz, A. (2019). Bayesian atmospheric tomography for detection and quantification of methane emissions: Application to data from the 2015 Ginninderra release experiment. *Atmospheric Measurement Techniques*, 12:4659–4676.
- Cressie, N. (2018). Mission CO₂ntrol: A statistical scientist’s role in remote sensing of atmospheric carbon dioxide. *Journal of the American Statistical Association*, 113:152–168.
- De Visscher, A. (2013). *Air Dispersion Modeling: Foundations and Applications*. Wiley, Hoboken, NJ.
- Feitz, A., Schroder, I., Phillips, F., Coates, T., Negandhi, K., Day, S., Luhar, A., Bhatia, S., Edwards, G., Hrabar, S., Hernandez, E., Wood, B., Naylor, T., Kennedy, M., Hamilton, M., Hatch, M., Malos, J., Kochanek, M., Reid, P., Wilson, J., Deutscher, N., Zegelin, S., Vincent, R., White, S., Ong, C., George, S., Maas, P., Towner, S., and Griffith, D. (2018). The Ginninderra CH₄ and CO₂ release experiment: an evaluation of gas detection and quantification techniques. *International Journal of Greenhouse Gas Control*, 70:202–224.

- Flannery, T. (2009). *Now or Never: Why We Must Act Now to End Climate Change and Create a Sustainable Future*. Grove Atlantic, New York, NY.
- Francom, D., Sansó, B., Bulaevskaya, V., Lucas, D., and Simpson, M. (2019). Inferring atmospheric release characteristics in a large computer experiment using Bayesian adaptive splines. *Journal of the American Statistical Association*, 114:1450–1465.
- Ganesan, A., Manning, A., Grant, A., Young, D., Oram, D., Sturges, W., Moncrieff, J., and O’Doherty, S. (2015). Quantifying methane and nitrous oxide emissions from the UK and Ireland using a national-scale monitoring network. *Atmospheric Chemistry and Physics*, 15:6393–6406.
- Goodfellow, I., Bengio, Y., and Courville, A. (2016). *Deep Learning*. MIT Press, Cambridge, MA.
- Gurney, K. (2003). *An Introduction to Neural Networks*. Taylor & Francis, London, UK, illustrated, reprint edition. (Original edition published 1997).
- Halldor, B. and Venegas, S. (1997). “A manual for EOF and SVD analyses of climate data”. CCGCR Report No. 97-1, McGill University, Montréal, Québec.
- Hannachi, A., Jolliffe, I. T., and Stephenson, D. B. (2007). Empirical orthogonal functions and related techniques in atmospheric science: A review. *International Journal of Climatology*, 27:1119–1152.
- Harvey, N., Huntley, N., Dacre, H., Goldstein, M., Thomson, D., and Webster, H. (2018). Multi-level emulation of a volcanic ash transport and dispersion model to quantify sensitivity to uncertain parameters. *Natural Hazards and Earth System Sciences*, 18:41–63.
- Higdon, D., Gattiker, J., Williams, B. J., and Rightley, M. (2008). Computer model calibration using high-dimensional output. *Journal of the American Statistical Association*, 103:570–583.
- Holden, P. B., Edwards, N. R., Garthwaite, P. H., and Wilkinson, R. D. (2015). Emulation and interpretation of high-dimensional climate model outputs. *Journal of Applied Statistics*, 42:2038–2055.
- Houweling, S., Bergamaschi, P., Chevallier, F., Heimann, M., Kaminski, T., Krol, M., Michalak, A. M., and Patra, P. (2017). Global inverse modeling of CH₄ sources and sinks: an overview of methods. *Atmospheric Chemistry and Physics*, 17:235–256.
- Humphries, R., Jenkins, C., Leuning, R., Zegelin, S., Griffith, D., Caldow, C., Berko, H., and Feitz, A. (2012). Atmospheric tomography: a Bayesian inversion technique for determining the rate and location of fugitive emissions. *Environmental Science and Technology*, 46:1739–1746.
- Jones, A., Thomson, D., Hort, M., and Devenish, B. (2007). The UK Met Office’s Next-Generation Atmospheric Dispersion Model, NAME III. In Borrego, C. and Norman, A.-L., editors, *Air Pollution Modeling and Its Application XVII*, pages 580–589. Banff, Canada. https://doi.org/10.1007/978-0-387-68854-1_62.

- Kingma, D. P. and Welling, M. (2014). Auto-Encoding Variational Bayes. In Bengio, Y. and LeCun, Y., editors, *2nd International Conference on Learning Representations*. Banff, Canada. <https://dblp.org/db/conf/iclr/iclr2014.html>.
- Li, K. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association*, 86:316–327.
- Lin, J. C., Gerbig, C., Wofsy, S. C., Andrews, A. E., Daube, B. C., Davis, K. J., and Grainger, C. A. (2003). A near-field tool for simulating the upstream influence of atmospheric observations: The Stochastic Time-Inverted Lagrangian Transport (STILT) model. *Journal of Geophysical Research: Atmospheres*, 108(D16).
- Liu, X. and Guillas, S. (2017). Dimension reduction for Gaussian Process emulation: An application to the influence of bathymetry on tsunami heights. *SIAM/ASA Journal on Uncertainty Quantification*, 5:787–812.
- Lucas, D. D., Simpson, M., Cameron-Smith, P., and Baskett, R. L. (2017). Bayesian inverse modeling of the atmospheric transport and emissions of a controlled tracer release from a nuclear power plant. *Atmospheric Chemistry and Physics*, 17:13521–13543.
- Luhar, A. K., Etheridge, D. M., Leuning, R., Loh, Z. M., Jenkins, C. R., and Yee, E. (2014). Locating and quantifying greenhouse gas emissions at a geological CO₂ storage site using atmospheric modeling and measurements. *Journal of Geophysical Research: Atmospheres*, 119:10959–10979.
- Myhre, G., Shindell, D., Bréon, F.-M., Collins, W., Fuglestad, J., Huang, J., Koch, D., Lamarque, J.-F., Lee, D., Mendoza, B., Nakajima, T., Robock, A., Stephens, G., Takemura, T., and Zhang, H. (2013). Anthropogenic and natural radiative forcing. In Stocker, T. F., Qin, D., Plattner, G.-K., Tignor, M., Allen, S. K., Doschung, J., Nauels, A., Xia, Y., Bex, V., and Midgley, P. M., editors, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, pages 659–740. Cambridge University Press, Cambridge, UK.
- Oney, B., Henne, S., Gruber, N., Leuenberger, M., Bamberger, I., Eugster, W., and Brunner, D. (2015). The CarboCount CH sites: characterization of a dense greenhouse gas observation network. *Atmospheric Chemistry and Physics*, 15:11147–11164.
- Pisso, I., Sollum, E., Grythe, H., Kristiansen, N. I., Cassiani, M., Eckhardt, S., Arnold, D., Morton, D., Thompson, R. L., Groot Zwaafink, C. D., Evangelizou, N., Sodemann, H., Haimberger, L., Henne, S., Brunner, D., Burkhart, J. F., Fouilloux, A., Brioude, J., Philipp, A., Seibert, P., and Stohl, A. (2019). The Lagrangian particle dispersion model FLEXPART version 10.4. *Geoscientific Model Development*, 12:4955–4997.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT Press, Cambridge, MA.
- Reddy, B., Assenza, G., Assenza, D., and Hasselmann, F. (2010). *Energy Efficiency and Climate Change: Conserving Power for a Sustainable Future*. SAGE Publications, New Delhi, India.

- Riddick, S. N., Connors, S., Robinson, A. D., Manning, A. J., Jones, P. S. D., Lowry, D., Nisbet, E., Skelton, R. L., Allen, G., Pitt, J., and Harris, N. R. P. (2017). Estimating the size of a methane emission point source at different scales: from local to landscape. *Atmospheric Chemistry and Physics*, 17:7839–7851.
- Seibert, P. and Frank, A. (2004). Source-receptor matrix calculation with a Lagrangian particle dispersion model in backward mode. *Atmospheric Chemistry and Physics*, 4:51–63.
- Song, T., Sun, J., Chen, B., Peng, W., and Song, J. (2019). Latent space expanded variational autoencoder for sentence generation. *IEEE Access*, 7:144618–144627.
- Soydaner, D. (2019). Rolling in the deep convolutional neural networks. *International Journal of Intelligent Systems and Applications in Engineering*, 7:222–226.
- Wikle, C., Zammit-Mangion, A., and Cressie, N. (2019). *Spatio-Temporal Statistics with R*. Chapman & Hall/CRC Press, Boca Raton, FL.
- Wolf, L. and Bileschi, S. (2005). Combining variable selection with dimensionality reduction. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. San Diego, CA.
- Xu, J. and Duraisamy, K. (2020). Multi-level convolutional autoencoder networks for parametric prediction of spatio-temporal dynamics. *Computer Methods in Applied Mechanics and Engineering*, 372(113379).
- Zammit-Mangion, A., Cressie, N., and Ganesan, A. L. (2016). Non-Gaussian bivariate modelling with application to atmospheric trace-gas inversion. *Spatial Statistics*, 18:194–220.
- Zammit-Mangion, A., Cressie, N., Ganesan, A. L., O’Doherty, S., and Manning, A. J. (2015). Spatio-temporal bivariate statistical models for atmospheric trace-gas inversion. *Chemometrics and Intelligent Laboratory Systems*, 15:227–241.
- Zhang, R., Lin, C. D., and Ranjan, P. (2018). Local Gaussian process model for large-scale dynamic computer experiments. *Journal of Computational and Graphical Statistics*, 27:798–807.
- Zidanšek, A., Blinc, R., Jeglič, A., Kabashi, S., Bekteshi, S., and Šlaus, I. (2009). Climate changes, biofuels and the sustainable future. *International Journal of Hydrogen Energy*, 34:6980–6983.

A Post-processing of FLEXPART output

In this appendix we describe the procedure used to convert the output obtained from backward-simulations using FLEXPART to time-integrated sensitivity plumes.

We first sum the output (in units of $\text{s m}^3 \text{ kg}^{-1}$) over all time points, which leads to an approximate integrated sensitivity, sometimes known as the “footprint”, of the trace gas (Oney et al., 2015). We next divide by the volume of the grid cell in each instance to remove the dependence on the size of the grid cell, thus yielding the sensitivities in units of s kg^{-1} . From here, we scale the sensitivities by 10^{-3} to make them have units of s g^{-1} .

When doing flux inversion, one would need to multiply these sensitivities by an emission rate. Doing this, we would obtain a mass mixing ratio (in units of g g^{-1}). To instead ensure we obtain a mole fraction after this multiplication, which is typically used for reporting measurements of atmospheric methane, we scale the sensitivities by the molar mass of “air” ($28.9644 \text{ g mol}^{-1}$) and divide by the molar mass of methane ($16.0425 \text{ g mol}^{-1}$). This results in sensitivities in units of $\text{s g}^{-1} \text{ mol}$ (of methane) mol^{-1} (of air), and a mole fraction in units of mol (methane) mol^{-1} (air) after multiplication by the emission rate. Finally, we scale our sensitivities by 10^9 , yielding units $\text{ns g}^{-1} \text{ mol}$ (methane) mol^{-1} (air), and so that the resulting value after multiplication by an emission rate is in units of nmol (methane) mol^{-1} (air), or ppbv.

For simplicity throughout this paper, we refer to the sensitivities as having units ns g^{-1} , rather than $\text{ns g}^{-1} \text{ mol mol}^{-1}$.

B Estimation of plume departure angle

This appendix describes the steps taken to estimate the departure angle for each plume. Figure 5 provides a visual representation of some of these steps.

First, we determine the approximate direction for the departure angle. This is done by constructing a small annulus around the plume’s point of origin, and then computing the angle, a , subtended from the plume’s origin to the grid cell with maximum sensitivity within the annulus. The reason an annulus was used, rather than a disc, was due to the occasional non-zero sensitivities located directly behind the plume’s point of origin, particularly when turbulent weather was present. In such cases, a search within a fixed radius may lead to an estimated departure angle that is in the opposite direction to the true departure angle.

Following this, we construct an “angle window” ($a - \pi/4, a + \pi/4$), and then compute a weighted average of the angles subtended from the plume’s origin, where the weights are all sensitivities that are within this window and up to two lat-lon degrees away from the plume’s origin. This weighted average is then our estimated plume departure angle.

C Additional emulation results

This appendix shows an additional three sets of emulated plumes together with associated prediction standard errors.

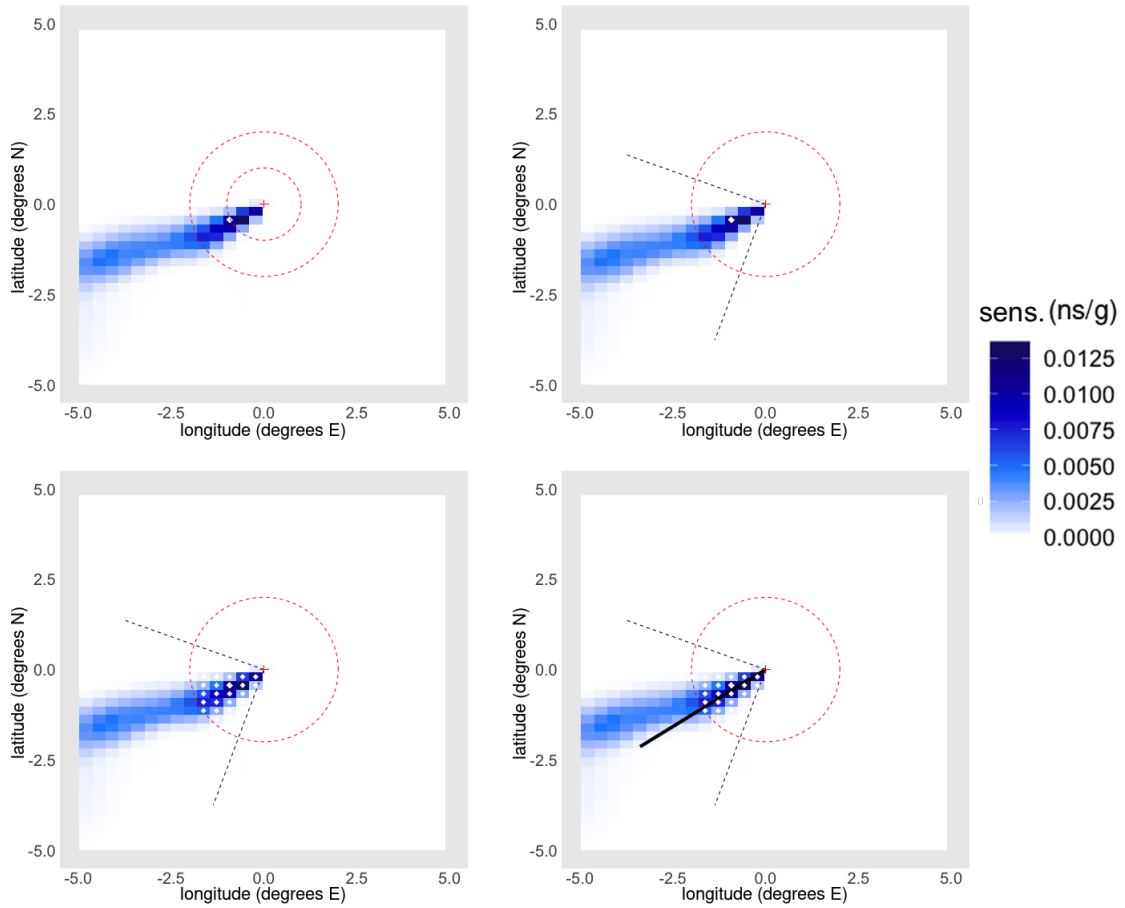


Figure 5: A visual representation of the departure-angle estimation process. Top-left: Step 1 – Locating the maximum sensitivity within the annulus. Top-right: Step 2 – Construction of the “angle window”. Bottom-left: Step 3 – Locating all points within the window, and within two lat-lon degrees of the plume’s origin. Bottom-right: Step 4 – Angle determined by computing a weighted average of the points identified in the bottom-left.

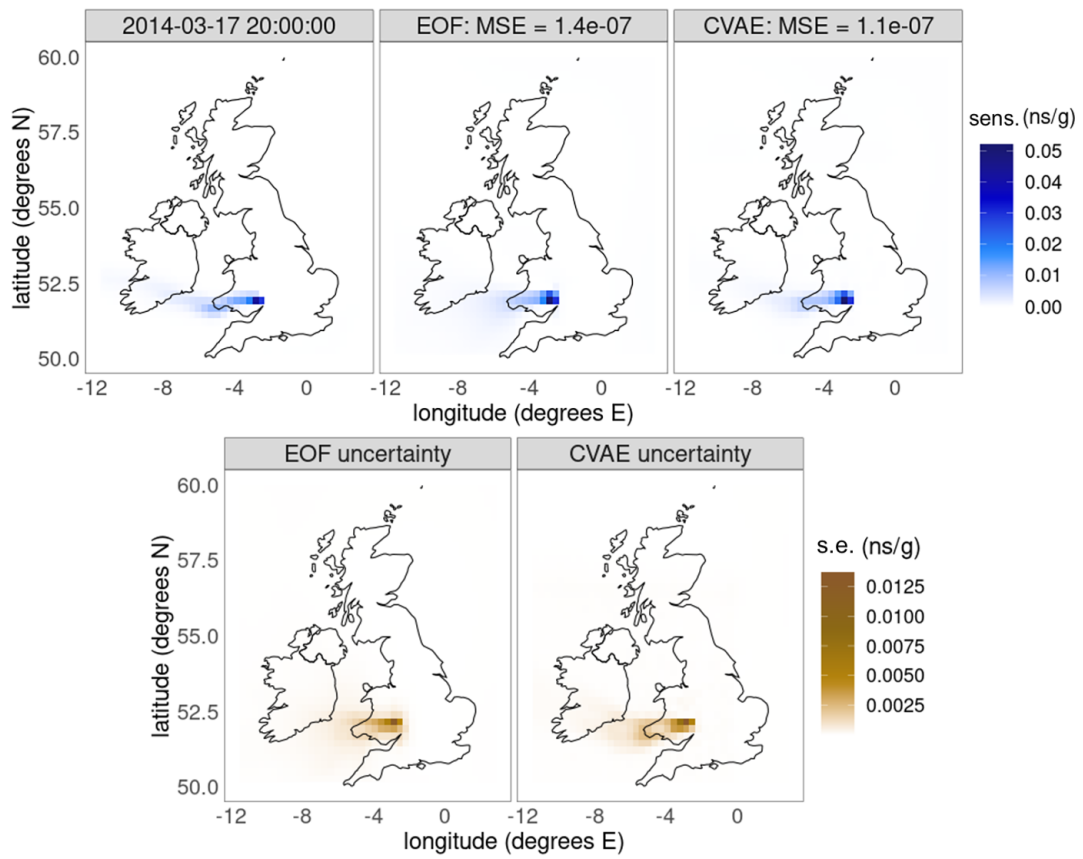


Figure 6: Same as Figure 4, but for 17 March 2014 at 20:00:00.




Figure 7: Same as Figure 4, but for 6 January 2014 at 08:00:00.




Figure 8: Same as Figure 4, but for 23 March 2014 at 16:00:00.