4DModeller (fdmr): A Comprehensive R Package for Spatio-Temporal Modelling *

Xueqing Yin[1[0000−0003−1103−8939]], John M. Aiken[2,3[0000−0003−2258−3836]], and Jonathan L. Bamber[1,4[0000−0002−2280−2819]]

[1] School of Geographical Sciences, University of Bristol, Bristol, UK
`de22002@bristol.ac.uk, J.Bamber@bristol.ac.uk`
[2] Expert Analytics, Oslo, Norway
`john@xal.no`
[3] Njord Centre, Departments of Physics and Geosciences, University of Oslo, Oslo, Norway
[4] Department of Aerospace and Geodesy, Technical University of Munich, Munich, Germany

Abstract. Spatio-temporal data analysis is crucial in many research fields. However, modelling large-scale spatio-temporal data presents challenges such as high computational demands, complex correlation structures, and the separation of mixed sources. To address these issues, we are developing 4DModeller (fdmr), a robust and user-friendly R package designed to model spatio-temporal data within a Bayesian framework. The software package offers a comprehensive solution for visualizing, analyzing and modelling different types of spatio-temporal data in various disciplines. By incorporating Bayesian hierarchical models, "fdmr" allows for the flexible integration of prior knowledge and data uncertainty into the modelling process. By utilizing the Integrated Nested Laplace Approximations (INLA) algorithm and the stochastic partial differential equations (SPDE) method for model inference, "fdmr" significantly reduces the computational complexity of handling high-resolution, high-dimensional spatio-temporal data. Furthermore, "fdmr" provides intuitive and interactive visual analytics tools that facilitate the exploration of data patterns across both space and time. This paper aims to introduce the "fdmr" package, and outline its core modelling framework through an example study on the spread of COVID-19 infection rates in England from 19 December, 2020 to 20 March, 2021.

Keywords: Spatio-temporal data analysis · fdmr package· Bayesian modelling framework.

DOI: https://doi.org/10.25436/E27C7F

* Supported by UKRI grant EP/X022641/1, and German BMBF grant 01DD20001.

1 Introduction

Spatio-temporal data analysis plays a crucial role across a wide range of research domains, such as earth and environmental sciences, ecosystem dynamics, epidemiology, crime analysis, and urban studies. It enables researchers and decision-makers to uncover hidden spatio-temporal patterns, identify spatial and temporal dependencies within the data, investigate the complex interplay between spatial and temporal controls, and make error-bounded predictions. With the increasing availability of large-scale spatio-temporal datasets, the need for advanced methodologies and tools to effectively handle and analyze such data has become paramount. However, the modelling of large-scale spatio-temporal data poses many challenges, such as intensive computational burden, complex spatio-temporal correlation structures, and the separation of mixed sources into their component processes. Furthermore, spatio-temporal modelling packages often require sophisticated statistical and discipline based knowledge and experience, which creates a further hurdle for users. To overcome these challenges, our team is currently engaged in the development of a robust and user-friendly R package known as "fdmr", which stands for "4D-Modeller". The "fdmr" package is designed to tackle spatio-temporal inference problems that occur in various disciplines, spanning from regional to global scales. This package extends the Bayesian hierarchical model developed during a successful ERC-funded project called GlobalMass (www.globamass.eu), where it was able to reliably handle large volumes of both satellite and in-situ data, and accurately decompose signals from the corresponding time series into the underlying physical processes that drive sea level rise [28]. The objective of the "fdmr" package is to provide researchers and practitioners a straightforward and efficient solution for handling and analyzing different types of spatio-temporal data, enabling a comprehensive analysis of the underlying spatio-temporal patterns and processes. There are three basic types of spatial data based on the characteristics of the research domain, including the geostatistical (point-referenced) data, areal data, and point patterns data [6]. The current development of the "fdmr" package supports geostatistical data and areal data over multiple time points. In geostatistical data, the study domain is a continuous fixed set, and the spatial data consist of observations measured at many fixed and precise locations. Examples of these type of data are air pollution, rainfall values or GPS ground displacements measured at several monitoring stations. In areal data, the domain is fixed (of regular or irregular shape) and partitioned into a finite number of areal units (e.g., polygons) with well-defined boundaries. Examples of areal data are attributes collected by ZIP code, census tract [17] or satellite imagery.

There are several existing R packages on CRAN for analyzing spatial or spatio-temporal data, either in a point-referenced or areal format. For point-referenced data, they include "fields" [18], "FRK" [33], "geoR" [21], "gstat" [19], "R-INLA" [14], and "spBayes" [9]. For areal data, they include "brms" [4], "CARBayes" [12], "CARBayes" [12], "CARBayesST" [13] and "hglm" [22]. These packages feature different degrees of flexibility, model inference methods as well as computational capacity. Unlike these aforementioned packages, the "fdmr" package is designed to analyze both point-referenced and areal data using a common Bayesian framework and syntax structure, and it is able to accommodate various response types, such as Gaussian, Poisson, and binomial data. "fdmr" also offers many features absent in the aforementioned packages - together within a single package, such as intuitive tools for mesh building and priors setting, interactive plotting tools for visualising and mapping the data in an intuitive manner, the option to use different model inference methods and high computational efficiency for "Big Data" problems. Thus, the "fdmr" package streamlines and simplifies the workflow, allowing users to use its comprehensive functionality without the need to rely on multiple packages.

In the "fdmr" package, spatio-temporal data are modelled within a Bayesian framework using Bayesian hierarchical models, which are a powerful tool for understanding intricate processes in space and time. These models integrate prior information and data observations to update process models, providing a posterior probabilistic understanding of the underlying spatio-temporal dynamics in the observed data. Additionally, the Bayesian hierarchical model accounts for spatial-temporal autocorrelation in the data by modelling a set of random effects via a spatio-temporal continuous Gaussian random field process. In comparison to existing frequentist statistical models for spatio-temporal analysis, such as the geographically weighted regression [3], spatial error model [8] and spatial lag model [30], our Bayesian modelling approach has several advantages in terms of computational accuracy and efficiency. Firstly, it allows for the incorporation of prior knowledge and a more robust assessment of the uncertainties in predictions by specifying prior distributions through a hierarchical modelling scheme [11, 32, 7]. Secondly, it is able to accommodate missing data by borrowing information from nearby spatial locations and time points based on the estimated spatio-temporal dependence structure, which enables predictions at any spatial location and time point, even in the case of missing data. Finally, our model inference is carried out using the INLA [26, 27] and SPDE method [15] via the R-inlabru package, which considerably reduces the computational burdens of handling high-resolution spatio-temporal data compared to existing packages [9, 12] that use Markov Chain Monte Carlo simulation to make model inference. Since the INLA-SPDE approach involves a finite element solution using a triangulated mesh of the study region that the processes are estimated on, the "fdmr" package also allows users to generate and customize a mesh based on their expert knowledge of the process they wish to infer. For example, high resolution for point sources of air pollution within urban areas and lower resolution in unpopulated regions. Another notable feature of "fdmr" is its functionality in data visualization. The package includes a set of intuitive and interactive visual analytics functions that enable users to visualize the data patterns across space and time, which aid in identifying spatial clusters, temporal trends, and anomalies. We have applied the package to several test case problems related to COVID-19 transmission, hydropower generation in Norway, and lake growth in the Tibetan Plateau. Thes examples were chosen to demonstrate the capabilities for different classes of problem and data. Detailed instructions on installing the "fdmr" package in R [20] and a list of worked-out tutorials can be found at `https://4dmodeller.github.io/fdmr/index.html`. This list is dynamic and being continuously added to as the user community grows.

This paper serves as a brief introduction to the "fdmr" package. Section 2 introduces the fundamental modelling framework utilized through the package, by presenting a case study that examines the spread of COVID-19 infection rate in England using the "fdmr". Section 3 summarises the primary application results, while the paper ends with a discussion in Section 4.

2 Spatio-temporal modelling framework

In this section, we demonstrate the spatio-temporal modelling framework employed in the package "fdmr" in the context of investigating the spread of COVID-19 infection rates in mainland England during the period from 19 December, 2020 to 20 March, 2021. Fig. 1 displays our study region, which is partitioned into $n = 6789$ neighbourhood units called Middle Layer Super Output Areas (MSOAs).

Fig. 1: A map of the study region. This map is made using the fdmr plotting function **plot_map()**.

The weekly number of reported COVID-19 cases between 19 December, 2020 and 20 March, 2021 for each MSOA was obtained from the UK COVID-19 dashboard (`https://coronavirus.data.gov.uk/details/download`). The aim of this study is to predict the COVID-19 infection rates across England and identify the temporal trends as weel as dominant risk factors. An infection rate is the probability or risk of an infection in a population, which in epidemiology is defined as the proportion of affected people in a population during a specific time period [24]. Here it is calculated as the ratio of the number of reported COVID-19 cases to the total population in each MSOA and week. However, naively using such infection rate data as a measure of the variations in COVID-19 transmission ignores the spatio-temporal autocorrelation inherent in the disease's dynamics. Moreover, it ignores the potential effects of risk factors on COVID-19 infections. Therefore, it is necessary to develop a model-based approach that can capture the spatio-temporal variations in disease spread, separate the variations from random noise and account for the spatio-temporal correlation structure in the data. To achieve this, we have developed a Bayesian hierarchical spatio-temporal modelling approach within the "fdmr" package, which is outlined in Section 2.1.

2.1 Bayesian hierarchical model

This section outlines the modelling specification in the context of analyzing the COVID-19 infections data described above.

Level 1 - Data likelihood model

Let $Y(\boldsymbol{s}_i, t)$ and $N(\boldsymbol{s}_i, t)$ be the number of reported COVID-19 cases and total population in MSOA $i \in (1, \ldots, n_t)$ during week $t \in (1, \ldots, T)$, respectively. Here $\boldsymbol{s}_i \in R^2$ denotes the geographical location for MSOA i, and n_t represents the number of MSOAs that have reported COVID-19 cases during week t. As the response variable is a count, the first level of our Bayesian hierarchical model is the Poisson log-linear specification given by

$$\begin{aligned} Y(\boldsymbol{s}_i, t) &\sim \text{Poisson}\left(N(\boldsymbol{s}_i, t)\theta(\boldsymbol{s}_i, t)\right), \quad i = 1, \ldots, n_t; \quad t = 1, \ldots, T, \\ \ln(\theta(\boldsymbol{s}_i, t)) &= \boldsymbol{x}(\boldsymbol{s}_i, t)^\top \boldsymbol{\beta} + \xi(\boldsymbol{s}_i, t), \end{aligned} \tag{1}$$

where $\theta(\boldsymbol{s}_i, t)$ is the COVID-19 infection rate in MSOA i and week t, which is modelled by two components. The first component is the vector of covariates (if needed) given by $\boldsymbol{x}(\boldsymbol{s}_i, t)$ for location $\boldsymbol{s}_i$ and time t, including an intercept term, with a vector of regression parameters $\boldsymbol{\beta}$. The second component is the spatio-temporal random effect $\xi(\boldsymbol{s}_i, t)$. Note that although the model introduced above is used to model count data , the "fdmr" package also accommodates Gaussian and binomial type responses.

Level 2 - Spatio-temporal random effects model

The spatio-temporal random effect $\xi(\boldsymbol{s}_i, t)$ represents the realization of a spatio-temporal process for the logarithm of the observed COVID-19 infection rates after covariate adjustment. The spatial correlation is modelled by location specific random effects through a Gaussian random field process, which captures the correlation via a covariance matrix expressed as a function of distance between locations. The temporal correlation is modelled by a first-order autoregressive

(AR1) process. Specifically, we have

$$\begin{aligned} \xi(\boldsymbol{s}_j, t) &= \alpha \times \xi(\boldsymbol{s}_j, t-1) + \omega(\boldsymbol{s}_j, t), \\ \xi(\boldsymbol{s}_j, 1) &\sim \mathrm{N}\left(0, \frac{\sigma_\omega^2}{1-\alpha^2}\right). \end{aligned} \tag{2}$$

Here α is a temporal dependence parameter such that $|\alpha| < 1$. $\boldsymbol{\omega}(\boldsymbol{s}_j, t) = (\omega(\boldsymbol{s}_1, t), \ldots, \omega(\boldsymbol{s}_{n_t}, t))^\top$ is a spatial random effect that is assumed to follow a multivariate Gaussian distribution and have $\boldsymbol{\omega}(\boldsymbol{s}_j, t) \sim \mathrm{N}(\boldsymbol{0}_{n_t}, \sigma_\omega^2 \boldsymbol{\Sigma}_\omega)$, where $\boldsymbol{0}_{n_t}$ is a $n_t \times 1$ vector of zeros, σ_ω^2 is the marginal variance of the spatial process, and $\boldsymbol{\Sigma}_\omega$ is the $n_t \times n_t$ covariance matrix with elements

$$(\boldsymbol{\Sigma}_\omega)_{ij} = C(||\boldsymbol{s}_i - \boldsymbol{s}_j||),$$

where $||\boldsymbol{s}_i - \boldsymbol{s}_j||$ is the distance between locations $(\boldsymbol{s}_i, \boldsymbol{s}_j)$, and $C(\cdot)$ is the Matern function [31] given by

$$C(||\boldsymbol{s}_i - \boldsymbol{s}_j||) = \frac{1}{2^{\nu-1}\Gamma(\nu)}(\kappa||\boldsymbol{s}_i - \boldsymbol{s}_j||)^\nu K_\nu(\kappa||\boldsymbol{s}_i - \boldsymbol{s}_j||), \tag{3}$$

where $K_\nu(\cdot)$ is the modified Bessel function of second kind, and $\Gamma(\nu)$ is the Gamma function. ν is the smoothness parameter of the Matern covariance function. κ is a scaling parameter controlling the spatial correlation range ρ, which is the distance at which the correlation function has fallen to about 0.13 and is given by $\rho = \sqrt{8\nu}/\kappa$.

Level 3- Prior distributions

The regression parameters $\boldsymbol{\beta}$ are assigned independent weakly informative zero-mean Gaussian prior distributions with a large variance, i.e., $\beta_j \sim \mathrm{N}(0, 1000)$ for $j = 0, \ldots, p$, to ensure their values are mainly informed by data. Penalized complexity priors [10] are specified for the spatial correlation range ρ and the marginal standard deviation parameter σ_ω. Specifically, a prior of $p(\rho < v_\rho) = p_\rho$ is specified for ρ, which means that the probability that ρ is smaller than v_ρ is p_ρ. A prior of $p(\sigma_\omega > v_{\sigma_\omega}) = p_\sigma$ is specified for σ_ω, indicating that the probability of σ_ω being greater than v_{σ_ω} is p_σ. The temporal autoregressive parameter α is also assigned a penalized complexity prior, with $p(\alpha > 0) = 0.9$. Note that the "fdmr" package provides default priors for the model parameters, but it also allows users to specify a range of different types of prior distributions based on the expert information in the research field.

2.2 Mesh construction

The model is implemented through the INLA-SPDE approach, which approximates a continuous spatial process, i.e., a Gaussian random field process, with the Matern covariance function defined in equation (3), by a discretely indexed spatial random process known as a Gaussian Markov random field (GMRF). The GMRF has zero mean and uses a sparse precision matrix which substantially reduces the computational cost in matrix algebra operations compared to using dense covariance matrices [25]. To represent the Matern field as a GMRF, the SPDE approach discretizes the space by defining a mesh composed of non-intersecting triangles that partition the domain of the study region [15]. These triangles allow the spatial autocorrelation between data observations to be calculated in the modelling process. For example, Fig. 2 displays the mesh we created for England, and the points in the figure represent the locations of MSOAs with COVID-19 reported cases. Then the INLA algorithm estimates the posterior distribution of the latent Gaussian process and hyperparameters using the Laplace approximation [27]. More details on this methodology can be found in the work of Blangiardo & Cameletti (2015) [1].

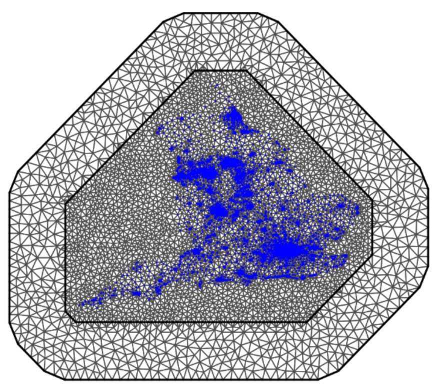

Fig. 2: A triangulated mesh over the study region.

The construction of a mesh has an impact on the performance of model inference and predictions. Therefore, it is important to develop a robust mesh that minimizes the sensitivity in the results. The choice of mesh varies depending on the case study. In general, high resolution meshes with large buffer regions tend to yield more accurate predictions [23], while such meshes may also increase computational costs. Lindgren and Rue (2015) [14] provided details on the construction of an optimal mesh. Our "fdmr" package will provide an intuitive mesh builder tool that allows users to interactively adjust the parameters related to mesh resolution and shape with ease, and then displays the resulting mesh in a few seconds in a user-friendly interface.

3 Results summary

The Bayesian spatio-temporal model outlined in Section 2.1 was fitted to the COVID-19 infection data in England. Although the model is outlined in its most general form that includes covariate information, covariates are not included in the COVID-19 application in this paper. Table 1 summarises the estimates for the spatial correlation range parameter ρ and temporal dependence parameter α from the model. It indicates that there is strong evidence of significant spatial and temporal autocorrelation in the spread of COVID-19 infections across England, with a spatial range of 0.295 and a high degree of temporal dependence of 0.836.

Our "fdmr" package provides automated plotting tools for time series visualization and spatial pattern mapping. For example, Fig. 3 illustrates the predicted COVID-19 infection rates for all MSOAs by week. The boxplots reveal the presence of health inequalities in COVID-19 infection rates across MSOAs in England, because there were substantial variations in the predicted infection rates, with values ranging between 0.0002 and 0.023 **(ADD UNITS)**. The average infection rate curve shows that there was a steady rise in recorded infection levels starting from 19 December 2020, which was driven by the emergence of the Alpha variant. This upward trend reached its peak in early January 2021, after which the recorded infection rate exhibited a gradual decrease until the end of the study. This reduction is result of the second national lockdown restriction implemented by the UK government, which limited the transmission of the virus. Fig. 4 displays the spatial pattern of the time averaged rates of infection in England. The high infection rates were mainly concentrated in the northeast of England, such as districts of Northumberland, Cumbria, and Lancashire, as well as in the southeast, particularly including London and some surrounding areas such as Kent, Essex, Surrey and Sussex. The goals here is not to present a detailed analysis of COVID infection dynamics but to illustrate the capabilities of fdmr for this example.

Table 1: Summary of the spatial correlation range parameter ρ and the temporal dependence parameter α.

Parameter	**mean**	**0.025quant**	**0.975quant**
Spatial range ρ	0.295	0.290	0.300
Temporal dependence α	0.836	0.832	0.839

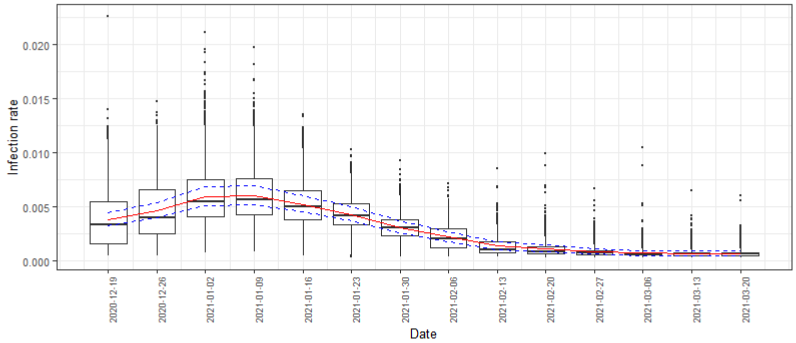

Fig. 3: Summary of the estimated COVID-19 infection rates across all MSOAs between 19 December, 2020 and 20 March, 2021. The red line displays the average estimated infection rate across mainland England in weeks and the dashed lines show the corresponding 95% credible intervals by week. This plot is made using the fdmr plotting functions **plot_boxplot()** and **plot_line_average()**.

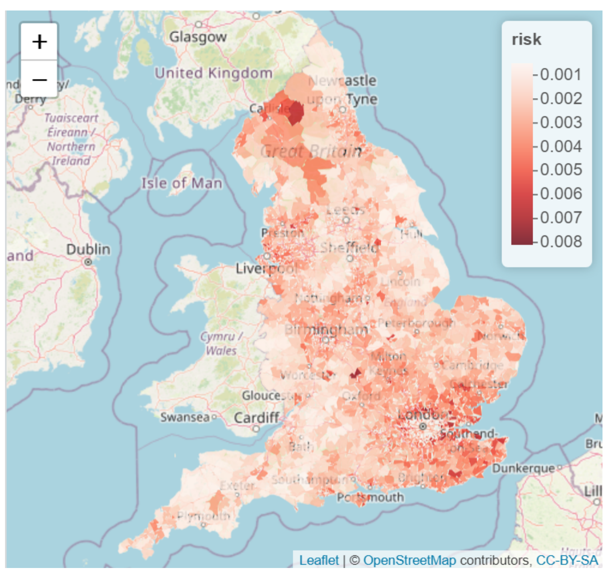

Fig. 4: Map of the average predicted infection rates in England over the study period. This map is made using the fdmr plotting function **plot_map()**.

4 Conclusion and discussion

In addition to the COVID-19 transmission case study outlined above, two other applications of the "fdmr" package include hydropower production in Norway, and the spatial dynamics of lakes on the Tibetan Plateau. Power in Norway comes primarily from hydropower stations throughout the country. Water runoff that feeds hydropower reservoirs is primarily due to snow melt. This snow melts in the spring, running into reservoirs that then feed hydropower stations. Gauging stations are placed throughout catchments (i.e., the area or region where water drains into a river), to monitor stream flow. The "fdmr" package is used to calculate the spatial and temporal flow rates per catchment, given precipitation data (ERA5-land precipitation [5]) as a fixed effect and an SPDE model which assumes there is some other correlation structure in the streamflow data due to unobserved variables, such as local elevation, temperature, soil permeability, or presence of vegetation. This study can provide data-driven expectations for reservoirs, aiding in effective hydropower management. In the Tibetan Plateau study, the "fdmr" package offers the capability to calculate lake change and investigate the underlying processes affecting the area of lakes. The Tibetan Plateau is a largely endorheic region, that is, water in this region does not travel outside the catchment. It is climatically linked to the larger High Mountain Asia region. It could be that glacier melt is the primary driver of lake growth in a high mountain context. However, glaciers seem to account for only a small amount of lake growth [2], while precipitation in the region seems to be increasing in the area [29], and ground thermal changes are likely also increasing permafrost thaw [16]. "fdmr" allows us to combine all these features into one model, thus providing insights into the dynamics of lake growth in the Tibetan Plateau and facilitating a better understanding of the region's hydrological system and its response to various environmental factors. New case studies are being added as the user community grows. A fdmr hackathon in November 2023 will further expand the portfolio of use cases and tutorials available.

The "fdmr" package aims to provide a robust and versatile toolkit for the exploration, analysis, and modelling of spatio-temporal data across diverse research domains. It has the benefits of supporting various data formats, efficiently handling large datasets, flexibility in specifying prior distributions, and accounting for data uncertainty during the modelling process. Furthermore, the package provides intuitive data visualization tools to aid in the interpretation of results. Despite being a work in progress, the package has demonstrated its efficacy and utility in several real-world applications. Future package development efforts will focus on expanding its capabilities and broadening its applicability. This includes, but is not limited to, automated mesh generation, the development of an R Shiny-based tool for customized meshes, the integration of additional model diagnostic tools for effective model evaluation and selection, an automated prior picking mechanism, and the enhancement of flexible prior distribution options for model parameters. Moreover, our team actively seeks for interdisciplinary collaborations to further expand the modelling framework and tailor it to the specific needs of diverse disciplines.

Acknowledgements The development of the "fdmr" package in R programming is supported by UKRI grant EP/X022641/1, and German BMBF grant 01DD20001.

References

1. Blangiardo, M., Cameletti, M.: Spatial and spatio-temporal Bayesian models with R-INLA. John Wiley & Sons (2015)
2. Brun, F., Treichler, D., Shean, D., Immerzeel, W.W.: Limited contribution of glacier mass loss to the recent increase in Tibetan Plateau lake volume. Frontiers in Earth Science **8**, 582060 (2020)
3. Brunsdon, C., Fotheringham, S., Charlton, M.: Geographically weighted regression. Journal of the Royal Statistical Society: Series D (The Statistician) **47**(3), 431–443 (1998)
4. Bürkner, P.C.: brms: An r package for bayesian multilevel models using stan. Journal of statistical software **80**, 1–28 (2017)
5. Copernicus Climate Change Service (C3S): ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate. Copernicus Climate Change Service Climate Data Store (CDS) (2017). https://doi.org/10.1002/qj.3803, `https://cds.climate.copernicus.eu/cdsapp#!/home`, Accessed: 2023-07-20
6. Cressie, N.: Statistics for spatial data. John Wiley & Sons (2015)
7. Cressie, N., Calder, C.A., Clark, J.S., Hoef, J.M.V., Wikle, C.K.: Accounting for uncertainty in ecological analysis: the strengths and limitations of hierarchical statistical modeling. Ecol Appl **19**(3), 553–570 (2009)
8. Darmofal, D.: Spatial analysis for the social sciences. Cambridge University Press (2015)
9. Finley, A.O., Banerjee, S., Carlin, B.P.: spbayes: an R package for univariate and multivariate hierarchical point-referenced spatial models. Journal of statistical software **19**(4), 1 (2007)
10. Fuglstad, G.A., Simpson, D., Lindgren, F., Rue, H.: Constructing priors that penalize the complexity of Gaussian random fields. J Am Stat Assoc **114**(525), 445–452 (2019)
11. Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: Bayesian data analysis. Chapman and Hall/CRC (1995)
12. Lee, D.: Carbayes: an r package for bayesian spatial modeling with conditional autoregressive priors. Journal of Statistical Software **55**(13), 1–24 (2013)
13. Lee, D., Rushworth, A., Napier, G.: Carbayesst: spatio-temporal generalised linear mixed models for areal unit data. R package version **2** (2015)
14. Lindgren, F., Rue, H.: Bayesian spatial modelling with r-inla. J Stat Softw **63**, 1–25 (2015)
15. Lindgren, F., Rue, H., Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. J R Stat Soc Series B Stat Methodo **73**(4), 423–498 (2011)
16. Martin, L.C., Westermann, S., Magni, M., Brun, F., Fiddes, J., Lei, Y., Kraaijenbrink, P., Mathys, T., Langer, M., Allen, S., et al.: Recent ground thermo-hydrological changes in a Tibetan endorheic catchment and implications for lake level changes. Hydrology and Earth System Sciences Discussions pp. 1–45 (2022)
17. Moraga, P.: Geospatial health data: Modeling and visualization with R-INLA and shiny. CRC Press (2019)

18. Nychka, D., Furrer, R., Paige, J., Sain, S.: fields: Tools for spatial data. Boulder, CO, USA: University Corporation for Atmospheric Research (2017)
19. Pebesma, E.J.: Multivariable geostatistics in s: the gstat package. Computers & geosciences **30**(7), 683–691 (2004)
20. R Core Team: R: A language and environment for statistical computing. Vienna, Austria (2013)
21. Ribeiro Jr, P.J., Diggle, P.J.: Analysis of geostatistical data. The geoR package, version pp. 1–6 (2006)
22. Rönnegård, L., Shen, X., Alam, M.: hglm: A package for fitting hierarchical generalized linear models. The R Journal **2**(2), 20–28 (2010)
23. Røste, J.: The Importance of Mesh Resolution When Using the SPDE Approach. Master's thesis, NTNU (2020)
24. Rothman, K.J., Greenland, S., Lash, T.L., et al.: Modern epidemiology, vol. 3. Wolters Kluwer Health/Lippincott Williams & Wilkins Philadelphia (2008)
25. Rue, H., Held, L.: Gaussian Markov random fields: theory and applications. Chapman and Hall/CRC (2005)
26. Rue, H., Martino, S.: Approximate Bayesian inference for hierarchical Gaussian Markov random field models. Journal of statistical planning and inference **137**(10), 3177–3192 (2007)
27. Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. J R Stat Soc Series B Stat Methodo **71**(2), 319–392 (2009)
28. Sha, Z., Rougier, J., Schumacher, M., Bamber, J.: Bayesian model–data synthesis with an application to global glacio-isostatic adjustment. Environmetrics **30**(1), e2530 (2019)
29. Treichler, D., Kääb, A., Salzmann, N., Xu, C.Y.: Recent glacier and lake changes in High Mountain Asia and their relation to precipitation changes. The Cryosphere **13**(11), 2977–3005 (2019)
30. Ward, M.D., Gleditsch, K.S.: Spatial regression models, vol. 155. Sage Publications (2018)
31. Williams, C.K., Rasmussen, C.E.: Gaussian processes for machine learning, vol. 2. MIT press Cambridge, MA (2006)
32. Wintle, B.A., McCarthy, M.A., Volinsky, C.T., Kavanagh, R.P.: The use of Bayesian model averaging to better represent uncertainty in ecological models. Conserv Biol **17**(6), 1579–1590 (2003)
33. Zammit-Mangion, A., Cressie, N.: Frk: An r package for spatial and spatio-temporal prediction with large datasets. arXiv preprint arXiv:1705.08105 (2017)