University of Wollongong

Research Online

National Institute for Applied Statistics
Research Australia Working Paper Series

Faculty of Engineering and Information
Sciences

2016

Non-gaussian bivariate modelling with application to atmospheric trace-gas inversion

Andrew Zammit-Mangion
University of Wollongong

Noel Cressie
University of Wollongong

Anita L. Ganesan
University of Bristol

Follow this and additional works at: <https://ro.uow.edu.au/niasrawp>

Recommended Citation

Zammit-Mangion, Andrew; Cressie, Noel; and Ganesan, Anita L., Non-gaussian bivariate modelling with application to atmospheric trace-gas inversion, National Institute for Applied Statistics Research Australia, University of Wollongong, Working Paper 06-16, 2016, 45.
<https://ro.uow.edu.au/niasrawp/39>

Research Online is the open access institutional repository for the University of Wollongong. For further information contact the UOW Library: research-pubs@uow.edu.au

Non-gaussian bivariate modelling with application to atmospheric trace-gas inversion

Abstract

Atmospheric trace-gas inversion is the procedure by which the sources and sinks of a trace gas are identified from observations of its mole fraction at isolated locations in space and time. This is inherently a spatio-temporal bivariate inversion problem, since the mole-fraction field evolves in space and time and the flux is also spatio-temporally distributed. Further, the bivariate model is likely to be non-Gaussian since the flux field is rarely Gaussian. Here, we use conditioning to construct a non-Gaussian bivariate model, and we describe some of its properties through auto- and cross-cumulant functions. A bivariate non-Gaussian, specifically trans-Gaussian, model is then achieved through the use of Box-Cox transformations, and we facilitate Bayesian inference by approximating the likelihood in a hierarchical framework. Trace-gas inversion, especially at high spatial resolution, is frequently highly sensitive to prior specification. Therefore, unlike conventional approaches, we assimilate trace-gas inventory information with the observational data at the parameter layer, thus shifting prior sensitivity from the inventory itself to its spatial characteristics (e.g., its spatial length scale). We demonstrate the approach in controlled-experiment studies of methane inversion, using fluxes extracted from inventories of the UK and Ireland and of Northern Australia.

NIASRA

NATIONAL INSTITUTE FOR APPLIED
STATISTICS RESEARCH AUSTRALIA

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

University of Wollongong, Australia

Working Paper

06-16

**Non-Gaussian Bivariate Modelling with Application to
Atmospheric Trace-gas Inversion**

Andrew Zammit-Mangion, Noel Cressie, and Anita L. Ganesan

*Copyright © 2016 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 Phone +61 2 4221 5435, Fax +61 2 4221 4998.
Email: karink@uow.edu.au

Non-Gaussian bivariate modelling with application to atmospheric trace-gas inversion

Andrew Zammit-Mangion^{a,*}, Noel Cressie^a, Anita L. Ganesan^b

^a*National Institute for Applied Statistics Research Australia (NIASRA), School of Mathematics and Applied Statistics (SMAS), University of Wollongong, Northfields Avenue, Wollongong, NSW 2522, Australia*

^b*School of Geographical Sciences, University of Bristol, University Road, Bristol, BS8 1SS, UK*

Abstract

Atmospheric trace-gas inversion is the procedure by which the sources and sinks of a trace gas are identified from observations of its mole fraction at isolated locations in space and time. This is inherently a spatio-temporal bivariate inversion problem, since the mole-fraction field evolves in space and time and the flux is also spatio-temporally distributed. Further, the bivariate model is likely to be non-Gaussian since the flux field is rarely Gaussian. Here, we use conditioning to construct a non-Gaussian bivariate model, and we describe some of its properties through auto- and cross-cumulant functions. A bivariate non-Gaussian, specifically trans-Gaussian, model is then achieved through the use of Box–Cox transformations, and we facilitate Bayesian inference by approximating the likelihood in a hierarchical framework. Trace-gas inversion, especially at high spatial resolution, is frequently highly sensitive to prior specification. Therefore, unlike conventional approaches, we assimilate trace-gas inventory information with the observational data at the parameter layer, thus shifting prior sensitivity from the inventory itself to its spatial characteristics (e.g., its spatial length scale). We demonstrate the approach in controlled-experiment studies of methane inversion, using fluxes extracted from inventories of the UK and Ireland and of Northern Australia.

Keywords: Bivariate spatial model, Conditional multivariate model, Methane emissions, Multivariate geostatistics, Trans-Gaussian model, Box–Cox transformation

1. Introduction

Atmospheric trace-gas inversion is the procedure by which flux fields (gas sinks and sources) are identified from gas mole-fraction observations. Unlike conventional problems in spatial

*Corresponding author

Email addresses: azm@uow.edu.au (Andrew Zammit-Mangion), ncressie@uow.edu.au (Noel Cressie), anita.ganesan@bristol.ac.uk (Anita L. Ganesan)

statistics, the spatial field of principal interest, the flux field, is rarely directly observed. Instead, satellite and surface gas-concentration instruments are sensitive to gas particles after they have been transported over potentially large distances over time. Spatio-temporal statistical methodology is well placed to obtain global space-time maps of mole fractions of an atmospheric constituent from irregular observations (see, for example, Cressie et al., 2010; Cameletti et al., 2013; Lindgren et al., 2011); however, the inversion of these maps in order to pinpoint the sources and sinks of the gas is a much more difficult problem. Its solution is key to effective policy implementation with regard to greenhouse gas emissions and climate change (Edenhofer et al., 2014).

Current approaches to trace-gas inversion build on the data-assimilation framework described, for instance, in Tarantola (2005), where prior beliefs on a flux field are updated with observations, to produce posterior beliefs (see, for example, Rigby et al., 2011; Stohl et al., 2009). The prior distribution is usually formulated to have prior expectation equal to values in an *inventory*, a flux database constructed from auxiliary activity data (e.g., vehicles per unit area) and emission factors associated with the emission sector (e.g., carbon dioxide emissions per vehicle), while the prior covariance is usually constructed using values calculated from the inventory. Reasonable prior marginal variances are typically assumed; for example, a prior density function at each location may be chosen such that the area under the curve between 0.5 and 1.5 times the value of the inventory at that location is around 68% (Ganesan et al., 2014). Frequently, a diagonal structure is imposed on the prior covariance for all locations. It is also usually assumed that the prior expectation and prior covariance completely specify the prior distribution.

These current approaches can be critiqued due to their use of inappropriate or inflexible models. First, atmospheric trace-gas inversion is inherently a bivariate spatio-temporal problem, where observations are not readings of fluxes, but rather readings of a second, mole-fraction, field. The mole-fraction field is generated by the underlying flux field and by meteorology, which determines transport of the trace gas. Making this distinction allows one to attribute uncertainties appropriately, either to instrumentation error or to imprecise mole-fraction field modelling (e.g., due to linearisation of the flux-mole-fraction mapping or imprecise specification of transport modelling/boundary conditions). Second, a Gaussian model is often inappropriate for the flux field (Ganesan et al., 2014; Miller et al., 2014), which can be inherently non-negative and which, as seen from inventories, may exhibit skewness and higher-order features.

Third, spatial correlations in prior error covariances, even when using regressors, should be assumed since errors in the inventories are likely to be spatially correlated. Fourth, several works claim that inventories are highly inaccurate in certain regions (e.g., Lunt et al., 2015). Consequently, there is a drive to divert from use of the inventories in the model, and to take a data-driven approach to flux inversion, even when sub-national resolution is required and so prior information is especially important.

These reservations were first discussed in Zammit-Mangion et al. (2015), where the authors constructed a hierarchical, lognormal bivariate spatio-temporal model for flux inversion of methane in the UK and Ireland. In that article, an empirical hierarchical modelling approach (Cressie and Wikle, 2011, Section 2.1) was adopted, where numerous parameters were estimated offline prior to drawing samples from the empirical posterior distribution of the flux field. In this article, we extend the modelling and inferential approach in Zammit-Mangion et al. (2015) in two ways. First, we relax the assumption of lognormality by defining a more general non-Gaussian model, and subsequently we model the flux field as a trans-Gaussian (here a Box–Cox) spatial process (De Oliveira et al., 1997). The class of Box–Cox spatial processes includes the lognormal spatial process as a special case. Second, following a likelihood approximation, we adopt a fully Bayesian approach to flux-field inversion that naturally propagates variability in the parameters (parameters that were estimated offline in Zammit-Mangion et al., 2015) to our inferences on the flux field. In order to reduce reliance on the flux inventory, which is based on emission factors that are not precisely known, we do not assume it to be the prior mean of the flux field. Instead, we take a new approach by assuming that it is an independent realisation of the flux field that we wish to infer. Hence, we assume that the inventory is informative of the spatial (prior) properties of the flux field, but not of the flux field itself, so that posterior inferences on fluxes are still predominantly data-driven. We demonstrate the value of this approach in a controlled experiment where we infer flux fields in the UK using the true flux as our inventory, an inventory with similar spatial properties (extracted from mainland Europe), and one with highly dissimilar spatial properties (from Northern Australia). In the experiment, inferred flux fields are compared to the known, true flux field.

The article is organised as follows. In Section 2 we motivate atmospheric trace-gas inversion as a problem in bivariate modelling (Cressie and Zammit-Mangion, 2015; Royle and Berliner, 1999) and outline a few of the proposed bivariate model’s theoretical properties. In Section 3 we focus on the bivariate trans-Gaussian model and place it within a hierarchical framework

appropriate for describing the problem of trace-gas inversion. The transformation that we feature here is the well known Box–Cox transformation. The hierarchical model we employ is summarised in Section 3.3; the succinct summary presented there can help to put the model descriptions of the preceding sections into perspective. In Section 4 we outline the Markov chain Monte Carlo (MCMC) scheme adopted and the approximations we use to facilitate its implementation. In Section 5 we apply the framework to a realistic simulation study of methane emissions in the UK and Ireland, where fluxes and synthetic mole-fraction observations are simulated using meteorology from a numerical model and UK/Ireland flux inventories. To show the flexibility of the model, we also apply it to the case when mole fractions are simulated from a Northern Australian inventory. The methane emissions in this region are considerably smoother (spatially) than those in the UK and Ireland. Conclusions are drawn in Section 6, and the article finishes with three technical appendices.

2. Application and modelling overview

In this section we motivate the problem of atmospheric trace-gas inversion as a problem in non-Gaussian bivariate modelling (Section 2.1), and discuss some properties of the non-Gaussian bivariate model (Section 2.2).

2.1. Atmospheric trace-gas inversion

Monitoring of greenhouse gases and air pollutants is a key priority for environmental agencies and institutions worldwide. One of its primary aims is to be able to determine where the gas is being produced (sources) and where it is being removed (sinks). This is not a straightforward task, since the net flux of the production and removal processes is related to the observed mole fraction at a particular time and location after transport by meteorology and alteration by chemical reactions. These processes can only be obtained from computationally intensive numerical models.

The numerical models employed are either Eulerian (e.g., the Comprehensive Air quality Model with eXtensions (CAMx); see Emery et al., 2012) or Lagrangian (e.g., the Numerical Atmospheric-dispersion Modelling Environment (NAME) or the FLEXible PARTicle dispersion model (FLEXPART); see Jones et al., 2007; Thompson and Stohl, 2014) in nature. Although these two classes of models are fundamentally different, they serve the same purpose in inversion, namely to establish the sensitivity of the mole fractions to the flux. In the application we

consider in Section 5, we shall use a Lagrangian model for regional inversion (specifically the model NAME), however the inversion framework we construct is applicable to both classes of models.

Now, consider a spatially referenced flux field, $Y_1(\cdot)$, and a spatially referenced mole-fraction field (at some given time instant), $Y_2(\cdot)$. Then meteorology and loss (chemical) processes imply that

$$Y_2(\cdot) = \mathcal{H}(Y_1(\cdot)), \quad (1)$$

where $\mathcal{H}(\cdot)$ is some nonlinear operator that Eulerian or Lagrangian models attempt to reconstruct. Following Cressie and Wikle (2011, Section 7.3), a stochastic version of this model can be written, conditional on $Y_1(\cdot)$, as

$$Y_2(\cdot) = \mathcal{H}(Y_1(\cdot); \xi(\cdot)), \quad (2)$$

where $\xi(\cdot)$ represents a process that expresses uncertainty in the physical relationship (1) that defines a model for $Y_2(\cdot)$ in terms of $Y_1(\cdot)$.

Consequently, (2) defines a bivariate stochastic process that incorporates not only knowledge about inventories and transport, but also uncertainty around how that knowledge results in the mole-fraction field $Y_2(\cdot)$. With a few exceptions (e.g., Peters et al., 2005), a linear mapping that maintains the physical relationship to leading order, $\mathcal{H}_L$, is substituted for $\mathcal{H}$, and a statistical model

$$Y_2(\cdot) = \mathcal{H}_L(Y_1(\cdot)) + \zeta(\cdot), \quad (3)$$

is assumed, conditional on $Y_1(\cdot)$. In (3), $\zeta(\cdot)$ collects discrepancies and uncertainties from (i) using leading-order, linear physical relationships and (ii) inaccuracies inherent to the chemical transport model. When analysing short-lived gases, such as carbon monoxide, the sink due to chemical reaction (a component of $\mathcal{H}$) can be nonlinearly dependent on the flux. However, for long-lived gases such as methane, considered in Section 5, this sink is approximately flux-independent. We therefore expect uncertainties arising from (i) to be dominated by (ii). When using $\mathcal{H}_L$ instead of $\mathcal{H}$, model interpretation and inference is greatly facilitated. In all that follows, we base our work on (3).

Using this conditional approach to building a bivariate process that respects the physical system expressed (to leading order) by (3), is very powerful, since it now only remains to define

a univariate spatial process for $Y_1(\cdot)$. That process need not be Gaussian, and indeed for the methane flux fields considered in Section 5, it will generally not be. Consequently, from (2) and (3) the joint process $(Y_1(\cdot), Y_2(\cdot))$ is tied to both the physics expressed in $\mathcal{H}_L$ and the stochastic (possibly non-Gaussian) properties of $Y_1(\cdot)$.

The joint statistical properties of $Y_1(\cdot)$ and $\zeta(\cdot)$ fully determine the statistical properties of $Y_2(\cdot)$. We briefly discuss these properties for general $Y_1(\cdot)$ and when $\mathcal{H}_L$ is an integral transform in Section 2.2. Unlike $Y_2(\cdot)$, observations of $Y_1(\cdot)$, when available, only capture the process' very fine scales. Hence one cannot enjoy the use of typical exploratory-data-analysis techniques to guide the choice of model for $Y_1(\cdot)$. The presence of *flux inventories* can help, although over-reliance on these inventories, which are hard to validate, is an ongoing concern (e.g., Berchet et al., 2015); we return to this point in Section 3.2.

This article is concerned with an appropriate model choice for $Y_1(\cdot)$ and $\zeta(\cdot)$. Both for analytical and computational convenience, $Y_1(\cdot)$ and $\zeta(\cdot)$ are frequently assumed to be Gaussian (e.g., Stohl et al., 2009) in which case $Y_2(\cdot)$ is Gaussian and $(Y_1(\cdot), Y_2(\cdot))$ is jointly Gaussian. In recent years, there has been a shift away from Gaussianity assumptions; for example, Rigby et al. (2011) considered exponential distributions for $Y_1(\cdot)$ evaluated at the grid-cell level, Ganesan et al. (2014) considered lognormal distributions, and Miller et al. (2014) considered a truncated spatial Gaussian process for $Y_1(\cdot)$. Zammit-Mangion et al. (2015) showed that both lognormality and spatial correlation in $Y_1(\cdot)$ are important for prediction in the UK and Ireland. However, it is also the case that lognormality may not be a suitable assumption everywhere and at all grid resolutions. A flexible class of models that can accommodate different degrees of non-Gaussianity and spatial correlation is the class of trans-Gaussian spatial processes (Cressie, 1993, p. 137); we use them within a trace-gas inversion framework in Section 3 and illustrate the importance of the added flexibility in Section 5.

2.2. Properties of the bivariate non-Gaussian model

If $Y_1(\cdot)$ and $\zeta(\cdot)$ are Gaussian processes then the properties of the joint process $(Y_1(\cdot), Y_2(\cdot))$ are straightforward to derive (Cressie and Zammit-Mangion, 2015). When $Y_1(\cdot)$ is a lognormal process, then the first-order and second-order moments of the finite-dimensional distributions of $(Y_1(\cdot), Y_2(\cdot))$ can be expressed in closed form (Zammit-Mangion et al., 2015). In this section we discuss the properties of the joint process $(Y_1(\cdot), Y_2(\cdot))$ when $Y_1(\cdot)$ is any (in general, non-Gaussian) process.

Consider a real-valued random field $\{Y_1(\mathbf{s}) : \mathbf{s} \in D \subset \mathbb{R}^d\}$, where typically $d = 2$ and D is a continuously indexed domain. Construct the real-valued field $\{Y_2(\mathbf{s}) : \mathbf{s} \in D \subset \mathbb{R}^d\}$ through

$$Y_2(\mathbf{s}) = \int_D b(\mathbf{s}, \mathbf{u}) Y_1(\mathbf{u}) d\mathbf{u} + \zeta(\mathbf{s}); \quad \mathbf{s} \in D, \quad (4)$$

where we term the integral kernel $b(\cdot, \cdot)$ an interaction function and $\zeta(\cdot)$ is a discrepancy term independent of $Y_1(\cdot)$. The integral in (4) exists for all $\mathbf{s} \in D$ as a quadratic mean limit of Riemann sums provided that

$$C_{Y_2 Y_2}(\mathbf{s}_1, \mathbf{s}_2) \equiv \int_D \int_D b(\mathbf{s}_1, \mathbf{u}_1) b(\mathbf{s}_2, \mathbf{u}_2) C_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2) d\mathbf{u}_1 d\mathbf{u}_2 < \infty; \quad \mathbf{s}_1, \mathbf{s}_2 \in D, \quad (5)$$

and $C_{Y_1 Y_1}(\cdot, \cdot) \equiv \text{cov}(Y_1(\cdot), Y_1(\cdot))$ is continuous on $D \times D$ (see Yaglom (1987), p. 68, and Lindgren (2012), Theorem 2.16). Thus, provided that $\zeta(\cdot)$ is well defined and (5) holds, $Y_2(\cdot)$ exists by construction. Further, from (5), if $\frac{\partial}{\partial \mathbf{s}} b(\mathbf{s}, \cdot)$ exists for all $\mathbf{s} \in D$, and if $\zeta(\cdot)$ is differentiable in quadratic mean, then $Y_2(\cdot)$ is differentiable in quadratic mean as well (Åström, 2006, p.37).

Without loss of generality, we consider the (joint) spatial cumulants of the processes evaluated at $n \geq 1$, possibly repeated, locations. Let $\varphi_{\mathbf{Y}_1}(\mathbf{t}_1)$ be the characteristic function of the vector $(Y_1(\mathbf{u}_1), \dots, Y_1(\mathbf{u}_n))'$, where it is understood that $\varphi_{\mathbf{Y}_1}(\mathbf{t}_1)$ is a function of the locations $\mathbf{u}_1, \dots, \mathbf{u}_n$. We define the n -th order *spatial cumulant function* of $\mathbf{u}_1, \dots, \mathbf{u}_n$ to be

$$\kappa_{Y_1 \dots Y_1}^n(\mathbf{u}_1, \dots, \mathbf{u}_n) \equiv \frac{1}{\iota^n} \frac{\partial^n}{\partial t_{11} \dots \partial t_{1n}} \ln \varphi_{\mathbf{Y}_1}(\mathbf{t}_1) \Big|_{\mathbf{t}_1 = \mathbf{0}}, \quad (6)$$

where the superscript of $\kappa(\cdot)$ denotes the spatial cumulant order, the subscript denotes the processes associated with the spatial locations, and ι is the imaginary unit. Since (6) is dependent only on Y_1 , it is an *auto-cumulant function*. *Cross-cumulant functions* can be constructed by assuming that a possibly different process is evaluated at each location, in which case

$$\kappa_{Y_{j_1} \dots Y_{j_n}}^n(\mathbf{s}_1, \dots, \mathbf{s}_n) \equiv \frac{1}{\iota^n} \frac{\partial^n}{\partial t_{j_1 1} \dots \partial t_{j_n n}} \ln \varphi_{\mathbf{Y}}(\mathbf{t}) \Big|_{\mathbf{t} = \mathbf{0}}, \quad (7)$$

where $j_k \in \{1, 2\}$, for $k = 1, \dots, n$; $\mathbf{Y} = (Y_{j_k}(\mathbf{s}_k) : k = 1, \dots, n)'$; and $\mathbf{t} = (t_{j_k k} : k = 1, \dots, n)'$.

We are interested in characterising the joint process $(Y_1(\cdot), Y_2(\cdot))$ when the auto-cumulant functions associated with $Y_1(\cdot)$ are known. We start by considering the auto-cumulant functions of $Y_2(\cdot)$. Since the cumulant of a sum of independent random variables is the sum of the

individual cumulants, the auto-cumulant functions associated with $Y_2(\cdot)$ are a sum of those associated with the processes, $\int_D b(\cdot, \mathbf{u})Y_1(\mathbf{u})d\mathbf{u}$ and $\zeta(\cdot)$. Hence, it can be shown through use of the characteristic function (Kuznetsov et al., 1965, Appendix III) that, provided all auto-cumulant functions associated with $Y_1(\cdot)$ are integrable,

$$\kappa_{Y_2}^1(\mathbf{s}_1) = \int_D b(\mathbf{s}_1, \mathbf{u}_1)\kappa_{Y_1}^1(\mathbf{u}_1)d\mathbf{u}_1 + \kappa_{\zeta}^1(\mathbf{s}_1); \quad \mathbf{s}_1 \in D, \quad (8)$$

$$\kappa_{Y_2 Y_2}^2(\mathbf{s}_1, \mathbf{s}_2) = \int_D \int_D b(\mathbf{s}_1, \mathbf{u}_1)b(\mathbf{s}_2, \mathbf{u}_2)\kappa_{Y_1 Y_1}^2(\mathbf{u}_1, \mathbf{u}_2)d\mathbf{u}_1 d\mathbf{u}_2 + \kappa_{\zeta \zeta}^2(\mathbf{s}_1, \mathbf{s}_2); \quad \mathbf{s}_1, \mathbf{s}_2 \in D, \quad (9)$$

and the n -th order auto-cumulant function associated with $Y_2(\cdot)$ is, for $n \geq 1$,

$$\begin{aligned} \kappa_{Y_2 \dots Y_2}^n(\mathbf{s}_1, \dots, \mathbf{s}_n) &= \int_D \dots \int_D b(\mathbf{s}_1, \mathbf{u}_1) \dots b(\mathbf{s}_n, \mathbf{u}_n) \kappa_{Y_1 \dots Y_1}^n(\mathbf{u}_1, \dots, \mathbf{u}_n) d\mathbf{u}_1 \dots d\mathbf{u}_n \\ &\quad + \kappa_{\zeta \dots \zeta}^n(\mathbf{s}_1, \dots, \mathbf{s}_n); \quad \mathbf{s}_1, \dots, \mathbf{s}_n \in D. \end{aligned} \quad (10)$$

When $\zeta(\cdot) = 0$ in (4), and hence $\kappa_{\zeta \dots \zeta}^n(\mathbf{s}_1, \dots, \mathbf{s}_n) \equiv 0$, we have

$$\kappa_{Y_2 \dots Y_2}^n(\mathbf{s}_1, \dots, \mathbf{s}_n) = \int_D \dots \int_D b(\mathbf{s}_1, \mathbf{u}_1) \dots b(\mathbf{s}_n, \mathbf{u}_n) \kappa_{Y_1 \dots Y_1}^n(\mathbf{u}_1, \dots, \mathbf{u}_n) d\mathbf{u}_1 \dots d\mathbf{u}_n,$$

which is a generalisation to cumulants of the bilinear property of covariances, namely

$$\text{cov}\left(\sum_{k=1}^{N_1} a_k U_{1k}, \sum_{l=1}^{N_2} b_l U_{2l}\right) = \sum_{k=1}^{N_1} \sum_{l=1}^{N_2} a_k b_l \text{cov}(U_{1k}, U_{2l}),$$

where $\mathbf{U}_1$ and $\mathbf{U}_2$ are random vectors of length N_1 and N_2 , respectively; $a_k \in \mathbb{R}$, for $k = 1, \dots, N_1$; and $b_l \in \mathbb{R}$, for $l = 1, \dots, N_2$.

In order to fully characterise the joint process $(Y_1(\cdot), Y_2(\cdot))$, we also need the cross-cumulant functions that involve both $Y_1(\cdot)$ and $Y_2(\cdot)$. This entails a lengthy but straightforward extension to the arguments of Kuznetsov et al. (1965), which we defer to Appendix A. We show that the cross-cumulant functions are rather simple in form; for example,

$$\kappa_{Y_1 Y_1 Y_2}^3(\mathbf{u}_1, \mathbf{u}_2, \mathbf{s}_3) = \int \kappa_{Y_1 Y_1 Y_1}^3(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3) b(\mathbf{s}_3, \mathbf{u}_3) d\mathbf{u}_3. \quad (11)$$

All cross-cumulant functions of $(Y_1(\cdot), Y_2(\cdot))$ do not involve the process $\zeta(\cdot)$ since $Y_1(\cdot)$ and $\zeta(\cdot)$ are independent. Of all the auto-cumulant and cross-cumulant functions, the second-order ones are arguably the most important since they are also the marginal and cross-covariance functions

of the joint process. It can be shown (Cressie and Zammit-Mangion, 2015, Section 3) that the 2×2 covariance matrix $(\kappa_{Y_i Y_j}^2(\mathbf{s}_1, \mathbf{s}_2) : i, j \in \{1, 2\})$ given by (9) is non-negative definite for all $\mathbf{s}_1, \mathbf{s}_2 \in D$.

In the special case where $b(\mathbf{s}, \mathbf{u})$; $\mathbf{s}, \mathbf{u} \in D$, is solely a function of $\mathbf{s} - \mathbf{u}$, the spectra of the auto-cumulants of $Y_2(\cdot)$ (known as polyspectra) can be found through a Fourier transform of (10). However, in our application, $b(\cdot, \cdot)$ describes a sensitivity of mole fraction to flux at a given spatial location at a single time instance. It is meteorology-driven, and hence it is highly spatially dependent and asymmetric. Wind induces directional sensitivity and, consequently, the function $b(\mathbf{s}, \cdot)$ frequently exhibits behaviour that is close to a step change at $\mathbf{s}$. As we show through a simple example in Appendix B, this can result in interesting correlation structures and higher-order properties that are not possible to obtain using Gaussian, symmetric, multivariate spatial models (e.g., as found in Gneiting et al., 2010).

3. Hierarchical trans-Gaussian bivariate models for atmospheric trace-gas inversion

In this section we discuss the construction of our framework for atmospheric trace-gas inversion. First, in Section 3.1 we develop the hierarchical model that uses the special class of non-Gaussian bivariate models, obtained by modelling $Y_1(\cdot)$ as a trans-Gaussian process with the Box-Cox spatial process as a leading case. In Section 3.2 we outline how to assimilate available flux inventories within the framework. In Section 3.3, we provide a succinct summary of the hierarchical model.

3.1. The hierarchical model

In Section 2.2 we considered the relationship of mole fraction to the flux field at a notional time instance. In practice, very little information on the flux field can be extracted from the mole-fraction field at a single time instance and, hence, considering its temporal evolution is important. In the following subsections we consider a spatio-temporal variant of the non-Gaussian model and place it within a Bayesian hierarchical framework appropriate for describing the problem of atmospheric trace-gas inversion.

3.1.1. The mole-fraction observations

Although the flux field Y_1 varies spatio-temporally, we treat it as a time-averaged spatial process. There are two reasons for this: First, the data available is not informative enough to infer flux evolution at such fine time scales and, second, we are primarily interested in the

average flux over a larger time window. In contrast, the mole-fraction field, which is directly observed with relatively low measurement error, is only informative of the flux field once the meteorology (that varies rapidly in space and time) and any chemistry is catered for. To account for time dependence, from now on we add the subscript t to both the mole-fraction field and the interaction function; these become $Y_{2,t}(\cdot)$ and $b_t(\cdot, \cdot)$, respectively.

The flux field is not directly observed; rather, the mole-fraction field is observed at different time points, $t \in \mathcal{T} \equiv \{1, 2, \dots, T\}$. Denote the observations at time point t as $\mathbf{Z}_{2,t} \equiv (Z_{2,t}(\mathbf{s}) : \mathbf{s} \in D_2^O)'$, where the finite set $D_2^O \subset D$. For clarity of exposition we omit the subscript t on D_2^O , however our model can readily accommodate time-dependent observation locations. It is reasonable to assume that the observations are conditionally independent and Gaussian, given the true mole-fraction field at the corresponding locations at time t . Denote these true mole fractions as $\mathbf{Y}_{2,t}^Z \equiv (Y_{2,t}(\mathbf{s}) : \mathbf{s} \in D_2^O)'$; then the observation model is:

$$(\mathbf{Z}_{2,t} \mid Y_{2,t}(\cdot)) \sim \mathcal{N}(\mathbf{Y}_{2,t}^Z, \mathbf{V}_t); \quad t \in \mathcal{T}, \quad (12)$$

where the covariance matrices $\{\mathbf{V}_t : t \in \mathcal{T}\}$ are diagonal and, in our case, assumed known.

3.1.2. The mole-fraction field

The mole-fraction field at time t , $Y_{2,t}(\cdot)$, is, to leading order, a linearly transformed version of the flux field $Y_1(\cdot)$, where the linear operator depends on the interaction function. In general, a model for the mole-fraction field follows through conditioning on the flux field:

$$(\{Y_{2,t}(\cdot) : t \in \mathcal{T}\} \mid Y_1(\cdot)) \sim \text{Dist}\left(\left\{\int_D b_t(\cdot, \mathbf{u})Y_1(\mathbf{u})d\mathbf{u} : t \in \mathcal{T}\right\}; \boldsymbol{\Theta}\right),$$

where $\text{Dist}(\{\mu_t(\cdot) : t \in \mathcal{T}\}; \boldsymbol{\Theta})$ is a possibly non-Gaussian, spatio-temporal (discrete-time) process with mean function $\{\mu_t(\cdot) : t \in \mathcal{T}\}$ and higher-order cumulant functions (Section 2.2) parameterised by $\boldsymbol{\Theta}$. This conditional distribution is, in reality, extremely complicated, but to date there has been very little effort to adequately characterise it. A natural way forward is to assume (4), adapted for the presence of the temporal index t :

$$Y_{2,t}(\mathbf{s}) = \int_D b_t(\mathbf{s}, \mathbf{u})Y_1(\mathbf{u})d\mathbf{u} + \zeta_t(\mathbf{s}); \quad \mathbf{s} \in D, t \in \mathcal{T}, \quad (13)$$

where $\zeta_t(\cdot)$, for $t \in \mathcal{T}$, has zero expectation. The term $\zeta_t(\cdot)$ in (13) is required to capture the spatio-temporal variability in $Y_{2,t}(\cdot)$ that is not explained through the integral operator.

Such discrepancy could be due to, for example, incorrect specification of $b_t(\cdot, \cdot)$ in (13) or its approximate numerical implementation.

Since $\{Y_{2,t}(\cdot) : t \in \mathcal{T}\}$, is a non-Gaussian spatio-temporal field, modelling and computation are critical considerations. Now, without compromising the non-Gaussian nature of the flux field $Y_1(\cdot)$ (and hence of $Y_{2,t}(\cdot)$, for $t \in \mathcal{T}$), we make the assumption that the additive discrepancy $\zeta_t(\cdot)$, for $t \in \mathcal{T}$, is Gaussian. For this model, the n -th order spatial cumulant functions associated with $Y_{2,t}(\cdot)$, for $t \in \mathcal{T}$ and $n \geq 3$ are

$$\kappa_{Y_{2,t} \dots Y_{2,t}}^n(\mathbf{s}_1, \dots, \mathbf{s}_n) = \int_D \dots \int_D b_t(\mathbf{s}_1, \mathbf{u}_1) \dots b_t(\mathbf{s}_n, \mathbf{u}_n) \kappa_{Y_1 \dots Y_1}^n(\mathbf{u}_1, \dots, \mathbf{u}_n) d\mathbf{u}_1 \dots d\mathbf{u}_n;$$

$$\mathbf{s}_1, \dots, \mathbf{s}_n \in D,$$

since all cumulants associated with a Gaussian process are zero for orders 3 and higher.

In practice, each $b_t(\cdot, \cdot)$ is evaluated at a finite set of pairwise discrete locations to yield the matrix $\tilde{\mathbf{B}}_t \equiv (b_t(\mathbf{s}, \mathbf{u}) : \mathbf{s} \in D_2^L, \mathbf{u} \in D_1^L)$, where D_1^L and D_2^L are finite, possibly different, subsets of D on which we wish to model (and predict) $Y_1(\cdot)$ and $Y_{2,t}(\cdot)$, respectively. Define $\mathbf{Y}_1 \equiv (Y_1(\mathbf{s}) : \mathbf{s} \in D_1^L)'$, $\mathbf{Y}_{2,t} \equiv (Y_{2,t}(\mathbf{s}) : \mathbf{s} \in D_2^L)'$, and $\boldsymbol{\zeta}_t \equiv (\zeta_t(\mathbf{s}) : \mathbf{s} \in D_2^L)'$. Then (13) can be approximated by, $\mathbf{Y}_{2,t} \approx \mathbf{B}_t \mathbf{Y}_1 + \boldsymbol{\zeta}_t$, where $\mathbf{B}_t \equiv \tilde{\mathbf{B}}_t \boldsymbol{\Delta}$ and $\boldsymbol{\Delta}$ is a diagonal matrix containing integration weights. Provided $D_2^O \subset D_2^L$, we can obtain $\mathbf{Y}_{2,t}^Z$ from $\mathbf{Y}_{2,t}$ through a known incidence matrix $\mathbf{C}_t$, so that $\mathbf{Y}_{2,t}^Z = \mathbf{C}_t \mathbf{Y}_{2,t}$ for each $t \in \mathcal{T}$. Henceforth, we assume that $D_2^O \subset D_2^L$.

In the absence of prior insight on the structure of $\zeta_t(\cdot)$, we further assume that its space-time covariance function is separable. Specifically, we assume equally spaced time points, and we let $\boldsymbol{\Sigma}_{\zeta; \tau_2, a, d} \equiv (\text{cov}(\zeta_t(\mathbf{s}_1), \zeta_{t'}(\mathbf{s}_2) \mid a, d) : t, t' \in \mathcal{T}, \mathbf{s}_1, \mathbf{s}_2 \in D_2^L) = \frac{1}{\tau_2} (\mathbf{Q}_{\zeta; a}^t)^{-1} \otimes \mathbf{R}_{\zeta; d}^s$, where τ_2 is a precision parameter; $\mathbf{Q}_{\zeta; a}^t$ is the precision matrix corresponding to a first-order auto-regressive discrete-time process (Rue and Held, 2005, Chapter 1) with auto-regressive parameter a ; and $\mathbf{R}_{\zeta; d}^s \equiv (\rho_\zeta(\|\mathbf{s}_1 - \mathbf{s}_2\| \mid d) : \mathbf{s}_1, \mathbf{s}_2 \in D_2^L)$, where $\rho_\zeta(u \mid d) \equiv \exp(-u/d)$, for $u \geq 0$, is the exponential correlation function. The motivation for this choice is primarily computational: In atmospheric trace-gas inversion with ground station data, $|\mathcal{T}|$ is large (several hundreds) while $|D_2^L|$ is typically small and frequently equal to $|D_2^O|$, which is generally less than 10. Hence, $\boldsymbol{\Sigma}_{\zeta; \tau_2, a, d}^{-1}$ will be sparse, although $\mathbf{R}_{\zeta; d}^s$ (and $(\mathbf{R}_{\zeta; d}^s)^{-1}$) is dense. Other approaches, such as covariance tapering (Kaufman et al., 2008) or dimensionality reduction (Cressie and Johannesson, 2008), might instead be employed to facilitate the computation.

Note that $\zeta_t(\cdot)$ induces an auto-regressive structure on $Y_{2,t}(\cdot)$. To see this, one can use the alternative representation (Storvik et al., 2002), $\zeta_t(\cdot) = a\zeta_{t-1}(\cdot) + e_t(\cdot)$, to re-write (13) as

$$Y_{2,t}(\mathbf{s}) = \int_D (b_t(\mathbf{s}, \mathbf{u}) - ab_{t-1}(\mathbf{s}, \mathbf{u}))Y_1(\mathbf{u})d\mathbf{u} + aY_{2,t-1}(\mathbf{s}) + e_t(\mathbf{s}); \quad \mathbf{s} \in D, \quad t \in \mathcal{T}, \quad (14)$$

where $e_t(\cdot)$ is temporally uncorrelated and has spatial correlation function $\rho_\zeta(u \mid d)$. In what follows, we collect the spatial and temporal correlation parameters into the parameter vector $\boldsymbol{\theta}_2 \equiv (a, d)'$.

3.1.3. The flux-field model

The focus now turns to flexible non-Gaussian modelling of $Y_1(\cdot)$, for which we employ a trans-Gaussian-process approach. For several trace gases, such as methane (considered in Section 5), it is reasonable to assume that the flux is only positive. To this end, we apply a monotonic function that has as domain the positive real line, such that $\tilde{Y}_1(\cdot) \equiv g_\lambda(Y_1(\cdot))$ is approximately a Gaussian process with mean $\tilde{\mu}_1(\cdot) \equiv \mathbf{x}(\cdot)'\boldsymbol{\beta}$, and covariance function $\tilde{C}_{Y_1Y_1}(\cdot, \cdot \mid \boldsymbol{\theta}_1)$, where $\mathbf{x}(\cdot)$ is a vector of covariate functions, $\boldsymbol{\beta} \in \mathbb{R}^p$, and λ is a parameter of the transformation $g_\lambda(\cdot)$.

A leading case of a trans-Gaussian process is the Box–Cox spatial process. Specifically, in line with De Oliveira et al. (1997), we consider the *power-normal* family of spatial processes for the flux field, obtained through application of the Box–Cox transformation

$$g_\lambda(y) = \begin{cases} \frac{y^\lambda - 1}{\lambda}; & \lambda \neq 0 \\ \ln y; & \lambda = 0 \end{cases}.$$

Define the correlation function $\tilde{R}_{Y_1Y_1}(\cdot, \cdot \mid \boldsymbol{\theta}_1) \equiv \tau_1 \tilde{C}_{Y_1Y_1}(\cdot, \cdot \mid \boldsymbol{\theta}_1)$, where τ_1 is a precision parameter. Further, define $\tilde{\mathbf{R}}_{Y_1Y_1;\boldsymbol{\theta}_1} \equiv (\tilde{R}_{Y_1Y_1}(\mathbf{u}_1, \mathbf{u}_2 \mid \boldsymbol{\theta}_1) : \mathbf{u}_1, \mathbf{u}_2 \in D_1^L)$, $\mathbf{X} \equiv (\mathbf{x}(\mathbf{u}_1) : \mathbf{u}_1 \in D_1^L)'$, and $\mathbf{Y}_1 \equiv (Y_1(\mathbf{u}_1) : \mathbf{u}_1 \in D_1^L)'$. In the case of the Box–Cox family of transformations, the positive-only domain of $g_\lambda(\cdot)$ also restricts the domain of $g_\lambda^{-1}(\cdot)$. In particular, since $g_\lambda^{-1}(\tilde{y}) = (\lambda\tilde{y} + 1)^{1/\lambda}$, then $\tilde{y} > -1/\lambda$ if $\lambda > 0$ and $\tilde{y} < -1/\lambda$ if $\lambda < 0$. Therefore, the process $\tilde{Y}_1(\cdot)$ that we specify in this case is a *truncated* Gaussian process. Then the probability density function of $\tilde{\mathbf{Y}}_1 \equiv g_\lambda(\mathbf{Y}_1)$ is

$$p(\tilde{\mathbf{Y}}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) = \frac{\left(\frac{\tau_1}{2\pi}\right)^{\frac{|D_1^L|}{2}} |\tilde{\mathbf{R}}_{Y_1Y_1;\boldsymbol{\theta}_1}|^{-\frac{1}{2}}}{K_{1;\boldsymbol{\beta},\tau_1,\boldsymbol{\theta}_1,\lambda}} \exp\left(-\frac{\tau_1}{2}(\tilde{\mathbf{Y}}_1 - \mathbf{X}\boldsymbol{\beta})'\tilde{\mathbf{R}}_{Y_1Y_1;\boldsymbol{\theta}_1}^{-1}(\tilde{\mathbf{Y}}_1 - \mathbf{X}\boldsymbol{\beta})\right) T_\lambda(\tilde{\mathbf{Y}}_1),$$

where the normalising factor $K_{1;\beta,\tau_1,\theta_1,\lambda} \leq 1$ is a function of all unknown parameters; it is understood that $g_\lambda(\mathbf{Y}_1)$ denotes a vector obtained by element-wise application of the Box–Cox transformation; and

$$T_\lambda(\tilde{\mathbf{Y}}_1) \equiv \begin{cases} \mathbb{1}(\tilde{\mathbf{Y}}_1 > -1/\lambda); & \lambda > 0 \\ \mathbb{1}(\tilde{\mathbf{Y}}_1 < -1/\lambda); & \lambda < 0 \end{cases}. \quad (15)$$

In (15), $\mathbb{1}(\cdot)$ is the indicator function as applied to vectors, that is, it returns a one if all the elements of the vector satisfy the condition and a zero otherwise.

From $p(\tilde{\mathbf{Y}}_1 \mid \beta, \tau_1, \theta_1, \lambda)$, we obtain $p(\mathbf{Y}_1 \mid \beta, \tau_1, \theta_1, \lambda)$ through the transformation $g_\lambda(\cdot)$ to yield the flux-process layer in our hierarchical model,

$$p(\mathbf{Y}_1 \mid \beta, \tau_1, \theta_1, \lambda) = \frac{\left(\frac{\tau_1}{2\pi}\right)^{\frac{|D_1^L|}{2}} |\tilde{\mathbf{R}}_{Y_1 Y_1; \theta_1}|^{-\frac{1}{2}}}{K_{1;\beta,\tau_1,\theta_1,\lambda}} \exp\left(-\frac{\tau_1}{2} (g_\lambda(\mathbf{Y}_1) - \mathbf{X}\beta)' \tilde{\mathbf{R}}_{Y_1 Y_1; \theta_1}^{-1} (g_\lambda(\mathbf{Y}_1) - \mathbf{X}\beta)\right) \\ \times J_\lambda \mathbb{1}(\mathbf{Y}_1 > 0), \quad (16)$$

where the Jacobian $J_\lambda \equiv \prod_{i=1}^{|D_1^L|} |Y_{1,i}^{\lambda-1}|$. Note that the support of $p(\mathbf{Y}_1 \mid \beta, \tau_1, \theta_1, \lambda)$ is positive, as required. We term (16) a Box–Cox spatial model and $Y_1(\cdot)$ a Box–Cox spatial process.

3.1.4. The parameter model

Finally, we need to specify prior distributions over the parameters. This should be done with care since, first, the range of the transformed observations changes with the transformation parameter and, second, posterior-distribution impropriety can easily follow from inadequate specification of non-informative prior distributions (Berger et al., 2001). In their original paper, Box and Cox (1964) tackled the first issue by considering the conditional prior distribution, $p(\beta, \tau_1, \theta_1 \mid \lambda) \propto h(\lambda)p(\theta_1)/\tau_1$, and deducing what $h(\lambda)$ should be for the prior probability measure to be approximately independent of λ . Their choice, $h(\lambda) = J_\lambda^{p/|D_1^L|}$, is not appropriate in our context where $\mathbf{Y}_1$ is not observed. We therefore adopt the alternative of Pericchi (1981), who noted that if the conditional prior distribution, $p(\beta, \tau_1, \theta_1 \mid \lambda) \propto h(\lambda)p(\theta_1)\tau_1^{\frac{p}{2}-1}$, then $h(\lambda)$ must be constant.

The prior distribution we choose for the flux-process layer is therefore,

$$p(\beta, \tau_1, \theta_1, \lambda) \propto p(\theta_1)p(\lambda)\tau_1^{\frac{p}{2}-1},$$

where $p(\theta_1)$ and $p(\lambda)$ are marginal prior distributions and, in order to ensure propriety of

the posterior distribution, the parameter spaces of $\boldsymbol{\theta}_1$ and λ are bounded (De Oliveira et al., 1997). For the remaining parameters appearing in the mole-fraction-process layer (τ_2, a and d), we use independent, proper, prior distributions, $p(\tau_2)$, $p(a)$, and $p(d)$, respectively.

For convenience, we specify bounded uniform distributions for $\boldsymbol{\theta}_1, \lambda, \ln \tau_2^{-1}, a$, and $\ln d$; see Section 3.3 for details.

3.2. Incorporation of the inventory

In atmospheric trace-gas inversion, one is usually supplied with a flux ‘inventory,’ which can be used to reduce uncertainty on inferences on the flux field, $Y_1(\cdot)$. We denote this inventory, available for spatial locations in D_1^L , as $\mathbf{W}_1$. Incorporation of $\mathbf{W}_1$ in the model is akin to the problem of *data assimilation*, where observational data is fused with computer-model output in order to obtain an estimate that is optimal in some sense (e.g., Wikle and Berliner, 2007). Standard data assimilation typically assumes that $E(\mathbf{Y}_1) = \mathbf{W}_1$. However, there are two concerns with data assimilation in this context. First, no guarantees are provided on the quality of $\mathbf{W}_1$, which has often been shown to be inaccurate (Lunt et al., 2015; Miller et al., 2013, amongst others). Second, it is unclear how to construct the covariance matrix $\text{var}(\mathbf{Y}_1)$; typically a diagonal structure is assumed with elements based on expert judgement (e.g., Ganesan et al., 2014), however this is likely to be overly simplistic. Another approach to assimilation is that of Fuentes and Raftery (2005), in which $\mathbf{W}_1$ is treated as a (possibly both additively and multiplicatively biased) observation of $\mathbf{Y}_1$. However, in their approach one also has a $\mathbf{Z}_1$ that provides a direct (unbiased) observation of $\mathbf{Y}_1$. In our context, where direct observations of $\mathbf{Y}_1$ are not available, data-driven inferences on the bias terms would be ill-constrained.

We adopt a different approach to the two outlined above, by assuming that $\mathbf{W}_1$ is only informative on the spatial properties of the process. This choice is motivated by the fact that inventories are typically constructed bottom-up from spatially referenced datasets on (say) agricultural productivity and transport infrastructure. The emission factors used in this construction are typically not well constrained, however the main spatial features are still adequately represented. We implement the approach by assuming that $\mathbf{W}_1$ is a noiseless observation of an independent realisation of $(\mathbf{Y}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)$. Under this modelling assumption,

$$p(\mathbf{Y}_1 \mid \mathbf{W}_1, \mathbf{Z}_2, \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) \equiv p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda),$$

due to the conditional independence of $\mathbf{Y}_1$ and $\mathbf{W}_1$ when conditioned on $(\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)$. Hence,

this modelling assumption allows us to use Bayesian learning to glean information on $(\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)$ from $\mathbf{W}_1$. At the same time, inferences on $\mathbf{Y}_1$, when conditioned on $\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1$, and λ , are purely a function of $\mathbf{Z}_2$ and are hence data-driven. We propose that the term *data-feature assimilation* be used to describe this methodology, to distinguish it from standard data assimilation.

In this application, we envision that the transformation parameter λ , and consequently all third-order and higher-order joint cumulants of $\mathbf{Y}_1$, will be highly sensitive to $\mathbf{W}_1$. Hence, this procedure is analogous to the use of a *training image* to estimate third-order and higher-order cumulants prior to carrying out spatial prediction (Dimitrakopoulos et al., 2010).

As a result of the modelled conditional independence between $\mathbf{Y}_1$ and $\mathbf{W}_1$, we augment the conditional distribution (16) as follows:

$$\begin{aligned} p(\mathbf{Y}_1, \mathbf{W}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) &= p(\mathbf{Y}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) p(\mathbf{W}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) \\ &= \frac{\left(\frac{\tau_1}{2\pi}\right)^{|D_1^L|} |\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}|^{-1}}{K_{2; \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda}} \exp \left(-\frac{\tau_1}{2} \sum_{j=1}^2 (\mathbf{G}_{\lambda, j} - \mathbf{X}\boldsymbol{\beta})' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} (\mathbf{G}_{\lambda, j} - \mathbf{X}\boldsymbol{\beta}) \right) \\ &\quad \times \left(\prod_{j=1}^2 J_{\lambda, j} \right) \mathbb{1}(\mathbf{Y}_1 > 0) \mathbb{1}(\mathbf{W}_1 > 0), \end{aligned} \quad (17)$$

where $K_{2; \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda}$ is a normalising factor, $\mathbf{G}_{\lambda, 1} \equiv g_\lambda(\mathbf{Y}_1)$, $\mathbf{G}_{\lambda, 2} \equiv g_\lambda(\mathbf{W}_1)$, $J_{\lambda, 1} \equiv J_\lambda$ in (16), and $J_{\lambda, 2} \equiv \prod_{i=1}^{|D_1^L|} |W_{1, i}^{\lambda-1}|$. For notational convenience, we frequently consider the joint vector $\underline{\mathbf{Y}}_1 \equiv (\mathbf{Y}_1', \mathbf{W}_1')'$; then the conditional distribution (17) becomes

$$\begin{aligned} p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) &= \frac{\left(\frac{\tau_1}{2\pi}\right)^{|D_1^L|} |\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}|^{-1/2}}{K_{2; \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda}} \exp \left(-\frac{\tau_1}{2} (\underline{\mathbf{G}}_\lambda - \underline{\mathbf{X}}\boldsymbol{\beta})' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} (\underline{\mathbf{G}}_\lambda - \underline{\mathbf{X}}\boldsymbol{\beta}) \right) \\ &\quad \times \underline{J}_\lambda \mathbb{1}(\underline{\mathbf{Y}}_1 > 0). \end{aligned} \quad (18)$$

In (18), $\underline{\mathbf{G}}_\lambda \equiv (\mathbf{G}'_{\lambda, 1}, \mathbf{G}'_{\lambda, 2})'$, $\underline{\mathbf{X}} \equiv (\mathbf{X}', \mathbf{X})'$, $\underline{J}_\lambda \equiv \prod_{i=1}^{2|D_1^L|} |Y_{-1, i}^{\lambda-1}|$, and $\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1} \equiv \text{bdiag}(\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}, \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1})$, where $\text{bdiag}(\cdot)$ creates a block-diagonal matrix from its arguments.

3.3. Summary of the hierarchical model

The graphical model that we construct is given in Fig. 1, where for simplicity we omit the discrepancies $\{\zeta_t : t \in \mathcal{T}\}$ and instead show the auto-regressive structure on $\{\mathbf{Y}_{2, t} : t \in \mathcal{T}\}$ that can be obtained from (14). In this graphical model we illustrate the time evolution of the mole-fraction field, however it is convenient, both for notational purposes and for inference in Section 4, to construct vectors and matrices that are blocked with time.

Fig. 1: Graphical representation of the hierarchical model, where $\{\mathbf{Z}_{2,t} : t \in \mathcal{T}\}$ are the mole-fraction observations, $\{\mathbf{Y}_{2,t} : t \in \mathcal{T}\}$ are the mole fractions, $\mathbf{Y}_1$ is the flux field, $\mathbf{W}_1$ is the flux-field inventory, $\boldsymbol{\theta}_2$ are mole-fraction-discrepancy spatio-temporal correlation parameters, τ_2 is the mole-fraction-discrepancy precision parameter, τ_1 is the flux precision parameter, $\boldsymbol{\beta}$ are the flux-field regression coefficients, $\boldsymbol{\theta}_1$ are the flux-field correlation parameters, and λ is the Box–Cox transformation parameter.

Define $\mathbf{Z}_2 \equiv (\mathbf{Z}'_{2,t} : t \in \mathcal{T})'$, $\mathbf{C} \equiv \text{bdiag}(\{\mathbf{C}_t : t \in \mathcal{T}\})$, $\mathbf{V} \equiv \text{bdiag}(\{\mathbf{V}_t : t \in \mathcal{T}\})$, $\mathbf{Y}_2 \equiv (\mathbf{Y}'_{2,t} : t \in \mathcal{T})'$, and $\mathbf{B} \equiv (\mathbf{B}'_t : t \in \mathcal{T})'$. Then the hierarchical model is:

Observation model (mole fraction): $(\mathbf{Z}_2 \mid \mathbf{Y}_2) \sim \mathcal{N}(\mathbf{C}\mathbf{Y}_2, \mathbf{V})$,

Process model 2 (mole fraction): $(\mathbf{Y}_2 \mid \mathbf{Y}_1, \boldsymbol{\theta}_2, \tau_2) \sim \mathcal{N}(\mathbf{B}\mathbf{Y}_1, \boldsymbol{\Sigma}_{\zeta; \tau_2, a, d})$,

Process model 1 (flux): $(\mathbf{Y}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) \sim \mathcal{BC}\left(\mathbf{X}\boldsymbol{\beta}, \frac{1}{\tau_1} \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}, \lambda\right)$,

Inventory model (flux): $(\mathbf{W}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) \sim \mathcal{BC}\left(\mathbf{X}\boldsymbol{\beta}, \frac{1}{\tau_1} \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}, \lambda\right)$,

independent of Process model 1,

Parameter model 2 (mole-fraction): $\ln(\tau_2^{-1}) \sim \mathcal{U}(\gamma_{\tau_2^{-1}}^l, \gamma_{\tau_2^{-1}}^u)$, $a \sim \mathcal{U}(\gamma_a^l, \gamma_a^u)$,

$\ln d \sim \mathcal{U}(\gamma_d^l, \gamma_d^u)$,

Parameter model 1 (flux): $p(\tau_1) \propto \tau_1^{\frac{p}{2}-1}$, $p(\boldsymbol{\beta}) \propto 1$, $\lambda \sim \mathcal{U}(\gamma_\lambda^l, \gamma_\lambda^u)$,

$\theta_{1i} \sim \mathcal{U}(\gamma_{\theta_{1i}}^l, \gamma_{\theta_{1i}}^u); \quad i = 1, 2, \dots, n_{\boldsymbol{\theta}_1}$,

where $\mathcal{BC}$ is an abbreviation for ‘Box–Cox’: If $\mathbf{D} \sim \mathcal{BC}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \lambda)$, then the probability density function of $g_\lambda(\mathbf{D})$ is proportional to the multivariate normal density function with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$, multiplied by $T_\lambda(\mathbf{D})$ given by (15). The distribution $\mathcal{U}(\gamma_\eta^l, \gamma_\eta^u)$ is the bounded uniform distribution over η with limits γ_η^l and γ_η^u , and $n_{\boldsymbol{\theta}_1}$ is the number of elements in the vector $\boldsymbol{\theta}_1$. For trace-gas inversion of methane, all hyper-parameters appearing in the parameter models can be elicited from physical considerations (e.g., Ganesan et al., 2015).

4. Inference

Since the normalising constant of the truncated multivariate normal density function is intractable, inference using the model in Section 3.3 is problematic. Such models, where the normalising constants of both the likelihood and the posterior distribution are intractable, have been called *doubly intractable* (Murray et al., 2006). The approach of De Oliveira et al. (1997) is to implicitly assume that the truncated volume is approximately 0. Then, in (16) and (18), the factor $K_{i;\beta,\tau_1,\theta_1,\lambda} \approx 1$, for $i = 1, 2$, independently of $\{\beta, \tau_1, \theta_1, \lambda\}$. From here on we assume that $K_{i;\beta,\tau_1,\theta_1,\lambda} = 1$, for $i = 1, 2$.

The Box–Cox transformation can only induce normality in $\tilde{\mathbf{Y}}_1$ if $\lambda = 0$ (Poirier, 1978). In the univariate case, truncation is kept to a minimum when the mean of $\tilde{Y}_1$ is large and the coefficient of variation of $\tilde{Y}_1$ is small (Freeman and Modarres, 2006). Draper and Cox (1969) conclude that even though the Box–Cox transformation does not achieve exact Gaussian multivariate normality, it has a beneficial effect on third-order cumulants (i.e., skewness); we believe that this is the main motivation for its continued use in applications, such as in this article.

Once we assert that the truncation effect in (18) is small, then the inferential approach can be based on the standard Gibbs sampler. We can then also analytically marginalise out the variables τ_1 , β , and $\mathbf{Y}_2$. There are several reasons why marginalisation is useful, especially when using improper prior distributions (see Berger et al., 1999, for a discussion), however our prime motivation here is to improve the convergence properties of the sampler (Van Dyk and Park, 2008).

The required full conditional distributions are

$$\begin{aligned} p(\tau_2, \theta_2 \mid \mathbf{Z}_2, \mathbf{Y}_1), \\ p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \mathbf{W}_1, \theta_2, \tau_2, \theta_1, \lambda), \\ p(\theta_1, \lambda \mid \mathbf{Y}_1), \end{aligned}$$

where in the Gibbs sampler we block sample $\{\tau_2, \theta_2\}$ and $\{\theta_1, \lambda\}$. Here we outline some features of these full conditional distributions; further details are given in Appendix C.

The first full conditional distribution, $p(\tau_2, \theta_2 \mid \mathbf{Z}_2, \mathbf{Y}_1) \propto p(\mathbf{Z}_2 \mid \tau_2, \theta_2, \mathbf{Y}_1)p(\tau_2)p(\theta_2)$, which

requires the marginalisation,

$$p(\mathbf{Z}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1) = \int p(\mathbf{Z}_2 \mid \mathbf{Y}_2) p(\mathbf{Y}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1) d\mathbf{Y}_2. \quad (19)$$

This is straightforward since both $p(\mathbf{Z}_2 \mid \mathbf{Y}_2)$ and $p(\mathbf{Y}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1)$ are multivariate Gaussian distributions. The resulting conditional distribution of $(\tau_2, \boldsymbol{\theta}_2)$ requires the Cholesky decomposition of $\boldsymbol{\Sigma}_{\zeta; \tau_2, a, d}^{-1}$ or other sparse variants of it that can easily be computed due to the imposed auto-regressive structure on $\zeta_t(\cdot), t \in \mathcal{T}$. The full conditional density function is given in (C.1) and we sampled from it using a slice sampler (Neal, 2003).

The second full conditional distribution,

$$p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \mathbf{W}_1, \boldsymbol{\theta}_2, \tau_2, \boldsymbol{\theta}_1, \lambda) \propto p(\mathbf{Z}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1) p(\mathbf{Y}_1 \mid \boldsymbol{\theta}_1, \lambda), \quad (20)$$

also requires marginalisation:

$$p(\mathbf{Y}_1 \mid \boldsymbol{\theta}_1, \lambda) = \int p(\mathbf{Y}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda) p(\boldsymbol{\beta}, \tau_1) d\boldsymbol{\beta} d\tau_1,$$

where the distribution of $(\mathbf{Y}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)$ is given by (18). This is the integrated likelihood function of $\{\lambda, \boldsymbol{\theta}_1\}$ (Berger et al., 2001), which partially depends on $\mathbf{Y}_1$ through the sum of squared residuals,

$$S_{\boldsymbol{\theta}_1, \lambda}^2 \equiv (\mathbf{G}_\lambda - \mathbf{X}\hat{\boldsymbol{\beta}}_{\boldsymbol{\theta}_1, \lambda})' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} (\mathbf{G}_\lambda - \mathbf{X}\hat{\boldsymbol{\beta}}_{\boldsymbol{\theta}_1, \lambda}). \quad (21)$$

In (21), $\hat{\boldsymbol{\beta}}_{\boldsymbol{\theta}_1, \lambda}$ is the generalised least squares estimate of $\boldsymbol{\beta}$ conditioned on $\boldsymbol{\theta}_1$ and λ ; that is,

$$\hat{\boldsymbol{\beta}}_{\boldsymbol{\theta}_1, \lambda} = (\mathbf{X}' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} \mathbf{X})^{-1} \mathbf{X}' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} \mathbf{G}_\lambda. \quad (22)$$

The sum of squared residuals, $S_{\boldsymbol{\theta}_1, \lambda}^2$, needs to be non-zero everywhere for (20) to be proper (see (C.6) in Appendix C). A sufficient condition for this is that there does not exist a $\boldsymbol{\beta}$ such that $\mathbf{W}_1 - \mathbf{X}\boldsymbol{\beta} = \mathbf{0}$. This condition is difficult to violate in practice (unless the inventory is used as a covariate) and, thus, propriety of the conditional distribution is practically guaranteed.

Since the gradient of this full conditional distribution can be easily found (see Appendix C), we implemented a Hamiltonian Monte Carlo (HMC) sampler for sampling from this conditional distribution (Neal, 2011). Previous work on this application (Zammit-Mangion et al., 2015) indicated an important benefit of HMC sampling. The gradient information of the conditional

density function of $\mathbf{Y}_1$ improved the mixing properties of the MCMC chains when compared to other samplers, such as slice samplers or Metropolis samplers. Since we analyse several models in this work, we auto-adapt the HMC sampler’s step size in the burn-in period, in order to achieve suitable acceptance ratios.

The third full conditional distribution we require is $p(\boldsymbol{\theta}_1, \lambda \mid \underline{\mathbf{Y}}_1)$, which recall is guaranteed to be proper under the chosen prior distribution. This conditional distribution can be obtained by following the approach in De Oliveira et al. (1997), after modifying it slightly to include $\mathbf{W}_1$ in addition to $\mathbf{Y}_1$. In summary, the distribution of $(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda)$ is multivariate Normal-Gamma and thus has a known normalising constant in terms of all the other parameters. Hence, the required conditional distribution can be found, up to a constant of proportionality, by writing it out as

$$p(\boldsymbol{\theta}_1, \lambda \mid \underline{\mathbf{Y}}_1) = \frac{p(\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda \mid \underline{\mathbf{Y}}_1)}{p(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda)}.$$

We sampled from this conditional distribution using a slice sampler (Neal, 2003).

5. Trace-gas inversion in the UK and Ireland

In this section we consider the problem of methane flux inversion in the UK and Ireland using the model of Section 3.3 and the Gibbs sampler of Section 4. In Sections 5.1 and 5.2 we describe the overall simulation setup and experimental conditions, while the results are discussed in Section 5.3.

5.1. Observation system simulation experiment (OSSE) setup

The main methane emissions inventory we use is based on the UK National Atmospheric and Emissions Inventory (NAEI, UK Department for Environment Food and Rural Affairs (DEFRA), 2014) and the Emissions Database for Global Atmospheric Research 4.2 (EDGAR, Joint Research Centre of the European Commission (JRC) and the Netherlands Environmental Assessment Agency (PBL), 2011); see Ganesan et al. (2015) for details.

Since methane is a long-lived gas, we expect that the relationship between the flux and the mole-fraction is linear (Section 2.1). The matrices used for the linear mapping, $\{\mathbf{B}_t : t \in \mathcal{T}\}$, were constructed from $b_t(\mathbf{s}, \mathbf{u})$ for $t \in \mathcal{T}$, $\mathbf{s} \in D_2^O$, and $\mathbf{u} \in D_1^L$, using the UK Met Office’s Lagrangian Particle Dispersion Model (LPDM), NAME. The domain D_1^L was defined as a lon-lat grid, with each grid cell of size $0.7^\circ \times 0.5^\circ$; here, $|D_1^L| = 122$. The domain D_2^L was chosen to be D_2^O , the set of four locations of ground stations recording methane mole fraction in the UK

and Ireland; hence $|D_2^L| = 4$. These four stations are located in Mace Head (Ireland), Ridge Hill (England), Tacolneston (England), and Angus (Scotland). The temporal domain $\mathcal{T}$ contains every 2 h interval between January 01 2014 at 00:00 and April 01 2014 at 00:00; a total of 1080 time intervals. Since the Scottish Highlands experience much lower methane emissions than the rest of the UK, we constructed $\mathbf{X}$ with two columns; the first column contains ones in rows corresponding to grid cells above 56.4° latitude (corresponding to the latitudes of the Scottish Highlands) and zeros otherwise, and the second column contains ones in rows corresponding to grid cells below 56.4° and zeros otherwise. The domain of interest D , station locations D_2^O , and the sum over rows of each column of $\mathbf{B}_1$, are shown in Fig. 2, left panel.

In order to mimic a typical regional inversion study, we carried out what is known as an *observing system simulation experiment* (OSSE), where observations (including measurement error) are forward-simulated using the bivariate model together with a known flux field. Initially, we assumed that the true flux $\mathbf{Y}_1$ was equal to the UK and Ireland inventory $\mathbf{W}_1$ (which therefore was not simulated from a Box–Cox spatial model). We further let $\tau_2 = 0.01 \text{ ppb}^{-2}$, $a = 0.9$, and $d = 2.5^\circ$, in line with what is expected in practice (Zammit-Mangion et al., 2015). We next simulated the mole-fraction field using a discrete approximation to (13) with the matrices $\{\mathbf{B}_t : t \in \mathcal{T}\}$ supplied from NAME. Finally, the observations were simulated from (12). We assumed that all the matrices in $\{\mathbf{V}_t : t \in \mathcal{T}\}$ were equal to the identity matrix and that data was missing in the OSSE at the same times and locations that data was indeed missing from the stations whose locations we use. The missingness was used to construct the incidence matrices $\{\mathbf{C}_t : t \in \mathcal{T}\}$.

In reality, the detected methane mole fraction is the sum of the contributions from regional emissions, and a background level (around 1800 ppb) that is not accounted for by the LPDM since it is run for a temporal horizon of only 30 days. This background level varies both in space (e.g., with latitude) and time (e.g., seasonally) and can be included in our model through $E(\zeta_t(\cdot))$. In this OSSE, we are implicitly assuming that the background is known and that the discrepancy has been adequately corrected for. Therefore, $E(\zeta_t(\cdot)) = 0$, and henceforth we refer to $Y_{2,t}(\mathbf{s})$ as a background-enhanced value of methane mole fraction. Estimation of background mole-fractions is discussed in Ganesan et al. (2015). The simulated observations from the four stations are depicted in Fig. 2, right panel. Negative measurements are present in these time series, primarily due to the Gaussian discrepancy $\{\zeta_t(\cdot) : t \in \mathcal{T}\}$. In practice, negative residual mole fractions can arise from incorrect characterisation of the background process. Note that

Fig. 2: Left panel: The region of interest (UK and Ireland) and the station locations considered (black dots), that is, Mace Head (MHD), Ridge Hill (RGL), Tacolneston (TAC), and Angus (TTA). The quantity depicted is the vector $\bar{\mathbf{B}}_1 = \mathbf{1}'\mathbf{B}_1$ mapped out geographically; this plot superimposes the sensitivity of mole fraction to the flux at the first time point at the four stations to give a visual summary of $\mathbf{B}_1$. Right panel: The four simulated measurement time series at the four station locations of background-enhanced methane mole fraction. The first time point corresponds to January 01 2014 at 00:00.

the 30-day temporal horizon used in NAME to obtain the required interaction functions is unrelated to the window we use for flux inference (namely, 3 months).

We considered six models for the flux field, namely a Box–Cox spatial process (Model 1), a lognormal spatial process (Box–Cox parameter λ fixed to 0; Model 2), a truncated Gaussian process (λ fixed to 1 but with positivity imposed; Model 3), and the same three sequences of models but with no spatial correlation assumed (Models 4–6). Specifically, for Models 1–3, we let the transformed flux-field covariance function, $\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2 | \boldsymbol{\theta}_1)$, be given by

$$\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2 | \boldsymbol{\theta}_1) = \frac{1}{\tau_1} \exp\left(-\theta_{11} \|\mathbf{u}_1 - \mathbf{u}_2\|^{\theta_{12}}\right),$$

where $\boldsymbol{\theta}_1 \equiv (\theta_{11}, \theta_{12})'$, and $\theta_{11} > 0, 0 < \theta_{12} < 2$; for Models 4–6, we let $\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2 | \boldsymbol{\theta}_1) = \frac{1}{\tau_1} \mathbb{1}(\mathbf{u}_1 = \mathbf{u}_2)$. Using prediction performance measures, these models will allow us to weigh the importance of modelling non-Gaussianity against that of spatial modelling. Note that Models 1–6 are all stationary models; while this may be a reasonable assumption when modelling the flux at this resolution on a domain of this size, more complex models will be needed for continental (and upwards) domains. In the general case, one might ensure that the employed model is representative of the flux by fitting the flux model to the inventory and looking at simple diagnostics (e.g., through leave-one-out cross-validation).

Fig. 3: Inventories used in the study both for OSSE-data generation and data-feature assimilation experiments: The correct inventory of the UK and Ireland (left panel), an incorrect inventory taken from mainland Europe but with similar spatial properties to the UK (centre panel), and an incorrect inventory taken from Northern Australia with highly dissimilar spatial properties (right panel).

In order to assess the role the inventory plays in the inference, we re-ran the algorithm with Model 1 using the same observations $\mathbf{Z}_2$, but we replaced $\mathbf{W}_1$ with one of two ‘incorrect’ inventories extracted from EDGAR in different world regions; see Fig. 3. The first inventory, $\mathbf{W}_1^*$, was taken from mainland Europe and has second-order properties that are qualitatively similar to those of the UK and Ireland. The second inventory was taken from Northern Australia. Since methane emissions are much lower in this latter region, we shifted and scaled the emissions in this inventory so that they have the same mean and variance as $\mathbf{W}_1$; we denote this inventory as $\mathbf{W}_1^{**}$. We denote the models utilising these two inventories by Model 1* and Model 1**, respectively.

Finally, since it is known that the lognormal spatial process is an adequate model for emissions in the UK and Ireland (Zammit-Mangion et al., 2015), we re-ran the OSSE with Models 1–6 assuming that the true emissions field is given by the inventory from Northern Australia (Fig. 3, right panel), while for simplicity keeping the measurement locations and meteorology unchanged. Upon comparing the panels of Fig. 3, it is readily apparent how different the right-hand panel is, and that a lognormal process will not be suitable for modelling the flux field. In this case, the added flexibility of the trans-Gaussian (here, Box–Cox) process is important.

5.2. MCMC details and diagnostics

We used bounded uniform prior distributions that encompass physically plausible values on parameter transformations. For the mole-fraction parameters, we used ‘Parameter model 2’ in Section 3.3, and we set $\gamma_{\tau_2-1}^l = -2$, $\gamma_{\tau_2-1}^u = 20$, $\gamma_a^l = -1$, $\gamma_a^u = 1$, $\gamma_d^l = \ln(0.1)$, and $\gamma_d^u = \ln(5)$. For the flux-field parameters, we used ‘Parameter model 1’, and we set $\gamma_{\theta_{11}}^l = 0$, $\gamma_{\theta_{11}}^u = 2$, $\gamma_{\theta_{12}}^l =$

0, $\gamma_{\theta_{12}}^u = 2$, $\gamma_{\lambda}^l = -3$, and $\gamma_{\lambda}^u = 3$ as in De Oliveira et al. (1997).

For each model, we ran 10 parallel MCMC chains with 12000 samples each. We used the first 1000 samples for HMC adaptation, during which the step size used for discretisation of the Hamiltonian dynamics was updated at a decreasing rate of adaptation in order to achieve an acceptance ratio between 30% and 80%. These samples, together with a further 7000, were discarded to ensure adequate burn-in. We then took every 10-th sample from the resulting chains, to give a total of 4000 posterior samples for each model. Acceptance ratios were checked and found to be reasonable for all the traces while mixing was verified by visual inspection of the concatenated/thinned groups of chains and the auto-correlation plots. For the flux field, we considered 10 of the spatial locations, chosen at random, out of the possible 122 locations. Since λ was well constrained by the inventory, there was no noticeable additional computational burden in sampling from its conditional distribution. The computational bottleneck for each model was the sampling of the flux field using HMC. Each study required between 5 and 7 hours of computation time and 10 Gb (1 Gb per thread) of memory to run using R (R Core Team, 2015) and OpenBLAS (Xianyi, 2016) on a computer with 64 AMD Opteron 6376 2.3 GHz processors.

In order to study model performance we considered the root-mean-squared prediction error (RMSPE) and the mean continuous rank probability score (MCRPS, Gneiting et al., 2005) that penalises both for location error in the posterior distribution and under/over-confidence. We also carried out a simple validation study on the mole fractions, where we diagnosed the posterior distributions at times when observations were missing, using the same diagnostics as for the flux field. Samples from the mole-fraction field, which was fully marginalised out in our sampling scheme, were obtained by using the samples from the grouped/thinned chains. Specifically, at the i -th iteration we sampled $\mathbf{Y}_2^{(i)}$ from $p(\mathbf{Y}_2 \mid \mathbf{Z}_2, \mathbf{Y}_1^{(i)}, \boldsymbol{\theta}_2^{(i)}, \tau_2^{(i)})$, where the superscript ‘ (i) ’ indicates the i -th sample.

5.3. Results

Flux prediction: Flux diagnostics in Table 1 indicate that there is no practical difference between the lognormal model (i.e., $\lambda = 0$) and the Box–Cox model when carrying out trace-gas inversion in the UK and Ireland; this is not surprising since both the posterior mean and posterior median of λ in Model 1 is 0.19. In separate simulation studies on the example illustrated in Appendix B, we found very little difference in the prediction performance between the log-

Table 1: Flux field and mole-fraction prediction diagnostics for Models 1–6, Model 1*, and Model 1**: Entries give the root-mean-squared prediction error (RMSPE) and the mean continuous rank probability score (MCRPS). Flux-field diagnostics are averages of scores at each $\mathbf{s} \in D_1^L$. Mole-fraction-field diagnostics are averages of scores at the space-time points at which observations are missing. Model 1* and Model 1** denote the cases when an inventory from mainland Europe and an inventory from Northern Australia are used for $\mathbf{W}_1$, respectively.

Model	Flux (g s^{-1})		Mole Fraction (ppb)	
	RMSPE	MCRPS	RMSPE	MCRPS
1	294.8	154.7	18.4	10.9
2	293.4	154.1	18.4	10.8
3	327.4	180.2	18.5	10.9
4	306.4	155.5	18.5	10.9
5	319.4	160.7	18.5	10.9
6	332.5	178.2	18.5	10.9
1*	308.8	165.6	18.4	10.8
1**	329.1	191.4	18.4	10.8

normal process and the Box–Cox process at these small values for λ . On the other hand, as seen from Table 1, the truncated Gaussian model (Model 3) performs considerably worse than Models 1 and 2. Further, the spatially uncorrelated Box–Cox and lognormal models perform considerably worse than their correlated counterparts, both in terms of RMSPE and CRPS.

In Fig. 4, we provide a violin plot (Hintze and Nelson, 1998), which shows where the true flux values (here, the UK and Ireland inventory) lie in relation to the posterior distribution of the flux at each spatial location for the Box–Cox process (Model 1). The colour of the violin is related to the posterior median, while the background colour of each grid cell denotes the true flux. Most of the time, the true (in this case the inventory) flux values coincide with regions of high posterior probability density. Although the width of the intervals containing high posterior density are large for some regions, the colour map indicates a broad agreement between the posterior median and the true flux. We can see however that the extremes, namely the low emissions in the Scottish Highlands and the high emissions in the English Midlands are not correctly captured. The lognormal process (Model 2) fared slightly better in this respect, while the poor diagnostics associated with the truncated Gaussian model (Models 3 and 6) are attributed to over- and under-estimating the flux in these regions, respectively.

Although the uncertainty on the flux in the *individual* grid cells is quite large (as can be seen from Fig. 4), the uncertainty on the *sum* of fluxes is not so poorly constrained. This is a natural consequence of averaging and of the linear operator obtained from the transport model, which integrates the flux over a large domain at each time point. We consider two levels of aggregation

Fig. 4: Violin plot summarising the posterior distributions of the fluxes (in g/s) at each grid cell contained in D_1^L for Model 1. Each violin is composed of a double-sided density plot obtained from the samples of $(\mathbf{Y}_1 | \mathbf{Z}_2, \mathbf{W}_1)$. In each grid cell the ‘x’ symbol and the background cell colour denote the true (unobserved) total flux at that grid cell. The fill colour of the density denotes the posterior median. The top and right labels on the region denote the cell centre’s longitude and latitude coordinate, respectively.

that are important, the entire geographic region (i.e., the UK and Ireland together), and the smaller, country level of aggregation (England, Northern Ireland, Scotland, Wales, and Ireland separately). Total flux for land territories at the country level for Model 1 are given in Fig. 5, left panel. These distributions are well constrained and, with the exception of Scotland, contain the true flux well within the 95% credible interval. The poorer performance in Scotland is due to the near-zero fluxes in the Highlands that take anomalously low values. Clearly, inferences on fluxes become more well constrained when spatially aggregating to the resolution of countries within the UK and Ireland.

Fig. 5: Box-and-whisker plots summarising the posterior distribution of country-level and whole-domain flux aggregates. Left panel: The total flux in the land territories of England, Northern Ireland, Scotland, Wales, and Ireland for Model 1. The triangle denotes the true (unobserved) flux. Right panel: The total flux in the land territories of the UK and Ireland for Models 1–6, Model 1*, and Model 1**. The triangle denotes the true (unobserved) total flux of 2.73 Tg yr^{-1} . The boxes denote the interquartile range, the whiskers extend to the last values that are within 1.5 times the interquartile range from the quartiles, and the dots show the samples that lie beyond the end of the whiskers.

The posterior distribution of total methane emissions in the entire geographic region are summarised in Fig. 5, right panel. The truncated Gaussian models (i.e., Models 3 and 6) do not contain the true value (2.73 Tg yr^{-1}) within the 95% credible interval. Interestingly, Model 1* (that uses the mainland Europe inventory) performed just as well as Model 1. Clearly, the sensitivity to the inventory when inferring over aggregated regions is less than when predicting at a high resolution, as expected. From Fig. 5, right panel, it is apparent that capturing non-Gaussianity is more important than capturing spatial correlations when predicting aggregates. On the other hand, models that capture spatial correlation provide considerably better point-level predictions (compare with Table 1).

Mole-fraction-field prediction: We carry out inference on the mole-fraction field at the space-time points for which no mole-fraction observation is available (note that $b_t(\mathbf{s}, \cdot)$ is still evaluated at these space-time points and, hence, the true mole fraction is known). From Table 1, we see that all methods have similar out-of-sample mole-fraction prediction performance. This is due to the ill-posed nature of the problem: There are a large number of plausible flux fields that yield similar observations. Hence, unless $b_t(\mathbf{s}, \cdot)$ evaluated at the validation space-time points is very different from all other evaluations (which is unlikely due to prevailing winds and

Fig. 6: Left panel: Posterior densities of the parameters θ_{11} , θ_{12} , and λ , from left to right, which appear in the flux field for Models 1–4. Note that the spatial correlation function is specified to be the indicator function in Models 4–6. Since the transformation is also assumed known for Models 5–6, no flux-field parameters need to be inferred for these models. Right panel: Box-and-whisker plots of samples from the marginal posterior distributions of the parameters τ_2 , a , and d , from top to bottom, which appear in the mole-fraction discrepancy term for Models 1–6. See Fig. 5 for a description of the box-and-whisker plot.

recurring meteorology patterns) mole-fraction prediction performance clearly cannot be used to draw any conclusions on the flux-field prediction performance.

Posterior distributions of parameters: Flux-parameter posterior distributions for Models 1–6 are summarised in Fig. 6, left panel, while mole-fraction-parameter posterior distributions are summarised in Fig. 6, right panel (recall that the true values are $\tau_2 = 0.01$ ppb $^{-2}$, $a = 0.9$, and $d = 2.5^\circ$). Irrespective of the model used to describe the flux field, posterior inferences on the parameters describing the discrepancy term $\zeta_t(\cdot)$ were practically identical for each case. This is reassuring, as it indicates that inferences over the discrepancy $\{\zeta_t : t \in \mathcal{T}\}$ are well constrained and, by implication, that inferences over $\{\mathbf{B}_t \mathbf{Y}_1 : t \in \mathcal{T}\}$ are as well. This does not imply that inferences over $\mathbf{Y}_1$ are well constrained, and posterior marginal uncertainty over the elements of $\mathbf{Y}_1$ clearly depends on the dimensionality of $\mathbf{Y}_1$ due to the aggregation induced by the matrices $\{\mathbf{B}_t : t \in \mathcal{T}\}$.

Inventory sensitivity: Over-reliance on the quality of the ‘prior inventory,’ and uncertainties associated with it, is an ongoing topic of concern in atmospheric trace-gas inversion (e.g., Berchet et al., 2015). In Table 1, we give diagnostics for two other models (Model 1* and Model 1**) that are identical to Model 1 but with inventories $\mathbf{W}_1^*$ (taken from mainland Europe containing parts of France, Switzerland, and Germany) and $\mathbf{W}_1^{**}$ (taken from Northern Australia) instead. The inventory $\mathbf{W}_1^*$ is, at first sight, entirely different from that attributed

to the UK and Ireland (see Fig. 3). However, it contains similar spatial properties that aid in the assimilation; in particular, the maximum likelihood estimate of λ obtained using `geoR` (Ribeiro Jr. and Diggle, 2015) with $\mathbf{W}_1^*$ as data, is 0.21. As seen from Table 1, the quality of flux prediction for Model 1* deteriorates but not drastically so.

On the other hand, the inventory $\mathbf{W}_1^{**}$ is much smoother (recall Fig. 3, right panel); the maximum likelihood estimate of λ obtained using `geoR` with $\mathbf{W}_1^{**}$ as data, is 0.83, and the posterior mean obtained using our hierarchical model with the Box–Cox process (Section 3.3), is 0.68. This value for λ is considerably higher than that obtained with the correct inventory and shows the impact that the inventory has on the inference of the transformation parameter. As expected, the use of this inventory considerably deteriorates the flux prediction (see Table 1).

From these two experiments with $\mathbf{W}_1^*$ and $\mathbf{W}_1^{**}$, we conclude that while the inventory remains important in trace-gas inversion in this framework, its effect on the flux-prediction performance may not dominate. Further, all that is needed is an inventory that has similar spatial characteristics as those of the true flux field; one can justifiably expect that any accepted methane inventory for the UK and Ireland will be reasonable for use within our framework. From a statistical perspective, prior sensitivity has been shifted from the inventory (process layer) to the spatial properties of the inventory (parameter layer).

Lognormal model may not always be adequate: From the results above, one might think that for our study there is no benefit in considering the Box–Cox trans-Gaussian model over the lognormal model. We provide a counter-example by simply replacing the inventory we used for the UK and Ireland with the one from Northern Australia (Figure 3, right panel), and re-running the entire OSSE with the same meteorology and station locations as those in the UK and Ireland for Models 1–6. Since the methane emissions in this region are lower than those in the UK and Ireland, we set $\tau_2 = 4 \text{ ppb}^{-2}$ so that the discrepancy does not overwhelm the signal of interest. Recall that the maximum likelihood estimate of λ obtained using `geoR` for this inventory is 0.83; the posterior mean obtained for Model 1 using the true inventory in the assimilation was 0.81. From these estimates, one might expect that a lognormal spatial process model for the flux field will not perform well in this region; indeed, in Table 2 we see that both the Box–Cox spatial process (Model 1) and the truncated Gaussian process (Model 3) considerably outperform the lognormal process (Model 2) in this case. Again, spatially correlated models are seen to outperform the spatially uncorrelated ones in flux prediction but,

Table 2: Flux field and mole-fraction prediction diagnostics for Models 1–6 when the mole-fraction data is simulated from an inventory from Northern Australia: Entries give the root-mean-squared prediction error (RMSPE) and the mean continuous rank probability score (MCRPS). Flux-field diagnostics are averages of scores at each $\mathbf{s} \in D_1^L$. Mole-fraction-field diagnostics are averages of scores at the space-time points at which observations are missing.

Model	Flux (g s^{-1})		Mole Fraction (ppb)	
	RMSPE	MCRPS	RMSPE	CRPS
1	13.7	7.1	0.8	0.5
2	17.8	10.3	0.9	0.5
3	13.9	7.2	0.8	0.5
4	18.8	10.4	0.9	0.5
5	24.5	14.3	0.9	0.5
6	19.1	10.5	0.9	0.5

more importantly, the Box–Cox model is able to adapt well in this very different scenario.

Reproducible code for this OSSE is available from <https://github.com/andrewzm/atminv>.

6. Discussion

Lognormality in flux-field modelling was considered for methane flux inversion by Ganesan et al. (2014) and Ganesan et al. (2015), however they did not make use of spatial models. Spatial lognormality of the flux field was subsequently considered in Zammit-Mangion et al. (2015), in which a two-stage inferential approach was adopted. In this article, we extend these works by proposing a new class of non-Gaussian spatio-temporal bivariate models for use in atmospheric trace-gas inversion. The bivariate model may be fully characterised through the use of cumulant functions (Appendix A). It is very flexible, and yet it requires only the specification of a univariate Box–Cox spatial process, a univariate Gaussian spatio-temporal field, and an interaction function. Most previous works in this area use models that contain only parts of the class of models we consider in this article.

In the hierarchical model of Section 3.3, the quantity of interest, $\mathbf{Y}_1$, is not directly observed. This implies that the mole-fraction data is *conditionally uninformative* of the flux field (Poirier, 1998) and, in this case, the quantity $\mathbf{Y}_1$ is frequently termed *unidentifiable*. Unidentifiability does not prohibit Bayesian learning on $\mathbf{Y}_1$, since $\mathbf{Y}_1$ and $\mathbf{Y}_2$ are never *a priori* independent in this problem. Xie and Carlin (2006) proposed a way to measure the permissible extent of learning obtained from the data in this context of “unidentifiability.” Such measures may prove useful for trace-gas-inversion experimental design through the use of simulation studies, such as

the one conducted in this article.

We implemented a fully Bayesian approach to flux inversion, in which a degree of non-Gaussianity was incorporated through the Box–Cox-transformation parameter. Our results provide further evidence that both non-Gaussianity *and* spatial correlation are particularly important features to model in trace-gas inversion of methane and, further, that the degree of non-Gaussianity is also important. The truncated Gaussian assumption (adopted in Miller et al., 2014) seems unsuitable for the UK and Ireland case, while the lognormal process (Zammit-Mangion et al., 2015) seems unsuitable for parts of Northern Australia. The Box–Cox process adapts well to both and more generally to largely different types of flux fields.

Covariates, such as population density, could be used to describe spatial features, and it is also possible that their inclusion could motivate the use of a Gaussian field for the stochastic component that remains. However, the use of covariates does not especially simplify the problem. In this example, covariates will almost certainly need to be spatially weighted and, further, their inclusion precludes the possibility for using the inventory in the assimilation, since these inventories are constructed using surrogate information such as population density. The advantage of only using second-order and higher-order prior descriptors is that the posterior distribution of the spatial pattern of fluxes is predominantly data-driven. Critically, in this study we have shown that usable flux inferences can be obtained at several resolutions of importance (e.g., the country level) with a reduced reliance on the inventory. This result may have implications for other geophysical fields of study that rely heavily on data assimilation.

Although transport models are in constant development and regularly validated (e.g., Ahmadov et al., 2009; Ryall and Maryon, 1998), they are by no means perfect. We acknowledge this with the use of a discrepancy term that we model as conditionally Gaussian and space-time separable (see (13)). This is overly simplistic; however, there is not enough information in the data to warrant a more complex model. For short-lived gases such as carbon monoxide, the discrepancy term also captures misspecification arising from linearisation. If nonlinearity of the flux-mole-fraction mapping is deemed important to model, then one might consider replacing the linear mapping with a higher-order one or even a stochastic one (using, for example, a Gaussian process emulator). In the latter case, computation and model interpretability are rendered more difficult. For example, the bivariate process of Section 2.2 would then need to be interpreted using *modified* cumulants (Schultz et al., 1978).

Finally, a comment is needed on inferring the transformation parameter λ within our frame-

work. With both the UK/Ireland and the Australia OSSE simulations, fitting a standard Box–Cox spatial model to the inventory data using `geoR` yielded maximum likelihood estimates within 0.02 of the posterior mean obtained using our fully Bayesian approach (on the models with the correct inventory). Thus, as expected, the data in trace-gas inversion is not informative of λ , although clearly this parameter plays a big role in the flux-prediction performance. Box and Cox (1964) suggested that a value of λ be selected in light of its posterior distribution, following which a standard analysis with fixed λ could be carried out. Christensen et al. (2001) further argue that only a select-few ‘interpretable’ values for λ should be chosen and that standard maximum likelihood can guide one’s choice. We do not oppose these statements and, in light of the insensitivity of the mole-fraction data on λ in this application, one might prefer to estimate λ offline and use it as a “plug-in.” This pragmatic approach is unlikely to adversely affect the results in any meaningful way in regional inversions.

Acknowledgements

As well as their University of Wollongong appointments, Andrew Zammit-Mangion is Honorary Research Fellow at the University of Bristol, and Noel Cressie is Distinguished Visiting Scientist at the Jet Propulsion Laboratory of the National Aeronautics and Space Administration (NASA). Noel Cressie’s research was partially supported by a 2015–2017 Australian Research Council Discovery Project, DP150104576, and partially supported by NASA grant NNH11ZDA001N-OCO2. Anita Ganesan was funded by the Natural Environment Research Council (NERC) Independent Research Fellowship NE/L010992/1. We would like to thank Alistair J. Manning for providing us with the NAME data used in this study. The UK National Atmospheric Emissions Inventory (NAEI) used here was funded by DECC, the Department for Environment, Food and Rural Affairs (DEFRA), the Scottish Government, the Welsh Government, and the Northern Ireland Department of Environment. We would also like to thank Bohai Zhang for providing some references, Jonathan Rougier for providing the *R* software for slice sampling, and Matt Rigby for ongoing discussions related to this application. Finally, our sincere appreciation goes to two referees and the Associate Editor for constructive comments on the initial submission.

References

References

- Ahmadvov, R., Gerbig, C., Kretschmer, R., Körner, S., Rödenbeck, C., Bousquet, P., Ramonet, M., 2009. Comparing high resolution WRF-VPRM simulations and two global CO₂ transport models with coastal tower measurements of CO₂. *Biogeosciences* 6, 807–817.
- Aitchison, J., Brown, J. A. C., 1957. *The Lognormal Distribution with Special Reference to its Uses in Economics*. Cambridge University Press, Cambridge, UK.
- Åström, K. J., 2006. *Introduction to Stochastic Control Theory*. Dover, New York, NY.
- Berchet, A., Pison, I., Chevallier, F., Bousquet, P., Bonne, J.-L., Paris, J.-D., 2015. Objectified quantification of uncertainties in Bayesian atmospheric inversions. *Geoscientific Model Development* 8, 1525–1546.
- Berger, J. O., De Oliveira, V., Sansó, B., 2001. Objective Bayesian analysis of spatially correlated data. *Journal of the American Statistical Association* 96, 1361–1374.
- Berger, J. O., Liseo, B., Wolpert, R. L., 1999. Integrated likelihood methods for eliminating nuisance parameters. *Statistical Science* 14, 1–28.
- Box, G. E., Cox, D. R., 1964. An analysis of transformations. *Journal of the Royal Statistical Society: Series B* 26, 211–252.
- Cameletti, M., Lindgren, F., Simpson, D., Rue, H., 2013. Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *Advances in Statistical Analysis* 97, 109–131.
- Christensen, O. F., Diggle, P. J., Ribeiro Jr, P. J., 2001. Analysing positive-valued spatial data: the transformed Gaussian model. In: Monestiez, P., Allard, D., Froidevaux, R. (Eds.), *geoENV III Geostatistics for Environmental Applications*. Kluwer Academic Publishers, Dordrecht, Netherlands, pp. 287–296.
- Cressie, N., 1993. *Statistics for Spatial Data*, rev. edn. John Wiley & Sons, New York, NY.
- Cressie, N., Johannesson, G., 2008. Fixed Rank Kriging for very large spatial data sets. *Journal of the Royal Statistical Society: Series B* 70, 209–226.

- Cressie, N., Shi, T., Kang, E. L., 2010. Fixed Rank Filtering for spatio-temporal data. *Journal of Computational and Graphical Statistics* 19, 724–745.
- Cressie, N., Wikle, C. K., 2011. *Statistics for Spatio-Temporal Data*. John Wiley and Sons, Hoboken, NJ.
- Cressie, N., Zammit-Mangion, A., 2015. Bivariate spatial covariance models: A conditional approach. <http://arxiv.org/abs/1504.01865>, [Online; accessed 09 November 2015].
- De Oliveira, V., Kede, B., Short, D. A., 1997. Bayesian prediction of transformed Gaussian random fields. *Journal of the American Statistical Association* 92, 1422–1433.
- Department for Environment Food and Rural Affairs (DEFRA), 2014. National Atmospheric Emissions Inventory. <http://naei.defra.gov.uk>, [Online; accessed 01 March 2015].
- Dimitrakopoulos, R., Mustapha, H., Gloaguen, E., 2010. High-order statistics of spatial random fields: exploring spatial cumulants for modeling complex non-Gaussian and non-linear phenomena. *Mathematical Geosciences* 42, 65–99.
- Draper, N. R., Cox, D. R., 1969. On distributions and their transformation to normality. *Journal of the Royal Statistical Society: Series B* 31, 472–476.
- Edenhofer, O., R., Pichs-Madruga, Y., Sokona, E., Farahani, S., Kadner, K., Seyboth, A., Adler, I., Baum, S., Brunner, P., Eickemeier, B., Kriemann, J., Savolainen, S., Schlmer, C., von Stechow, T., Zwickel, Minx, J., 2014. IPCC, 2014: Summary for Policymakers. In: *Climate Change 2014: Mitigation of Climate Change. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK.
- Emery, C., Jung, J., Downey, N., Johnson, J., Jimenez, M., Yarwood, G., Morris, R., 2012. Regional and global modeling estimates of policy relevant background ozone over the United States. *Atmospheric Environment* 47, 206–217.
- Freeman, J., Modarres, R., 2006. Inverse Box–Cox: the power-normal distribution. *Statistics & Probability Letters* 76, 764–772.
- Fuentes, M., Raftery, A. E., 2005. Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics* 61, 36–45.

- Ganesan, A. L., Manning, A. J., Grant, A., Young, D., Oram, D. E., Sturges, W. T., Moncrieff, J. B., 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.
- 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., Weiss, R. F., 2014. Characterization of uncertainties in atmospheric trace gas inversions using hierarchical Bayesian methods. *Atmospheric Chemistry and Physics* 14, 3855–3864.
- Gneiting, T., Kleiber, W., Schlather, M., 2010. Matérn cross-covariance functions for multivariate random fields. *Journal of the American Statistical Association* 105, 1167–1177.
- Gneiting, T., Raftery, A. E., Westveld III, A. H., Goldman, T., 2005. Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation. *Monthly Weather Review* 133, 1098–1118.
- Hintze, J. L., Nelson, R. D., 1998. Violin plots: a box plot-density trace synergism. *The American Statistician* 52, 181–184.
- Joint Research Centre of the European Commission (JRC) and the Netherlands Environmental Assessment Agency (PBL), 2011. Emission Database for Global Atmospheric Research (EDGAR), release version 4.2. <http://edgar.jrc.ec.europa.eu>, [Online; accessed 12 May 2014].
- Jones, A., Thomson, D. J., Hort, M., Devenish, B., 2007. The UK Met Office's next-generation atmospheric dispersion model, NAME III. In: Borrego, C., Norman, A.-L. (Eds.), *Air Pollution Modeling and its Application XVII*. Springer, New York, NY, pp. 580–589.
- Kaufman, C. G., Schervish, M. J., Nychka, D. W., 2008. Covariance tapering for likelihood-based estimation in large spatial data sets. *Journal of the American Statistical Association* 103, 1545–1555.
- Kotz, S., Balakrishnan, N., Johnson, N. L., 2000. *Continuous Multivariate Distributions*. Volume 1: Models and Applications, 2nd Edition. John Wiley & Sons, New York, NY.

- Kuznetsov, P. I., Stratonovich, R. L., Tikhonov, V. I., 1965. The transmission of certain random functions through linear systems. In: *Non-linear Transformations of Stochastic Processes*. Pergamon Press Ltd., Oxford, UK, pp. 3–28.
- Lindgren, F., Rue, H., Lindström, J., 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B* 73, 423–498.
- Lindgren, G., 2012. *Stationary Stochastic Processes: Theory and Applications*. CRC Press, Boca Raton, FL.
- 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., Simmonds, P. G., 2015. Reconciling reported and unreported HFC emissions with atmospheric observations. *Proceedings of the National Academy of Sciences* 112, 5927–5931.
- Miller, S. M., Michalak, A. M., Levi, P. J., 2014. Atmospheric inverse modeling with known physical bounds: an example from trace gas emissions. *Geoscientific Model Development* 7, 303–315.
- Miller, S. M., Wofsy, S. C., Michalak, A. M., Kort, E. A., Andrews, A. E., Biraud, S. C., Dlugokencky, E. J., Eluszkiewicz, J., Fischer, M. L., Janssens-Maenhout, G., Miller, B. R., Miller, J. B., Montzka, S. A., Nehrkorn, T., Sweeney, C., 2013. Anthropogenic emissions of methane in the United States. *Proceedings of the National Academy of Sciences* 110, 20018–20022.
- Murray, I., Ghahramani, Z., MacKay, D., 2006. MCMC for doubly-intractable distributions. In: *Proceedings of the 22nd Annual Conference on Uncertainty in Artificial Intelligence*. pp. 359–366.
- Neal, R. M., 2003. Slice sampling. *Annals of Statistics* 31, 705–741.
- Neal, R. M., 2011. MCMC using Hamiltonian dynamics. In: Brooks, S., Gelman, A., Jones, G. L., Meng, X.-L. (Eds.), *Handbook of Markov Chain Monte Carlo*. Chapman & Hall/CRC Press, Boca Raton, FL, pp. 113–162.

- O'Hagan, A., 2000. Kendall's Advanced Theory of Statistics, Volume 2B: Bayesian Analysis, 2nd Edition. Arnold, London, UK.
- Pericchi, L., 1981. A Bayesian approach to transformations to normality. *Biometrika* 68, 35–43.
- Peters, W., Miller, J., Whitaker, J., Denning, A., Hirsch, A., Krol, M., Zupanski, D., Bruhwiler, L., Tans, P., 2005. An ensemble data assimilation system to estimate CO₂ surface fluxes from atmospheric trace gas observations. *Journal of Geophysical Research*, D24304.
- Poirier, D. J., 1978. The use of the Box-Cox transformation in limited dependent variable models. *Journal of the American Statistical Association* 73, 284–287.
- Poirier, D. J., 1998. Revising beliefs in nonidentified models. *Econometric Theory* 14, 483–509.
- R Core Team, 2015. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria.
- Ribeiro Jr., P. J., Diggle, P. J., 2015. *geoR: Analysis of Geostatistical Data*. <http://CRAN.R-project.org/package=geoR>, R package version 1.7-5.1, [Online; accessed 01 December 2015].
- Rigby, M., Manning, A. J., Prinn, R. G., 2011. Inversion of long-lived trace gas emissions using combined Eulerian and Lagrangian chemical transport models. *Atmospheric Chemistry and Physics* 11, 9887–9898.
- Royle, J. A., Berliner, L. M., 1999. A hierarchical approach to multivariate spatial modeling and prediction. *Journal of Agricultural, Biological, and Environmental Statistics* 4, 29–56.
- Rue, H., Held, L., 2005. *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC Press, Boca Raton, FL.
- Ryall, D., Maryon, R., 1998. Validation of the UK Met. Office's NAME model against the ETEX dataset. *Atmospheric Environment* 32, 4265–4276.
- Schultz, G. W., Tsokos, C. P., Rao, A., 1978. Statistical properties of a linear stochastic system. *Information and Control* 39, 92–117.
- Severini, T. A., 2005. *Elements of Distribution Theory*. Cambridge University Press, Cambridge, UK.

- Stohl, A., Seibert, P., Arduini, J., Eckhardt, S., Fraser, P., Grealley, 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., Yokouchi, Y., 2009. An analytical inversion method for determining regional and global emissions of greenhouse gases: sensitivity studies and application to halocarbons. *Atmospheric Chemistry and Physics* 9, 1597–1620.
- Storvik, G., Frigessi, A., Hirst, D., 2002. Stationary space-time Gaussian fields and their time autoregressive representation. *Statistical Modelling* 2, 139–161.
- Tarantola, A., 2005. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, Philadelphia, PA.
- Thompson, R. L., Stohl, A., 2014. FLEXINVERT: an atmospheric Bayesian inversion framework for determining surface fluxes of trace species using an optimized grid. *Geoscientific Model Development* 7, 2223–2242.
- Van Dyk, D. A., Park, T., 2008. Partially collapsed Gibbs samplers: theory and methods. *Journal of the American Statistical Association* 103, 790–796.
- Wikle, C. K., Berliner, L. M., 2007. A Bayesian tutorial for data assimilation. *Physica D: Nonlinear Phenomena* 230, 1–16.
- Xianyi, Z., 2016. OpenBLAS: an optimized BLAS library. <http://www.openblas.net/>, [Online; accessed 01 May 2016].
- Xie, Y., Carlin, B. P., 2006. Measures of Bayesian learning and identifiability in hierarchical models. *Journal of Statistical Planning and Inference* 136, 3458–3477.
- Yaglom, A. M., 1987. *Correlation Theory of Stationary and Related Random Functions Volume I: Basic Results*. Springer-Verlag, New York, NY.
- Zammit-Mangion, A., Cressie, N., Ganesan, A., O'Doherty, S., Manning, A., 2015. Spatio-temporal bivariate statistical models for atmospheric trace-gas inversion. *Chemometrics and Intelligent Laboratory Systems* 149, 227–241.

Appendix A. Derivation of cross-cumulant functions

Since all joint cumulants associated with independent variables are zero (e.g., Severini, 2005, Theorem 4.16), we focus here on the linear system (4) in which the discrepancy term is zero.

Consider two random vectors $\mathbf{Y}_1 \in \mathbb{R}^N$ and $\mathbf{Y}_2 \in \mathbb{R}^M$ such that $\mathbf{Y}_2 = \mathbf{A}\mathbf{Y}_1$, and $\mathbf{A}$ is a real-valued $M \times N$ matrix. Let the columns of $\mathbf{A}$ be denoted as $\mathbf{a}_1, \dots, \mathbf{a}_N$, and denote the characteristic functions of $\mathbf{Y}_1$ and $\mathbf{Y}_2$ as

$$\varphi_{\mathbf{Y}_1}(\mathbf{t}_1) = \mathbb{E}(\exp(\iota \mathbf{t}_1' \mathbf{Y}_1)) = \int_{\mathbb{R}^N} \exp(\iota \mathbf{t}_1' \mathbf{Y}_1) p(\mathbf{Y}_1) d\mathbf{Y}_1,$$

and

$$\varphi_{\mathbf{Y}_2}(\mathbf{t}_2) = \mathbb{E}(\exp(\iota \mathbf{t}_2' \mathbf{Y}_2)) = \int_{\mathbb{R}^M} \exp(\iota \mathbf{t}_2' \mathbf{Y}_2) p(\mathbf{Y}_2) d\mathbf{Y}_2,$$

respectively. The distribution of $\mathbf{Y}_1$ is obtained from

$$p(\mathbf{Y}_1) = \int_{\mathbb{R}^M} p(\mathbf{Y}_2 | \mathbf{Y}_1) p(\mathbf{Y}_1) d\mathbf{Y}_2 = \int_{\mathbb{R}^M} \delta(\mathbf{Y}_2 - \mathbf{A}\mathbf{Y}_1) p(\mathbf{Y}_1) d\mathbf{Y}_2,$$

where $\delta(\cdot)$ is the Dirac delta function. Let $\mathbf{Y} \equiv (\mathbf{Y}_1', \mathbf{Y}_2')'$ and $\mathbf{t}_\mathbf{Y} \equiv (\mathbf{t}_1', \mathbf{t}_2')'$. Then the characteristic function of $\mathbf{Y}$, $\varphi_\mathbf{Y}(\mathbf{t}_\mathbf{Y}) = \mathbb{E}(\exp(\iota \mathbf{t}_\mathbf{Y}' \mathbf{Y}))$, is given by

$$\begin{aligned} \varphi_\mathbf{Y}(\mathbf{t}_\mathbf{Y}) &= \int_{\mathbb{R}^N} \int_{\mathbb{R}^M} \exp\left(\iota [\mathbf{t}_1', \mathbf{t}_2'] \begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{Y}_2 \end{bmatrix}\right) p(\mathbf{Y}_2 | \mathbf{Y}_1) p(\mathbf{Y}_1) d\mathbf{Y}_2 d\mathbf{Y}_1 \\ &= \int_{\mathbb{R}^N} \exp\left(\iota [\mathbf{t}_1', \mathbf{t}_2'] \begin{bmatrix} \mathbf{Y}_1 \\ \mathbf{A}\mathbf{Y}_1 \end{bmatrix}\right) p(\mathbf{Y}_1) d\mathbf{Y}_1 \\ &= \int_{\mathbb{R}^N} \exp\left(\iota [\mathbf{t}_1', \mathbf{t}_2'] \begin{bmatrix} \mathbf{I} \\ \mathbf{A} \end{bmatrix} \mathbf{Y}_1\right) p(\mathbf{Y}_1) d\mathbf{Y}_1. \end{aligned}$$

We therefore obtain the result,

$$\varphi_\mathbf{Y}(\mathbf{t}_\mathbf{Y}) = \varphi_{\mathbf{Y}_1}(t_{11} + \mathbf{t}_2' \mathbf{a}_1, \dots, t_{1N} + \mathbf{t}_2' \mathbf{a}_N). \quad (\text{A.1})$$

Equation (A.1) is a natural extension to the result of Kuznetsov et al. (1965, Appendix I), who showed that

$$\varphi_{\mathbf{Y}_2}(\mathbf{t}_2) = \varphi_{\mathbf{Y}_1}(\mathbf{t}_2' \mathbf{a}_1, \dots, \mathbf{t}_2' \mathbf{a}_N), \quad (\text{A.2})$$

but (A.1) allows us to derive *all* cross-cumulant functions as limits of Riemann sums.

Returning to the original application, consider a random field $Y_2(\mathbf{s}) = \int_D b(\mathbf{s}, \mathbf{u}) Y_1(\mathbf{u}) d\mathbf{u}$ (i.e., let $\zeta(\cdot) = 0$). If $Y_1(\cdot)$ has continuous realisations with probability one, and the function

$b(\mathbf{s}, \mathbf{u})$ is continuous, then this integral can be represented approximately as,

$$Y_2(\mathbf{s}_i) = \sum_{j=1}^N Y_1(\mathbf{u}_j) b(\mathbf{s}_i, \mathbf{u}_j) \Delta_{\mathbf{u}_j}; \quad \mathbf{s}_i \in D,$$

where N is large and $\Delta_{\mathbf{u}_j}$ is a small region in D centred at $\mathbf{u}_j$. Now, let $\mathbf{Y}_1 \equiv (Y_1(\mathbf{u}_j) : j = 1, \dots, N)'$ and $\mathbf{Y}_2 \equiv (Y_2(\mathbf{s}_i) : i = 1, \dots, M)'$. Further, recall that $\varphi_{\mathbf{Y}_1}(\mathbf{t}_1)$ and $\varphi_{\mathbf{Y}_2}(\mathbf{t}_2)$ are the characteristic functions of $\mathbf{Y}_1$ and $\mathbf{Y}_2$, respectively. We saw that $\mathbf{Y} \equiv (\mathbf{Y}_1', \mathbf{Y}_2')'$ has characteristic function $\varphi_{\mathbf{Y}}(\mathbf{t}_{\mathbf{Y}})$, where $\mathbf{t}_{\mathbf{Y}} \equiv (\mathbf{t}_1', \mathbf{t}_2')'$. Then, from (A.1), we have the approximate representation,

$$\varphi_{\mathbf{Y}}(\mathbf{t}_{\mathbf{Y}}) = \varphi_{\mathbf{Y}_1} \left(t_{11} + \sum_{i=1}^M t_{2i} b(\mathbf{s}_i, \mathbf{u}_1) \Delta_{\mathbf{u}_1}, \dots, t_{1N} + \sum_{i=1}^M t_{2i} b(\mathbf{s}_i, \mathbf{u}_N) \Delta_{\mathbf{u}_N} \right). \quad (\text{A.3})$$

Under conditions outlined in Kuznetsov et al. (1965), the characteristic function of the process $Y_1(\cdot)$ evaluated at $\{\mathbf{u}_j : j = 1, \dots, N\}$ can be expanded as a series

$$\varphi_{\mathbf{Y}_1}(\mathbf{t}_1) = \exp \left(\iota \sum_{j=1}^N \kappa_{Y_1}^1(\mathbf{u}_j) t_{1j} + \frac{\iota^2}{2} \sum_{j,j'=1}^N \kappa_{Y_1 Y_1}^2(\mathbf{u}_j, \mathbf{u}_{j'}) t_{1j} t_{1j'} + \dots \right).$$

Using this series definition for $\varphi_{\mathbf{Y}_1}(\mathbf{t}_1)$ in (A.3), we obtain

$$\begin{aligned} \varphi_{\mathbf{Y}}(\mathbf{t}_{\mathbf{Y}}) = & \exp \left(\iota \sum_{j=1}^N \kappa_{Y_1}^1(\mathbf{u}_j) \left[t_{1j} + \sum_{i=1}^M t_{2i} b(\mathbf{s}_i, \mathbf{u}_j) \Delta_{\mathbf{u}_j} \right] \right. \\ & + \frac{\iota^2}{2} \sum_{j,j'=1}^N \kappa_{Y_1 Y_1}^2(\mathbf{u}_j, \mathbf{u}_{j'}) \left[t_{1j} + \sum_{i=1}^M t_{2i} b(\mathbf{s}_i, \mathbf{u}_j) \Delta_{\mathbf{u}_j} \right] \left[t_{1j'} + \sum_{i'=1}^M t_{2i'} b(\mathbf{s}_{i'}, \mathbf{u}_{j'}) \Delta_{\mathbf{u}_{j'}} \right] \\ & \left. + \dots \right). \end{aligned} \quad (\text{A.4})$$

Now $\varphi_{\mathbf{Y}}(\mathbf{t}_{\mathbf{Y}})$ may also be expanded as a series:

$$\begin{aligned} \varphi_{\mathbf{Y}}(\mathbf{t}_{\mathbf{Y}}) = & \exp \left(\iota \sum_{j=1}^N \kappa_{Y_1}^1(\mathbf{u}_j) t_{1j} + \iota \sum_{i=1}^M \kappa_{Y_2}^1(\mathbf{s}_i) t_{2i} \right. \\ & + \frac{\iota^2}{2} \sum_{i,i'=1}^M \kappa_{Y_2 Y_2}^2(\mathbf{s}_i, \mathbf{s}_{i'}) t_{2i} t_{2i'} + \frac{\iota^2}{2} \sum_{j,j'=1}^N \kappa_{Y_1 Y_1}^2(\mathbf{u}_j, \mathbf{u}_{j'}) t_{1j} t_{1j'} \\ & \left. + \frac{\iota^2}{2} \sum_{i,j'=1}^{M,N} \kappa_{Y_2 Y_1}^2(\mathbf{s}_i, \mathbf{u}_{j'}) t_{2i} t_{1j'} + \frac{\iota^2}{2} \sum_{i',j=1}^{M,N} \kappa_{Y_1 Y_2}^2(\mathbf{u}_j, \mathbf{s}_{i'}) t_{2i'} t_{1j} + \dots \right). \end{aligned} \quad (\text{A.5})$$

Expanding (A.4) and comparing coefficients of the products of elements in $\mathbf{t}_Y$ to those in (A.5), we deduce that

$$\begin{aligned}\kappa_{Y_2}^1(\mathbf{s}_i) &= \sum_{j=1}^N b(\mathbf{s}_i, \mathbf{u}_j) \kappa_{Y_1}^1(\mathbf{u}_j) \Delta_{\mathbf{u}_j}, \\ \kappa_{Y_1 Y_2}^2(\mathbf{u}_j, \mathbf{s}'_i) &= \sum_{j'=1}^N \kappa_{Y_1 Y_1}^2(\mathbf{u}_j, \mathbf{u}_{j'}) b(\mathbf{s}'_i, \mathbf{u}_{j'}) \Delta_{\mathbf{u}_{j'}}, \\ \kappa_{Y_2 Y_2}^2(\mathbf{s}_i, \mathbf{s}'_i) &= \sum_{j, j'=1}^N b(\mathbf{s}_i, \mathbf{u}_j) b(\mathbf{s}'_i, \mathbf{u}_{j'}) \kappa_{Y_1 Y_1}^2(\mathbf{u}_j, \mathbf{u}_{j'}) \Delta_{\mathbf{u}_j} \Delta_{\mathbf{u}_{j'}}, \\ &\vdots\end{aligned}$$

These sums tend to Riemann integrals, such as those shown in (8) and (9), provided $b(\cdot, \cdot)$, $\kappa_{Y_1}^1(\cdot)$, $\kappa_{Y_1 Y_1}^2(\cdot, \cdot)$, $\dots$ are Riemann-integrable. Similar arguments can be used to derive third-order cross-cumulants. For example, it can be shown that

$$\kappa_{Y_1 Y_1 Y_2}^3(\mathbf{u}_j, \mathbf{u}_{j'}, \mathbf{s}''_i) = \int \kappa_{Y_1 Y_1 Y_1}^3(\mathbf{u}_j, \mathbf{u}_{j'}, \mathbf{u}_{j''}) b(\mathbf{s}''_i, \mathbf{u}_{j''}) d\mathbf{u}_{j''},$$

and thus we can formulate all the cross-cumulants of the joint process $(Y_1(\cdot), Y_2(\cdot))$ from those of $Y_1(\cdot)$ through the interaction function $b(\cdot, \cdot)$.

These results can be extended to (4), where the discrepancy term is included (and is independent of $Y_1(\cdot)$) by taking advantage of the additive property of cumulants and the fact that all joint cumulants associated with independent variables are zero (Severini, 2005, Theorem 4.16).

Appendix B. Example illustrating higher-order cross-cumulant functions for bivariate processes

From the relationship between moments and cumulants, it can be shown that $\kappa_{Y_i}^1(\cdot) = E(Y_i(\cdot))$, $\kappa_{Y_i Y_j}^2(\cdot, \cdot) = C_{Y_i Y_j}(\cdot, \cdot) \equiv \text{cov}(Y_i(\cdot), Y_j(\cdot))$ and $\kappa_{Y_i Y_j Y_k}^3(\cdot, \cdot, \cdot) = E((Y_i(\cdot) - \mu_i(\cdot))(Y_j(\cdot) - \mu_j(\cdot))(Y_k(\cdot) - \mu_k(\cdot)))$. These relationships hold even if the moment generating function does not exist (as is the case, for example, with the lognormal process; see Severini, 2005, p. 116). Assume, for illustration, that $Y_1(\cdot)$ is a spatial lognormal process such that $\ln Y_1(\cdot)$ is a Gaussian process with mean function $\tilde{\mu}_1(\cdot)$ and covariance function $\tilde{C}_{Y_1 Y_1}(\cdot, \cdot)$; then the first two spatial cumulant functions are the well known mean and covariance functions of the lognormal process (e.g., Aitchison and Brown, 1957), while the third-order spatial cumulant function can be obtained

by repeated application of the standard expression used to find the moments of a multivariate lognormal distribution (e.g., Kotz et al., 2000, p. 219):

$$\begin{aligned} \kappa_{Y_1 Y_1 Y_1}^3(\mathbf{u}_1, \mathbf{u}_2, \mathbf{u}_3) = & \kappa_{Y_1}^1(\mathbf{u}_1) \kappa_{Y_1}^1(\mathbf{u}_2) \kappa_{Y_1}^1(\mathbf{u}_3) \Big(\exp(\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2)) + \tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_3) + \tilde{C}_{Y_1 Y_1}(\mathbf{u}_2, \mathbf{u}_3) \\ & - \exp(\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_2)) - \exp(\tilde{C}_{Y_1 Y_1}(\mathbf{u}_1, \mathbf{u}_3)) - \exp(\tilde{C}_{Y_1 Y_1}(\mathbf{u}_2, \mathbf{u}_3)) + 2 \Big). \end{aligned}$$

We use these expressions to illustrate second-order and third-order properties of the joint process $(Y_1(\cdot), Y_2(\cdot))$ in a simple scenario with $\zeta(s) = 0; s \in D = [-10, 10] \subset \mathbb{R}$ (i.e., $d = 1$). Let $Y_1(\cdot)$ be a lognormal process on D , with $\ln Y_1(\cdot)$ having mean function $\tilde{\mu}_1(s) = -2$ and covariance function

$$\tilde{C}_{Y_1 Y_1}(u_1, u_2 \mid \boldsymbol{\theta}_1) = \frac{1}{\tau_1} \exp\left(-\theta_{11}|u_1 - u_2|^{\theta_{12}}\right), \quad (\text{B.1})$$

where $\boldsymbol{\theta}_1 \equiv (\theta_{11}, \theta_{12})'$, and $\theta_{11} > 0, 0 < \theta_{12} < 2$. Here, we set the scale parameter $\theta_{11} = 0.8$, the smoothness parameter $\theta_{12} = 1.7$, and the precision parameter $\tau_1 = 1$. We let $b(s, u); s, u \in D$, be a truncated Gaussian density function centred and truncated at s with variance varying smoothly with u , representing a typical highly directional interaction function in u ; see Fig. B.7, top-left panel. We now concentrate on the spatial location $s = 0$, and we carry out the integrations required to find the spatial cumulant functions, such as those in (8)–(11), on a discrete grid $D^L = \{-9.9, -9.7, \dots, 9.9\}$. In the bottom-left, top-right, and bottom-right panels of Fig. B.7, we show the second-order cross-cumulant function (i.e., cross-covariance function) $\kappa_{Y_2 Y_1}^2(0, u_2)$, the third-order cross-cumulant function $\kappa_{Y_2 Y_1 Y_1}^3(0, u_2, u_3)$, and the third-order auto-cumulant function $\kappa_{Y_2 Y_2 Y_2}^3(0, s_2, s_3)$, respectively, for $u_2, u_3, s_2, s_3 \in D^L$. It is clear from these panels that asymmetry is naturally present in the cross-covariance functions and the third-order cross-cumulant functions. Also, the two third-order cross-cumulant functions are markedly different from each other; in contrast, if $Y_1(\cdot)$ were Gaussian the third-order cumulants would have been zero everywhere. Of course, $Y_1(\cdot)$ need not be lognormal, but the spatial cumulant functions of $Y_2(\cdot)$ and those of the joint process $(Y_1(\cdot), Y_2(\cdot))$ depend critically on the spatial cumulant functions of $Y_1(\cdot)$.

Fig. B.7: Properties of a non-Gaussian process constructed from (4) with $\zeta(\cdot) = 0$ and $Y_1(\cdot)$ a lognormal process. Top-left panel: The interaction function $b(s, u); s, u \in D^L$. Bottom-left panel: The interaction function $b(0, u_2); u_2 \in D^L$, and the second-order cross-cumulant function $\kappa^2_{Y_2 Y_1}(0, u_2); u_2 \in D^L$. Top-right panel: The third-order cross-cumulant function $\kappa^3_{Y_2 Y_1 Y_1}(0, u_2, u_3); u_2, u_3 \in D^L$. Bottom-right panel: The third-order auto-cumulant function $\kappa^3_{Y_2 Y_2 Y_2}(0, s_2, s_3); s_2, s_3 \in D^L$. In all panels, the diamond symbol denotes the origin $u_2 = u_3 = s_2 = s_3 = 0$.

Appendix C. Derivation of conditional distributions

Appendix C.1. The discrepancy parameters

Consider the conditional distribution, $p(\tau_2, \boldsymbol{\theta}_2 \mid \mathbf{Z}_2, \mathbf{Y}_1) \propto p(\mathbf{Z}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1)p(\tau_2)p(\boldsymbol{\theta}_2)$. This is given by

$$\begin{aligned}
 p(\tau_2, \boldsymbol{\theta}_2 \mid \mathbf{Z}_2, \mathbf{Y}_1) &\propto \left(\int p(\mathbf{Z}_2 \mid \mathbf{Y}_2)p(\mathbf{Y}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1)d\mathbf{Y}_2 \right) p(\tau_2)p(\boldsymbol{\theta}_2) \\
 &\propto |\mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}|^{1/2} \left[\int \exp \left(-\frac{1}{2}(\mathbf{Z}_2 - \mathbf{C}\mathbf{Y}_2)' \mathbf{V}^{-1}(\mathbf{Z}_2 - \mathbf{C}\mathbf{Y}_2) \right. \right. \\
 &\quad \left. \left. - \frac{1}{2}(\mathbf{Y}_2 - \mathbf{B}\mathbf{Y}_1)' \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}(\mathbf{Y}_2 - \mathbf{B}\mathbf{Y}_1) \right) d\mathbf{Y}_2 \right] p(\tau_2)p(\boldsymbol{\theta}_2) \\
 &\propto \frac{p(\tau_2)p(\boldsymbol{\theta}_2)|\mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}|^{1/2}}{|\mathbf{C}'\mathbf{V}^{-1}\mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}|^{1/2}} \\
 &\quad \times \exp \left(-\frac{1}{2}\mathbf{Y}_1' \mathbf{B}' \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} \mathbf{B} \mathbf{Y}_1 + \frac{1}{2}\mathbf{D}'_{\boldsymbol{\theta}_2}(\mathbf{C}'\mathbf{V}^{-1}\mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2})^{-1} \mathbf{D}_{\boldsymbol{\theta}_2} \right), \tag{C.1}
 \end{aligned}$$

where $\mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} \equiv \boldsymbol{\Sigma}_{\zeta; \tau_2, \boldsymbol{\theta}_2}^{-1}$, $\mathbf{C} \equiv \text{bdiag}(\{\mathbf{C}_t : t \in \mathcal{T}\})$, $\mathbf{V} \equiv \text{bdiag}(\{\mathbf{V}_t : t \in \mathcal{T}\})$, $\mathbf{B} \equiv (\mathbf{B}'_t : t \in \mathcal{T})'$ and

$$\mathbf{D}_{\boldsymbol{\theta}_2} \equiv (\mathbf{C}'\mathbf{V}^{-1}\mathbf{Z}_2 + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}\mathbf{B}\mathbf{Y}_1).$$

The prior distributions $p(\tau_2)$ and $p(\boldsymbol{\theta}_2)$ are bounded uniform distributions, with limits as described in Section 5.

Appendix C.2. The flux field

It is straightforward to show that using the Pericchi (1981) prior distribution under the assumption that the transformed field is multivariate normal, $p(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda)$ is multivariate Normal-Gamma (O'Hagan, 2000, Chapter 9); that is,

$$(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda) \sim \mathcal{NGa}\left(\hat{\boldsymbol{\beta}}_{\boldsymbol{\theta}_1, \lambda}, \underline{\mathbf{X}}'\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1}\underline{\mathbf{X}}, |D_1^L|, \frac{S_{\boldsymbol{\theta}_1, \lambda}^2}{2}\right), \quad (\text{C.2})$$

where $\mathcal{NGa}(\boldsymbol{\mu}, \mathbf{Q}, \alpha_1, \alpha_2)$ denotes the multivariate Normal-Gamma distribution with location $\boldsymbol{\mu}$, scale $\mathbf{Q}$, shape α_1 and rate α_2 . Writing out

$$p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\theta}_1, \lambda) = \frac{p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)p(\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1 \mid \lambda)p(\lambda)}{p(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda)p(\boldsymbol{\theta}_1, \lambda)}, \quad (\text{C.3})$$

we deduce that

$$p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\theta}_1, \lambda) \propto \left(\frac{S_{\boldsymbol{\theta}_1, \lambda}^2}{2}\right)^{-|D_1^L|} \underline{J}_\lambda. \quad (\text{C.4})$$

Note that $(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\theta}_1, \lambda)$ depends on $\underline{\mathbf{Y}}_1$ through $S_{\boldsymbol{\theta}_1, \lambda}^2$ as in normal-linear models (O'Hagan, 2000, Section 9.29) and, in our case, also on the Jacobian, $\underline{J}_\lambda$.

Now, the required full conditional distribution is

$$p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \mathbf{W}_1, \boldsymbol{\theta}_2, \tau_2, \boldsymbol{\theta}_1, \lambda) \propto p(\mathbf{Z}_2 \mid \tau_2, \boldsymbol{\theta}_2, \mathbf{Y}_1)p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\theta}_1, \lambda). \quad (\text{C.5})$$

The second term in the product on the right-hand side of (C.5) is given by (C.4), while the first

term is obtained through marginalisation of $\mathbf{Y}_2$, as in (C.1). Hence,

$$\begin{aligned}
& \ln p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \mathbf{W}_1, \boldsymbol{\theta}_2, \tau_2, \boldsymbol{\theta}_1, \lambda) \\
&= \text{const} - \frac{1}{2} \mathbf{Y}_1' \mathbf{B}' (\mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} - \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} (\mathbf{C}' \mathbf{V}^{-1} \mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2})^{-1} \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}) \mathbf{B} \mathbf{Y}_1 \\
&\quad + \mathbf{Y}_1' \mathbf{B}' \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} (\mathbf{C}' \mathbf{V}^{-1} \mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2})^{-1} \mathbf{C}' \mathbf{V}^{-1} \mathbf{Z}_2 \\
&\quad - |D_1^L| \ln \left(\frac{S_{\boldsymbol{\theta}_1, \lambda}^2}{2} \right) + \ln J_\lambda,
\end{aligned} \tag{C.6}$$

where ‘const’ collects all the terms that are not a function of $\mathbf{Y}_1$.

For Hamiltonian Monte Carlo, we also need the derivative of (C.6). By substituting (22) into (21) we can re-write $S_{\boldsymbol{\theta}_1, \lambda}^2$ as $S_{\boldsymbol{\theta}_1, \lambda}^2 = \underline{\mathbf{G}}_\lambda' \boldsymbol{\Psi}_{\boldsymbol{\theta}_1} \underline{\mathbf{G}}_\lambda$, where

$$\boldsymbol{\Psi}_{\boldsymbol{\theta}_1} = \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} - \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} \mathbf{X} (\mathbf{X}' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} \mathbf{X})^{-1} \mathbf{X}' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1}.$$

Then, the required derivative is

$$\begin{aligned}
\frac{\partial \ln p(\mathbf{Y}_1 \mid \mathbf{Z}_2, \mathbf{W}_1, \boldsymbol{\theta}_2, \tau_2, \boldsymbol{\theta}_1, \lambda)}{\partial \mathbf{Y}_1'} &= -\mathbf{B}' (\mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} - \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} (\mathbf{C}' \mathbf{V}^{-1} \mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2})^{-1} \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2}) \mathbf{B} \mathbf{Y}_1 \\
&\quad + \mathbf{B}' \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2} (\mathbf{C}' \mathbf{V}^{-1} \mathbf{C} + \mathbf{Q}_{\zeta; \tau_2, \boldsymbol{\theta}_2})^{-1} \mathbf{C}' \mathbf{V}^{-1} \mathbf{Z}_2 \\
&\quad - \frac{2|D_1^L|}{S_{\boldsymbol{\theta}_1, \lambda}^2} \frac{\partial \underline{\mathbf{G}}_\lambda'}{\partial \mathbf{Y}_1'} \boldsymbol{\Psi}_{\boldsymbol{\theta}_1} \underline{\mathbf{G}}_\lambda + \mathbf{G}_{\lambda, 1}^{(2)} \oslash \mathbf{G}_{\lambda, 1}^{(1)},
\end{aligned}$$

where $\mathbf{G}_{\lambda, 1}^{(1)} \equiv (\partial g_\lambda(Y_{1,i}) / \partial Y_{1,i} : i = 1, \dots, |D_1^L|)$, $\mathbf{G}_{\lambda, 1}^{(2)} \equiv (\partial^2 g_\lambda(Y_{1,i}) / \partial Y_{1,i}^2 : i = 1, \dots, |D_1^L|)$, the operator $\oslash$ denotes element-wise division,

$$\frac{\partial \underline{\mathbf{G}}_\lambda'}{\partial \mathbf{Y}_1'} = \begin{pmatrix} \frac{\partial \mathbf{G}_{\lambda, 1}'}{\partial \mathbf{Y}_1'} & \mathbf{0} \end{pmatrix} = \begin{pmatrix} \text{diag}(\mathbf{G}_{\lambda, 1}^{(1)}) & \mathbf{0} \end{pmatrix}, \tag{C.7}$$

and we have used the fact that

$$\frac{\partial S_{\boldsymbol{\theta}_1, \lambda}^2}{\partial \mathbf{Y}_1'} = \frac{\partial \underline{\mathbf{G}}_\lambda' \boldsymbol{\Psi}_{\boldsymbol{\theta}_1} \underline{\mathbf{G}}_\lambda}{\partial \mathbf{Y}_1'} = 2 \left(\frac{\partial \underline{\mathbf{G}}_\lambda'}{\partial \mathbf{Y}_1'} \right) \boldsymbol{\Psi}_{\boldsymbol{\theta}_1} \underline{\mathbf{G}}_\lambda. \tag{C.8}$$

Appendix C.3. The flux-field parameters

From (C.3), we have

$$p(\boldsymbol{\theta}_1, \lambda \mid \underline{\mathbf{Y}}_1) = \frac{p(\underline{\mathbf{Y}}_1 \mid \boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1, \lambda)}{p(\boldsymbol{\beta}, \tau_1 \mid \underline{\mathbf{Y}}_1, \boldsymbol{\theta}_1, \lambda)} \frac{p(\boldsymbol{\beta}, \tau_1, \boldsymbol{\theta}_1 \mid \lambda) p(\lambda)}{p(\underline{\mathbf{Y}}_1)},$$

from which we deduce that

$$p(\boldsymbol{\theta}_1, \lambda \mid \underline{\mathbf{Y}}_1) \propto |\tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}|^{-1/2} |\underline{\mathbf{X}}' \tilde{\mathbf{R}}_{Y_1 Y_1; \boldsymbol{\theta}_1}^{-1} \underline{\mathbf{X}}|^{-1/2} (S_{\boldsymbol{\theta}_1, \lambda}^2)^{-|D_1^L|} \underline{J}_\lambda p(\boldsymbol{\theta}_1) p(\lambda). \quad (\text{C.9})$$

This is similar to Equation 8 in De Oliveira et al. (1997); however they used the prior distribution given by Box and Cox (1964) instead, which results in different powers of $S_{\boldsymbol{\theta}_1, \lambda}^2$ and $\underline{J}_\lambda$ in (C.9).