Bayesian spatio-temporal inference of trace gas emissions using an integrated nested Laplacian approximation and Gaussian Markov random fields

Luke M. Western[1], **Zhe Sha**[2,3], **Matthew Rigby**[1], **Anita L. Ganesan**[2], **Alistair J. Manning**[4], **Kieran M. Stanley**[1,5],
Simon J. O'Doherty[1], **Dickon Young**[1], and **Jonathan Rougier**[3]

[1]School of Chemistry, University of Bristol, Bristol, UK
[2]School of Geographical Sciences, University of Bristol, Bristol, UK
[3]School of Mathematics, University of Bristol, Bristol, UK
[4]Hadley Centre, Met Office, Exeter, UK
[5]Institute for Atmospheric and Environmental Science, Goethe University Frankfurt, Frankfurt am Main, Germany

Correspondence: Luke M. Western (luke.western@bristol.ac.uk)

Abstract. We present a method to infer spatially and spatio-temporally correlated emissions of greenhouse gases from atmospheric measurements and a chemical transport model. The method allows fast computation of spatial emissions using a hierarchical Bayesian framework as an alternative to Markov chain Monte Carlo algorithms. The spatial emissions follow a Gaussian process with a Matérn correlation structure which can be represented by a Gaussian Markov random field through a stochastic partial differential equation approach. The inference is based on an integrated nested Laplacian approximation (INLA) for hierarchical models with Gaussian latent fields. Combining an autoregressive temporal correlation and the Matérn field provides a full spatio-temporal correlation structure. We first demonstrate the method on a synthetic data example and follow this using a well-studied test case of inferring UK methane emissions from tall tower measurements of atmospheric mole fraction. Results from these two test cases show that this method can accurately estimate regional greenhouse gas emissions, accounting for spatio-temporal uncertainties that have traditionally been neglected in atmospheric inverse modelling.

1 Introduction

Emissions of greenhouse gases, ozone-depleting substances and air pollutants are increasingly inferred indirectly from atmospheric trace gas concentration observations and chemical transport models. These "top-down" or "inverse" methods complement inventory- or process-model-based "bottom-up" techniques that are used, for example, in national reporting of greenhouse gas emissions to the United Nations Framework Convention on Climate Change (UNFCCC; e.g. Leip et al., 2018).

Top-down methods rely on some form of statistical inference, or inverse theory, to infer emissions at global (e.g. Saunois et al., 2016) to regional (e.g. Brunner et al., 2017) scales. They also require a chemical transport model to provide the relationship between atmospheric mole fraction and emissions. The most common type of inverse method uses Bayesian inference. Bayesian inference relies on the information about the uncertainties in the measurement, transport model and prior probability of the parameters to inform the emissions estimate. Inverse methods often assume that uncertainties in the likelihood and prior probabilities are known exactly and Gaussian (e.g. Stohl et al., 2009; Brioude et al., 2013). These assumptions allow large inverse problems to be solved efficiently. However, when uncertainties are poorly understood, Bayesian methods have been shown to lead to

posterior solutions that are highly dependent on these assumptions (e.g. Ganesan et al., 2014).

Recently, hierarchical Bayesian schemes have been developed to infer unknown uncertainties in the inversion framework (e.g. Ganesan et al., 2014; Jeong et al., 2016; Lunt et al., 2016). These hierarchical Bayesian inference schemes use Markov chain Monte Carlo (MCMC) methods that are well-suited to small-dimensional problems. However, they can suffer from issues with convergence, especially as the dimension of the problem increases (Cowles and Carlin, 1996). Hierarchical inference of spatially and spatio-temporally correlated emissions is computationally difficult and currently relies on methods such as Kalman filters with assumed covariance structures (e.g. Brunner et al., 2012) or an empirical hierarchical framework, where unknown hyperparameters are not integrated out during inference (e.g. Michalak et al., 2005; Berchet et al., 2015). Previous work shows, however, that including spatial correlation improves the fit between modelled and observed data (Zammit-Mangion et al., 2016), while temporal correlation is important to represent the dependence of emissions between time periods. The size and spatial coverage of measurement networks available for inferring gas emissions are growing, particularly through satellite observations (e.g. Bergamaschi et al., 2007; Veefkind et al., 2012; Wecht et al., 2014; Ganesan et al., 2017). Therefore, there is a need to develop methods that can utilise these big data sets whilst maintaining the benefits of uncertainty quantification in hierarchical methods, ideally extending to spatio-temporal inference.

This work presents a computationally efficient hierarchical Bayesian framework for inferring spatio-temporally correlated trace gas emissions in a widely used regional atmospheric chemical transport modelling framework. We use an integrated nested Laplacian approximation (INLA) for the Bayesian inference. The spatial correlation structure with spatial Markov properties results from the Gaussian random field being a solution to a particular stochastic partial differential equation. Kronecker product algebra allows efficient extension to spatio-temporal correlation. Section 2 presents the construction of the hierarchical problem, describes the formation of the correlation structure and introduces INLA. Section 3 contains two case studies as a proof of concept for the method with a discussion of their implementation. The first is an example using four consecutive periods of pseudo-data observations of methane from four UK monitoring stations. This then extends to using real observations from the four measurement stations to infer UK emissions for 2014, split into four 3-month periods. Section 4 discusses the results and computational performance of the method, and Sect. 5 presents the conclusions of the study.

2 Methods

This section details an efficient approach to forming spatial and spatio-temporal correlation functions and outlines how this applies to the inference of regional trace gas emissions from measurements using fast inference for hierarchical models. We limit the scope of this paper to the well-established problem of regional inference of long-lived trace gas emissions using a backward-running Lagrangian particle dispersion model (Stohl et al., 2009; Manning et al., 2011; Henne et al., 2016). In this framework, the model directly calculates the sensitivity of the measurements to emissions from each grid cell in the domain. Extension to other systems (e.g. global models) would require modification to provide sensitivity on a global scale with substantially different temporal and spatial emissions. We begin by introducing the model and the inferred latent parameters (the emission fluxes and boundary conditions), followed by an introduction to Gaussian Markov random fields and how they are useful for efficient calculation of spatial and spatio-temporal correlation structures. All together, this forms a Bayesian hierarchical model, from which emissions can be inferred.

2.1 Model framework

The aim is to infer some parameters of interest, here a spatial field of a priori emissions scaled by some factor, x, from some measurement. The a priori emissions at a given location are generally informed by spatially resolved bottom-up inventories (as in Sect. 3.1.3) or extrapolation from some reported emissions value. The value x is a multiplicative scaling of this a priori value for emissions, most generally expressed as a quantity of gas per unit time per unit area. The approach taken in this work uses a Gaussian Markov random field for fast and efficient calculation of spatial correlation for x. This means that the emissions field is required to be a latent Gaussian field, which will be discussed further in Sect. 2.2. Net surface fluxes of many greenhouse gases are positive at the scales resolved by the model. In this work, due to the usage of a latent Gaussian field, which must be defined over both positive and negative values, we choose to look at the deviations of emissions from the prior mean emissions field. This is in an effort to fit the physical model to the imposed statistical model that fast computation requires. Taking the approach that, for many regional inverse problems involving longer-lived trace gases, there is a linear relationship between emissions which are constant in time and observed atmospheric concentration, the relationship between measurements and emissions is

$$y = \mathbf{H}x + \mathbf{K}u + \epsilon, \tag{1}$$

where y is a vector of the residual between the measured and a priori predicted measurement; $\mathbf{H}$ is a Jacobian (or sensitivity) matrix, which maps the surface emissions to the measurements; u is a vector of independent and identically dis-

tributed variables containing the contribution to the measurement of mole fractions at the boundary of the domain minus the prior mean contribution, with an associated sensitivity matrix $\mathbf{K}$; and ϵ is some stochastic error. The variable x has a Gaussian prior probability with zero expectation and a covariance described by a Gaussian random field. This will be solved using a hierarchical framework (Sect. 2.4).

2.2 Gaussian Markov random fields

The emissions scaling from its a priori value x is spatially continuous over the domain of interest. We assume that it exhibits a spatial correlation structure because emissions at one location are generally not independent from all other locations in the same field. We choose to model the covariance in this field using a Matérn covariance function, which Stein (1999) shows is well-suited to natural systems due to its flexibility, with other correlation structures (e.g. exponential) being special cases of the Matérn family (Guttorp and Gneiting, 2006). A stationary Gaussian random field with a Matérn covariance is the solution to a particular stochastic partial differential equation (Whittle, 1954, 1963), given by

$$(\kappa^2 - \Delta)^{\frac{\alpha}{2}}(\tau x(s)) = \mathcal{W}(s), \quad s \in \Omega, \tag{2}$$

where κ is the spatial scale parameter, τ influences the variance, Δ is the Laplacian operator, Ω is the spatial domain and $\mathcal{W}(s)$ is a Gaussian stochastic process for locations s, noting that the dependence is only on the Euclidean distance; the process is therefore isotropic. For an example of nonstationary and anisotropic fields see Marques et al. (2019). The smoothness parameter α gives a continuous domain Markov field for integer values and is set to $\alpha = 2$ (see Whittle, 1954). Smaller values will give more short-scale variability and can be difficult to differentiate from noise. Lindgren et al. (2011) show that if the field is represented using a Gaussian Markov random field (see Rue and Held, 2005), then it is very efficient to directly construct the precision, or inverse covariance, matrix using a basis function representation

$$x(s) = \sum_k \psi_k(s) x_k, \tag{3}$$

where $\psi_k(\cdot)$ is a piecewise linear function in each triangle of a mesh construction of the spatial domain, where $\psi_k(\cdot)$ is 1 at vertex k and 0 at all other vertices. This mesh is created using constrained refined Delaunay triangulation (Shewchuk, 2002), placing nodes at the main points of interests and infilling the rest of the space using some condition of minimum and maximum length of the vertices. In this work we choose to represent the UK and Ireland with an evenly spaced denser mesh with a coarser mesh outside of this region. A mesh could be further refined, for example by creating an even denser mesh close to the measurement site where sensitivity is higher. It is important to extend the mesh beyond the region where the measurements are sensitive to emissions


Figure 1. A mesh constructed using constrained refined Delaunay triangulation, where the distribution of nodes is denser around the United Kingdom and Ireland.

as the mesh is constructed using Neumann boundary conditions, which trigger reflection and therefore overestimation close to the boundaries. This has no effect on the inferred emissions, provided that the mesh is extended far enough around the domain of interest. Figure 1 shows the mesh used in this work, excluding the extended outer mesh region.

2.3 Extension to spatio-temporal correlation

A spatio-temporal extension to the forward model (Eq. 1) is possible by including the spatial correlation structure introduced in Sect. 2.2 in a temporal framework (Cameletti et al., 2013). Similar to Eq. (1), the deviation from the prior mean measurement at site l made at time t in its simplest form is

$$y_{tl} = \sum_{i=1}^{n} h_{itl} x_{it} + \sum_{j=1}^{4} k_{jtl} u_{jt} + \epsilon_{tl}, \tag{4}$$

where i represents each of the total n nodes and j represents the edge of the domain at each of the four cardinal directions. We make the assumption that measurements made at a given time are independent, giving a vectorised observation vector y_t at each time. This can be generalised to include correlated measurements, where the covariance function should result in a sparse matrix to retain computational efficiency. Following this, the matrix $\mathbf{H}$ from Eq. (1) becomes the sparse block diagonal matrix

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_1 & 0 & 0 & \dots & 0 \\ 0 & \mathbf{H}_2 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \\ 0 & 0 & 0 & & \mathbf{H}_m \end{bmatrix}, \tag{5}$$

which operates on the vectorised spatio-temporal scaling of the emissions field $x = [x_1, x_2, \dots x_m]$ to model the obser-

vations $\boldsymbol{y} = [\boldsymbol{y}_1, \boldsymbol{y}_2, \ldots \boldsymbol{y}_m]$. The time varying structure of $\mathbf{H}_t$ and $\boldsymbol{x}$ applies also to $\mathbf{K}_t$, and $\boldsymbol{u}_t$. We impose the temporal correlation structure between emissions over time to be a first-order autoregressive model, where

$$\boldsymbol{x}_t = \phi \boldsymbol{x}_{t-1} + \delta_t; \quad \delta_t \sim \mathcal{N}\left(0, \mathbf{Q}_S^{-1}\right) \quad t = 1, \ldots, m, \quad (6a)$$

and

$$\boldsymbol{x}_1 \sim \mathcal{N}\left(0, \left(\frac{\mathbf{Q}_S}{1 - \phi^2}\right)^{-1}\right), \quad (6b)$$

where ϕ is the temporal correlation and $\mathbf{Q}_S^{-1}$ is the spatial correlation structure described by a Matérn field using the stochastic partial differential equation approach for a Gaussian Markov random field. We have chosen a first-order autoregressive model as we believe that emissions are generally similar to those at the previous time step. In the given model, ϕ is the correlation between the previous time step and the current time step. This means that the emissions at time t have a similarity of ϕ to emissions at $t-1$ plus some spatially correlated random effect δ_t. The vectorisation of $\boldsymbol{x}$ allows a separable covariance structure for the temporal and spatial covariances, which means that the spatio-temporal precision matrix can be expressed using a Kronecker product as follows (Mardia et al., 1979):

$$\mathbf{Q} = \mathbf{Q}_T \otimes \mathbf{Q}_S. \quad (7)$$

Estimating hourly emissions at each time t soon makes inference prohibitive due to burden. Instead we make the assumption that emissions are constant over a 3-month period to reduce the computational size. We continue with the assumption that measurements within a single time period are independent, although this can be generalised if required. In this experiment we make the assumption that emissions are constant over a 3-month period and that the correlation between these 3-month periods is autoregressive of the order of 1.

2.4 Hierarchical model

Inferring the emissions and the related uncertainties requires a hierarchical model to infer the quantities of interest from measurements while estimating some unknown parameters which are necessary for inference. The main focus of this work is to estimate the posterior distribution of the emissions field $\boldsymbol{x}$ based on observations $\boldsymbol{y}$. We follow a typical Bayesian hierarchical framework

$$p(\boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta} \mid \boldsymbol{y}) \propto p(\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta}) p(\boldsymbol{x} \mid \boldsymbol{u}, \boldsymbol{\theta}) p(\boldsymbol{u} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}), \quad (8)$$

where $\boldsymbol{\theta}$ is vector of hyperparameters describing the variances and covariances in $\boldsymbol{x}$, $\boldsymbol{u}$ and $\boldsymbol{y}$, noting that $\boldsymbol{x} \perp\!\!\!\perp \boldsymbol{y} \mid \boldsymbol{\theta}$.

We assume

$$\boldsymbol{y} \mid \boldsymbol{x}, \boldsymbol{u}, \boldsymbol{\theta} \sim \mathcal{N}(\mathbf{H}\boldsymbol{x} + \mathbf{K}\boldsymbol{u}, \mathbf{Q}_y(\boldsymbol{\theta})^{-1}),$$
$$\boldsymbol{x} \mid \boldsymbol{u}, \boldsymbol{\theta} \sim \mathcal{N}(0, \mathbf{Q}(\boldsymbol{\theta})^{-1}),$$
$$\boldsymbol{u} \mid \boldsymbol{\theta} \sim \mathcal{N}(0, \mathbf{I}\sigma_{\text{BC}}^2),$$
$$\boldsymbol{\theta} \sim p(\boldsymbol{\theta}),$$

where the precision matrix of the model–measurement uncertainty $\mathbf{Q}_y$ contains the hyperparameter for the standard deviation of the model–measurement standard deviation σ_y, σ_{BC} is the standard deviation of the prior for $\boldsymbol{u}$ and $\mathbf{I}$ is the identity matrix. Together the hyperparameters make the vector $\boldsymbol{\theta} = (\rho, \sigma, \phi, \sigma_{\text{BC}}, \sigma_y)$, which have independent prior distributions. The hyperparameters for the spatial precision matrix $\mathbf{Q}(\boldsymbol{\theta})$ are transformations of the variables in Eq. (2),

$$\rho = \frac{\sqrt{8}}{\kappa}, \quad (9a)$$

$$\sigma = \left(\frac{1}{\sqrt{4\pi\kappa\tau}}\right)^{-\frac{1}{2}}, \quad (9b)$$

where ρ is the range parameter and σ is the marginal standard deviation of the latent field (Lindgren et al., 2011). We use penalised complexity priors to define the prior probabilities for these parameters (Simpson et al., 2017; Fuglstad et al., 2018). Penalised complexity priors allow the formation of priors when there is only a vague understanding of their true values while enforcing more constraint than using a broad uniform or a Jeffreys prior. This uses the information loss of deviating from some baseline estimate of the parameter. Penalised priors do not increase the computational speed for the given case and are chosen instead for their intuitiveness. Penalised complexity priors are specified by the probability of the parameters being less than or greater than some baseline value

$$p(\rho > \rho_0) = p_\rho,$$
$$p(\sigma < \sigma_0) = p_\sigma,$$

where ρ_0 and σ_0 are the baseline values and p_ρ and p_σ are the associated probabilities defined by the user. The hyperparameter ϕ controls the temporal correlation between latent variables, and its prior probability is defined on a beta distribution scaled between -1 and 1 as

$$\phi \sim \text{Beta}(a, b),$$

where a and b are assumed coefficients. The matrix $\mathbf{Q}_y$ is the precision matrix for the combined measurement and model errors. The diagonal of $\mathbf{Q}_y$ contains the square of the hyperparameter σ_y, where the prior probability is defined on $\log \frac{1}{\sigma_y^2}$ as

$$\log \frac{1}{\sigma_y^2} \sim \mathcal{N}\left(\mu_{\sigma_y}, \sigma_{\sigma_y}^2\right).$$

The marginal standard deviation of the a priori boundary conditions σ_{BC} is also constructed this way, giving

$$\log \frac{1}{\sigma_{BC}^2} \sim \mathcal{N}\left(\mu_{\sigma_{BC}}, \sigma_{\sigma_{BC}}^2\right).$$

A Bayesian hierarchical model requires a method of inference to estimate the parameters of interest and any parameters that are not of direct interest but required, and uncertain, in order to infer the other parameters in the hierarchy. Many methods of inference exist and have been applied to the problem of estimating emissions of trace gases (see references in Sect. 1). A promising approach using the correlation structure in Sect. 2.2 and 2.3 is INLA (Lindgren and Rue, 2015). Section 2.5 outlines this approach to inference.

2.5 Inference using an integrated nested Laplacian approximation

An integrated nested Laplacian approximation (Rue et al., 2009) provides a fast and efficient framework to infer the latent variable x and hyperparameters θ from measurements y. The calculation of the INLA is possible using the R-INLA package (Lindgren and Rue, 2015). In this work we use R-INLA version 17.06.20. The speed in this approach comes from solving the marginal posteriors for x_i, i.e. each element i of the latent field, through the numerical integration

$$p(\theta_j \mid y) = \int p(\theta \mid y)\mathrm{d}\theta_{-j}, \tag{11}$$

$$p(x_i \mid y) = \int p(x_i \mid y, \theta) p(\theta \mid y)\mathrm{d}\theta, \tag{12}$$

where j is the jth element in θ and $-j$ indicates all by the element j. Equations (11) and (12) make use of a Laplace approximation by approximating $p(\theta \mid y)$ by

$$p(\theta \mid y) \propto \left.\frac{p(x, \theta, y)}{p_G(x \mid \theta, y)}\right|_{x = x^*(\theta)}, \tag{13}$$

where $p_G(x \mid \theta, y)$ is a Gaussian approximation to the full conditional of x and the right-hand side is evaluated at $x^*(\theta)$, which is the modal probability of x for a given θ. This approximation is exact if $p(\theta \mid y)$ is Gaussian and gives a good approximation for log concave problems (Tierney and Kadane, 1986). Then, it is possible to approximate $p(x_i \mid y, \theta)$ using another Laplace approximation

$$p_{LA}(x_i \mid \theta, y) \propto \left.\frac{p(x, \theta, y)}{p_G(x_{-i} \mid x_i, \theta, y)}\right|_{x_{-i} = x^*_{-i}(x_i, \theta)}, \tag{14}$$

where $p_G(x_{-i} \mid x_i, \theta, y)$ is the Gaussian approximation to $x_{-i} \mid x_i, \theta, y$ evaluated at the mode $x^*_{-i}(x_i, \theta)$. See Rue et al. (2009) and Martins et al. (2013) for a more in-depth description.

While this method relies on calculating only the marginal posterior distributions of x_i, it is still possible to predict a linear combination of the field to provide regional emissions totals (e.g. country totals). We define a linear predictor of emissions for a given region η^*, defined using the basis function representation of the mesh where each node contains information on the spatial area represented by that node and its connecting vertices and zeros for all nodes that are not in the region of interest. Then we can approximate the linear combination of the parameters of interest, giving a combined emissions total η^*, using Eqs. (11) and (12) and transforming the predicted latent field by the weightings containing the area information.

3 Case studies

This section presents two case studies to demonstrate how the method applies to inferring trace gas emissions. The first uses simulated methane observations from four tall-tower measurement sites to infer simulated spatio-temporal emissions from the UK. The second case study expands on the first case study by using real observations from the four tall towers to infer emissions of methane from the UK over four 3-month periods in 2014. While the size of this problem is not particularly large, we demonstrate the method using UK methane emissions as a proof of concept as it is a well-studied test case (Manning et al., 2011; Ganesan et al., 2015; Lunt et al., 2015), which should exhibit a spatio-temporal correlation structure. The method can be extended to larger spatial domains or data set sizes as required.

3.1 Transport model and data

3.1.1 Measurement data

The case study observations are from four measurement sites: three in the UK and one in Ireland, which are part of the UK Deriving Emissions related to Climate Change (DECC) network (Stanley et al., 2018). Figure 2 shows the location of these four measurement stations: Ridge Hill in the west of England, Angus on the east coast of Scotland, Tacolneston on the east coast of England and Mace Head on the west coast of Ireland. Measurements are made quasi-continuously throughout this network but are averaged into hourly values here, consistent with the time step of the meteorology that drives the atmospheric transport model (Sect. 3.1.2). The data set contains $\sim 10\,000$ measurements to demonstrate the capabilities of the method at handling moderate data volumes. We consider the scalability of this method in the discussion (Sect. 4).

3.1.2 Transport model

An atmospheric transport model calculates the sensitivity of hourly measurements to the emissions or boundary condi-
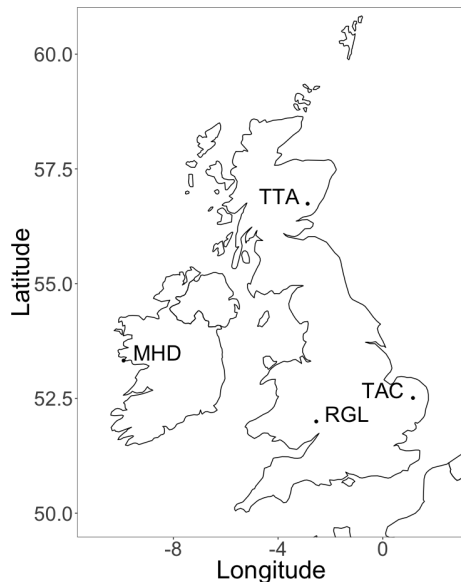
Figure 2. The case studies use the four UK DECC network measurement stations located on this map. RGL is Ridge Hill, TAC is Tacolneston, TTA is Angus and MHD is Mace Head station.

tions, from which the matrices **H** and **K** can be formed. This work uses the NAME III (Numerical Atmospheric-dispersion Modelling Environment) version 7.2 Lagrangian particle dispersion model (Jones et al., 2006) to simulate the transport of methane in the atmosphere. For each measurement, NAME tracks 20 000 gas particles, released over a 1 h period, backward in time for 30 d from the measurement site. We record the times and locations that the particles drop below 40 m a.g.l. and reach the computational domain boundary, which is at 5° S, 74° N, 55 and 192° E, to calculate the sensitivity of the data to emissions from the surface or to the mole fraction at the domain edge. NAME was driven by the Met Office's Unified Model UK Variable (UKV) three-hourly meteorological analysis (Cullen, 1993).

3.1.3 Prior emissions inventory

Inventory data from the Emissions Database for Global Atmospheric Research (EDGAR) v4.3.2. provide the mean prior emissions (Janssens-Maenhout et al., 2017) using the most recent emissions map from 2012. This database contains anthropogenic emissions, which are the dominant methane sources in the UK (Manning et al., 2011), and is deemed appropriate for the application. The EDGAR emissions inventory provides an a priori estimate of UK methane emissions of $2.46 \, \text{Tg} \, \text{yr}^{-1}$. The MOZART-4 global chemistry model (Emmons et al., 2010), driven by global fields of natural and anthropogenic emissions and sink terms (Ganesan et al., 2017), provides the prior mean estimate of methane mole fraction at the boundaries of the inversion region.

3.2 Pseudo-data experiment

We test the method by performing an inversion using pseudo-data for four consecutive time periods of 1 month. By creating a known-emissions field we are able to validate the method through comparing the inferred emissions to the known emissions, which is not possible in the real world. We form a synthetic-emissions field by allowing the emissions to deviate from the prior mean emissions according to a Matérn field (see Sect. 2.2). We choose to simulate the data using $\sigma = 0.5$ as the uncertainty in the EDGAR v4.3.2. inventory is around 50 % for methane emissions and $\rho = 3.25$, which is similar to the correlation length scale in UK emissions in the EDGAR v4.3.2. inventory estimated using a variogram (Cressie, 2015), although the correlation length scale in the uncertainty is unknown. We use $\phi = 0.8$ as expert experience suggests that UK emissions of methane are generally highly correlated in time.

To create the synthetic-emissions field, the NAME sensitivities for each measurement at each grid cell, detailed in Sect. 3.1.1, are multiplied by the corresponding EDGAR inventory emissions in that grid cell, which are then transformed into the triangulation nodes in Fig. 1. This forms the matrix **H**. We randomly sample the full spatio-temporal precision matrix for the latent field (Sect. 2.3) using the GMR-FLib library (Rue and Follestad, 2001) to generate the latent field x. In this experiment we treat the boundary conditions as known as in practice these are generally well-constrained (or treated as known) during an inversion. The observations are simulated using the simulated latent field and sensitivities following Sect. 2.1 with additive Gaussian noise with a standard deviation of 15 % of y. The synthetic observations contain a total of 11 520 measurement points, which are used to infer 646 emissions nodes for each time period, i.e. the nodes of the mesh in Fig. 1. Figure 3a–d show the synthetic-emissions field in terms of deviation from some prior emissions field. The synthetic total deviation in emissions from the prior mean for the UK for each period is 0.07, 1.29, 1.01 and $-0.33 \, \text{Tg} \, \text{yr}^{-1}$.

The inference needs prior probabilities for the hyperparameters, which are known exactly here, but we set them to be deliberately incorrect, but feasible based on true prior knowledge, to check that the inversion method can still recover the correct emissions. For the inversion we assign a prior probability for ϕ using $a = 6.5$ and $b = 0.1$, for σ using $\sigma_0 = 0.1$ and $p_\sigma = 0.01$ and for ρ using $\rho_0 = 5$ and $p_\rho = 0.5$. We base the constraint on the spatio-temporal emissions on the assumption that methane emissions in the UK are likely to be strongly correlated between time periods, that they are unlikely to vary by more than 10 % of the a priori emissions from the previous time step and that there is little knowledge of the spatial correlation structure. For the model measurement error we assign the prior probability on the log precision as $\log \frac{1}{\sigma_{\tilde{y}}^2} \sim \mathcal{N}(-5, 1)$, which represents approxi-
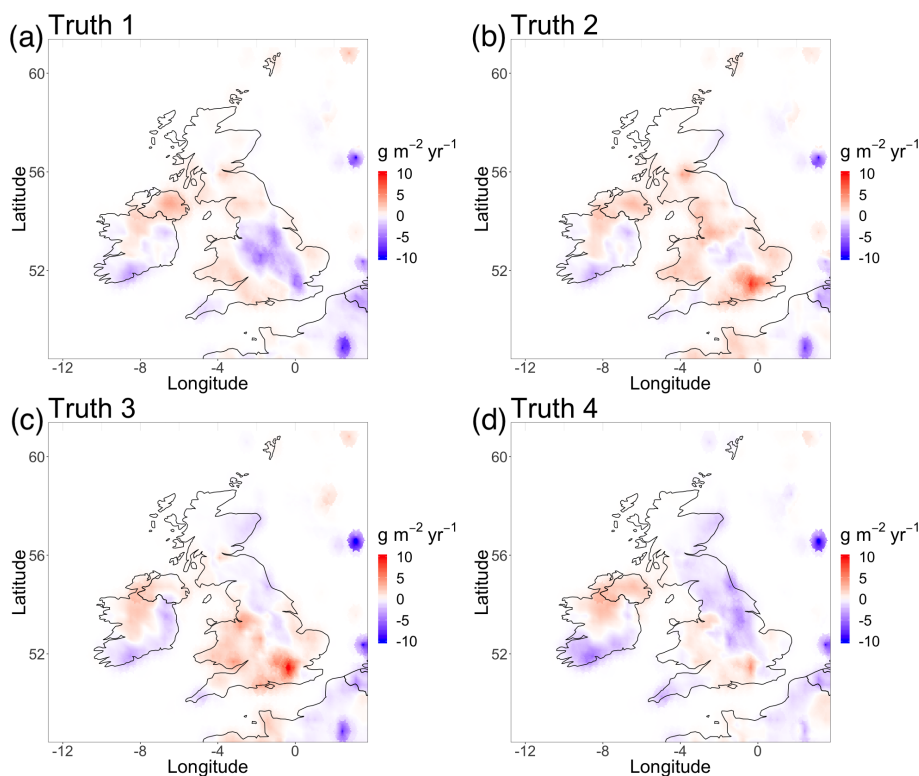
Figure 3. Time-correlated deviation of emissions from the prior mean, simulated using a Matérn field for **(a)** the first time period, **(b)** the second time period, **(c)** the third time period and **(d)** the fourth time period.

mately 68 % probability that the error falls between 8 and 20 ppb, in line with previous experience.

Figure 4 shows the resultant inferred mean changes in the emissions field for each time period using INLA (Sect. 2.5) and the hierarchical model in Sect. 2.4. The change in the emissions field is the inferred latent field scaling parameter multiplied by the inventory emissions in Sect. 3.1.3. Figure S1 shows maps of the 95 % uncertainty in the difference in emissions. A linear combination of the posterior latent field multiplied by the inventory emissions field gives total emissions for the UK. The inferred posterior mean change in total emissions for the UK from the prior mean with their associated 95 % uncertainties for the four time periods are 0.01 $[-0.10, 0.12]$, 1.23 $[1.10, 1.36]$, 1.07 $[0.97, 1.18]$ and -0.30 $[-0.41, -0.18]$ Tg yr^{-1}, which are shown in Fig. 5. The inferred emissions for the UK as a whole agree well with the synthetic emissions, with all true synthetic emissions totals falling within the 95 % uncertainty of the inferred total emissions. These synthetic tests cover both small and large differences between the a priori estimate of methane emissions in the UK and their true value. The larger differences are most likely an overestimate of what would be expected in practice, although this provides a good test case. The posterior mean spatial distribution of emissions in Fig. 4 is qualitatively similar to those in Fig. 3, although we avoid reading heavily into

comparisons of the mean estimates plotted on spatial maps (see Gelman and Price, 1999).

Hyperparameter estimation is less accurate than for the latent field. The estimation of the noise σ_y generally captures the imposed noise well, which had a mean value of 5.5 ppm and an estimated value of 6.6 [6.5, 6.8] ppm. The correlation structures are less well-captured. The temporal correlation ϕ was estimated with a mean value of 0.7 [0.6, 0.8]; the range ρ has a value of 1.7 [1.3, 2.1] and the marginal standard deviation of the latent field σ has a value of 0.8 [0.7,0.9]. It is promising that the posterior mean estimates of all hyperparameters show an improvement on their prior mean or baseline values, although only the true value for ρ falls within the estimated 95 % uncertainty.

3.3 Real-data experiment

This section presents methane emissions estimates for the UK in 2014. The year is split into four time periods: January to March, April to June, July to September and October to December. Based on the synthetic data set in Sect. 3.2, we use the same prior probabilities for the hyperparameters as the inversion. The prior boundary conditions are distributed with $\mu_{\sigma_{BC}} = 3.2$ and $\sigma_{\sigma_{BC}} = 0.4$, based on expert judgement. The a priori values for these boundary conditions come from the MOZART-4 model as in Sect. 3.1.3. The linear mapping
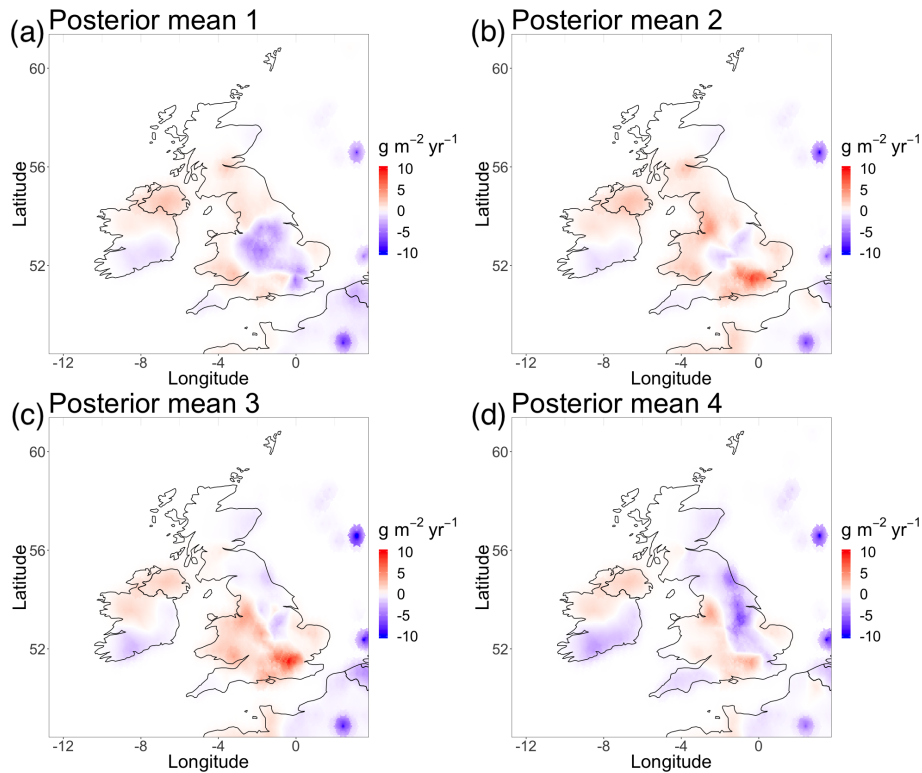
Figure 4. The inferred mean deviation of emissions from the prior mean for **(a)** the first time period, **(b)** the second time period, **(c)** the third time period and **(d)** the fourth time period.
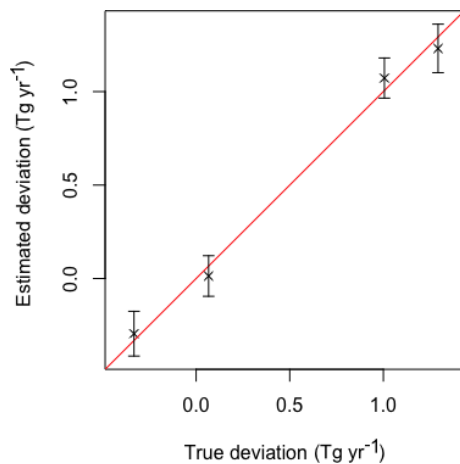

Figure 5. The inferred and true difference between UK emissions and a synthetic inventory value for each of the four time periods. The crosses show the inferred mean difference in total emissions with their associated 95 % uncertainty. The red line indicates one-to-one agreement.

from the latent field to the measurements is generated using the NAME-derived sensitivities described in Sect. 3.1.2 multiplied by the inventory emissions detailed in Sect. 3.1.3.

Figure 6 shows maps of the inferred mean difference in emissions from the a priori inventory for the four time periods using INLA as in Sect. 2.5. This result is the mean posterior scaling for the latent field multiplied by the inventory value (see Sect. 3.1.3). Figure S2 shows maps of the 95 % uncertainty in the difference in emissions. A linear combination of the posterior latent field multiplied by the inventory emissions field gives total emissions for the UK. We estimate total UK methane emissions in 2014, with their associated 95 % uncertainty, of 2.14 [1.90, 2.38], 2.28 [2.06, 2.51], 2.62 [2.39, 2.84] and 2.09 [1.87, 2.30] Tg yr^{-1} for the periods January to March, April to June, July to September and October to December, respectively. These are plotted along with the inventory estimate in Fig. 7. This emissions trend suggests that, for 2014, there was an increase in methane emissions in the UK during the summer months compared to the winter months. The uncertainties, however, are large, meaning that this increase may not be as stark as suggested by the mean estimates. Combining these emissions into a mean annual emission for 2014 with its associated 2-standard-deviation uncertainty assuming that the time periods are correlated with the modal posterior value of ϕ gives 2.28 ± 0.33 Tg yr^{-1}. The emissions are similar to mean UK estimates from previous hierarchical inversions using NAME of 2.09 [1.65, 2.67] Tg yr^{-1} by Ganesan et al. (2015) and 2.28 [2.04, 2.52] Tg yr^{-1} by Lunt et al. (2016). All of these
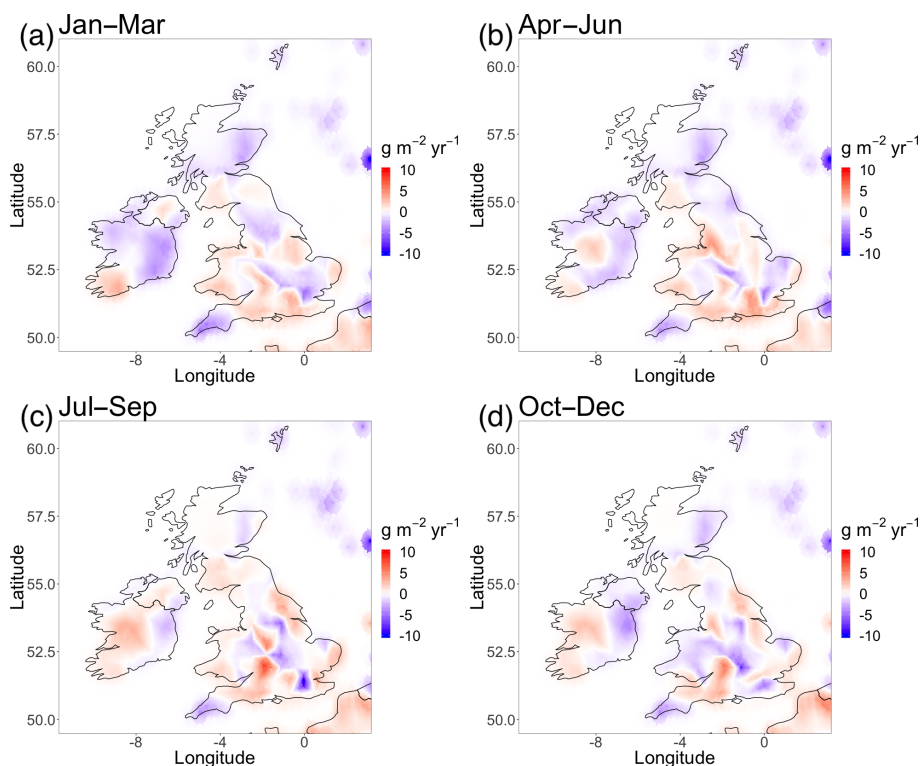
Figure 6. Inferred mean difference in methane emissions for the UK in 2014 compared to the EDGAR inventory for **(a)** January to March, **(b)** April to June, **(c)** July to September and **(d)** October to December.

estimates are broadly in line with the United Kingdom's 2014 emissions estimate of $2.13\,\mathrm{Tg\,yr^{-1}}$ reported in its national inventory (Department of Business, Energy and Industrial Strategy, 2019).

4 Discussion

The real benefit of the presented inversion method is speed, while still maintaining the idea that, in this application, uncertainties exist within the uncertainties that are inherent to hierarchical Bayesian inverse methods. The computation of the marginal posterior distribution of the latent field is readily suited to run in parallel across multiple cores, making this approach scalable to problems with a larger parameter space. This, however, requires that sufficient memory allocation is available. The inference for the experiment in Sect. 3.3 took around 2 h wall-clock time using a Quad-Core MacBook Pro with 8 GB RAM. For comparison, we were unable to get the problem in Sect. 3.3 to converge using 22 million iterations of a Metropolis-adjusted Langevin diffusion MCMC algorithm (Roberts and Rosenthal, 1998), taking on the order of days to complete and whose convergence is difficult to diagnose (Cowles and Carlin, 1996). Reducing the problem to purely spatial inference reduces the run time down to around 10 min wall-clock time for a single 3-month period. Inference using INLA may have smaller computational gains over


Figure 7. Estimated total UK methane emissions for 2014. The prior mean comes from EDGAR v4.32 inventory data (red dashed). For comparison, the figure also shows the posterior means for UK methane emissions from Ganesan et al. (2015) (black dotted) and Lunt et al. (2016) (black dash-dotted).

hierarchical methods other than MCMC, for example empirical hierarchical inference (e.g. Michalak et al., 2005). Such methods, however, would benefit from using the stochastic

partial differential equation approach to Gaussian Markov random fields for spatial correlation in Sect. 2.2, vastly reducing the cost of inverting dense covariance matrices.

The latent Gaussian field, crucial to this method, has the problem that it does not restrict emissions to strictly positive values, which is a physical requirement for many gas emissions. In addition, the INLA method relies on the assumption of approximate multivariate normality of the posterior linear predictors. MCMC algorithms do not suffer from this issue as any prior probability density function can be chosen, and no assumption is made about the posterior linear predictors. This has to be traded off, however, with the speed and ease of implementation of the method and the scale of the problem that the user wishes to solve. If strictly positive posterior emissions are an absolute requirement, another possible modification to the approach is to use a Taylor expansion around the nonlinear model $Z = He^x + \epsilon$, where x is a multiplicative scaling of the prior mode which follows a Matérn field. This Taylor expansion could either be around zero if the prior inventory is a good estimate of the true emissions or around the posterior mode itself, which can be found through iterations of the inverse method at the posterior mode of the previous iteration. An alternative route would be to instead use a non-Gaussian Matérn field, and this will likely prove a promising future avenue for geostatistical inference (Wallin and Bolin, 2015).

A potential extension to this work is global-scale modelling, for example a global study of methane emission from global satellite measurements. Spatio-temporal estimation of global CO_2 fluxes using Gaussian Markov random fields and a global measurement network gives promising results (Dahlén et al., 2019). The nature of the method makes it well-suited to parallel implementation and thus upscaling. Upscaling to global studies may introduce additional difficulties, which have been otherwise ignored in the regional case. For example it may no longer be possible to assume a spatial correlation structure controlled by only two hyperparameters, i.e. the correlation structure may vary in space. The stochastic partial differential equation approach to Gaussian Markov random fields can handle this by allowing the hyperparameters ρ and σ to become vectors or a non-stationary covariance (Lindgren et al., 2011) or by subsetting the space controlled by the different covariance structures (Sha et al., 2019). The problem of non-stationary covariances may also be present for inference of other spatial scales. For example natural features, such as lakes, may cause abrupt changes in correlation structures. Emissions of anthropogenic greenhouse gases may exhibit no spatial correlation structure (e.g. Mühle et al., 2019; Rigby et al., 2019). In this case a Matérn field would be inappropriate, but an autoregressive model may still be applicable for time-varying emissions. The inclusion of a non-stationary covariance makes the problem much more computationally expensive, and a more parsimonious approach may often suffice (Fuglstad et al., 2015).

5 Conclusions

This work presents a fast and efficient method using an integrated nested Laplacian approximation for hierarchical inference of trace gas emissions. This method is particularly well-suited to assimilating large data sets. We show that INLA with a stochastic partial differential equation approach for spatial correlation can reproduce synthetic emissions from pseudo-observations and benchmarked emissions using real data.

A real advantage over other hierarchical Bayesian inversion methods is the attractive convergence properties, which can be difficult to obtain using methods such as Markov chain Monte Carlo algorithms. As the method computes the marginal variance for each node, this allows for efficient parallel implementation and significant computational savings compared to other hierarchical methods. Computational speed will become increasingly important as more data from space-borne sensors become available, which will offer more measurements and increased spatial coverage.

References

Berchet, A., Pison, I., Chevallier, F., Bousquet, P., Bonne, J.-L., and Paris, J.-D.: Objectified quantification of uncertainties in Bayesian atmospheric inversions, Geosci. Model Dev., 8, 1525–1546, https://doi.org/10.5194/gmd-8-1525-2015, 2015.

Bergamaschi, P., Frankenberg, C., Meirink, J. F., Krol, M., Dentener, F., Wagner, T., Platt, U., Kaplan, J. O., Körner, S., Heimann, M., Dlugokencky, E. J., and Goede, A.: Satellite chartography of atmospheric methane from SCIA-MACHY on board ENVISAT: 2. Evaluation based on inverse model simulations, J. Geophys. Res.-Atmos., D02304, 112, https://doi.org/10.1029/2006JD007268, 2007.

Brioude, J., Angevine, W. M., Ahmadov, R., Kim, S.-W., Evan, S., McKeen, S. A., Hsie, E.-Y., Frost, G. J., Neuman, J. A., Pollack, I. B., Peischl, J., Ryerson, T. B., Holloway, J., Brown, S. S., Nowak, J. B., Roberts, J. M., Wofsy, S. C., Santoni, G. W., Oda, T., and Trainer, M.: Top-down estimate of surface flux in the Los Angeles Basin using a mesoscale inverse modeling technique: assessing anthropogenic emissions of CO, NOx and CO2 and their impacts, Atmos. Chem. Phys., 13, 3661–3677, https://doi.org/10.5194/acp-13-3661-2013, 2013.

Brunner, D., Henne, S., Keller, C. A., Reimann, S., Vollmer, M. K., O'Doherty, S., and Maione, M.: An extended Kalman-filter for regional scale inverse emission estimation, Atmos. Chem. Phys., 12, 3455–3478, https://doi.org/10.5194/acp-12-3455-2012, 2012.

Brunner, D., Arnold, T., Henne, S., Manning, A., Thompson, R. L., Maione, M., O'Doherty, S., and Reimann, S.: Comparison of four inverse modelling systems applied to the estimation of HFC-125, HFC-134a, and SF6 emissions over Europe, Atmos. Chem.

Phys., 17, 10651–10674, https://doi.org/10.5194/acp-17-10651-2017, 2017.

Cameletti, M., Lindgren, F., Simpson, D., and Rue, H.: Spatio-temporal modeling of particulate matter concentration through the SPDE approach, ASTA-Adv. Stat. Anal., 97, 109–131, https://doi.org/10.1007/s10182-012-0196-3, 2013.

Cowles, M. K. and Carlin, B. P.: Markov Chain Monte Carlo Convergence Diagnostics: A Comparative Review, J. Am. Stat. Assoc., 91, 883–904, https://doi.org/10.2307/2291683, 1996.

Cressie, N. A. C.: Statistics for spatial data, John Wiley & Sons, Inc, Hoboken, NJ, revised edition, 2015.

Cullen, M. J. P.: The Unified Forecast/Climate Model, Meteorol. Mag., 122, 81–94, 1993.

Dahlén, U., Lindström, J., and Scholze, M.: Spatiotemporal reconstructions of global CO_2-fluxes using Gaussian Markov random fields, Environmetrics, 2019, e2610, https://doi.org/10.1002/env.2610, 2019.

Department of Business, Energy and Industrial Strategy: 2017 UK Greenhouse Gas Emissions, Final Figures, 2019.

Emmons, L. K., Walters, S., Hess, P. G., Lamarque, J.-F., Pfister, G. G., Fillmore, D., Granier, C., Guenther, A., Kinnison, D., Laepple, T., Orlando, J., Tie, X., Tyndall, G., Wiedinmyer, C., Baughcum, S. L., and Kloster, S.: Description and evaluation of the Model for Ozone and Related chemical Tracers, version 4 (MOZART-4), Geosci. Model Dev., 3, 43–67, https://doi.org/10.5194/gmd-3-43-2010, 2010.

Fuglstad, G.-A., Simpson, D., Lindgren, F., and Rue, H.: Does non-stationary spatial data always require non-stationary random fields?, Spat. Stat., 14, 505–531, https://doi.org/10.1016/j.spasta.2015.10.001, 2015.

Fuglstad, G.-A., Simpson, D., Lindgren, F., and Rue, H.: Constructing Priors that Penalize the Complexity of Gaussian Random Fields, J. Am. Stat. Assoc., 114, 445–452, https://doi.org/10.1080/01621459.2017.1415907, 2018.

Ganesan, A. L., Rigby, M., Zammit-Mangion, A., Manning, A. J., Prinn, R. G., Fraser, P. J., Harth, C. M., Kim, K.-R., Krummel, P. B., Li, S., Mühle, J., O'Doherty, S. J., Park, S., Salameh, P. K., Steele, L. P., and Weiss, R. F.: Characterization of uncertainties in atmospheric trace gas inversions using hierarchical Bayesian methods, Atmos. Chem. Phys., 14, 3855–3864, https://doi.org/10.5194/acp-14-3855-2014, 2014.

Ganesan, A. L., Manning, A. J., Grant, A., Young, D., Oram, D. E., Sturges, W. T., Moncrieff, J. B., and O'Doherty, S.: Quantifying methane and nitrous oxide emissions from the UK and Ireland using a national-scale monitoring network, Atmos. Chem. Phys., 15, 6393–6406, https://doi.org/10.5194/acp-15-6393-2015, 2015.

Ganesan, A. L., Rigby, M., Lunt, M. F., Parker, R. J., Boesch, H., Goulding, N., Umezawa, T., Zahn, A., Chatterjee, A., Prinn, R. G., Tiwari, Y. K., van der Schoot, M., and Krummel, P. B.: Atmospheric observations show accurate reporting and little growth in India's methane emissions, Nat. Commun., 8, 836, https://doi.org/10.1038/s41467-017-00994-7, 2017.

Gelman, A. and Price, P. N.: All maps of parameter estimates are misleading, Stat. Med., 18, 3221–3234, https://doi.org/10.1002/(SICI)1097-0258(19991215)18:23<3221::AID-SIM312>3.0.CO;2-M, 1999.

Guttorp, P. and Gneiting, T.: Studies in the history of probability and statistics XLIX On the Matérn correlation family, Biometrika, 93, 989–995, https://doi.org/10.1093/biomet/93.4.989, 2006.

Henne, S., Brunner, D., Oney, B., Leuenberger, M., Eugster, W., Bamberger, I., Meinhardt, F., Steinbacher, M., and Emmenegger, L.: Validation of the Swiss methane emission inventory by atmospheric observations and inverse modelling, Atmos. Chem. Phys., 16, 3683–3710, https://doi.org/10.5194/acp-16-3683-2016, 2016.

Janssens-Maenhout, G., Crippa, M., Guizzardi, D., Muntean, M., Schaaf, E., Dentener, F., Bergamaschi, P., and Pagliari, V.: EDGAR v4.3.2 Global Atlas of the three major Greenhouse Gas Emissions for the period 1970-2012, Open Access, p. 55, 2017.

Jeong, S., Newman, S., Zhang, J., Andrews, A. E., Bianco, L., Bagley, J., Cui, X., Graven, H., Kim, J., Salameh, P., LaFranchi, B. W., Priest, C., Campos-Pineda, M., Novakovskaia, E., Sloop, C. D., Michelsen, H. A., Bambha, R. P., Weiss, R. F., Keeling, R., and Fischer, M. L.: Estimating methane emissions in California's urban and rural regions using multitower observations: Methane Emissions in California, J. Geophys. Res.-Atmos., 121, 13031–13049, https://doi.org/10.1002/2016JD025404, 2016.

Jones, A., Thomson, D., Hort, M., and Devenish, B.: The U.K. Met Office's Next-Generation Atmospheric Dispersion Model, NAME III, in: Air Pollution Modeling and Its Application XVII, edited by Borrego, C. and Norman, A.-L., Springer US, Boston, MA, 580–589, https://doi.org/10.1007/978-0-387-68854-1_62, 2006.

Leip, A., Skiba, U., Vermeulen, A., and Thompson, R. L.: A complete rethink is needed on how greenhouse gas emissions are quantified for national reporting, Atmos. Environ., 174, 237–240, https://doi.org/10.1016/j.atmosenv.2017.12.006, 2018.

Lindgren, F. and Rue, H.: Bayesian Spatial Modelling with R-INLA, J. Stat. Softw., 63, https://doi.org/10.18637/jss.v063.i19, 2015.

Lindgren, F., Rue, H., and Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach: Link between Gaussian Fields and Gaussian Markov Random Fields, J. Roy. Stat. Soc. B, 73, 423–498, https://doi.org/10.1111/j.1467-9868.2011.00777.x, 2011.

Lunt, M. F., Rigby, M., Ganesan, A. L., Manning, A. J., Prinn, R. G., O'Doherty, S., Mühle, J., Harth, C. M., Salameh, P. K., Arnold, T., Weiss, R. F., Saito, T., Yokouchi, Y., Krummel, P. B., Steele, L. P., Fraser, P. J., Li, S., Park, S., Reimann, S., Vollmer, M. K., Lunder, C., Hermansen, O., Schmidbauer, N., Maione, M., Arduini, J., Young, D., and Simmonds, P. G.: Reconciling reported and unreported HFC emissions with atmospheric observations, P. Natl. Acad. Sci. USA, 112, 5927–5931, https://doi.org/10.1073/pnas.1420247112, 2015.

Lunt, M. F., Rigby, M., Ganesan, A. L., and Manning, A. J.: Estimation of trace gas fluxes with objectively determined basis functions using reversible-jump Markov chain Monte Carlo, Geosci. Model Dev., 9, 3213–3229, https://doi.org/10.5194/gmd-9-3213-2016, 2016.

Manning, A. J., O'Doherty, S., Jones, A. R., Simmonds, P. G., and Derwent, R. G.: Estimating UK methane and nitrous oxide emissions from 1990 to 2007 using an inversion modeling approach, J. Geophys. Res., 116, D02305, https://doi.org/10.1029/2010JD014763, 2011.

Mardia, K. V., Kent, J. T., and Bibby, J. M.: Multivariate analysis, Probability and mathematical statistics, Academic Press, London, New York, 1979.

Marques, I., Klein, N., and Kneib, T.: Non-stationary spatial regression for modelling monthly precipitation in Germany, Spat. Stat., 100386, https://doi.org/10.1016/j.spasta.2019.100386, 2019.

Martins, T. G., Simpson, D., Lindgren, F., and Rue, H.: Bayesian computing with INLA: New features, Comput. Stat. Data An., 67, 68–83, https://doi.org/10.1016/j.csda.2013.04.014, 2013.

Met Office: Operational Numerical Weather Prediction (NWP) Output from the UK Variable (UKV) Resolution Part of the Met Office Unified Model (UM), NCAS British Atmospheric Data Centre, available at: http://catalogue.ceda.ac.uk/uuid/292da1ccfebd650f6d123e53270016a8 (last access: 13 February 2020), 2013

Michalak, A. M., Hirsch, A., Bruhwiler, L., Gurney, K. R., Peters, W., and Tans, P. P.: Maximum likelihood estimation of covariance parameters for Bayesian atmospheric trace gas surface flux inversions, J. Geophys. Res.-Atmos., 110, D24107, https://doi.org/10.1029/2005JD005970, 2005.

Mühle, J., Trudinger, C. M., Western, L. M., Rigby, M., Vollmer, M. K., Park, S., Manning, A. J., Say, D., Ganesan, A., Steele, L. P., Ivy, D. J., Arnold, T., Li, S., Stohl, A., Harth, C. M., Salameh, P. K., McCulloch, A., O'Doherty, S., Park, M.-K., Jo, C. O., Young, D., Stanley, K. M., Krummel, P. B., Mitrevski, B., Hermansen, O., Lunder, C., Evangeliou, N., Yao, B., Kim, J., Hmiel, B., Buizert, C., Petrenko, V. V., Arduini, J., Maione, M., Etheridge, D. M., Michalopoulou, E., Czerniak, M., Severinghaus, J. P., Reimann, S., Simmonds, P. G., Fraser, P. J., Prinn, R. G., and Weiss, R. F.: Perfluorocyclobutane (PFC-318, c-C4F8) in the global atmosphere, Atmos. Chem. Phys., 19, 10335–10359, https://doi.org/10.5194/acp-19-10335-2019, 2019.

Prinn, R. G., Weiss, R. F., Arduini, J., Arnold, T., Fraser, P. J., Ganesan, A. L., Gasore, J., Harth, C. M., Hermansen, O., Kim, J., Krummel, P. B., Li, S., Loh, Z. M., Lunder, C. R., Maione, M., Manning, A. J., Miller, B. R., Mitrevski, B., Mühle, J., O'Doherty, S., Park, S., Reimann, S., Rigby, M., Salameh, P. K., Schmidt, R., Simmonds, P.G., Steele, L. P., Vollmer, M. K., Wang, R. H., and Young, D.: The ALE/GAGE/AGAGE Data Base, available at: http://agage.mit.edu/data, last access: 13 February 2020.

Rigby, M., Park, S., Saito, T., Western, L. M., Redington, A. L., Fang, X., Henne, S., Manning, A. J., Prinn, R. G., Dutton, G. S., Fraser, P. J., Ganesan, A. L., Hall, B. D., Harth, C. M., Kim, J., Kim, K.-R., Krummel, P. B., Lee, T., Li, S., Liang, Q., Lunt, M. F., Montzka, S. A., Mühle, J., O'Doherty, S., Park, M.-K., Reimann, S., Salameh, P. K., Simmonds, P., Tunnicliffe, R. L., Weiss, R. F., Yokouchi, Y., and Young, D.: Increase in CFC-11 emissions from eastern China based on atmospheric observations, Nature, 569, 546–550, https://doi.org/10.1038/s41586-019-1193-4, 2019.

Roberts, G. O. and Rosenthal, J. S.: Optimal scaling of discrete approximations to Langevin diffusions, J. Roy. Stat. Soc. B, 60, 255–268, https://doi.org/10.1111/1467-9868.00123, 1998.

Rue, H. and Follestad, T.: GMRFLib: a C-library for fast and exact simulation of Gaussian Markov random fields, Tech. rep., SIS-2002-236, 2001.

Rue, H. and Held, L.: Gaussian Markov random fields: theory and applications, Chapman & Hall/CRC, Boca Raton, oCLC: 839047556, 2005.

Rue, H., Martino, S., and Chopin, N.: Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations, J. Roy. Stat. Soc. B, 71, 319–392, https://doi.org/10.1111/j.1467-9868.2008.00700.x, 2009.

Saunois, M., Bousquet, P., Poulter, B., Peregon, A., Ciais, P., Canadell, J. G., Dlugokencky, E. J., Etiope, G., Bastviken, D., Houweling, S., Janssens-Maenhout, G., Tubiello, F. N., Castaldi, S., Jackson, R. B., Alexe, M., Arora, V. K., Beerling, D. J., Bergamaschi, P., Blake, D. R., Brailsford, G., Brovkin, V., Bruhwiler, L., Crevoisier, C., Crill, P., Covey, K., Curry, C., Frankenberg, C., Gedney, N., Höglund-Isaksson, L., Ishizawa, M., Ito, A., Joos, F., Kim, H.-S., Kleinen, T., Krummel, P., Lamarque, J.-F., Langenfelds, R., Locatelli, R., Machida, T., Maksyutov, S., McDonald, K. C., Marshall, J., Melton, J. R., Morino, I., Naik, V., O'Doherty, S., Parmentier, F.-J. W., Patra, P. K., Peng, C., Peng, S., Peters, G. P., Pison, I., Prigent, C., Prinn, R., Ramonet, M., Riley, W. J., Saito, M., Santini, M., Schroeder, R., Simpson, I. J., Spahni, R., Steele, P., Takizawa, A., Thornton, B. F., Tian, H., Tohjima, Y., Viovy, N., Voulgarakis, A., van Weele, M., van der Werf, G. R., Weiss, R., Wiedinmyer, C., Wilton, D. J., Wiltshire, A., Worthy, D., Wunch, D., Xu, X., Yoshida, Y., Zhang, B., Zhang, Z., and Zhu, Q.: The global methane budget 2000–2012, Earth Syst. Sci. Data, 8, 697–751, https://doi.org/10.5194/essd-8-697-2016, 2016.

Sha, Z., Rougier, J. C., Schumacher, M., and Bamber, J. L.: Bayesian model–data synthesis with an application to global glacio-isostatic adjustment, Environmetrics, 30, e2530, https://doi.org/10.1002/env.2530, 2019.

Shewchuk, J. R.: Delaunay refinement algorithms for triangular mesh generation, Comp. Geom., 22, 21–74, https://doi.org/10.1016/S0925-7721(01)00047-5, 2002.

Simpson, D., Rue, H., Riebler, A., Martins, T. G., and Sørbye, S. H.: Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors, Stat. Sci., 32, 1–28, https://doi.org/10.1214/16-STS576, 2017.

Stanley, K. M., Grant, A., O'Doherty, S., Young, D., Manning, A. J., Stavert, A. R., Spain, T. G., Salameh, P. K., Harth, C. M., Simmonds, P. G., Sturges, W. T., Oram, D. E., and Derwent, R. G.: Greenhouse gas measurements from a UK network of tall towers: technical description and first results, Atmos. Meas. Tech., 11, 1437–1458, https://doi.org/10.5194/amt-11-1437-2018, 2018.

Stein, M. L.: Interpolation of Spatial Data: Some Theory for Kriging, Springer Series in Statistics, Springer, New York, 1999.

Stohl, A., Seibert, P., Arduini, J., Eckhardt, S., Fraser, P., Greally, B. R., Lunder, C., Maione, M., Mühle, J., O'Doherty, S., Prinn, R. G., Reimann, S., Saito, T., Schmidbauer, N., Simmonds, P. G., Vollmer, M. K., Weiss, R. F., and Yokouchi, Y.: An analytical inversion method for determining regional and global emissions of greenhouse gases: Sensitivity studies and application to halocarbons, Atmos. Chem. Phys., 9, 1597–1620, https://doi.org/10.5194/acp-9-1597-2009, 2009.

Tierney, L. and Kadane, J. B.: Accurate Approximations for Posterior Moments and Marginal Densities, J. Am. Stat. Assoc., 81, 82–86, https://doi.org/10.2307/2287970, 1986.

Tørseth, K., Aas, W., Breivik, K., Fjæraa, A. M., Fiebig, M., Hjellbrekke, A. G., Lund Myhre, C., Solberg, S., and Yttri, K. E.: Introduction to the European Monitoring and Evaluation Programme (EMEP) and observed atmospheric composition change during 1972–2009, Atmos. Chem. Phys., 12, 5447–5481, https://doi.org/10.5194/acp-12-5447-2012, 2012.

Veefkind, J., Aben, I., McMullan, K., Förster, H., de Vries, J., Otter, G., Claas, J., Eskes, H., de Haan, J., Kleipool, Q., van Weele, M., Hasekamp, O., Hoogeveen, R., Landgraf, J., Snel, R., Tol, P., Ingmann, P., Voors, R., Kruizinga, B., Vink, R., Visser, H., and Levelt, P.: TROPOMI on the ESA Sentinel-5 Precursor: A GMES mission for global observations of the atmospheric composition for climate, air quality and ozone layer applications, Remote Sens. Environ., 120, 70–83, https://doi.org/10.1016/j.rse.2011.09.027, 2012.

Wallin, J. and Bolin, D.: Geostatistical Modelling Using Non-Gaussian Matérn Fields, Scand. J. Stat., 42, 872–890, https://doi.org/10.1111/sjos.12141, 2015.

Wecht, K. J., Jacob, D. J., Sulprizio, M. P., Santoni, G. W., Wofsy, S. C., Parker, R., Bösch, H., and Worden, J.: Spatially resolving methane emissions in California: constraints from the CalNex aircraft campaign and from present (GOSAT, TES) and future (TROPOMI, geostationary) satellite observations, Atmos. Chem. Phys., 14, 8173–8184, https://doi.org/10.5194/acp-14-8173-2014, 2014.

Western, L.: INLA_GHG_GMD, OSF, https://doi.org/10.17605/OSF.IO/53W96, 2019.

Whittle, P.: On Stationary Processes in the Plane, Biometrika, 434–449, 1954.

Whittle, P.: Stochastic-processes in several dimensions, Int Statistical Institude 428 Prinses Beatrixlaen, Voorburg, the Netherlands, 40, 974–994, 1963.

Zammit-Mangion, A., Cressie, N., and Ganesan, A. L.: Non-Gaussian bivariate modelling with application to atmospheric trace-gas inversion, Spat. Stat., 18, 194–220, https://doi.org/10.1016/j.spasta.2016.06.005, 2016.