

Interactive comment on “Bayesian spatiotemporal inference of trace gas emissions using an integrated nested Laplacian approximation and Gaussian Markov random fields” by Luke M. Western et al.

Luke M. Western et al.

luke.western@bristol.ac.uk

Received and published: 14 October 2019

We would like to thank the reviewers for their helpful and supportive comments on our discussion article ‘Bayesian spatiotemporal inference of trace gas emissions using an integrated nested Laplacian approximation and Gaussian Markov random fields’ (gmd-2019-66). The revised manuscript, along with new supplementary figures, are attached with revised sections indicated in red. Below we address each reviewer’s comments in turn, where the comment has been italicised.

C1

Reviewer 1

Page 3: In Section 2.1 the linear model for the spatial field is briefly described and derived. This section can be improved without loss of generality by starting from equation 3 that contains all relevant quantities of interest expressed as deviations from prior mean values. In addition, the physics behind factor x or some interpretation about its connection with the measurements could be added in this section to help the reader better understand the model.

We have addressed this comment by reworking most of section 2.1 (P3 L4-23). A physical interpretation of x has been added and the linear model starts from the previous equation 3 (now equation 1), which removes the need for the tilde notation in the remainder of the manuscript. The main additions are:

(L6-8) ‘The a priori emissions at given location are generally informed by spatially resolved bottom-up inventories (as in section 3.1.3) or extrapolation from some reported emissions value. The value x is a multiplicative scaling of this a priori value for emissions, most generally expressed as a quantity of gas per unit time per unit area.’

(L13-23) ‘This is in an effort to fit the physical model to the imposed statistical model that fast computation requires. Taking the approach that, for many regional inverse problems involving longer-lived trace gases, there is a linear relationship between emissions which are constant in time and observed atmospheric concentration, the relationship between measurements and emissions is,

$$\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{K}\mathbf{u} + \epsilon, \quad (1)$$

where \mathbf{y} is a vector of the residual between the measured and a priori predicted measurement, \mathbf{H} is a Jacobian (or sensitivity) matrix, which maps the surface emissions to the measurements, \mathbf{u} is a vector of independent and identically distributed variables containing the contribution to the measurement of mole fractions

C2

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

Page 4, Line 2: The solution to the stochastic differential is a stationary Gaussian Field with Matern covariance structure. The fact that the process is stationary should be mentioned. Moreover, if the covariance function does not depend on the direction but just on the Euclidean distances between s_i and s_j , then it should also be mentioned that the process is isotropic.

The stationarity has been mentioned (P3 L28) as well as elaboration of the process being isotropic, with further reference to information on non-stationary and anisotropic process (P4 L3-4) '...noting that the dependence is only on the Euclidean distance and so the process is isotropic. For an example of non-stationary and anisotropic fields see e.g. Marques et al. (2019).'

Page 5, equation 6: The indices and ranges should be explained. For instance, $i=1,\dots,n$. Does n represent the number of nodes in the mesh? Also, index j goes from 1 to 4. Does this represent the sides of the rectangular region?

This notation has now been clarified (P5 L6) ' i represents each of the total n nodes and j represents the edge of the domain at each of the four cardinal directions'.

Page 7, Line 15: Is there any advantage regarding the computational cost or convergence for using penalized complexity priors instead of a vague prior?

C3

Penalised priors do not increase the computational speed for the given case and are chosen instead for their intuitiveness. This has been added to the manuscript (P7 L 19).

Pages 10-13, Section 3.2: More information about the simulation study should be included in the main text or supplementary materials. Would be nice to see the relationship between the true parameters from the simulated data and the estimated values along with their corresponding uncertainty.

We have added a discussion on the estimated hyperparameters for the simulated data (P14 L4-9), 'Hyperparameter estimation is less accurate than for the latent field. The estimation of the noise σ_y generally captures the imposed noise well, which had a mean value of 5.5 ppm, with an estimated value of 6.6 [6.5, 6.8] ppm. The correlation structures are less well captured. The temporal correlation ϕ was estimated with a mean value 0.7 [0.6, 0.8], the range ρ has value 1.7 [1.3, 2.1] and the marginal standard deviation of the latent field σ with the value 0.8 [0.7,0.9]. It is promising that the posterior mean estimates of all hyperparameters show an improvement on their prior mean or baseline values, although only the true value for ρ falls within the estimated 95% uncertainty.'

Also, would be interesting to complement the posterior mean estimate plots as function of time with analogous plots showing some measure of the uncertainty on the estimates (e.g. SE or credible set).

These plots, showing the 95% uncertainty, have been added as supplementary plots (P25-26).

Last, the value of this manuscript would be greatly improved if the R code used for the

C4

numerical experiment could be shared. However, if the complexity and scale of the problem makes this option not feasible, then at least a toy example to illustrate the concept would be useful.

We have made the code and data for the simulated data example freely available via the Open Science Foundation (osf.io/53w96, DOI 10.17605/OSF.IO/53W96), which has been noted in the Code and data availability section of the manuscript.

Pages 15-16: The discussion section could mention some limitations of the method. First, the stationarity assumption may not be suitable when modelling some environmental phenomena. Particularly in atmospheric phenomena it may inappropriate to assume that the spatial correlation is the same throughout the domain as topographical variables (e.g. mountains, lakes, etc.) might have an influence on the spatial dependence. Second, INLA method relies on the assumption of approximate multivariate normality of the posterior linear predictors.

We have added additional discussion on these topics in section 4, in particular P18, L16-20. 'The problem of non-stationary covariances may also be present for inference of on other spatial scales. For example natural features, such as lakes, may cause abrupt changes in correlation structures. Emissions of anthropogenic greenhouse gases may exhibit no spatial correlation structure (e.g. Mühle et al., 2019; Rigby et al., 2019). In this case a Matérn field would be inappropriate, but an autoregressive model may still be applicable for time-varying emissions.'

Page 4, Line 19: provided instead of providing.

This has been changed.

Page 6, Line 12: consider rephrasing: ' . . . which means that the spatiotemporal

C5

precision matrix can be expressed using a Kronecker product as follows...', and delete Line 15.

This has been changed.

Page 6, Line 17: burden instead of size.

This has been changed.

Page 6, Line 19: consider rephrasing: '... constant over a three-month period ...'

This has been changed.

Page 6, Line 19-20: Later in the text you use four three-month periods, not three.

This has been changed to three-month

Page 7, Line 10: consider rephrasing '... or the spatial precision matrix $Q(\theta)^{-1}$, are ...'

This has been changed.

Page 16, Line 13: rephrase to '... correlation structure may vary in space.

This has been changed.

Reviewer 2

P3, L9: '... which maps the surface emissions (x) to the measurements ...'. The variable x was not defined here.

This is now defined in section 2.1, L7-8. 'The value x is a multiplicative scaling of this a priori value for emissions, most generally expressed as a quantity of gas per unit time per unit area.'

P3, L16: I do not understand the discussion about the change of variable in x (which is simply a shift in the distribution), and in particular how this relates to the positiveness of the fluxes. It seems one issue here could be that a log-normal distribution may

C6

be more suitable than a Gaussian one (if positiveness needs to be enforced). Apart from a re-centering of the prior pdf around zero, I do not see what the transformation achieves. Please clarify.

The use of GMRFs are a requirement for the computational efficiency of the method. Therefore, we have attempted to fit a statistic model to the physical model. We have clarified this point in the reworked section 2.1, in particular P3 L13. 'This is in an effort to fit the physical model to the imposed statistical model that fast computation requires.'

P6, L1: '*...measurements made at a given time are independent ...*'. Would it be easy to generalize the approach for spatially correlated observations at a given time? This, for instance, would be useful for inversions based on satellite measurements.

This could be generalised but this correlation would have to result in a sparse covariance matrix (e.g. AR1, Matern GMRF) to keep the computational gains. We have noted this on P6 L1-2. 'This can be generalised to include correlated measurements, where the covariance function should result in a sparse matrix to retain computational efficiency.'

P6, L6-9: Do you mean 'emissions' instead of 'measurements' here? You mention measurements (y) but then refer to emissions (x) in the equation. Please clarify.

Thank you for noticing this. The sentence has now been reworded to 'The time varying structure of \mathbf{H}_t and \mathbf{x} applies also to \mathbf{K}_t , and \mathbf{u}_t .' We have elaborated on the use of an autoregressive model and the parameter ϕ on P6 L11-14. 'We have chosen an autoregressive model of order one as we believe that emissions are generally similar to those at previous time step. In the given model, ϕ is the correlation between the

C7

previous time step and the current time step. This means that the emissions at time t have a similarity of ϕ to emissions at $t - 1$ plus some spatially correlated random effect δ_t .'

Section 2.5: This section should be entirely rewritten and much more details have to be provided. For instance, please explain the posterior sampling approach adopted here as well as the principle of the integrated nested Laplacian approximation used. This seems to be the core of the methodological contribution of this paper.

The section 2.5, P6, has been entirely rewritten with much more detail about the posterior sampling approach, outlining the principles of INLA.

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

$$(\theta_j | \mathbf{y}) = \int (\theta | \mathbf{y}) d\theta_{-j}, \quad (2)$$

$$(x_i | \mathbf{y}) = \int (x_i | \mathbf{y}, \theta) (\theta | \mathbf{y}) d\theta. \quad (3)$$

Equations 2 and 2 make use of a Laplace approximation, by approximating $(\theta | \mathbf{y})$ by

$$(\theta | \mathbf{y}) \propto \frac{G(\mathbf{x}, \theta, \mathbf{y})}{G(\mathbf{x} | \theta, \mathbf{y})} \Big|_{\mathbf{x}=\mathbf{x}^*(\theta)}, \quad (4)$$

where $G(\mathbf{x} | \theta, \mathbf{y})$ is a Gaussian approximation to the full conditional of \mathbf{x} and $\mathbf{x}^*(\theta)$ is the modal probability of \mathbf{x} for a given θ . This approximation is exact if $(\theta | \mathbf{y})$ is Gaussian,

C8

and gives a good approximation for log concave problems (Tierney and Kadane, 1986). Then, it is possible to approximate $(x_i | \mathbf{y}, \theta)$ using another Laplace approximation

$$LA(x_i | \theta, \mathbf{y}) \propto \frac{(\mathbf{x}, \theta, \mathbf{y})}{G(\mathbf{x}_{-i} | x_i, \theta, \mathbf{y})} \Big|_{\mathbf{x}_{-i} = \mathbf{x}_{-i}^*(x_i, \theta)}, \quad (5)$$

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

While this method relies on calculating only the marginal posterior distributions of x_i , it is still possible to predict a linear combination of the field to provide regional emissions totals (e.g. country totals). We define a linear predictor of emissions for a given region η^* , defined using the basis function representation of the mesh where each node contains information on the spatial area represented by that node and its connecting vertices and zeros for all nodes that are not in the region of interest. Then we can approximate the linear combination of the parameters of interest, giving a combined emissions total η^* , using equations 12 and 13 and transforming the predicted latent field by the weightings containing the area information.

Please also note the supplement to this comment:

<https://www.geosci-model-dev-discuss.net/gmd-2019-66/gmd-2019-66-AC1-supplement.pdf>

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2019-66>, 2019.